

Durham Research Online

Deposited in DRO:

03 September 2019

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Bodirsky, Manuel and Mottet, Antoine and Olsak, Miroslav and Oprsal, Jakub and Pinsker, Michael and Willard, Ross (2019) 'Topology is relevant (in a dichotomy conjecture for infinite-domain constraint satisfaction problems).', in 2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). .

Further information on publisher's website:

<https://doi.org/10.1109/LICS.2019.8785883>

Publisher's copyright statement:

© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

TOPOLOGY IS RELEVANT (IN A DICHOTOMY CONJECTURE FOR INFINITE-DOMAIN CONSTRAINT SATISFACTION PROBLEMS)

MANUEL BODIRSKY, ANTOINE MOTTET, MIROSLAV OLŠÁK, JAKUB OPRŠAL,
MICHAEL PINSKER, AND ROSS WILLARD

ABSTRACT. The algebraic dichotomy conjecture for Constraint Satisfaction Problems (CSPs) of reducts of (infinite) finitely bounded homogeneous structures states that such CSPs are polynomial-time tractable when the model-complete core of the template has a pseudo-Siggers polymorphism, and NP-complete otherwise.

One of the important questions related to this conjecture is whether, similarly to the case of finite structures, the condition of having a pseudo-Siggers polymorphism can be replaced by the condition of having polymorphisms satisfying a fixed set of identities of height 1, i.e., identities which do not contain any nesting of functional symbols. We provide a negative answer to this question by constructing for each non-trivial set of height 1 identities a structure whose polymorphisms do not satisfy these identities, but whose CSP is tractable nevertheless.

An equivalent formulation of the dichotomy conjecture characterizes tractability of the CSP via the local satisfaction of non-trivial height 1 identities by polymorphisms of the structure. We show that local satisfaction and global satisfaction of non-trivial height 1 identities differ for ω -categorical structures with less than double exponential orbit growth, thereby resolving one of the main open problems in the algebraic theory of such structures.

1. INTRODUCTION

Many computational problems in theoretical computer science can be phrased as *constraint satisfaction problems (CSPs)*: in such a problem, we are given a finite set of variables and a finite set of constraints that are imposed on the variables, and the task is to find values for the variables that satisfy all the given constraints. The computational complexity of a CSP depends on the language that we allow when formulating the constraints in the input. By appropriately choosing this language, many computational problems in optimisation, artificial intelligence, computational biology, verification, and many other areas can be precisely expressed as a CSP.

Manuel Bodirsky and Jakub Opršal were supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No 681988, CSP-Infinity). Jakub Opršal has also received funding from the UK EPSRC (grant No EP/R034516/1). Antoine Mottet has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No 771005, CoCoSym). Miroslav Olšák and Michael Pinsker have received funding from the Czech Science Foundation (grant No 13-01832S). Michael Pinsker has received funding from the Austrian Science Fund (FWF) through project No P27600.

Formally, we fix a structure \mathbb{B} (also called the *template* or *constraint language*). The problem $\text{CSP}(\mathbb{B})$ is the computational problem of deciding whether a given conjunction of atomic formulas over the signature of \mathbb{B} is satisfiable in \mathbb{B} . For example, if the domain of \mathbb{B} is the Boolean domain $\{0, 1\}$, and \mathbb{B} contains all binary Boolean relations, then $\text{CSP}(\mathbb{B})$ is precisely the 2-SAT problem, which can be solved in polynomial time, and if the structure \mathbb{B} is the complete graph \mathbb{K}_3 on three vertices (without loops), then $\text{CSP}(\mathbb{B})$ is precisely the graph 3-coloring problem, which is NP-complete. Note that it is not necessary for this definition that the domain of \mathbb{B} is finite, and indeed, many problems can only be expressed when the domain of \mathbb{B} is infinite. For example the satisfiability of a system of polynomial equations over the rational numbers can be formulated using a structure \mathbb{B} whose domain is the rationals, but certainly not with a structure \mathbb{B} that has a finite domain.

