

TOPOLOGY IS IRRELEVANT
(IN THE INFINITE DOMAIN DICHOTOMY CONJECTURE
FOR CONSTRAINT SATISFACTION PROBLEMS)

LIBOR BARTO AND MICHAEL PINSKER

ABSTRACT. We prove that an ω -categorical core structure primitively positively interprets all finite structures with parameters if and only if some stabilizer of its polymorphism clone has a homomorphism to the clone of projections, and that this happens if and only if its polymorphism clone does not contain operations α, β, s satisfying the identity $\alpha s(x, y, x, z, y, z) \approx \beta s(y, x, z, x, z, y)$.

This establishes an algebraic criterion equivalent to the conjectured borderline between P and NP-complete CSPs over reducts of finitely bounded homogenous structures, and accomplishes one of the steps of a proposed strategy for reducing the infinite domain CSP dichotomy conjecture to the finite case.

Our theorem is also of independent mathematical interest, characterizing a topological property of any ω -categorical core structure (the existence of a continuous homomorphism of a stabilizer of its polymorphism clone to the projections) in purely algebraic terms (the failure of an identity as above).

1. INTRODUCTION AND MAIN RESULTS

The Constraint Satisfaction Problem (CSP) over a relational structure \mathbb{A} in a finite language, denoted by $\text{CSP}(\mathbb{A})$, is the problem of deciding whether or not a given primitive positive (pp-) sentence in the language of \mathbb{A} holds in \mathbb{A} . An alternative, combinatorial definition of the CSP is also popular in the literature: $\text{CSP}(\mathbb{A})$ is the problem of deciding whether a given finite relational structure in the same language as \mathbb{A} maps homomorphically into \mathbb{A} .

For CSPs over certain structures, including all finite ones, a computational complexity classification has been conjectured, separating NP-hard problems from polynomial-time solvable ones. In the following, we shall state and discuss this conjecture, and subsequently present an improvement thereof which follows from our results.

In order to keep the presentation compact, we postpone most definitions to Section 2, and refer also to the survey [6] as well as to the shorter [1] for the basics of the finite domain CSP and to the monograph [12] and the shorter [38] for the infinite. As a reference for standard notions from model theory and universal algebra, we point to the textbooks [34, 9].

All structures in the present article are implicitly assumed to be finite or countable.

Libor Barto has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No 771005), and from the Czech Science Foundation (grant No 13-01832S). Michael Pinsker has received funding from the Austrian Science Fund (FWF) through project No P27600, and from the Czech Science Foundation (grant No 13-01832S).

Part of this work appeared in the Proceedings of the 31th Annual IEEE Symposium on Logic in Computer Science (LICS'16) – see [8].

1.1. The tractability conjecture. The CSP over a structure with finite domain is clearly contained in the class NP. Some well-known NP-complete problems, such as variants of 3-SAT or 3-COLORING, can be formulated as CSPs over suitable finite structures, as well as some problems solvable in polynomial time, such as 2-SAT, HORN-SAT, or systems of linear equations over finite fields. In fact, it was conjectured in [29] that CSPs over finite structures enjoy a dichotomy in the sense that every such CSP is either NP-complete, or tractable, i.e., solvable in polynomial time. A large amount of attention was brought to confirming or refuting this conjecture; a precise borderline between NP-complete and tractable CSPs was delineated [28] and thenceforth referred to as the *tractability conjecture* or also the *algebraic dichotomy conjecture*, since most of the equivalent formulations are algebraic. In a recent turn of events, the tractability conjecture was confirmed independently by Andrei Bulatov [27] and Dmitriy Zhuk [41].

When we allow the domain of \mathbb{A} to be infinite, the situation changes drastically: every computational decision problem is polynomial-time equivalent to $\text{CSP}(\mathbb{A})$ for some \mathbb{A} [14]! A reasonable assumption on \mathbb{A} which sends the CSP back to the class NP, and which still allows to cover many interesting computational problems which cannot be modeled as the CSP of a finite structure, is that \mathbb{A} is a reduct of a finitely bounded homogeneous structure. Substantial results for such CSPs include the full complexity classification of the CSPs over the reducts of $(\mathbb{Q}; <)$ in [16] (classifying the complexity of problems previously called temporal constraint satisfaction problems), the reducts of the random graph [21] (generalizing Schaefer's theorem for Boolean CSPs to what can be called the propositional logic for graphs), and the reducts of the binary branching C -relation [15] (classifying the complexity of problems known as phylogeny CSPs). The methods here include the algebraic methods from the finite, but in addition tools from model theory and Ramsey theory [20]. Moreover, topological considerations have played a significant role in the development of the theory [22], and indeed seem inevitable in a sense, although paradoxically it was believed or at least hoped that they would ultimately turn out inutile in a general complexity classification. On the other hand, due to the fact that the investigation of infinite domain CSPs is more recent, and the additional technical complications which are to be expected when passing from the finite to the infinite, the purely algebraic theory as known in the finite is still quite undeveloped in the infinite; the present work can be seen as the first purely algebraic result for such CSPs.

A generalization of the finite domain tractability conjecture has been formulated by Manuel Bodirsky and the second author of the present article. To state it, we first recall several basic facts. When \mathbb{B} has a primitive positive (pp-)interpretation without parameters in \mathbb{A} , then $\text{CSP}(\mathbb{B})$ reduces to $\text{CSP}(\mathbb{A})$. When \mathbb{A} is an ω -categorical core, then this statement is even true for pp-interpretations with parameters. By [11], every ω -categorical structure, in particular every reduct of a finitely bounded homogeneous structure, is homomorphically equivalent to an ω -categorical core, which is unique up to isomorphism. Moreover, the CSPs over any two structures which are homomorphically equivalent are equal, and so passing from an ω -categorical structure to its core does not result in any loss of information concerning the CSP.

These facts imply that the CSP over an ω -categorical structure is NP-hard whenever its core pp-interprets with parameters some structure whose CSP is NP-hard, such as

$$\mathbb{K}_3 = (\{1, 2, 3\}; \neq)$$

whose CSP is the 3-coloring problem, or

$$\mathbb{L} = (\{0, 1\}; R_{000}, R_{001}, R_{011}, R_{111}), R_{abc} = \{0, 1\}^3 \setminus \{(a, b, c)\}$$

whose CSP is the 3-SAT problem, or

$$\mathbb{M} = (\{0, 1\}; \{(0, 0, 1), (0, 1, 0), (1, 0, 0)\})$$

whose CSP is the positive 1-in-3-SAT problem. In fact, these three structures not only pp-interpret each other, they pp-interpret all finite structures. The infinite domain tractability conjecture postulates, as does the corresponding conjecture for finite structures, that pp-interpreting all finite structures with parameters in the core is the only source of hardness for CSPs over reducts of finitely bounded homogeneous structures.

Conjecture 1.1 (Bodirsky + Pinsker 2011; cf. [23]). *Let \mathbb{B} be a reduct of a finitely bounded homogeneous structure and let \mathbb{A} be the core of \mathbb{B} . Then*

- \mathbb{A} pp-interprets all finite structures with parameters (and thus $\text{CSP}(\mathbb{B})$ is NP-complete), or
- $\text{CSP}(\mathbb{B})$ is solvable in polynomial time.

In the present article, we show that the failure of the first condition of this conjecture is witnessed by a certain algebraic fact that could, similarly to the finite setting, potentially be exploited for proving tractability of the CSP. We now make the notion of an algebraic witness more precise.

1.2. The algebraic approach. The algebraic approach to finite domain CSPs is based on the fact that pp-interpretability strength of a finite structure \mathbb{A} is determined by its set of compatible operations, the so-called *polymorphism clone* of \mathbb{A} , denoted by $\text{Pol}(\mathbb{A})$. Namely, by classical universal algebraic results, a finite structure \mathbb{A} pp-interprets a finite structure \mathbb{B} if and only if there exists a clone homomorphism from $\text{Pol}(\mathbb{A})$ to $\text{Pol}(\mathbb{B})$, that is, a mapping which preserves arities and *identities* (universally quantified equations). This fact implies that the complexity of $\text{CSP}(\mathbb{A})$ only depends on the identities satisfied by operations in $\text{Pol}(\mathbb{A})$ and leads to an algebraic reformulation of the first item of Conjecture 1.1 for finite structures. The following notation is used: the *stabilizer* of $\text{Pol}(\mathbb{A})$ by finitely many constants c_1, \dots, c_n is denoted $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$; its elements are those polymorphisms of \mathbb{A} which preserve all unary relations $\{c_i\}$. The clones $\text{Pol}(\mathbb{L})$, $\text{Pol}(\mathbb{M})$, as well as $\text{Pol}(\mathbb{K}, 0, 1, 2)$ are *trivial*, i.e., they contain only projections. Let us denote the clone of projections on a 2-element set by \mathcal{P} . The clone of projections on any other set of at least 2 elements is isomorphic to \mathcal{P} .

Theorem 1.2 ([31, 25, 10], cf. [12]). *The following are equivalent for a finite relational structure \mathbb{A} with domain $A = \{c_1, \dots, c_n\}$.*

- \mathbb{A} pp-interprets all finite structures with parameters.
- There exists a clone homomorphism from $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ to \mathcal{P} .

For the second, algebraic statement of Theorem 1.2 numerous equivalent algebraic criteria have been obtained within the setting of finite structures [40, 33, 28, 36, 39, 3], making in particular the failure of the condition more easily verifiable: this failure is then usually witnessed by the satisfaction of particular identities in $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ which cannot be satisfied in \mathcal{P} .

Some of the above-mentioned facts about finite domain CSPs have analogues for ω -categorical structures. The complexity of $\text{CSP}(\mathbb{A})$ still only depends on the polymorphism clone $\text{Pol}(\mathbb{A})$

[19], and there is an analogue of Theorem 1.2, which however takes into consideration the natural topological structure of $\text{Pol}(\mathbb{A})$.

Theorem 1.3 ([22]). *The following are equivalent for an ω -categorical structure \mathbb{A} with domain A .*

- \mathbb{A} *pp*-interprets all finite structures with parameters.
- There exists a continuous clone homomorphism from $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ to \mathcal{P} , for some elements $c_1, \dots, c_n \in A$.

More generally, the complexity of $\text{CSP}(\mathbb{A})$ for ω -categorical structures provably only depends on the structure of $\text{Pol}(\mathbb{A})$ as a topological clone [22]. A natural, yet unresolved problem when comparing the finite with the ω -categorical setting then is whether the topological structure of the polymorphism clone is really essential in the infinite, or whether the abstract algebraic structure, i.e., the identities that hold in $\text{Pol}(\mathbb{A})$, is sufficient to determine the complexity of the CSP. This problem motivates, in particular, the related concept of *reconstruction* of the topology of a clone from its algebraic structure introduced in [24], which has its own purely mathematical interest.

1.3. The result. We show that the borderline proposed in Conjecture 1.1 is purely algebraic. In particular, if the conjecture is true, then the complexity of CSPs over structures concerned by the conjecture only depends on the identities which hold in the polymorphism clone of their core, rather than the additional topological structure thereof. Moreover, the borderline is characterized by a single simple identity generalizing that of [39]. We show the following.

Theorem 1.4. *The following are equivalent for an ω -categorical core structure \mathbb{A} with domain A .*

- (i) *There exists no continuous clone homomorphism $\text{Pol}(\mathbb{A}, c_1, \dots, c_n) \rightarrow \mathcal{P}$, for any $c_1, \dots, c_n \in A$.*
- (ii) *There exists no clone homomorphism $\text{Pol}(\mathbb{A}, c_1, \dots, c_n) \rightarrow \mathcal{P}$, for any $c_1, \dots, c_n \in A$.*
- (iii) *$\text{Pol}(\mathbb{A})$ contains a pseudo-Siggers operation, i.e., a 6-ary operation s such that*

$$\alpha s(x, y, x, z, y, z) \approx \beta s(y, x, z, x, z, y)$$

for some unary operations $\alpha, \beta \in \text{Pol}(\mathbb{A})$.

Consequently, the missing piece for proving Conjecture 1.1 can now be stated in purely algebraic terms.

Conjecture 1.5. *Let \mathbb{A} be the core of a reduct of a finitely bounded homogeneous structure. If $\text{Pol}(\mathbb{A})$ contains a pseudo-Siggers operation, then $\text{CSP}(\mathbb{A})$ is solvable in polynomial time.*

In a proposed strategy [38] for solving Conjecture 1.1, the first step asked to prove that for an ω -categorical structure \mathbb{A} , the existence of a clone homomorphism $\text{Pol}(\mathbb{A}) \rightarrow \mathcal{P}$ implies the existence of a continuous such homomorphism (cf. [23]). If this was true, then the failure of the first item of Conjecture 1.1 would have an algebraic witness, i.e., a non-trivial system of identities holding in some polymorphism clone $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$. The idea then is, roughly speaking, to use Ramsey theory to “lift” the algorithm for finite structures whose polymorphism clone satisfy this identity to show that $\text{CSP}(\mathbb{A})$ is tractable.

