

THE ALGEBRAIC DICHOTOMY CONJECTURE FOR INFINITE DOMAIN 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 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 monograph [Bod12] as well as to the shorter [Pin15]. As a reference for standard notions from model theory, we point to the textbook [Hod97].

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

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 has been conjectured [FV99] that CSPs over finite structures enjoy a dichotomy in the sense that every such CSP is either NP-complete,

Date: January 18, 2016.

Libor Barto was supported by the the Grant Agency of the Czech Republic, grant GAČR 13-01832S. Michael Pinsker has been funded through project P27600 of the Austrian Science Fund (FWF).

or tractable, i.e., solvable in polynomial time. A large amount of attention has been brought to confirming or refuting this conjecture, resulting in considerable progress; see [Bar15] for a recent brief introduction and survey. In particular, a precise borderline between NP-complete and tractable CSPs has been delineated [BKJ05] and is now referred to as the *tractability conjecture* or also the *algebraic dichotomy conjecture*, since most of the equivalent formulations are algebraic.

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} [BG08]! A reasonable assumption on \mathbb{A} which sends the CSP back to the class NP and 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 [BK09] (classifying the complexity of problems previously called temporal constraint satisfaction problems), the reducts of the random graph (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 [BJP16] (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 [BP11]. Moreover, topological considerations have played a significant role in the development of the theory [BP15], 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. 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 [Bod07], 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. [BPP14]). *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 what is the hope in 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 ([Gei68, BKKR69, Bir35], cf. [Bod12]). *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 [Tay77, HM88, BKJ05, MM08, Sig10, BK12b], 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 $\text{Pol}(\mathbb{A})$ [BN06], 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 ([BP15]). *The following are equivalent for an ω -categorical structure \mathbb{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 $c_1, \dots, c_n \in A$.

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 [BP15]. 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.

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 [Sig10]. We show the following.

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 $c_1, \dots, c_n \in \mathbb{A}$.*
- (ii) *There exists no clone homomorphism $\text{Pol}(\mathbb{A}, c_1, \dots, c_n) \rightarrow \mathcal{P}$, for any $c_1, \dots, c_n \in \mathbb{A}$.*
- (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 [Pin15] 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. [BPP14]). 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 identity holding in some polymorphism clone $\text{Pol}(\mathbb{A}, c_1, \dots, c_n)$. The idea is to then, roughly speaking, “lift” the algorithm for finite structures whose polymorphism clone satisfies this identity (assuming the finite tractability conjecture is true) via Ramsey theory 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 $\text{Pol}(\mathbb{A}) \rightarrow \mathcal{P}$ [BPP14], for an 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 [BEKP15], it might very well turn out that the answer to the original question is negative, but, as we would then see 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 [Sig10] (see also [KMM14]). Siggers’s reasoning is based on a “loop lemma” from Bulatov’s paper [Bul05] that refines the dichotomy theorem for finite undirected graphs [HN90].

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 [Bul05]. Instead of finite graphs we work with infinite objects which we call graph-group-systems, and which can be imagined as an infinite permutation group acting on an infinite fuzzy graph, which is a finite graph modulo the action. Theorem 1.4 is then derived from the pseudoloop lemma in Section 4 using a standard universal algebra technique adapted to the ω -categorical setting via a compactness argument.

Section 5 at the end of the article contains further discussion of our results in the light of other recent results, in particular from the wonderland of reflections [BOP15], as well as inspiration for future work.

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 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* can contain, in addition, elements of the domain.

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} 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} . 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 the structure. 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 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 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 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. Clone homomorphisms preserve all identities which hold in a function clone, as defined in the introduction.

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.

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

The function clones which are closed in the topology of pointwise convergence are precisely the polymorphism clones of relational structures.

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 $\alpha s(x, y, x, z, y, z) \approx \beta s(y, x, z, x, z, y)$ holds (where \approx means that 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, in 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.*

For the proof of Lemma 3.2, we need the following auxiliary definitions.

Definition 3.3. *Let $(\mathbb{G}, \mathcal{G})$ be a gg -system. The support of $(\mathbb{G}, \mathcal{G})$ are those elements of its domain which are contained in an edge.*

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} .

The support of a gg -system $(\mathbb{G}, \mathcal{G})$ is a union of \mathcal{G} -orbits by the remark below Definition 3.1. This justifies the following definition.