The class of CSPs is a large class which allows for a uniform mathematical approach to the question which interests us for computational problems in general: What kind of structure makes a problem easy (i.e., polynomial-time tractable), and what makes such a problem hard (i.e., NP-hard)? Evidence for the possibility of a clear structural characterization of tractability within the realm of large classes of CSPs has been found in the fact that finite-domain CSPs exhibit a P/NP-complete dichotomy, i.e., $\text{CSP}(\mathbb{B})$ is for every finite structure \mathbb{B} , in P or NP-complete. This was conjectured by Feder and Vardi [FV98], and recently proved by Bulatov [Bul17] and, independently, by Zhuk [Zhu17]. Both proofs rely on the universal-algebraic approach and recent developments in universal algebra. In fact, they prove a strengthening of the conjecture which in addition provides a precise condition that implies NP-completeness of a finite-domain CSP. This strengthening provided by Bulatov, Jeavons, and Krokhin [BJK05] uses algebraic language, in particular the notion of *polymorphisms*, which are (structure preserving) finitary functions on a structure \mathbb{B} . Such functions can be viewed as ‘higher-order symmetries’, and in particular they form a certain generalization of the automorphisms of \mathbb{B} . The essence of the algebraic approach is that the complexity of $\text{CSP}(\mathbb{B})$ is determined up to log-space reductions by the polymorphisms of \mathbb{B} : few polymorphisms imply hardness of the CSP, while interesting polymorphisms are meant to imply better algorithmic properties.

Before we move to the infinite case, let us first describe the situation in the finite case. We denote by $\text{Pol}(\mathbb{B})$ the set of all polymorphisms of \mathbb{B} , and by \mathcal{P} the set of projections, i.e., trivial polymorphisms (these are precisely the polymorphisms of 3-SAT). Using a result of Siggers [Sig10], the finite-domain CSP dichotomy can then be formulated as follows (see Section 2 for the definitions of the concepts that appear in the statement).

Theorem 1.1 (Bulatov-Zhuk [Bul17, Zhu17]). *Let \mathbb{B} be a finite structure. Exactly one of the following holds:*

- (1) *There exists a minion homomorphism $\text{Pol}(\mathbb{B}) \rightarrow \mathcal{P}$, and $\text{CSP}(\mathbb{B})$ is NP-complete,*
- (2) *$\text{Pol}(\mathbb{B})$ contains a function s satisfying the identity*

$$(\diamond) \quad \forall x, y, z \in B. s(x, y, x, z, y, z) = s(y, x, z, x, z, y),$$
and $\text{CSP}(\mathbb{B})$ is in P.

In particular, $\text{CSP}(\mathbb{B})$ is in P or NP-complete.

One astonishing fact is that the condition for tractability in this dichotomy has an elegant formulation using a single function satisfying a single simple identity.

For templates with an infinite domain, it is known that no such dichotomy exists in general and that CSPs exhaust all possible complexity classes, up to polynomial-time Turing reductions [BG08]. However, large classes of infinite templates have been proved to exhibit a P/NP-complete dichotomy. One of the largest and most robust classes that have been conjectured to have such a dichotomy is the class of so-called *first-order reducts of finitely bounded homogeneous structures*. This class is important for several reasons:

- it is a vast generalisation of the class of finite structures where it is possible to investigate deep questions about the nature of computation. Many problems studied e.g. in temporal and spatial reasoning can be formulated as CSPs for such structures [BJ17].
- the algebraic methods of finite-domain constraint satisfaction can still be used in this class, and due to its relative tameness it is an important framework where the tractability of large classes of computational problems can be tied to algebraic and topological properties of mathematical objects.

The dichotomy conjecture for these structures has been verified in numerous special cases, for example for all CSPs in the complexity class MMSNP [BMM18]; also see [BK09, BP15, BJP17, KP17, BMPP16]. There are various equivalent formulations of the infinite-domain tractability conjecture originally formulated in [BPP] (Conjecture 6.1). The most recent one, proposed by Barto, Opršal, and Pinsker [BOP18] and later proved to be equivalent to the original one [BKO⁺17, BKO⁺], is now considered the most satisfactory formulation both esthetically and practically:

Conjecture 1.2. *Let \mathbb{B} be a reduct of a finitely bounded homogeneous structure. Exactly one of the following holds:*

- (1) *There exists a uniformly continuous minion homomorphism from $\text{Pol}(\mathbb{B})$ to \mathcal{P} , and $\text{CSP}(\mathbb{B})$ is NP-complete,*
- (2) *$\text{Pol}(\mathbb{B})$ does not have a uniformly continuous minion homomorphism to \mathcal{P} , and $\text{CSP}(\mathbb{B})$ is in P.*

Here, uniform continuity is meant with respect to the pointwise convergence topology on the space of all maps from finite powers of \mathbb{B} into \mathbb{B} , where \mathbb{B} carries the discrete topology. It is known that if there exists a uniformly continuous minion homomorphism $\text{Pol}(\mathbb{B}) \rightarrow \mathcal{P}$, then $\text{CSP}(\mathbb{B})$ is NP-complete [BOP18].