While we do not answer this question, Theorem 1.4 gives an answer for the variant which is actually relevant for the CSP: for an ω -categorical core structure \mathbb{A} , the existence of a clone homomorphism $\text{Pol}(\mathbb{A}) \rightarrow \mathcal{P}$ implies the existence of a continuous clone homomorphism from

some stabilizer of $\text{Pol}(\mathbb{A})$ to \mathcal{P} . Taking into account the existence of non-continuous clone homomorphisms $\text{Pol}(\mathbb{A}) \rightarrow \mathcal{P}$ [23], even for ω -categorical \mathbb{A} , as well as the recent discovery of ω -categorical structures \mathbb{A}, \mathbb{A}' whose polymorphism clones are isomorphic algebraically, but not topologically [13], it might very well turn out that the answer to the original question is negative, but, as we could then conclude a posteriori, irrelevant for CSPs.

Let us also remark that Theorem 1.4 is, by the fact that every ω -categorical structure has a unique ω -categorical core, a statement about all ω -categorical structures, rather than only the structures concerned by Conjecture 1.1. Theorem 1.4 is therefore remarkable in that non-trivial statements about the class of all ω -categorical structures, other than the fundamental theorem of Ryll-Nardzewski, Engeler, and Svenonius characterizing them, are practically non-existent.

1.4. Outline and proof strategy. The strategy for proving Theorem 1.4 is similar to the finite analogue of Theorem 1.4 proved in [39] (see also [35]). Siggers’s reasoning is based on a “loop lemma” from Bulatov’s work in [26] that refines the CSP dichotomy theorem for finite undirected graphs [32].

After providing definitions and notation in Section 2, we start our proof in Section 3 with a generalization of the loop lemma, the *pseudoloop lemma*, using some of the ideas from [26]. Instead of finite graphs we work with infinite objects which we call graph-group-systems, and which consist of a permutation group acting on the vertex set of an infinite graph while preserving its edge relation. In our case, the action of the group will have finitely many orbits in its componentwise action on finite tuples of fixed length, so that in some sense the graph becomes finite modulo the group action. It might be some people’s cup of tea to imagine such systems as fuzzy finite graphs, whereas others will be inclined to draw lines between “potatoes” in order to achieve an appropriate visualization. In the second part of Section 3, we discuss possible generalizations of our pseudoloop lemma inspired by facts about finite graphs, and provide some evidence for the possibility of such generalizations.

Theorem 1.4 is derived from the pseudoloop lemma in Section 4 basically using a standard universal algebra technique, albeit adapted to the ω -categorical setting via a compactness argument. Again, we discuss possible extensions of our main theorem.

Section 5 contains further discussion on clone morphisms in the light of other recent results, in particular from the wonderland of reflections [7]. In particular, one of our examples there provides the answer to a question in [23].

We conclude our work with suggestions for research in Section 6.

2. DEFINITIONS AND NOTATION

Relational structures are denoted by blackboard bold letters, such as \mathbb{A} , and their domain by the same letter in the plain font, such as A . By a *graph* we mean a relational structure with a single symmetric binary relation.

2.1. The range of the infinite CSP conjecture. A relational structure \mathbb{B} is *homogeneous* if every isomorphism between finite induced substructures extends to an automorphism of the entire structure \mathbb{B} . In that case, \mathbb{B} is uniquely determined, up to isomorphism, by its *age*, i.e., the class of its finite induced substructures up to isomorphism. \mathbb{B} is *finitely bounded* if its signature is finite and its age is given by a finite set \mathcal{F} of forbidden finite substructures, i.e., the age consists precisely of those finite structures in its signature which do not (isomorphically) embed any member of \mathcal{F} . A *reduct* of a structure \mathbb{B} is a structure \mathbb{A} on the same domain which

is first-order definable without parameters in \mathbb{B} . Reducts \mathbb{A} of finitely bounded homogeneous structures are ω -categorical, i.e., the up to isomorphism unique countable model of their first-order theory. Equivalently, their automorphism groups are *oligomorphic*: they have finitely many orbits in their action on n -tuples over \mathbb{A} , for every finite $n \geq 1$.

2.2. pp-formulas and interpretations. A formula is *primitive positive*, in short *pp*, if it contains only equalities, existential quantifiers, conjunctions, and atomic formulas – in our case, relational symbols. A pp-formula *with parameters* over a structure \mathbb{A} can contain, in addition, elements of the domain of \mathbb{A} .

A *pp-interpretation* is a first-order interpretation in the sense of model theory where all the involved formulas are primitive positive: a structure \mathbb{A} *pp-interprets* \mathbb{B} (with parameters) if there exists a partial mapping f from a finite power A^n to B such that the domain of f , the f -preimage of the equality relation and the f -preimage of every relation in \mathbb{B} is pp-definable in \mathbb{A} (with parameters). In particular, \mathbb{A} pp-interprets its substructures induced by pp-definable subsets and also its quotients modulo a pp-definable equivalence relation.

2.3. Cores. An ω -categorical structure \mathbb{A} is a *core*, also called *model-complete core*, if all of its endomorphisms are elementary self-embeddings, i.e., preserve all first-order formulas over \mathbb{A} . This is the case if and only if its automorphism group is dense in its endomorphism monoid with respect to the pointwise convergence topology on functions on A ; cf. Section 2.5 for a description of the latter. Two structures \mathbb{A}, \mathbb{B} are *homomorphically equivalent* if there exist homomorphisms from \mathbb{A} into \mathbb{B} and vice-versa.

2.4. Clones. A *function clone* \mathcal{C} is a set of finitary operations (also called functions) on a fixed set C which contains all projections and which is closed under composition. A *polymorphism* of a relational structure \mathbb{A} is a finitary operation $f(x_1, \dots, x_n)$ on A which *preserves* all relations R of \mathbb{A} : this means that for all $\mathbf{r}_1, \dots, \mathbf{r}_n \in R$ we have that $f(\mathbf{r}_1, \dots, \mathbf{r}_n)$, calculated componentwise, is again contained in R . The *polymorphism clone* of \mathbb{A} , denoted by $\text{Pol}(\mathbb{A})$, consists of all polymorphisms of \mathbb{A} , and is always a function clone. Its unary operations are precisely the endomorphisms of \mathbb{A} , and its *invertible* unary operations (i.e., those which are bijections and whose inverse is also a polymorphism) are precisely the automorphisms of \mathbb{A} .

A *clone homomorphism* is a mapping from one function clone to another which preserves arities, composition, and which sends every projection of its domain to the corresponding projection of its co-domain. Equivalently, a clone homomorphism is a mapping ξ that preserves arities and *identities*, i.e., universally quantified equations over \mathcal{C} : more precisely, whenever an identity $t \approx s$ holds in \mathcal{C} , where t and s are terms over the signature which has one functional symbol for every element in \mathcal{C} , then the identity obtained from $t \approx s$ by replacing each occurrence of $f \in \mathcal{C}$ by $\xi(f)$ holds in \mathcal{D} . See [7] for a more detailed exposition of connections between various variants of clone homomorphisms and identities.

2.5. Topology. Function clones carry a natural topology, the topology of *pointwise convergence*, for which a subbasis is given by sets of functions which agree on a fixed finite tuple; the functions of a fixed arity in a function clone form a clopen set. Equivalently, the domain of a function clone is taken to be discrete, and the n -ary functions in the clone equipped with the product topology, for every $n \geq 1$; the whole clone is then the sum space of the spaces of n -ary functions. The function clones which are topologically closed within the space of all functions on their domain are precisely the polymorphism clones of relational structures.

We always understand continuity of clone homomorphisms with respect to this topology.

2.6. Core clones and oligomorphicity. We say that a closed function clone is a *core* if it is the polymorphism clone of a core.

A function clone \mathcal{C} is *oligomorphic* if and only if the permutation group $\text{Gr}(\mathcal{C})$ of unary invertible elements of \mathcal{C} is oligomorphic. When \mathcal{C} is closed, then this is the case if and only if it is the polymorphism clone of an ω -categorical structure, in which case $\text{Gr}(\mathcal{C})$ consists of the automorphisms of that structure.

Note that when \mathcal{C} is a core, then the set of its unary operations is the closure of $\text{Gr}(\mathcal{C})$.

2.7. Pseudo-Siggers operations. A 6-ary operation s in a function clone \mathcal{C} is a *pseudo-Siggers operation* if there exist unary $\alpha, \beta \in \mathcal{C}$ such that the identity $\alpha s(x, y, x, z, y, z) \approx \beta s(y, x, z, x, z, y)$ holds in \mathcal{C} (i.e., equality holds for all values for the variables in \mathcal{C}). We then also say that s *satisfies* the pseudo-Siggers identity.

3. THE PSEUDOLOOP LEMMA

The following definition is a generalization of finite graphs to the ω -categorical which is suitable for our purposes.

Definition 3.1. A graph-group-system, for short *gg-system*, is a pair $(\mathbb{G}, \mathcal{G})$, where \mathcal{G} is a permutation group on a set G , and $\mathbb{G} = (G; R)$ an (undirected) graph which is invariant under \mathcal{G} . We also write (R, \mathcal{G}) for the same *gg-system*.

The system is called *oligomorphic* if \mathcal{G} is; in that case, \mathbb{G} is ω -categorical, since its automorphism group contains \mathcal{G} and hence is oligomorphic.

The system *pp-interprets* (pp-defines) a structure \mathbb{B} if \mathbb{G} together with the orbits of \mathcal{G} on finite tuples does.

A *pseudoloop* of a *gg-system* $(\mathbb{G}, \mathcal{G})$ is an edge of \mathbb{G} of the form $(a, \alpha(a))$, where $\alpha \in \mathcal{G}$.

Note that R , as well as any relation that is first-order definable from a *gg-system* $(\mathbb{G}, \mathcal{G})$, is invariant under the natural action of \mathcal{G} on tuples. In particular, such relations are unions of orbits of the action of \mathcal{G} on tuples, and when \mathcal{G} is oligomorphic, then there are only finite many first-order definable relations of any fixed arity.

We are now ready to state our pseudoloop lemma for *gg-systems*.

Lemma 3.2 (The pseudoloop lemma). *Let $(\mathbb{G}, \mathcal{G})$ be an oligomorphic *gg-system*, where \mathbb{G} has a subgraph isomorphic to \mathbb{K}_3 . Then either it *pp-interprets* \mathbb{K}_3 with parameters, or it contains a pseudoloop.*

3.1. Proof of the pseudoloop lemma. We need the following auxiliary notation and definitions.

Definition 3.3. *Let $(\mathbb{G}, \mathcal{G})$ be a *gg-system*. For $a_1, \dots, a_n \in G$, we denote by $O(a_1, \dots, a_n)$ the orbit of the tuple (a_1, \dots, a_n) under \mathcal{G} .*

Definition 3.4. *A pseudoloop-free *gg-system* $(\mathbb{G}, \mathcal{G})$ which that \mathbb{G} contains (an isomorphic copy of) \mathbb{K}_3 is minimal if*

- *it does not *pp-define* any proper subset $S \subsetneq G$ and a binary symmetric relation R' on S such that R' is pseudoloop-free and contains a \mathbb{K}_3 , and*
- *it does not *pp-define* any non-trivial equivalence relation $\sim \subseteq G^2$ and a binary symmetric relation R' on the set G/\sim of its equivalence classes such that R' is pseudoloop-free and contains a \mathbb{K}_3 .*

We can now prove the pseudoloop lemma.

Proof of Lemma 3.2. Assuming that a gg-system $(\mathbb{G}, \mathcal{G})$, where $\mathbb{G} = (G; R)$ contains a \mathbb{K}_3 , has no pseudoloop, we show that it pp-interprets \mathbb{K}_3 with parameters.

Step 0: If $(\mathbb{G}, \mathcal{G})$ is not minimal, then we can replace it by a minimal gg-system as follows. When $(\mathbb{G}, \mathcal{G})$ pp-defines a relation R' on a pp-definable subset $S \subsetneq G$ which contains no pseudoloop but a \mathbb{K}_3 , then we replace $(\mathbb{G}, \mathcal{G})$ by the system thus obtained on S , consisting of R' and the restriction of the action of \mathcal{G} to S . If this is not the case, but $(\mathbb{G}, \mathcal{G})$ pp-defines a relation on a proper pp-definable factor of G which contains no pseudoloop but a \mathbb{K}_3 , then we proceed similarly, replacing the original system by the system obtained on the factor set. Iterating this procedure, it can only happen a finite number of times that we move to a system on a subset, since this step strictly decreases the finite number of orbits of the system, and since in factoring steps the number of orbits of the system does not increase either. Hence, from some point in the iteration, only factoring steps occur. But two factoring steps could be performed in a single step by combining the pp-definitions, and thus, since there exists only a finite number of pp-definable equivalence relations in the system, all factoring steps can be combined to a single step, after which the system cannot be further factored. The gg-system thus obtained must therefore be minimal.

Notice that minimality implies that every vertex is contained in an edge: otherwise restriction of R to the (pp-definable) subset of those vertices which are contained in an edge would yield a smaller system.

Step 1: R pp-defines a symmetric binary relation R' with the property that every edge of R' is contained in a \mathbb{K}_3 , i.e., every element of R' is contained in an induced subgraph of $(G; R')$ isomorphic to \mathbb{K}_3 , and which still shares our assumptions on R :

$$R'(x, y) :\leftrightarrow \exists z R(x, y) \wedge R(x, z) \wedge R(y, z).$$

Hence, replacing R by R' , we henceforth assume that every edge of R is contained in a \mathbb{K}_3 .