Definition 3.4. *A pseudoloop-free gg -system $(\mathbb{G}, \mathcal{G})$ containing a \mathbb{K}_3 is minimal if it does not pp -define a symmetric relation R' on G such that (R', \mathcal{G}) is a pseudoloop-free gg -system containing a \mathbb{K}_3 whose support consists of fewer orbits.*

We can now prove the pseudoloop lemma.

Proof of Lemma 3.2. Assuming that a gg -system $(\mathbb{G}, \mathcal{G})$, where $\mathbb{G} = (G; R)$, 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. We thus henceforth assume that it is minimal.

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 , $x = a_1$, $d_1 = a_2$, $d_2 = a_3$, \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 the support of R : it is clearly transitive and symmetric, and it is reflexive since on the support every vertex is contained in a \mathbb{K}_3 , by Step 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), it contains x , for otherwise we could intersect R with S^2 and obtain a relation whose support consists of a smaller number of orbits. 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 similarly we infer $R'(y, z')$.

Therefore, we may replace R by R' . If this replacement changes the equivalence \sim , we repeat Step 3. Since the relation R gets bigger after each step and \mathcal{G} is oligomorphic, this process stabilizes after finitely many steps.

Step 4: Now 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 neither pseudoloops (by Step 3), nor *diamonds*, that is, there do not exist distinct $a, b, c, d \in G^q$ such that $\{a, b, c\}$ and $\{b, c, d\}$ both induce a \mathbb{K}_3 . Moreover, every edge of \mathbb{G}^q is still contained in a \mathbb{K}_3 . We replace the gg-system $(\mathbb{G}, \mathcal{G})$ by $(\mathbb{G}^q, \mathcal{G}^q)$.

If this new gg-system was not minimal, we could repeat the whole proof starting from Step 0. In each reiteration, Step 0 decreases the number of orbits in the support, while no other step increases it, so this process will terminate after a finite number of steps, and we will end up with a minimal gg-system after Step 4 eventually. The first author insists to remark that this process is in fact unnecessary, as the gg-system is already minimal after the first round

of Steps 0 to 4: if after factoring it did pp-define a relation whose support consists of less \mathcal{G}^q -orbits, then the syntactically same pp-definition would show that the original gg-system was not minimal.

Whichever solution we prefer, summarizing we end up with a gg-system $(\mathbb{G}, \mathcal{G})$ which is minimal, pseudoloop-free, and diamond-free; still, every edge 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, there exists an $m \geq 1$ such that, for any $k \geq m$, \mathbb{G} has no induced subgraph isomorphic to \mathbb{T}_k .

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 $l = |\mathbb{T}_k| = 3^k$, 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 [BN06], 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 $|\mathbb{T}_k|$. [Bul05, Claim 3, Subsection 3.2] shows that the image of f' induces a graph isomorphic to \mathbb{T}_m for some $m > k$, a contradiction.

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 [Bul05]. \square

Lemma 3.5. *Let $(\mathbb{G}, \mathcal{G})$ be an oligomorphic gg-system containing a \mathbb{K}_3 and having no pseudoloops, and assume the system is minimal. Then \mathbb{G} has no induced subgraph isomorphic to \mathbb{T}_k for any $k \geq 1$ where $|\mathbb{T}_k| = 3^k$ exceeds the number of orbits of \mathcal{G} .*

Proof. Suppose there exists a counterexample $(\mathbb{G}, \mathcal{G})$, where $\mathbb{G} = (G; R)$. We may assume without loss of generality that every edge of \mathbb{G} is contained in a \mathbb{K}_3 by keeping only those which are, as in Step 1 of the proof of Lemma 3.2. Under those assumptions, we pp-define in $(\mathbb{G}, \mathcal{G})$ a symmetric relation $R' \supseteq R$ without pseudoloops in which every edge is still contained in a \mathbb{K}_3 , and such that the support of (R', \mathcal{G}) equals the support of (R, \mathcal{G}) . Repeating this process, by oligomorphicity we must after finitely many steps arrive at a gg-system (R', \mathcal{G}) which is not minimal. Hence, $(\mathbb{G}, \mathcal{G})$ was not minimal in the first place, a contradiction.