In the following, for $n \geq 1$ we say that $x, y \in G$ are *n-diamond-connected*, denoted by $x \sim_n y$, if there exist $a_1, b_1, c_1, d_1, \dots, a_n, b_n, c_n, d_n \in G$ such that, for every $1 \leq i \leq n$, both a_i, b_i, c_i and b_i, c_i, d_i induce \mathbb{K}_3 in \mathbb{G} , $x = a_1, d_1 = a_2, d_2 = a_3, \dots, d_{n-1} = a_n$, and $d_n = y$. They are *diamond-connected*, denoted by $x \sim y$, if they are *n-diamond-connected* for some $n \geq 1$.

Observe that \sim_n is a pp-definable relation from R (since our definition is in fact a pp-definition). Also recall that there are only finitely many binary relations first-order definable from R , and note that if x, y are *n-diamond-connected*, then they are *m-diamond-connected* for all $m \geq n$. Therefore, there exists an $n \geq 1$ such that x, y are diamond-connected if and only if they are *n-diamond connected*. In particular, the relation $x \sim y$ is pp-definable in \mathbb{G} . Note also that it is an equivalence relation on G : it is clearly transitive and symmetric, and it is reflexive since every vertex is contained in a \mathbb{K}_3 , by Steps 0 and 1.

Step 2: We claim that if $x, y \in G$ are *n-diamond-connected* for some $n \geq 1$, then $\neg R(x, y')$ for all $y' \in O(y)$. Otherwise, pick a counterexample x, y, y' with minimal $n \geq 1$.

Suppose first that n is odd and set $k := \frac{n-1}{2}$. Let a be the a_{k+1} from the chain of diamonds witnessing $x \sim_n y$. Consider the following pp-definition over $(\mathbb{G}, \mathcal{G})$:

$$S(w) :\leftrightarrow \exists u, v (u \in O(a) \wedge u \sim_k v \wedge R(v, w)) ;$$

in case that $k = 0$ we replace \sim_k by the equality relation. Then clearly $S(b_n)$ and $S(c_n)$. But we also have $S(y)$, since $S(y')$ holds by virtue of $a \sim_k x$ and $R(x, y')$ and since y is in the same

orbit as y' . Hence, since $d_n = y$, we have $S(d_n)$ and so S contains a \mathbb{K}_3 . By the minimality of $(\mathbb{G}, \mathcal{G})$ (see Step 0), we must have that S holds for all elements of G . Let $u, v \in G$ as in the definition of S witness that $S(x)$ holds. Then $u \sim_k v$, but also $u \sim_k x'$ for some $x' \in O(x)$, as $a \sim_k x$, $u \in O(a)$, and \sim_k is invariant under \mathcal{G} . Therefore, $v \sim_{n-1} x'$, which together with $R(v, x)$ contradicts the minimality of n when $n \geq 3$; when $n = 1$, this means that we have discovered a pseudoloop of $(\mathbb{G}, \mathcal{G})$, again a contradiction.

Suppose now that n is even and denote $k := \frac{n}{2} - 1$; the argument is similar. Let b, c be the b_{k+1}, c_{k+1} from the chain of diamonds witnessing $x \sim_n y$. Consider the following pp-definition:

$$S(w) : \leftrightarrow \exists u_b, u_c, u, v ((u_b, u_c) \in O(b, c) \wedge R(u, u_c) \wedge R(u, u_b) \wedge u \sim_k v \wedge R(v, w)) .$$

Then as in the odd case, $S(b_n), S(c_n), S(d_n)$, and so the set defined by S contains a \mathbb{K}_3 . By the minimality of $(\mathbb{G}, \mathcal{G})$, it contains x ; let $u_b, u_c, u, v \in G$ as in the definition of S witness this. Then $u \sim_k v$, but also $u \sim_{k+1} x'$ for some $x' \in O(y)$. Hence, $v \sim_{n-1} x'$ and $R(v, x)$ contradict the minimality of n .

Step 3: Defining

$$R'(x, y) : \leftrightarrow \exists x', y' (x \sim x' \wedge y \sim y' \wedge R(x', y'))$$

we obtain a relation $R' \supseteq R$ which does not contain a pseudoloop. Indeed, if $R'(x, y)$ is witnessed by x', y' and x and y are in the same orbit, then $x \sim y''$ for some $y'' \in O(y')$ since $y \in O(x)$ and since \sim is invariant under \mathcal{G} . Thus $x' \sim y''$ and $R(x', y')$, a contradiction with Step 2. Moreover, every edge in R' is contained in a \mathbb{K}_3 : if z' is so that $\{x', y', z'\}$ induce a \mathbb{K}_3 in R , then $\{x, y, z'\}$ induce a \mathbb{K}_3 in R' , for $z' \sim z'$ and $x \sim x'$ imply $R'(x, z')$, and $R'(y, z')$ can be inferred similarly.

Step 4: Note that R' is in fact a relation between equivalence classes of \sim and the naturally defined quotient gg-system $(\mathbb{G}^q, \mathcal{G}^q)$ on $G^q = G / \sim$ contains no pseudoloops (by Step 3). Moreover, every edge of \mathbb{G}^q is still contained in a \mathbb{K}_3 . By minimality, \sim must therefore be the equality relation. This means that R contains no *diamonds*, that is, there do not exist distinct $a, b, c, d \in G$ such that $\{a, b, c\}$ and $\{b, c, d\}$ both induce a \mathbb{K}_3 .

Summarizing we now know that our gg-system $(\mathbb{G}, \mathcal{G})$ is minimal, pseudoloop-free, and diamond-free; moreover, every edge and every vertex of \mathbb{G} is contained in a \mathbb{K}_3 .

Step 5: For $k \geq 1$, we denote the k -th power of \mathbb{K}_3 by \mathbb{T}_k . By Lemma 3.5 shown below, if \mathbb{G} contains an induced subgraph isomorphic to \mathbb{T}_k , where $|T_k|$ exceeds the number of orbits of \mathcal{G} , then $(\mathbb{G}, \mathcal{G})$ pp-defines a symmetric relation $R' \supseteq R$ without pseudoloops such that every edge of R' is still contained in a \mathbb{K}_3 . Replacing R by R' , we obtain a minimal gg-system which enjoys the properties summarized at the end of Step 4.

Iterating this process, by oligomorphicity we arrive after a finite number of steps at a gg-system $(\mathbb{G}, \mathcal{G})$ with the additional property that \mathbb{G} contains no induced subgraph isomorphic to \mathbb{T}_k for any $k \geq 1$ such that $|T_k|$ exceeds the number of orbits of \mathcal{G} .

Step 6: Recall that \mathbb{G} contains $\mathbb{T}_1 = \mathbb{K}_3$. By Step 5, there exists a maximal $k \geq 1$ such that \mathbb{G} contains an induced subgraph isomorphic to \mathbb{T}_k . Let k be that number and let a_1, \dots, a_l , where $l := |T_k| = 3^k$, denote the vertices of such an induced subgraph. We show that \mathbb{G} pp-defines the set $A = \{a_1, \dots, a_l\}$ with parameters a_1, \dots, a_l .

By [19], this is the case if each l -ary operation f in $\text{Pol}(\mathbb{G}, a_1, \dots, a_l)$ preserves A . So, suppose that such a function f does not preserve A . Now, f is a homomorphism $\mathbb{G}^l \rightarrow \mathbb{G}$ and its restriction to A is a homomorphism f' from \mathbb{T}_k to the diamond-free graph \mathbb{G} whose image, which contains A because f stabilizes each a_i , is strictly larger than $|\mathbb{T}_k|$. Then [26, Claim 3, Subsection 3.2] shows that the image of f' induces a graph isomorphic to \mathbb{T}_m for some $m > k$, contradicting the maximality of k .

Step 7: Step 6 implies that \mathbb{G} pp-interprets \mathbb{T}_k with parameters. But \mathbb{K}_3 can be pp-interpreted in \mathbb{T}_k with parameters by the final sentence of [26]. \square

It remains to prove the lemma referred to in Step 5 of the proof of Lemma 3.2 above. It might be worth noting that the essential conclusion of Step 5, namely that \mathbb{G} can be assumed not to contain an isomorphic copy of \mathbb{T}_k for k large enough, is trivial in the proof of the finite loop lemma in [26] for cardinality reasons: it is certainly true for all k such that $|\mathbb{T}_k| > |G|$. Lifting this statement to the ω -categorical setting seems to require a non-trivial combinatorial argument rather than, for example, a simple compactness argument. This is accomplished in the following lemma.

Lemma 3.5. *Let $(\mathbb{G}, \mathcal{G})$, where $\mathbb{G} = (G; R)$, be a minimal (hence pseudoloop-free) oligomorphic gg-system. Assume moreover that*

- every edge of \mathbb{G} is contained in a \mathbb{K}_3 , and
- \mathbb{G} contains an induced subgraph isomorphic to \mathbb{T}_k for some $k \geq 1$ such that $|\mathbb{T}_k| = 3^k$ exceeds the number of orbits of \mathcal{G} .

Then $(\mathbb{G}, \mathcal{G})$ pp-defines a symmetric relation $R' \supseteq R$ without pseudoloops such that every edge of R' is contained in a \mathbb{K}_3 .

Proof. Fix a copy of \mathbb{T}_k in \mathbb{G} , the elements of which we denote by tuples in $\{1, 2, 3\}^k$. So, two vertices in $\{1, 2, 3\}^k$ are adjacent if and only if they differ in every coordinate. From the cardinality assumption, we can pick two elements \mathbf{a}, \mathbf{a}' of the copy that belong to the same orbit A . Let \mathbf{b}, \mathbf{c} in the copy be so that $\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$ induce a \mathbb{K}_3 , and let B, C be their orbits. Since $(\mathbb{G}, \mathcal{G})$ has no pseudoloops, the three orbits A, B, C are distinct. Without loss of generality, assume $\mathbf{a} = 1^k$ (i.e., the k tuple all of whose entries equal 1), $\mathbf{b} = 2^k$, and $\mathbf{c} = 3^k$.

Define a relation

$$\begin{aligned} S(u, v) :& \leftrightarrow \exists a'', b'', c'', n_A, n_B, n_C \\ & (R(u, n_A) \wedge R(v, n_A) \wedge R(n_A, a'') \wedge a'' \in A \wedge \\ & R(u, n_B) \wedge R(v, n_B) \wedge R(n_B, b'') \wedge b'' \in B \wedge \\ & R(u, n_C) \wedge R(v, n_C) \wedge R(n_C, c'') \wedge c'' \in C) . \end{aligned}$$

In words, u, v have common neighbors adjacent to elements in A, B , and C .

The relation S is obviously symmetric. It is also reflexive. To see this, it suffices to observe that every element in G is a neighbor of a neighbor of an element in A , and similarly in B and C . But the latter follows from the minimality of $(\mathbb{G}, \mathcal{G})$: otherwise, we could restrict R to neighbors of neighbors of A , a set which contains $A \cup B \cup C$, and thus obtain a smaller gg-system containing a \mathbb{K}_3 , namely the one induced by $\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$.

Observe that whenever $S(u, v)$ holds, then every element of G is adjacent to a common neighbor of $O(u)$ and $O(v)$: this follows as above from the minimality of $(\mathbb{G}, \mathcal{G})$ since the elements of $A \cup B \cup C$ are adjacent to a common neighbor of $O(u)$ and $O(v)$.

Set

$$Q(u, v) := \exists s (R(u, s) \wedge S(s, v)) \wedge \exists t (S(u, t) \wedge R(t, v)).$$

Then $Q \supseteq R$: since S is reflexive, setting $s = v$ and $t = u$ in the above definition shows that $R(u, v)$ implies $Q(u, v)$. Moreover, Q is symmetric by definition. Let R' consist of those edges of Q which are contained in a \mathbb{K}_3 with respect to Q . We still have that $R' \supseteq R$.

We now show that (Q, \mathcal{G}) , and thus (R', \mathcal{G}) , has no pseudoloop. To this end, it suffices to show that whenever $R(u, v)$ holds, then we cannot have $S(u, v')$ for any $v' \in O(v)$. Suppose to the contrary that there exist such elements. The R -edge (u, v) is contained in a \mathbb{K}_3 , induced by $\{u, v, w\}$, for some $w \in G$. As observed above, each vertex, in particular the vertex w , is adjacent to a common neighbor of $O(u)$ and $O(v') = O(v)$. Therefore, there exists a common neighbor z of $O(u)$, $O(v)$ and $O(w)$. The set of neighbors of $O(z)$ contains $O(u)$, $O(v)$, and $O(w)$; it is a proper subset of G since $(\mathbb{G}, \mathcal{G})$ has no pseudoloops; it is pp-definable in $(\mathbb{G}, \mathcal{G})$; and finally, it contains a \mathbb{K}_3 , contradicting the minimality of $(\mathbb{G}, \mathcal{G})$.

Using for the first time the copy of \mathbb{T}_k in \mathbb{G} , we now show that R is properly contained in R' by showing that \mathbf{a}' , the second element of the copy of \mathbb{T}_k in the orbit A of \mathbf{a} , is related to \mathbf{b} and \mathbf{c} via R' . Note that this is sufficient since in \mathbb{T}_k , no two distinct elements are related to both \mathbf{b} and \mathbf{c} . We show only $R'(\mathbf{a}', \mathbf{b})$, the second claim is analogous. Reordering the tuples when necessary, we may assume that $a'_i \neq 2$ for all $1 \leq i \leq j$, and $a'_i = 2$ for all $j < i \leq k$. Since $\mathbf{a}' \neq \mathbf{b}$, we have $j \geq 1$. Observe that whenever $\mathbf{u}, \mathbf{v} \in \{1, 2, 3\}^k$ are of the form $(x, \dots, x, 2, \dots, 2)$ and $(x, \dots, x, 3, \dots, 3)$, respectively, where the number of occurrences of x equals j , then $S(\mathbf{u}, \mathbf{v})$: this is witnessed by their common neighbor $(y_1, \dots, y_j, 1, \dots, 1)$, where $y_i \notin \{a'_i, x\}$ for all $1 \leq i \leq j$, which is R -related to $\mathbf{a}' \in A$; their common neighbor $(z, \dots, z, 1, \dots, 1)$, starting with j occurrences of $z \notin \{2, x\}$, which is R -related to $\mathbf{b} \in B$; and their common neighbor $(w, \dots, w, 1, \dots, 1)$, starting with j occurrences of $w \notin \{3, x\}$, which is R -related to $\mathbf{c} \in C$. But now we see that $Q(\mathbf{a}', \mathbf{b})$ holds: setting $\mathbf{t} = (a'_1, \dots, a'_j, 3, \dots, 3)$, we have $S(\mathbf{a}', \mathbf{t})$ and $R(\mathbf{t}, \mathbf{b})$; on the other hand, setting $\mathbf{s} := (2, \dots, 2, 3, \dots, 3)$, with j occurrences of 2, we have $R(\mathbf{a}', \mathbf{s})$ and $S(\mathbf{s}, \mathbf{b})$. We can then conclude that $R'(\mathbf{a}', \mathbf{b})$ holds, since any two elements of $\{1, 2, 3\}^k$, in particular \mathbf{a}' and \mathbf{b} , have a common neighbor with respect to R , and hence also with respect to Q , showing that the Q -edge $(\mathbf{a}', \mathbf{b})$ is contained in a \mathbb{K}_3 with respect to Q . □

3.2. Discussion of the pseudoloop lemma. The proof of the pseudoloop lemma presented here draws on ideas from the proof in [26] of the loop lemma for finite undirected graphs [32, 26], which states that every finite non-bipartite graph either contains a loop (i.e., an edge of the form (a, a)) or pp-interprets \mathbb{K}_3 with parameters. A generalization of the latter result, which is useful in the theory of finite algebras and their applications to finite domain CSPs, is its expansion from non-bipartite undirected graphs to certain directed graphs (digraphs). We shall now state this generalization.

Recall that a digraph is *smooth* if each vertex has an incoming and an outgoing edge, and that a digraph has *algebraic length 1* if it contains a closed walk whose number of forward edges exceeds its number of backward edges by 1. Note that graphs containing an induced copy of \mathbb{K}_3 have algebraic length 1.

Theorem 3.6 ([5, 3]). *Let $\mathbb{G} = (G; R)$ be a smooth digraph of algebraic length 1 without a loop. Then \mathbb{G} pp-interprets \mathbb{K}_3 with parameters.*

We conjecture that the pseudoloop lemma can be generalized to such digraphs as well.

Conjecture 3.7. *Let $(\mathbb{G}, \mathcal{G})$ be an oligomorphic gg -system, where $\mathbb{G} = (G; R)$ is a smooth digraph of algebraic length 1 without a pseudoloop. Then $(\mathbb{G}, \mathcal{G})$ pp-interprets \mathbb{K}_3 with parameters.*

It is easy to reduce the loop lemma for finite non-bipartite undirected graphs to finite undirected graphs containing a \mathbb{K}_3 (see [32]). For the infinite pseudo-version, even this special case of Conjecture 3.7 is open. On the other hand, the generalization of the loop lemma for finite non-bipartite undirected graphs to the digraphs of Theorem 3.6 requires completely different techniques. It might be possible to circumvent the additional complications in the infinite by reducing to the finite, as exemplified in the proof of the following proposition.

Proposition 3.8. *Let $(\mathbb{G}, \mathcal{G})$ be an oligomorphic gg -system, where $\mathbb{G} = (G; R)$ is a smooth connected digraph of algebraic length 1. Moreover, assume that*

- \mathcal{G} acts with three orbits on G , and
- \mathbb{G} modulo the orbit equivalence on G is isomorphic to \mathbb{K}_3 (i.e., when $U, V \subseteq G$ are orbits of \mathcal{G} , then there is an edge from some $u \in U$ to some $v \in V$ if and only if $U \neq V$).

Then $(\mathbb{G}, \mathcal{G})$ pp-interprets \mathbb{K}_3 with parameters.

Proof. We prove the claim in three steps. In Step 1, we construct an equivalence relation \sim on G which is pp-definable from $(\mathbb{G}, \mathcal{G})$ and has the following properties.

- (i) If $a \sim b$, then $O(a) = O(b)$.
- (ii) If $a \sim b$, $R(a, a')$, $R(b, b')$, and $O(a') = O(b')$, then $a' \sim b'$.
- (iii) If $a \sim b$, $R(a', a)$, $R(b', b)$, and $O(a') = O(b')$, then $a' \sim b'$.

In Step 2, we observe that \mathbb{G} , factored by \sim , is a finite digraph and then, in Step 3, we use Theorem 3.6 to finish the proof.

Before we start, we make a simple observation which will be used throughout the proof. If $A, B \subseteq G$ are distinct orbits of \mathcal{G} and $a \in A$, then there exists $b \in B$ such that $R(a, b)$. Indeed, since \mathbb{G} modulo the orbits is isomorphic to \mathbb{K}_3 , there are $a' \in A$ and $b' \in B$ such that $R(a', b')$, and so we can set $b = \alpha(a)$ for any $\alpha \in \mathcal{G}$ with $\alpha(a') = a$. A similar statement holds, of course, with the role of a and b reversed.

Step 1: For each pp-definable equivalence relation \sim on G satisfying (i) and elements $a, a', b, b' \in G$ falsifying (ii) (or, similarly, (iii)) we are going to pp-define an equivalence relation \sim' containing \sim which still satisfies (i), and such that $a' \sim' b'$. This is enough for achieving the goal of Step 1 since the process of enlarging \sim (starting with the smallest equivalence relation on G) must stop after a finite number of steps due to oligomorphicity (recall the remark after Definition 3.1).

We will only consider the situation where (ii) does not yet hold for \sim , the situation for (iii) being completely analogous. So, let \sim satisfy (i) and let $a, a', b, b' \in G$ be as in (ii). Denote $U := O(a) = O(b)$ and $V := O(a') = O(b')$, and the remaining orbit of \mathcal{G} by W . Note that $U \cup V$ can be pp-defined since $x \in U \cup V$ if and only if $\exists y R(x, y) \wedge y \in W$ holds. Similarly, $U \cup W$ and $V \cup W$ are pp-definable in $(\mathbb{G}, \mathcal{G})$.

We first construct a reflexive binary relation $S := S_1 \cap S_2 \cap S_3 \cap S_4$ on G which contains both (a', b') and (b', a') and satisfies (i). Using S we define a symmetric and reflexive relation Q which contains $\sim \cup \{(a', b')\}$ and which still satisfies (i), and finally set \sim' to be the transitive closure of Q .

The relations S_i are pp-defined as follows:

$$\begin{aligned} S_1(x, y) &:\leftrightarrow \exists z_1, z_2 R(z_1, x) \wedge z_1 \sim z_2 \wedge R(z_2, y) \wedge z_1 \in U \cup V \\ S_2(x, y) &:\leftrightarrow \exists z_1, z_2 R(z_1, x) \wedge z_1 \sim z_2 \wedge R(z_2, y) \wedge z_1 \in U \cup W \\ S_3(x, y) &:\leftrightarrow \exists z_1, z_2, z_3, z_4 R(z_1, x) \wedge z_1 \sim z_2 \wedge R(z_2, z_3) \wedge R(z_4, z_3) \wedge R(z_4, y) \\ &\quad \wedge z_1 \in U \cup V \wedge z_3 \in U \cup V \wedge z_4 \in V \cup W \\ S_4(x, y) &:\leftrightarrow S_3(y, x) . \end{aligned}$$

We claim that each S_i is reflexive. Each $c \in G$ has an incoming R -edge from some $d \in U \cup V$. The choice $z_1 = z_2 = d$ then shows $S_1(c, c)$, proving that S_1 is reflexive. Similarly, S_2 is reflexive as well. To prove the claim for S_3 , let $c \in G$ arbitrary. If $c \in U \cup V$ then, for any $d \in U \cup V, e \in V \cup W$ with $R(d, c)$ and $R(e, c)$, the choice $(z_1, z_2, z_3, z_4) := (d, d, c, e)$ shows $S_3(c, c)$; if $c \in W$, then we pick $d \in V, e \in U$ such that $R(d, c)$, $R(d, e)$ and the choice $(z_1, z_2, z_3, z_4) = (d, d, e, d)$ works.

It is also easy to verify that (a', b') is contained in each S_i : $(a', b') \in S_1 \cap S_2$ follows from the choice $(z_1, z_2) = (a, b)$. To see that $(a', b') \in S_3$, set $(z_1, z_2, z_3, z_4) = (a, b, b', c)$, where $c \in W$ is such that $R(c, b')$. Finally, setting $(z_1, z_2, z_3, z_4) = (b, a, a', c)$ shows that $(b', a') \in S_3$, and $(a', b') \in S_4$. Similarly, (b', a') is also contained in each S_i .

Now we claim that S_1 does not intersect $U \times V$. Indeed, if $x \in U, y \in G$, and $S_1(x, y)$ is witnessed by $z_1, z_2 \in G$ via the defining formula, then $z_1 \in V$, since $x \in U, z_1 \in U \cup V, R(x, z_1)$, and \mathbb{G} has no pseudoloops. Thus, since \sim satisfies (i) and $z_1 \sim z_2$, we must also have $z_2 \in V$. Hence, we cannot have $y \in V$, since $R(z_2, y)$ and since \mathbb{G} has no pseudoloops. The same reasoning shows that S_1 does not intersect $V \times U$, that S_2 intersects neither $U \times W$ nor $W \times U$, that S_3 does not intersect $V \times W$, and that S_4 does not intersect $W \times V$. It follows that $S = S_1 \cap S_2 \cap S_3 \cap S_4$ satisfies (i).

In summary, S is reflexive, satisfies (i), and contains both (a', b') and (b', a') . We now define the relation Q by

$$Q(x, y) :\leftrightarrow \exists s (x \sim s \wedge S(s, y)) \wedge \exists t (S(t, x) \wedge t \sim y).$$

By definition, Q is symmetric, and it satisfies (i) as both \sim and S do. From the reflexivity of S , it follows that Q contains \sim ; in particular, it is itself reflexive. Moreover, the reflexivity of \sim together with $(a', b'), (b', a') \in S$ imply that $(a', b') \in Q$.

Finally, by repeatedly replacing Q by $Q'(x, y) :\leftrightarrow \exists z Q(x, z) \wedge Q(z, y)$ we obtain an increasing chain of reflexive symmetric relations. Using again the oligomorphicity of \mathcal{G} , this chain stabilizes after finitely many steps and we get the transitive closure of the original Q – the sought after equivalence relation \sim' .

Step 2: Let $(\mathbb{G}^q, \mathcal{G}^q)$, where $\mathbb{G}^q = (G^q; R^q)$, be the quotient gg-system modulo \sim : that is, G^q consists of the equivalence classes of \sim , \mathcal{G}^q is the natural action of \mathcal{G} on these classes, and $R^q(U, V)$ holds for classes U, V if $R(u, v)$ holds for some $u \in U$ and $v \in V$.

Note that $(\mathbb{G}, \mathcal{G})$ pp-interprets \mathbb{G}^q and that \mathbb{G}^q is a connected smooth digraph of algebraic length 1. Moreover, \mathbb{G}^q has no pseudoloops since \sim satisfies (i). From (ii) (or (iii)) it follows that if $R^q(U, V)$ and $R^q(U, W)$ (or $R^q(V, U)$ and $R^q(W, U)$) where V is in the same \mathcal{G}^q -orbit as W , then $V = W$.

We show that G^q is finite. Fix any vertex $U \in G^q$ and consider any $V, W \in G^q$ such that (U, V) is in the same \mathcal{G}^q -orbit of pairs as (U, W) , say $\alpha(U, V) = (U, W)$, where $\alpha \in \mathcal{G}^q$. Recall that \mathbb{G}^q is connected and consider any oriented path $U = U_0, U_1, \dots, U_k = V$ and its

α -image $U = \alpha(U_0), \alpha(U_1), \dots, \alpha(U_k) = W$. Clearly, each U_i is in the same orbit as $\alpha(U_i)$ and then, using the property of R^q from the previous paragraph, we obtain $\alpha(U_1) = U_1, \dots, \alpha(U_k) = U_k$, so $V = W$. We have proved that (U, V) and (U, W) are in different orbits of pairs whenever $V \neq W$. Since \mathcal{G}^q is oligomorphic, we conclude that G^q must be finite.

Step 3: Now \mathbb{G}^q is a finite smooth digraph of algebraic length 1 without a loop (in fact, it does not even have a pseudoloop). By Theorem 3.6, the quotient \mathbb{G}^q pp-interprets \mathbb{K}_3 with parameters. Since $(\mathbb{G}, \mathcal{G})$ pp-interprets \mathbb{G}^q , the proof is complete. \square

Some assumptions in Proposition 3.8 can be weakened. For example, the connectivity assumption can be omitted and more quotient digraphs (other than \mathbb{K}_3) can be allowed. We have decided to present the simplest version just to illustrate the general idea.