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 on the support of $(\mathbb{G}, \mathcal{G})$: every element of the support is a neighbor of a neighbor of an element in A , and similarly in B and C , by the minimality of $(\mathbb{G}, \mathcal{G})$; otherwise, we could restrict R to neighbors of neighbors of A , a set which contains $A \cup B \cup C$; we would therefore obtain a smaller support 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 the support of $(\mathbb{G}, \mathcal{G})$ 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

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

Then $Q \supseteq R$: since S is reflexive on the support of $(\mathbb{G}, \mathcal{G})$, 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 . □

4. PROOF OF THE MAIN RESULT

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 above, there exist unary functions $\alpha, \beta \in \mathcal{C}$ such that $\alpha s(x, y, x, z, y, z) = \beta s(y, x, z, x, z, y)$ 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 \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 this 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 [BN06]. Moreover, since \mathbb{A} is a core, the orbits of \mathcal{G} are pp-definable from \mathbb{A} as well by [Bod12]. It follows that \mathbb{A} pp-interprets \mathbb{K}_3 with parameters, as required. \square

We are now ready to prove Theorem 1.4.

Proof of Theorem 1.4. 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 homomorphism to \mathcal{P} . Then no such stabilizer has a continuous clone homomorphism to $\text{Pol}(\mathbb{K}_3)$ since it is well-known that the latter clone has a continuous homomorphism to \mathcal{P} . 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

5. DISCUSSION

Our main theorem can be used as a tool for proving hardness: If an ω -categorical core structure \mathbb{A} does not have a pseudo-Siggers polymorphism, then \mathbb{A} interprets all finite structures with parameters by the combination of Theorem 1.4 and Theorem 1.3, and therefore $\text{CSP}(\mathbb{A})$ is NP-hard. The pseudoloop lemma itself can be regarded as a hardness result as well:

Corollary 5.1. *Let \mathbb{A} be an ω -categorical core structure. If \mathbb{A} pp-defines with parameters a pseudoloop-free graph containing a \mathbb{K}_3 , then $\text{CSP}(\mathbb{A})$ is NP-hard.*

Recall that a digraph is *smooth* if each vertex has an incoming and an outgoing edge, and a digraph has *algebraic length 1* if it contains a closed walk with one more forward edges than backward edges. The finite loop lemma for graphs [HN90, Bul05] has been generalized to finite smooth digraphs with algebraic length 1 in [BKN09, BK12b]. We conjecture that the pseudoloop lemma can be generalized to such digraphs as well.

Conjecture 5.2. *Let \mathcal{G} be an oligomorphic permutation group on G and let \mathbb{G} be a countable smooth digraph of algebraic length 1 on G which is invariant under \mathcal{G} . Then either \mathbb{G} contains a pseudoloop, or \mathbb{G} together with the orbits of \mathcal{G} on finite tuples pp-interprets \mathbb{K}_3 with parameters.*

As intermediate steps we suggest the following stronger assumptions on \mathbb{G} .

- (a) \mathbb{A} is a non-bipartite graph;
- (b) \mathbb{A} is a digraph containing

$$(\{a, b, c\}; \{(a, b), (b, c), (c, a), (b, a)\})$$

as a subgraph (not necessarily induced);

- (c) \mathbb{A} is a strongly connected digraph of algebraic length 1 (equivalently, the GCD of the length of cycles is 1).

A positive answer to Conjecture 5.2 under the assumption (b) or (c) would allow a strengthening of item (iii) of Theorem 1.4 to a 4-variable pseudo-Siggers operation $\alpha s(r, a, r, e) \approx \beta s(a, r, e, a)$ (see [Sig10, KMM14]). Another open problem is whether it is possible to replace item (iii) of Theorem 1.4 by pseudo-weak-near-unanimity operations (see [MM08]). On the negative side, it has been observed that the CSP classification for the reducts of $(\mathbb{Q}; <)$ shows that the syntactically strongest characterization of (iii) in Theorem 1.4 in the finite case by means of cyclic operations [BK12b] cannot be lifted to the infinite, at least not in the straightforward way of adding unary functions.

Our final remarks concern a new dichotomy conjecture that has been formulated for CSPs of reducts of finitely bounded homogeneous structures [BOP15]. 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). When the polymorphism clone of an ω -categorical structure has a uniformly continuous h1 clone homomorphism to \mathcal{P} , then its CSP is NP-hard. The conjecture states that the reducts of finitely bounded homogeneous structures whose polymorphism clone does not have such a mapping have polynomial-time solvable CSP. The following corollary summarizes all we know so far about various clone homomorphisms to the clone of projections. In particular, it shows that the new conjecture is “better” in that it is implied by Conjecture 1.1 (but not necessarily vice-versa).