A theorem recently announced by Marcin Kozik suggests that even the following strengthening of Conjecture 3.7 might be true.

Conjecture 3.9. *Let $(\mathbb{G}, \mathcal{G})$ be an oligomorphic gg-system such that the quotient of \mathbb{G} modulo the \mathcal{G} -orbit equivalence is a smooth digraph of algebraic length 1 without a loop. Then $(\mathbb{G}, \mathcal{G})$ pp-interprets \mathbb{K}_3 with parameters.*

Kozik proved this conjecture in the finite case. The infinite version may be of use for, e.g., answering the original question about (continuous) homomorphisms to projections discussed after Conjecture 1.5 — the implication (2) \Rightarrow (1) in Corollary 5.1.

4. THE MAIN RESULT

We first prove Theorem 1.4 using the pseudoloop lemma, and then discuss the result and related results about finite algebras.

4.1. Proof of the pseudo-Siggers theorem. In order to derive Theorem 1.4, we will produce pseudo-Siggers operations locally using the pseudoloop lemma, and then derive a global pseudo-Siggers operation via a compactness argument.

Definition 4.1. *We say that a function clone \mathcal{C} has local pseudo-Siggers operations if for every finite $A \subseteq C$ there exists a 6-ary $s \in \mathcal{C}$ and unary $\alpha, \beta \in \mathcal{C}$ satisfying*

$$\alpha s(x, y, x, z, y, z) = \beta s(y, x, z, x, z, y)$$

for all $x, y, z \in A$.

Lemma 4.2. *Let \mathcal{C} be a closed oligomorphic function clone. If it has local pseudo-Siggers operations, then it has a pseudo-Siggers operation.*

Proof. Let $A_0 \subseteq A_1 \subseteq \dots$ be a sequence of finite subsets of C whose union equals C , and pick for every $i \in \omega$ a 6-ary operation $s_i \in \mathcal{C}$ witnessing the definition of local pseudo-Siggers operations on A_i , i.e., there exist unary $\alpha_i, \beta_i \in \mathcal{C}$ such that $\alpha_i s_i(x, y, x, z, y, z) = \beta_i s_i(y, x, z, x, z, y)$ for all $x, y, z \in A_i$. Note that if s_i is such a witness for A_i , then so is γs_i , for all $\gamma \in \text{Gr}(\mathcal{C})$. Hence, because $\text{Gr}(\mathcal{C})$ is oligomorphic, we may thin out the sequence in such a way that s_j agrees with s_i on A_i , for all $j > i \geq 0$. We briefly describe this standard compactness argument for the convenience of the reader: there exists a smallest $j_0 \geq 0$ such that for infinitely many $k \geq j_0$ there exists $\gamma_k \in \text{Gr}(\mathcal{C})$ such that $\gamma_k s_k$ agrees with s_{j_0} on A_0 , by oligomorphicity. Replace s_0 by s_{j_0} , all s_k as above by $\gamma_k s_k$, and remove all other $s_{k'}$ where $k' \geq 0$ from the sequence. Next repeat this process picking $j_1 \geq 1$ for A_1 , and so on. This completes the argument.

Since the elements of the sequence $(s_i)_{i \in \omega}$ agree on every fixed A_i eventually, and since \mathcal{C} is closed, they converge to a function $s \in \mathcal{C}$. The function s , restricted to any A_i , witnesses local pseudo-Siggers operations on A_i , i.e., there exist unary $\alpha_i, \beta_i \in \mathcal{C}$ such that $\alpha_i s(x, y, x, z, y, z) = \beta_i s(y, x, z, x, z, y)$ for all $x, y, z \in A_i$. By a similar compactness argument as the one above for the functions s_i , this time applied to the pairs (α_i, β_i) of functions, we may assume that these pairs converge to a pair (α, β) of functions $\alpha, \beta \in \mathcal{C}$. We then have $\alpha s(x, y, x, z, y, z) = \beta s(y, x, z, x, z, y)$ for all $x, y, z \in C$. \square

We now consider gg-systems where the group $\text{Gr}(\mathcal{C})$ of a closed oligomorphic function clone \mathcal{C} acts on finite powers of its domain.

Lemma 4.3. *Let \mathcal{C} be a closed oligomorphic function clone. Suppose that every gg-system $(\mathbb{G}, \mathcal{G})$ where*

- $\mathbb{G} = (C^k; R)$ for some $k \geq 1$,
- \mathcal{G} corresponds to the componentwise action of $\text{Gr}(\mathcal{C})$ on C^k ,
- \mathbb{G} contains a \mathbb{K}_3 , and
- $R \subseteq C^{2k}$ is invariant under \mathcal{C}

has a pseudoloop. Then \mathcal{C} has a pseudo-Siggers operation.

Proof. We show that \mathcal{C} has local pseudo-Siggers operations and apply Lemma 4.2. Let $A \subseteq C$ be finite, and pick $k \geq 1$ and $\mathbf{a}^x, \mathbf{a}^y, \mathbf{a}^z \in A^k$ such that the rows of the $(k \times 3)$ -matrix $(\mathbf{a}^x, \mathbf{a}^y, \mathbf{a}^z)$ form an enumeration of A^3 . Let R be the binary relation on C^k where tuples $\mathbf{b}, \mathbf{c} \in C^k$ are related via R if there exists a 6-ary $s \in \mathcal{C}$ such that $\mathbf{b} = s(\mathbf{a}^x, \mathbf{a}^y, \mathbf{a}^x, \mathbf{a}^z, \mathbf{a}^y, \mathbf{a}^z)$ and $\mathbf{c} = s(\mathbf{a}^y, \mathbf{a}^x, \mathbf{a}^z, \mathbf{a}^x, \mathbf{a}^z, \mathbf{a}^y)$. In other words, it is the \mathcal{C} -invariant subset of $(2k)$ -tuples generated by the six vectors obtained by concatenating \mathbf{a}^u and \mathbf{a}^v , where $u, v \in \{x, y, z\}$ are distinct. The latter description reveals that R is a symmetric relation on C^k invariant under \mathcal{C} and containing \mathbb{K}_3 , therefore the gg-system (R, \mathcal{G}) , where \mathcal{G} is the componentwise action of $\text{Gr}(\mathcal{C})$ on C^k , has a pseudoloop (\mathbf{b}, \mathbf{c}) . That means that there exists a 6-ary $s \in \mathcal{C}$ and $\alpha \in \text{Gr}(\mathcal{C})$ such that $s(\mathbf{a}^x, \mathbf{a}^y, \mathbf{a}^x, \mathbf{a}^z, \mathbf{a}^y, \mathbf{a}^z) = \alpha s(\mathbf{a}^y, \mathbf{a}^x, \mathbf{a}^z, \mathbf{a}^x, \mathbf{a}^z, \mathbf{a}^y)$, proving the claim. \square

Corollary 4.4. *Let \mathbb{A} be an ω -categorical core. Then either it pp-interprets \mathbb{K}_3 with parameters, or $\text{Pol}(\mathbb{A})$ has a pseudo-Siggers operation.*

Proof. We apply Lemma 4.3 to the clone $\mathcal{C} := \text{Pol}(\mathbb{A})$; then $\text{Gr}(\mathcal{C})$ consists precisely of the automorphisms of \mathbb{A} . If the assumptions of that lemma are satisfied, then \mathcal{C} has a pseudo-Siggers operation. Otherwise, there exists a pseudoloop-free gg-system $((C^k; R), \mathcal{G})$ satisfying the four conditions. By Lemma 3.2, this gg-system pp-interprets \mathbb{K}_3 with parameters. Since R is invariant under \mathcal{C} , it is pp-definable from \mathbb{A} by [19]. Moreover, since \mathbb{A} is a core, the orbits of \mathcal{G} are pp-definable from \mathbb{A} as well by [12]. It follows that \mathbb{A} pp-interprets \mathbb{K}_3 with parameters, as required. \square

We are now ready to prove Theorem 1.4. For the convenience of the reader, we restate it here.

Theorem 4.5 (Theorem 1.4). *The following are equivalent for an ω -categorical core structure \mathbb{A} .*

- (i) *There exists no continuous clone homomorphism $\text{Pol}(\mathbb{A}, c_1, \dots, c_n) \rightarrow \mathcal{P}$, for any $c_1, \dots, c_n \in A$.*
- (ii) *There exists no clone homomorphism $\text{Pol}(\mathbb{A}, c_1, \dots, c_n) \rightarrow \mathcal{P}$, for any $c_1, \dots, c_n \in A$.*

(iii) $\text{Pol}(\mathbb{A})$ contains a pseudo-Siggers operation, i.e., a 6-ary operation s such that

$$\alpha s(x, y, x, z, y, z) \approx \beta s(y, x, z, x, z, y)$$

for some unary operations $\alpha, \beta \in \text{Pol}(\mathbb{A})$.

Proof. We first prove that (iii) implies (ii). Take $\alpha, \beta, s \in \text{Pol}(\mathbb{A})$ satisfying the pseudo-Siggers identity. We claim that every stabilizer $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ has a pseudo-Siggers operation. To see that, consider the endomorphisms γ, δ of \mathbb{A} defined by $\gamma(x) := s(x, \dots, x)$, $\delta(x) := \alpha\gamma(x)$ ($= \beta\gamma(x)$ by the pseudo-Siggers identity). Because \mathbb{A} is a core, its automorphisms are dense in endomorphisms, thus there exist automorphisms ϵ, θ of \mathbb{A} such that $\epsilon(c_i) = \gamma(c_i)$ and $\theta(c_i) = \delta(c_i)$ for every i . But then $\theta^{-1}\alpha\epsilon$, $\theta^{-1}\beta\epsilon$ and $\epsilon^{-1}s$ are contained in $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ and satisfy

$$(\theta^{-1}\alpha\epsilon)(\epsilon^{-1}s)(x, y, x, z, y, z) \approx (\theta^{-1}\beta\epsilon)(\epsilon^{-1}s)(y, x, z, x, z, y).$$

The implication from (ii) to (i) is trivial.

Finally, assume that no stabilizer of $\text{Pol}(\mathbb{A})$ has a continuous clone homomorphism to \mathcal{P} . Then no such stabilizer has a continuous clone homomorphism to $\text{Pol}(\mathbb{K}_3)$ either, since it is well-known that the latter clone has a continuous clone homomorphism to \mathcal{P} : it is a folklore fact that all operations of $\text{Pol}(\mathbb{K}_3)$ are *essentially unary*, i.e., depend on at most one variable. By Theorem 1.3, \mathbb{A} does not pp-interpret \mathbb{K}_3 with parameters. Corollary 4.4 then tells us that $\text{Pol}(\mathbb{A})$ has a pseudo-Siggers operation. \square

4.2. Discussion of the pseudo-Siggers theorem. A positive answer to Conjecture 3.7, at least under the assumption that R contains edges (r, a) , (a, r) , (r, e) , (e, a) for some $a, e, r \in G$, would allow a strengthening of Item (iii) of Theorem 1.4 to a 4-variable pseudo-Siggers operation s satisfying $\alpha s(r, a, r, e) \approx \beta s(a, r, e, a)$ for some unary functions α, β (which works in the finite case, see [39, 35]). Another open problem is whether it is possible to replace item (iii) of Theorem 1.4 by a pseudo-weak-near-unanimity operation, i.e., an operation $w(x_1, \dots, x_n)$, for some $n \geq 2$, satisfying $\alpha_1 w(x, \dots, x, y) \approx \alpha_2 w(x, \dots, x, y, x) \approx \dots \approx \alpha_n w(y, x, \dots, x)$, for unary functions $\alpha_1, \dots, \alpha_n$ (which again works in the finite, see [36]).

On the negative side, it has been observed that the CSP classification for the reducts of $(\mathbb{Q}; <)$ in [16] shows that the syntactically strongest characterization of (iii) in Theorem 1.4 in the finite case by means of *cyclic operations* (see [3]) cannot be lifted to the infinite, at least not in the straightforward way of adding unary functions. For a well-known example, the clone of those functions in $\text{Pol}(\mathbb{Q}; <)$ which are injective (up to dummy variables) contains a pseudo-Siggers operation [16] (and the CSP of the corresponding structure is in P), but no operation $c(x_1, \dots, x_n)$, where $n \geq 2$, satisfying $\alpha c(x_1, \dots, x_n) \approx \beta c(x_2, \dots, x_n, x_1)$: if there was such an operation, then picking any non-constant tuple $(a_1, \dots, a_n) \in \mathbb{Q}^n$, there would be an automorphism γ of $(\mathbb{Q}; <)$ such that $c(b_1, \dots, b_n) = \gamma c(b_2, \dots, b_n, b_1)$ for all $b_1, \dots, b_n \in \{a_1, \dots, a_n\}$, since $(\mathbb{Q}; <)$ is a core structure. But then shifting n times we would get $c(a_1, \dots, a_n) = \gamma^n c(a_1, \dots, a_n)$, and so γ would fix $c(a_1, \dots, a_n)$, contradicting the injectivity of c .

Finally, we remark that our main theorem as well as the pseudoloop lemma can be used as tools for proving hardness: if an ω -categorical core structure \mathbb{A} does not have a pseudo-Siggers polymorphism (or \mathbb{A} pp-defines with parameters a pseudoloop-free graph containing a \mathbb{K}_3), then \mathbb{A} interprets all finite structures with parameters by the combination of Theorem 1.4 and Theorem 1.3 (or Lemma 3.2), and therefore $\text{CSP}(\mathbb{A})$ is NP-hard.

5. MAPPINGS TO THE PROJECTION CLONE

A recent revision of the basic reductions between CSPs in [7] has revealed the importance of a different kind of mapping – *uniformly continuous h1 clone homomorphisms* – which is obtained from continuous clone homomorphisms by strengthening the topological requirement and weakening the algebraic one. We now briefly introduce these concepts and refer to [7] for more details.

The pointwise convergence topology of a function clone is induced by a uniformity [30, 24] which is, in the case of a countable domain, induced by a metric. An arity-preserving mapping ξ from a clone \mathcal{C} on a set C to a clone \mathcal{D} on a finite set D is uniformly continuous with respect to this uniformity if there exists a finite subset $C' \subseteq C$ such that $\xi(f) = \xi(g)$ whenever f and g agree on C' . An *h1 clone homomorphism* is a mapping from one function clone to another which preserves arities and composition with projections; equivalently, preserves identities of height 1, i.e., identities of the form $f(\text{variables}) \approx g(\text{variables})$. For example, the pseudo-Siggers identity of our main theorem is not of height 1, but becomes height 1 when the outer unary functions are removed.

These notions give us $12=3 \times 2 \times 2$ different meaningful types of morphisms from a clone to the clone of projections: we may require that the mapping is uniformly continuous, or continuous, or we do not impose any topological condition at all; we may require that the mapping is a clone homomorphism or only an h1 clone homomorphism; finally, we can ask for a total mapping from the clone itself, or only for a mapping from some stabilizer. The trivial implications between the existence of such morphisms form a $3 \times 2 \times 2$ grid. However, as it turns out, for polymorphism clones of ω -categorical core structures there are at most six non-equivalent conditions and the implications among them form a chain.

Corollary 5.1. *Consider the following statements for an ω -categorical core \mathbb{A} .*

- (1) $\text{Pol}(\mathbb{A})$ has a uniformly continuous clone homomorphism to \mathcal{P} .
- (1') $\text{Pol}(\mathbb{A})$ has a continuous clone homomorphism to \mathcal{P} .
- (2) $\text{Pol}(\mathbb{A})$ has a clone homomorphism to \mathcal{P} .
- (3) Some $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ has a clone homomorphism to \mathcal{P} .
- (3') Some $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ has a continuous clone homomorphism to \mathcal{P} .
- (3'') Some $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ has a uniformly continuous clone homomorphism to \mathcal{P} .
- (4) Some $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ has a uniformly continuous h1 clone homomorphism to \mathcal{P} .
- (4') $\text{Pol}(\mathbb{A})$ has a uniformly continuous h1 clone homomorphism to \mathcal{P} .
- (5) $\text{Pol}(\mathbb{A})$ has a continuous h1 clone homomorphism to \mathcal{P} .
- (5') Some $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ has a continuous h1 clone homomorphism to \mathcal{P} .
- (6) $\text{Pol}(\mathbb{A})$ has an h1 clone homomorphism to \mathcal{P} .
- (6') Some $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ has an h1 clone homomorphism to \mathcal{P} .

Then all statements with equal number are equivalent, and (i) implies (j) for all $1 \leq i \leq j \leq 6$.

Proof. Items (1) and (1') are equivalent by the proof of Lemma 20 in [22] (cf. [30] for an explicit proof thereof), (3) and (3') are equivalent by Theorem 1.4, and (3') and (3'') again by the proof in [22]. Items (4) and (4'), (5) and (5'), as well as (6) and (6') are equivalent by Corollary 8.1 in [7] which implies that there is a uniformly continuous h1 clone homomorphism from $\text{Pol}(\mathbb{A})$ to $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$ (although the argument presented there does not work; a correct argument is obtained by combining Lemma 3.9 with Corollary 4.7 of that article). Given these equivalences it is clear that the strength of the statements is decreasing with increasing number. \square

In [7] it was shown that (4') still implies NP-hardness of the CSP. This challenged the original dichotomy conjecture of Bodirsky and Pinsker, Conjecture 1.1, which identified NP-hardness with (3'). More on the eventual resolution of this matter in the next subsection.

Trung Van Pham has observed that our main theorem implies that the failure of (4') is witnessed by the satisfaction of the pseudo-Siggers identity even for ω -categorical structures which are not cores. In the following proof, we assume that the reader is familiar with the basic facts about the core of an ω -categorical structure, and refer to [2] for more details.

Corollary 5.2. *Let \mathbb{B} be ω -categorical. If $\text{Pol}(\mathbb{B})$ has no uniformly continuous h1 clone homomorphism to \mathcal{P} , then $\text{Pol}(\mathbb{B})$ contains a pseudo-Siggers operation.*

Proof. Let \mathbb{A} be the core of \mathbb{B} . Since $\text{Pol}(\mathbb{B})$ has no uniformly continuous h1 clone homomorphism to \mathcal{P} , neither does $\text{Pol}(\mathbb{A})$, by the results from [7]. Hence, by Corollary 5.1 and Theorem 1.4, $\text{Pol}(\mathbb{A})$ has a pseudo-Siggers operation $s(x_1, \dots, x_6)$ whose pseudo-Siggers identity is witnessed by unary functions $\alpha, \beta \in \text{Pol}(\mathbb{A})$.

To show that $\text{Pol}(\mathbb{B})$ contains a pseudo-Siggers operation as well, we show that it contains local pseudo-Siggers operations and then refer to Lemma 4.2. So let $F \subseteq B$ be finite. We may assume that $A \subseteq B$ and that \mathbb{A} is an induced substructure of \mathbb{B} . Let h be a homomorphism from \mathbb{B} into \mathbb{A} . Let $F' \subseteq A$ be the image of F^6 under the operation $s'(x_1, \dots, x_6) := s(h(x_1), \dots, h(x_6)) \in \text{Pol}(\mathbb{B})$, and let γ be an automorphism of \mathbb{A} which agrees with h on F' ; this automorphism exists since \mathbb{A} is a core and since h restricted to A is an endomorphism of \mathbb{A} . Then for all $x_1, \dots, x_6 \in F$ we have

$$\alpha\gamma^{-1}hs'(x, y, x, z, y, z) = \beta\gamma^{-1}hs'(y, x, z, x, z, y).$$

Since $\alpha\gamma^{-1}h$ and $\beta\gamma^{-1}h$ are contained in $\text{Pol}(\mathbb{B})$, we see that the pseudo-Siggers identity is indeed satisfied by s' on F . \square

5.1. Converse implications. The implication (2) \Rightarrow (1) is the original question from [23] about clone homomorphisms to projections, which we already shortly discussed in the present paper after Conjecture 1.5. It remains open.

A counterexample to (3) \Rightarrow (2) can be found among the first-order reducts of $(\mathbb{Q}; <)$, investigated in [16]: it is the structure of *shuffle closed temporal relations*, or equivalently, those relations preserved by a binary operation called *pp*, which essentially behaves like one projection on one half plane, and like the other projection on the remaining half plane. It is not hard to see that it satisfies (3) but not (2).

The following example, a simplified version of a construction communicated to us by Ross Willard, shows that the implication from (2) to (3) cannot be reversed even for finite structures. The example is a disjoint union of two copies of \mathbb{K}_3 with an additional relation ensuring that the structure is a core.

Example 5.3. *Consider $\mathbb{A} = (A; R, S)$, where*

$$\begin{aligned} A &:= \{1_0, 2_0, 3_0, 1_1, 2_1, 3_1\}, \\ R(a_i, b_j) &:\leftrightarrow (i = j) \wedge (a \neq b), \text{ and} \\ S(a_i, b_j) &:\leftrightarrow i \neq j. \end{aligned}$$

The structure \mathbb{A} is a core. Indeed, for each endomorphism α of \mathbb{A} and each pair of elements $a_i, b_j \in \mathbb{A}$ with $\alpha(a_i) = \alpha(b_j)$ we have $i = j$ since α preserves S , and then $a = b$ since α preserves R . Therefore, each endomorphism of \mathbb{A} is injective, and hence an automorphism as \mathbb{A} is finite.

Next observe that, for an arbitrary chosen $c \in A$, the substructure of $(A; R)$ induced by the set $\{x \mid S(x, c)\}$, which is obviously pp-definable with a parameter, is isomorphic to \mathbb{K}_3 . It follows that \mathbb{A} pp-interprets \mathbb{K}_3 with parameters and therefore a stabilizer of $\text{Pol}(\mathbb{A})$ has a clone homomorphism to \mathcal{P} .

Finally, in order to show that $\text{Pol}(\mathbb{A})$ does not have a homomorphism to \mathcal{P} , we define polymorphisms α and s , where α is unary and s ternary, satisfying a non-trivial identity: that is, an identity not satisfiable by members of \mathcal{P} . Set

$$\alpha(a_i) := a_{1-i} \quad \text{and} \quad s(a_i, b_j, c_k) := \begin{cases} c_k & \text{if } i = j \\ a_i & \text{if } i \neq j \end{cases} .$$

Now α, s are clearly polymorphisms and they satisfy the nontrivial identity

$$s(x, x, y) \approx s(y, \alpha(y), x)$$

since both sides are equal to y by the definitions.

A counterexample to $(4) \Rightarrow (3)$ is given in [2]. However, it is also proved in the same paper that each such counterexample must have at least double exponential growth (in n) of the number of orbits of n -tuples. This is never the case for reducts of finitely bounded homogeneous structures, so that (3) and (4) are equivalent in the range of Conjectures 1.1 and 1.5, and the conjectures remain plausible even in the light of the above-mentioned results from [7] which show that (4) still implies NP-hardness.

No counterexamples are known for $(5) \Rightarrow (4)$ or $(6) \Rightarrow (5)$. We suspect that $(6) \Rightarrow (4)$ at least for reducts of finitely bounded homogeneous structures. Note that if Conjecture 1.1 is true and (6) does not imply (4), then we would be in a rather peculiar situation: The complexity of a CSP would depend only on the height 1 identities together with the uniform structure of the polymorphism clone (this fact follows from [7]), but height 1 identities per se would be insufficient to decide the complexity. On the other hand, the topology would be irrelevant when all the identities (or at least the pseudo-identities as in our main theorem) are taken into account.

We now provide an example showing that the implication $(6) \Rightarrow (5)$ does not hold in general for polymorphism clones. Our example answers Question 7.2 in [23] to the negative, by the results in [7, 30, 22].

Example 5.4. Define for every $n \geq 1$ a binary function $g_n: \omega^2 \rightarrow \omega$ as follows:

$$g_n(x, y) := \begin{cases} n & \text{if } x < n \\ x & \text{otherwise.} \end{cases}$$

Let \mathcal{C} be the smallest function clone of the domain ω which contains the set $\{g_n \mid n \geq 1\}$; then \mathcal{C} consists of all term functions over this set.

Let $f \in \mathcal{C}$, and denote its arity by m . We claim that there exists $k_f \geq 0$ such that $f(x_1, \dots, x_m) \geq k_f$ for all $x_1, \dots, x_m \in [0, k_f)$, and such that f restricted to $[k_f, \infty)^m$ equals a projection. To see this, we use induction over terms. Clearly, the claim is true for the functions g_n (set $k_{g_n} := n$) and all projections. Now let

$$f(x_1, \dots, x_m) = g_n(s(x_1, \dots, x_m), t(x_1, \dots, x_m)),$$

where $n \geq 1$, and $s, t \in \mathcal{C}$ satisfy the claim.

- If $x_1, \dots, x_m \in [0, k_s)$, then $f(x_1, \dots, x_m) \geq s(x_1, \dots, x_m) \geq k_s$ since $g_n(x, y) \geq x$ for all $x, y \in \omega$ by definition.

- If $x_1, \dots, x_m \in [0, n)$, then either $s(x_1, \dots, x_m) \geq n$ and hence

$$f(x_1, \dots, x_m) \geq s(x_1, \dots, x_m) \geq n,$$

or $s(x_1, \dots, x_m) < n$, in which case we have $g_n(s(x_1, \dots, x_m), t(x_1, \dots, x_m)) = n$ by the definition of g_n .

- On $[\max\{k_s, n\}, \infty)^m$ we have that $s(x_1, \dots, x_m)$ behaves like a projection, by the definition of k_s . Therefore, for $x_1, \dots, x_m \in [\max\{k_s, n\}, \infty)$ we have $s(x_1, \dots, x_m) \in \{x_1, \dots, x_m\}$, and in particular $s(x_1, \dots, x_m) \geq n$. Thus,

$$g_n(s(x_1, \dots, x_m), t(x_1, \dots, x_m)) = s(x_1, \dots, x_m).$$

Consequently, f behaves like the same projection as s on $[\max\{k_s, n\}, \infty)^m$.

Summarizing, we obtain that setting $k_f := \max\{k_s, n\}$ proves the claim.

The mapping ξ which sends every function $f(x_1, \dots, x_m) \in \mathcal{C}$ to the projection which it equals on $[k_f, \infty)^m$ clearly is a clone homomorphism from \mathcal{C} to \mathcal{P} .