Corollary 5.3. *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. It is clear that the strength of the statements is decreasing; the only non-trivial part are the equivalences. (1) and (1') are equivalent by the proof in [BP15] (cf. [GP] for an explicit proof thereof). (3) and (3') are equivalent by Theorem 1.4, and (3') and (3'') again by the proof in [BP15]. (4) and (4'), (5) and (5'), as well as (6) and (6') are equivalent by [BOP15]. \square

We remark that (4) is the weakest condition known to imply NP-hardness of the CSP.

An example communicated to us by Ross Willard shows that the implication from (2) to (3) cannot be reversed, even for finite \mathbb{A} ; another example is known among the reducts of $(\mathbb{Q}; <)$. For all remaining implications, no counterexamples are known. The implication from (2) to (1) is conjectured in [BPP14], the implication from (4) to (3) would, if true, imply that the two conjectures are equivalent. The most optimistic speculation would be that (6) implies (3). A positive answer would show that both conjectures can be true and yield a particularly esthetically pleasing — core-free, topology-free, and without identities of height greater than 1 — formulation of the tractability conjecture.

REFERENCES

- [Bar15] Libor Barto. The constraint satisfaction problem and universal algebra. *The Bulletin of Symbolic Logic*, 21:319–337, 2015.
- [BEKP15] Manuel Bodirsky, David Evans, Michael Kompatscher, and Michael Pinsker. A counterexample to the reconstruction of omega-categorical structures from their endomorphism monoids. Preprint arXiv:1510.00356, 2015.
- [BG08] 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.
- [Bir35] Garrett Birkhoff. On the structure of abstract algebras. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):433–454, 1935.
- [BJP16] Manuel Bodirsky, Peter Jonsson, and Trung Van Pham. The complexity of phylogeny constraint satisfaction. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS)*, 2016. Preprint arXiv:1503.07310.
- [BK09] 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).
- [BK12a] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem. *Logical Methods in Computer Science*, 8(1), 2012.
- [BK12b] 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.
- [BKJ05] 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.
- [BKKR69] 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.
- [BKN09] 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.
- [BN06] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.
- [Bod07] Manuel Bodirsky. Cores of countably categorical structures. *Logical Methods in Computer Science*, 3(1):1–16, 2007.
- [Bod12] 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.
- [BOP15] Libor Barto, Jakub Opršal, and Michael Pinsker. The wonderland of reflections. Preprint arXiv:1510.04521, 2015.
- [BP11] Manuel Bodirsky and Michael Pinsker. Reducts of Ramsey structures. *AMS Contemporary Mathematics, vol. 558 (Model Theoretic Methods in Finite Combinatorics)*, pages 489–519, 2011.
- [BP15] Manuel Bodirsky and Michael Pinsker. Topological Birkhoff. *Transactions of the American Mathematical Society*, 367:2527–2549, 2015.

- [BPP14] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Projective clone homomorphisms. Preprint arXiv:1409.4601, 2014.
- [Bul05] Andrei A. Bulatov. H -coloring dichotomy revisited. *Theoretical Computer Science*, 349(1):31–39, 2005.
- [FV99] Tomás 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.
- [Gei68] David Geiger. Closed systems of functions and predicates. *Pacific Journal of Mathematics*, 27:95–100, 1968.
- [GP] Mai Gehrke and Michael Pinsker. Uniform Birkhoff. In preparation.
- [HM88] David Hobby and Ralph McKenzie. *The structure of finite algebras*, volume 76 of *Contemporary Mathematics*. American Mathematical Society, Providence, RI, 1988.
- [HN90] Pavol Hell and Jaroslav Nešetřil. On the complexity of H -coloring. *Journal of Combinatorial Theory, Series B*, 48:92–110, 1990.
- [Hod97] Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.
- [KMM14] 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.
- [MM08] Miklós Maróti and Ralph McKenzie. Existence theorems for weakly symmetric operations. *Algebra Universalis*, 59(3-4):463–489, 2008.
- [Pin15] Michael Pinsker. Algebraic and model theoretic methods in constraint satisfaction. arXiv:1406.7870, 2015.
- [Sig10] Mark H. Siggers. A strong Mal’cev condition for locally finite varieties omitting the unary type. *Algebra universalis*, 64(1-2):15–20, 2010.
- [Tay77] Walter Taylor. Varieties obeying homotopy laws. *Canadian Journal of Mathematics*, 29(3):498–527, 1977.

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/>

DEPARTMENT OF ALGEBRA, MFF UK, SOKOLOVSKA 83, 186 00 PRAHA 8, CZECH REPUBLIC
E-mail address: marula@gmx.at
URL: <http://dmg.tuwien.ac.at/pinsker/>