However, the clone \mathcal{C} does not enjoy any uniformly continuous h1 clone homomorphism ξ' to \mathcal{P} : uniform continuity would imply that there exists a finite set $F \subseteq \omega$ such that the value of every function in \mathcal{C} under ξ' only depends on its restriction to F . But then $\xi'(g_n(x, y)) = \xi'(g_n(y, x))$ for any $n \geq 1$ such that $F \subseteq [0, n)$, a contradiction since $\xi'(g_n(y, x))$ is, by the definition of an h1 clone homomorphism, also the opposite projection of $\xi'(g_n(x, y))$ (as such homomorphisms preserve the switching of variables).

The topological closure $\bar{\mathcal{C}}$ of \mathcal{C} in the space of all finitary functions on ω is a polymorphism clone. We argue that every $f \in \bar{\mathcal{C}}$ still enjoys the property of the claim above: for any sequence $(f_n)_{n \in \omega}$ in \mathcal{C} converging to f we must have that $(k_{f_n})_{n \in \omega}$ is eventually constant, and so setting k_f to that value works. Hence, $\bar{\mathcal{C}}$ has a clone homomorphism to \mathcal{P} as well (but still none which is uniformly continuous).

We remark that the mapping ξ is continuous: if a sequence $(f_n)_{n \in \omega}$ in $\bar{\mathcal{C}}$ converges to $f \in \bar{\mathcal{C}}$, then we must have $k_f = k_{f_n}$ eventually, and that all f_n behave like the same projection on $[k_f, \infty)^m$ eventually.

5.2. Taylor equations without idempotency. Long before the CSP motivated the intensive investigation of function clones that do not have a clone homomorphism to \mathcal{P} , such clones were characterized by Walter Taylor by means of certain identities [40]. His theorem does not require any finiteness or topological assumptions. On the other hand, it only concerns *idempotent* clones, that is, clones whose all operations f satisfy the identity $f(x, x, \dots, x) \approx x$ (or, in other words, all operations are polymorphisms of all singleton unary relations).

Theorem 5.5 ([40]). *Let \mathcal{C} be an idempotent function clone. Then \mathcal{C} does not have a clone homomorphism to \mathcal{P} if and only if it contains an operation $t(x_1, \dots, x_n)$ for some $n \geq 1$ such that, for any $1 \leq i \leq n$, t satisfies an identity of the form $t(\dots, x, \dots) \approx t(\dots, y, \dots)$, where x and y are at the i -th position (and the remaining positions contain some variables which are not further specified).*

The standard proof (see, e.g., [33]) of the non-trivial implication (i.e., from left to right) can be split into two parts. First, every h1 clone homomorphism from an idempotent clone to \mathcal{P} is, in fact, a clone homomorphism. Second, the absence of t as in Theorem 5.5 implies the existence of an h1 clone homomorphism by a compactness argument. While the former part can be generalized to the non-idempotent situation as proved in the following proposition, Example 5.7 below shows that the latter part does not generalize in a straightforward fashion.

For the purpose of the next proposition, we say that a mapping $\xi: \mathcal{C} \rightarrow \mathcal{D}$, where \mathcal{C}, \mathcal{D} are function clones, is an *almost clone homomorphism* if it is an h1 clone homomorphism and it preserves the composition with unary members both from inside and outside; equivalently, ξ preserves arities and all identities of the form

$$\alpha(f(\beta_1(x_{i_1}), \dots, \beta_n(x_{i_n}))) \approx \gamma(g(\delta_1(x_{j_1}), \dots, \delta_n(x_{j_n}))),$$

where some of the unary functions of the identity can be missing (so that in particular, h1 identities are of this form).

Proposition 5.6. *Let \mathcal{C} be a function clone and let $\xi: \mathcal{C} \rightarrow \mathcal{P}$ be a mapping. Then ξ is a clone homomorphism if and only if it is an almost clone homomorphism.*

Proof. Assume that ξ is an almost clone homomorphism. We will show that ξ is a clone homomorphism in a sequence of steps. The first step is well-known.

Step 1: It is enough to show that ξ preserves the projections and identities of the form

$$f(x_1, \dots, x_k, g(x_{k+1}, \dots, x_l), x_{l+1}, \dots, x_m) \approx h(x_1, x_2, \dots, x_m).$$

This follows from the fact that each term can be obtained by repeatedly using the type of composition on the left hand side and then merging the variables.

Step 2: Whenever an operation $f(x_1, \dots, x_n) \in \mathcal{C}$ depends only on its i -th argument, then it is mapped via ξ to the i -th n -ary projection; in particular, ξ preserves projections. Indeed, such f satisfies the identity

$$f(x_1, \dots, x_n) \approx f(y_1, \dots, y_{i-1}, x_i, y_{i+1}, \dots, y_n)$$

and the only projection satisfying this identity is the projection onto the i -th coordinate. Since ξ is an h1 clone homomorphism, the claim follows.

Step 3: The mapping ξ preserves identities of the form

$$f(g(x_{1,1}, \dots, x_{1,m}), g(x_{2,1}, \dots, x_{2,m}), \dots, g(x_{n,1}, \dots, x_{n,m})) \approx h(x_{1,1}, \dots, x_{1,m}, x_{2,1}, \dots, \dots, x_{n,m}).$$

Suppose $\xi(f)$ and $\xi(g)$ are the i -th and j -th projection, respectively. We need to show that the above identity holds in \mathcal{P} after application of ξ , or in other words, that

$$x_{i,j} \approx \xi(h)(x_{1,1}, \dots, x_{1,m}, x_{2,1}, \dots, \dots, x_{n,m})$$

holds in \mathcal{P} .

Denoting $\alpha(x) := g(x, \dots, x)$, the following identity holds in \mathcal{C} :

$$h(\underbrace{x_1, \dots, x_1}_{m \times}, \underbrace{x_1, \dots, x_1}_{m \times}, \dots, \underbrace{x_n, \dots, x_n}_{m \times}) \approx f(\alpha(x_1), \alpha(x_2), \dots, \alpha(x_n)).$$

Since ξ preserves this identity and $\xi(\alpha)$ is the identity mapping, we obtain

$$\xi(h)(\underbrace{x_1, \dots, x_1}_{m \times}, \underbrace{x_2, \dots, x_2}_{m \times}, \dots, \underbrace{x_n, \dots, x_n}_{m \times}) \approx \xi(f)(x_1, x_2, \dots, x_n) \approx x_i$$

and it follows that

$$x_{i,\ell} \approx \xi(h)(x_{1,1}, \dots, x_{1,m}, x_{2,1}, \dots, \dots, x_{n,m})$$

for some $1 \leq \ell \leq m$.

Similarly, from

$$h(x_1, \dots, x_m, x_1, \dots, x_m, \dots) \approx \beta(g(x_1, \dots, x_m)),$$

where $\beta(x) := f(x, \dots, x)$, we get that

$$x_{k,j} \approx \xi(h)(x_{1,1}, \dots, x_{1,m}, x_{2,1}, \dots, \dots, x_{n,m})$$

for some $1 \leq k \leq n$, proving the claim.

Step 4: Now we can finish the proof using Step 1. For simplicity, consider the case where g appears at the first coordinate. So, suppose the following identity holds in \mathcal{C} .

$$f(g(x_1, \dots, x_m), x_{m+1}, \dots, x_n) \approx h(x_1, x_2, \dots, x_n).$$

Set

$$t(x_{1,1}, \dots, x_{1,m}, x_{2,1}, \dots, \dots, x_{n,m}) := f(g(x_{1,1}, \dots, x_{1,m}), g(x_{2,1}, \dots), \dots).$$

From Step 3 we get $\xi(t)(x_{1,1}, \dots, x_{n,m}) \approx \xi(f)(\xi(g)(x_{1,1}, \dots), \dots)$, in particular,

$$\begin{aligned} & \xi(t)(x_1, \dots, x_m, \underbrace{x_{m+1}, \dots, x_{m+1}}_{m \times}, \dots, \underbrace{x_n, \dots, x_n}_{m \times}) \approx \\ & \xi(f)(\xi(g)(x_1, \dots, x_m), \xi(g)(x_{m+1}, \dots, x_{m+1}), \dots, \xi(g)(x_n, \dots, x_n)). \end{aligned}$$

Since $\xi(g)$ is a projection, we have $\xi(g)(x, \dots, x) \approx x$. Finally, from the definition of t it follows that

$$t(x_1, \dots, x_m, x_{m+1}, \dots, x_{m+1}, \dots, x_n, \dots, x_n) \approx h(x_1, \dots, x_m, \alpha(x_{m+1}), \dots, \alpha(x_n)),$$

where $\alpha(x) := g(x, \dots, x)$, so

$$\xi(t)(x_1, \dots, x_m, x_{m+1}, \dots, x_{m+1}, \dots, x_n, \dots, x_n) \approx \xi(h)(x_1, \dots, x_m, x_{m+1}, \dots, x_n).$$

Combining these facts we obtain

$$\begin{aligned} & \xi(f)(\xi(g)(x_1, \dots, x_m), x_{m+1}, \dots, x_n) \\ & \approx \xi(f)(\xi(g)(x_1, \dots, x_m), \xi(g)(x_{m+1}, \dots, x_{m+1}), \dots, \xi(g)(x_n, \dots, x_n)) \\ & \approx \xi(t)(x_1, \dots, x_m, x_{m+1}, \dots, x_{m+1}, \dots, x_n, \dots, x_n) \\ & \approx \xi(h)(x_1, \dots, x_n), \end{aligned}$$

which finishes the proof. □

The following example due to Miroslav Olšák shows that a direct analogue of Theorem 5.5, or more precisely the above-mentioned second part of its standard proof, does not hold.

Example 5.7. *Let \mathcal{C} be the clone of all term operations of the free algebra over countably many generators in the signature consisting of unary operations p, q and binary operations r, s modulo the identities*

$$pr(x, y) \approx ps(x, y), \quad qr(x, y) \approx qs(y, x).$$

Since this system of identities is not satisfiable by projections, the clone \mathcal{C} does not have a clone homomorphism to \mathcal{P} .

Let $t \in \mathcal{C}$ be an arbitrary n -ary function. We claim that t has a coordinate $1 \leq k \leq n$ such that no identity of the following form is satisfied:

$$\alpha(t(\beta_1(x_{i_1}), \dots, \beta_n(x_{i_n}))) \approx \gamma(t(\delta_1(x_{j_1}), \dots, \delta_n(x_{j_n}))),$$

where $i_k \neq j_k$, and some of the unary functions can be missing in the identity.

Striving for a contradiction, suppose there was an n -ary $t \in \mathcal{C}$ satisfying such an identity for every $1 \leq k \leq n$. Fix a term $t'(x_1, \dots, x_n)$ over the variables x_1, \dots, x_n in the signature $\{p, q, r, s\}$ which evaluates in \mathcal{C} as t .

Assume first that the outermost symbol of t' is not r . We inductively define a mapping θ from the set of all terms over the variables x_1, \dots, x_n to the set $\{1, \dots, n\}$ as follows:

$$\theta(z) = \begin{cases} i & \text{if } z = x_i \\ \theta(z_2) & \text{if } z = qr(z_1, z_2) \\ \theta(z_1) & \text{otherwise, and } z = f(z_1, \dots, z_l) \text{ for some } f \in \{p, q, r, s\} \end{cases} .$$

In other words, $\theta(z)$ is the index of its leftmost variable after rewriting it using the rule which replaces each subterm of the form $qr(z_1, z_2)$ by $qs(z_2, z_1)$. It is clear that θ preserves composition with unary operations symbols from the inside and, for terms z whose outermost symbol is not r , also from the outside.

It is apparent from the defining identities of \mathcal{C} that $\theta(z_1) = \theta(z_2)$ whenever z_1 and z_2 evaluate to the same function in \mathcal{C} . On the other hand, setting $k := \theta(t')$ we obtain

$$\begin{aligned} \theta(\alpha(t'(\beta_1(x_{i_1}), \dots, \beta_n(x_{i_n})))) &= \theta(t'(x_{i_1}, \dots, x_{i_n})) = i_k, \text{ and} \\ \theta(\gamma(t'(\delta_1(x_{j_1}), \dots, \delta_n(x_{j_n})))) &= j_k, \end{aligned}$$

a contradiction.

If the outermost symbol of t' is r , then we adjust the definition of θ by swapping the roles of r and s , and argue until a contradiction in the same fashion.

Note that any pseudo-Siggers operation satisfies the modification of Taylor identities (from Theorem 5.5) which is described in the second paragraph of Example 5.7 (even without the inner unary functions). Hence, the non-satisfaction of identities of that form does, in the case of polymorphism clones of ω -categorical core structures, imply the existence of a uniformly continuous clone homomorphism from some stabilizer to \mathcal{P} , and then also the existence of an h1 clone homomorphism to \mathcal{P} by Corollary 5.1. The non-satisfaction of Taylor identities (without inner and outer unary functions) is not sufficient to this end: the function clone of all (up to dummy variables) injective functions on a countable set has no Taylor term (whose identities are incompatible with injectivity), but satisfies non-trivial h1 identities, e.g., those implicitly described in the statement of Lemma IV.3 in [2].

5.3. Ternary identities. Perhaps surprisingly in the light of our main theorem, it has been observed by Jakub Opršal and independently by Ross Willard (and perhaps might have been known to others as well) that the existence of clone homomorphisms to \mathcal{P} only depends on the structure of the ternary operations in the clone.

Proposition 5.8. *Let \mathcal{C} be a function clone, and let ξ be a mapping from the ternary operations in \mathcal{C} to \mathcal{P} which preserves arities and identities (h1 identities). Then ξ has a unique extension to a clone homomorphism (h1 clone homomorphism) from \mathcal{C} to \mathcal{P} .*

Proof. We give a self-contained proof, although the claim essentially follows from the results and arguments in [7], in particular, Section 7 thereof.

Let $f(x_1, \dots, x_n) \in \mathcal{C}$, where $n \geq 1$, and let $a = (a_1, \dots, a_n) \in \{0, 1\}^n$. We define $f_a(x, y) \in \mathcal{C}$ to be the binary function obtained by identifying all variables x_i of f with x whenever $a_i = 0$, and all other variables with y . We further define an n -ary operation $\xi'(f)$ on $\{0, 1\}$ by setting $\xi'(f)(a) := \xi(f_a)(0, 1)$ for all $a \in \{0, 1\}^n$.

It is easy to verify that ξ' is an extension of ξ , since ξ preserves h1 identities on the ternary functions. Moreover, it is clear for the same reason that any extension of ξ to an h1 clone homomorphism into \mathcal{P} must be compatible with the above definition, and hence be equal to ξ' , proving uniqueness of the extension.

We claim that ξ' maps into \mathcal{P} . To show this, we use the fact that $\mathcal{P} = \text{Pol}(\mathbb{M})$, where

$$\mathbb{M} = (\{0, 1\}; \{(0, 0, 1), (0, 1, 0), (1, 0, 0)\}),$$

as in the introduction. If $\xi'(f)$ was not a projection for some $f(x_1, \dots, x_n) \in \mathcal{C}$, then there would exist triples r_1, \dots, r_n in the relation of \mathbb{M} such that $\xi'(f)(r_1, \dots, r_n)$, calculated componentwise, is not contained in that relation. Since the relation contains only three triples, we can define a ternary operation $g(x, y, z)$ from f by replacing x_i by x whenever r_i equals $(0, 0, 1)$, similarly for y and $(0, 1, 0)$, and for z and $(1, 0, 0)$. Then $\xi'(g)$ still violates \mathbb{M} , a contradiction since $\xi(g) = \xi'(g)$ and since $\xi(g)$ is a projection.

To see that ξ' is an h1 clone homomorphism, consider an identity

$$f(y_{i_1}, \dots, y_{i_n}) \approx s(y_1, \dots, y_m),$$

and say that $\xi'(f)$ projects onto the j -th coordinate, where $1 \leq j \leq n$. Then

$$\xi'(f)(y_{i_1}, \dots, y_{i_n}) \approx y_{i_j}$$

holds in \mathcal{P} ; we wish to show that $\xi'(s)(y_1, \dots, y_m) \approx y_{i_j}$. Let $a = (a_1, \dots, a_m) \in \{0, 1\}^m$ be the tuple in which $a_{i_j} = 1$, and all other $a_i = 0$; since $\xi'(s) \in \mathcal{P}$, it suffices to show that $\xi'(s)(a) = 1$. By the definition of ξ' and the first identity above, we have that $\xi'(s)(a) = \xi(g)(0, 1)$, where $g(x, y)$ is the binary function obtained from $f(x_1, \dots, x_n)$ by inserting y for x_k whenever $i_k = i_j$, and x otherwise. Again by the definition of ξ' and the first identity, this value is equal to $\xi'(f)(b)$, where $b = (b_1, \dots, b_n) \in \{0, 1\}^n$ is the tuple where $b_k = 1$ if and only if $i_k = i_j$. By the second identity above, that value equals 1, as desired.

When ξ preserves identities, then it is not hard to show that ξ' is a clone homomorphism: one argues that it is an almost clone homomorphism by reducing to ternary functions, and then refers to Proposition 5.6. \square

A standard application of the compactness theorem of logic yields the following corollary.

Corollary 5.9. *Let \mathcal{C} be a function clone. Then there is no clone homomorphism (h1 clone homomorphism) $\mathcal{C} \rightarrow \mathcal{P}$ if and only if \mathcal{C} satisfies a non-trivial finite system of identities (h1 identities) of ternary functions.*

In particular, our pseudo-Siggers identity, which involves a 6-ary function, always implies a non-trivial system of ternary identities – but note that there is no bound on the number of identities in the system.

6. CONCLUSION

The presented results show the possibility of a purely algebraic theory for function clones that are concerned by the infinite-domain tractability conjecture. They have already found several applications: in the classification of CSPs over reducts of unary structures [18], and indirectly (via results in [2] inspired by the present work) in the algebraic dichotomy for CSPs in MMSNP [17]. Nevertheless, the theory is presently only in its infancy compared to its finite counterpart. We now briefly discuss three of the research directions that call for further exploration.

Firstly, the technical core of our main result is the pseudoloop lemma which generalizes a loop lemma for finite undirected graphs. The extension of the latter result to directed graphs stated in Theorem 3.6 has led to the development of absorption theory [3, 4], which in turn proved to be a powerful tool in the context of finite domain CSPs and universal algebra. While it might be possible to lift some results from the finite to the infinite, e.g., in a manner akin to the one shown in Proposition 3.8, it seems more valuable to find a genuine generalization of absorption theory itself. Some ideas from this theory and, in particular, a loop lemma was used in a different infinite setting, namely that of idempotent clones [37]. Is it possible to merge these approaches?

Secondly, as discussed in Section 5, there are several meaningful types of morphisms to the clone of projections and their exact relation in various settings (e.g., polymorphism clones of structures, polymorphism clones of ω -categorical structures, polymorphism clones of reducts of homogeneous structures) requires further investigation. For finite core clones all the items (2) – (6) in Corollary 5.1 are equivalent, and many alternative characterizations are available: for example by means of identities, the structure of invariant relations, absorption properties, the Tame Congruence Theory types, Bulatov’s local algebraic properties, etc. Are there infinite analogues to these characterizations?

Thirdly, clone homomorphisms to other clones rather than only the projection clone \mathcal{P} are significant for CSPs as well as universal algebra. For example, the clones of modules play a similar role to \mathcal{P} in the context of solvability of finite-domain CSPs by consistency methods. Again, are there infinite analogues?

REFERENCES

- [1] Libor Barto. The constraint satisfaction problem and universal algebra. *The Bulletin of Symbolic Logic*, 21:319–337, 2015.
- [2] Libor Barto, Michael Kompatscher, Miroslav Olšák, Michael Pinsker, and Trung Van Pham. The equivalence of two dichotomy conjectures for infinite domain constraint satisfaction problems. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science – LICS’17*, 2017. Preprint arXiv:1612.07551.
- [3] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms and the constraint satisfaction problem. *Logical Methods in Computer Science*, 8/1(07):1–26, 2012.
- [4] Libor Barto and Marcin Kozik. Absorption in universal algebra and CSP. In *The constraint satisfaction problem: complexity and approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 45–77. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2017.
- [5] Libor Barto, Marcin Kozik, and Todd Niven. The CSP dichotomy holds for digraphs with no sources and no sinks (a positive answer to a conjecture of Bang-Jensen and Hell). *SIAM Journal on Computing*, 38(5), 2009.
- [6] Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them. In *The constraint satisfaction problem: complexity and approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 1–44. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2017.
- [7] Libor Barto, Jakub Opršal, and Michael Pinsker. The wonderland of reflections. *Israel Journal of Mathematics*, 223(1):363–398, 2018.
- [8] Libor Barto and Michael Pinsker. The algebraic dichotomy conjecture for infinite domain constraint satisfaction problems. In *Proceedings of the 31th Annual IEEE Symposium on Logic in Computer Science – LICS’16*, pages 615–622, 2016.
- [9] Clifford Bergman. *Universal algebra*, volume 301 of *Pure and Applied Mathematics (Boca Raton)*. CRC Press, Boca Raton, FL, 2012. Fundamentals and selected topics.
- [10] Garrett Birkhoff. On the structure of abstract algebras. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):433–454, 1935.
- [11] Manuel Bodirsky. Cores of countably categorical structures. *Logical Methods in Computer Science*, 3(1):1–16, 2007.

- [12] Manuel Bodirsky. Complexity classification in infinite-domain constraint satisfaction. Mémoire d'habilitation à diriger des recherches, Université Diderot – Paris 7. Available at arXiv:1201.0856, 2012.
- [13] Manuel Bodirsky, David Evans, Michael Kompatscher, and Michael Pinsker. A counterexample to the reconstruction of ω -categorical structures from their endomorphism monoids. *Israel Journal of Mathematics*, 224(1):57–82, 2018.
- [14] Manuel Bodirsky and Martin Grohe. Non-dichotomies in constraint satisfaction complexity. In Luca Aceto, Ivan Damgard, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, Lecture Notes in Computer Science, pages 184–196. Springer Verlag, July 2008.
- [15] Manuel Bodirsky, Peter Jonsson, and Trung Van Pham. The Complexity of Phylogeny Constraint Satisfaction Problems. *ACM Transactions on Computational Logic (TOCL)*, 18(3), 2017. An extended abstract appeared in the Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS) 2016.
- [16] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2):1–41, 2009. An extended abstract appeared in the Proceedings of the Symposium on Theory of Computing (STOC'08).
- [17] Manuel Bodirsky, Florent R. Madelaine, and Antoine Mottet. A universal-algebraic proof of the complexity dichotomy for monotone monadic SNP. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 105–114, 2018.
- [18] Manuel Bodirsky and Antoine Mottet. A dichotomy for first-order reducts of unary structures. *Logical Methods in Computer Science*, 14(2), 2018.
- [19] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.
- [20] Manuel Bodirsky and Michael Pinsker. Reducts of Ramsey structures. *AMS Contemporary Mathematics, vol. 558 (Model Theoretic Methods in Finite Combinatorics)*, pages 489–519, 2011.
- [21] Manuel Bodirsky and Michael Pinsker. Schaefer's theorem for graphs. *Journal of the ACM*, 62(3):52 pages (article number 19), 2015. A conference version appeared in the Proceedings of STOC 2011, pages 655–664.
- [22] Manuel Bodirsky and Michael Pinsker. Topological Birkhoff. *Transactions of the American Mathematical Society*, 367:2527–2549, 2015.
- [23] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Projective clone homomorphisms. Preprint arXiv:1409.4601, 2014.
- [24] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Reconstructing the topology of clones. *Transactions of the American Mathematical Society*, 369:3707–3740, 2017.
- [25] 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.
- [26] Andrei A. Bulatov. H -coloring dichotomy revisited. *Theoretical Computer Science*, 349(1):31–39, 2005.
- [27] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017*, pages 319–330. IEEE Computer Soc., Los Alamitos, CA, 2017.
- [28] Andrei A. Bulatov, Andrei A. Krokhin, and Peter G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.
- [29] Tomáš Feder and Moshe Y. 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.
- [30] Mai Gehrke and Michael Pinsker. Uniform Birkhoff. *Journal of Pure and Applied Algebra*, 222(5):1242–1250, 2018.
- [31] David Geiger. Closed systems of functions and predicates. *Pacific Journal of Mathematics*, 27:95–100, 1968.
- [32] Pavol Hell and Jaroslav Nešetřil. On the complexity of H -coloring. *Journal of Combinatorial Theory, Series B*, 48:92–110, 1990.
- [33] David Hobby and Ralph McKenzie. *The structure of finite algebras*, volume 76 of *Contemporary Mathematics*. American Mathematical Society, Providence, RI, 1988.
- [34] Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.
- [35] Keith Kearnes, Petar Marković, and Ralph McKenzie. Optimal strong Malcev conditions for omitting type 1 in locally finite varieties. *Algebra universalis*, 72(1):91–100, 2014.

- [36] Miklós Maróti and Ralph McKenzie. Existence theorems for weakly symmetric operations. *Algebra Universalis*, 59(3-4):463–489, 2008.
- [37] Miroslav Olsák. The weakest nontrivial idempotent equations. *Bull. Lond. Math. Soc.*, 49(6):1028–1047, 2017.
- [38] Michael Pinsker. Algebraic and model theoretic methods in constraint satisfaction. arXiv:1406.7870, 2015.
- [39] Mark H. Siggers. A strong Mal'cev condition for locally finite varieties omitting the unary type. *Algebra universalis*, 64(1-2):15–20, 2010.
- [40] Walter Taylor. Varieties obeying homotopy laws. *Canadian Journal of Mathematics*, 29(3):498–527, 1977.
- [41] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In *58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017*, pages 331–342. IEEE Computer Soc., Los Alamitos, CA, 2017.

DEPARTMENT OF ALGEBRA, MFF UK, SOKOLOVSKA 83, 186 00 PRAHA 8, CZECH REPUBLIC
E-mail address: `libor.barto@gmail.com`
URL: <http://www.karlin.mff.cuni.cz/~barto/>

INSTITUT FÜR DISKRETE MATHEMATIK UND GEOMETRIE, FG ALGEBRA, TU WIEN, AUSTRIA, AND DEPARTMENT OF ALGEBRA, CHARLES UNIVERSITY, CZECH REPUBLIC
E-mail address: `marula@gmx.at`
URL: <http://dmg.tuwien.ac.at/pinsker/>