There are two major differences between the above conjecture and the finite-domain CSP dichotomy as phrased in Theorem 1.1. First of all, there is a topological content in Conjecture 1.2. This topological nature can loosely be explained in the following terms: the non-existence of a uniformly continuous minion homomorphism $\text{Pol}(\mathbb{B}) \rightarrow \mathcal{P}$ can be characterised by the fact that non-trivial identities are satisfied on every finite subset of \mathbb{B} , while the non-existence of a minion homomorphism $\text{Pol}(\mathbb{B}) \rightarrow \mathcal{P}$ is characterised by the fact that some non-trivial identities are satisfied on the whole structure \mathbb{B} . This local/global distinction evidently only arises when \mathbb{B} is an infinite structure, and can be understood as one of the major obstacles towards solving Conjecture 1.2.

Second, even when ignoring the topology, it is not known whether the second item can be expressed by the satisfaction of some fixed height 1 identities in $\text{Pol}(\mathbb{B})$

as is the case in Theorem 1.1(2). This naturally raises the following questions, which were also asked in [BP16, BOP18, BP18]:

- (1) Does the existence of a minion homomorphism $\text{Pol}(\mathbb{B})$ to \mathcal{P} imply the existence of a uniformly continuous one? In other words, can the requirement of uniform continuity in Conjecture 1.2 be dropped?
- (2) Can the non-existence of a minion homomorphism to \mathcal{P} be replaced by a statement positing that some fixed set of height 1 identities holds in $\text{Pol}(\mathbb{B})$?

We note that a positive answer to the second question would have certain implications which make a positive answer to the first question more likely. Moreover, the corresponding natural questions were asked about the historically first conjecture (see [BPP] and Conjecture 6.1) and were proved to have positive answers [BP16, BP18], thus showing that topology was irrelevant in that formulation of the conjecture.

The second question is purely algebraic, and therefore of interest to universal algebra as well. Similar questions have been asked about various properties of algebras, e.g. [Tay88, Olš17]. One notable open problem in this field is whether the algebraic condition describing structures whose CSP can be solved by a Datalog program [BK14] can be described by a single fixed set of identities as well. As in Theorem 1.1, it is known that there is such a set of identities when we restrict to finite domains.

Contributions

In the present paper, we give a negative answer to the second question, proving that no system of height 1 identities (also called *height 1 condition*) can be used as a replacement for the condition in the second item of Conjecture 1.2. Our result is formalized as follows:

Theorem 1.3. *For every non-trivial height 1 condition Σ there exists a structure \mathbb{B} such that*

- \mathbb{B} is a first-order reduct of a finitely bounded homogeneous structure;
- $\text{Pol}(\mathbb{B})$ does not satisfy Σ ;
- $\text{Pol}(\mathbb{B})$ satisfies some other non-trivial height 1 condition (consequently, there is no minion homomorphism to \mathcal{P});
- $\text{CSP}(\mathbb{B})$ is in P .

Nevertheless, we give a partially positive answer: we prove that there exists an infinite chain of weaker and weaker systems of height 1 identities¹ that can be used to describe the non-existence of a minion homomorphism to \mathcal{P} :

Theorem 1.4. *There is a countable strictly decreasing chain of height 1 conditions such that every clone that does not have a minion homomorphism to \mathcal{P} satisfies a condition from the chain.*

We note that in both proofs of Theorem 1.1, such a decreasing chain of height 1 conditions was enough to prove the dichotomy (i.e., the authors do not rely on the satisfaction of the Siggers identity).

¹Such chains are commonly studied in universal algebra since many interesting properties of algebras can be expressed by such a condition, e.g. congruence distributivity [Jón67].

This still leaves the possibility that Question (1) above has a positive answer. However, we use the chain from the previous statement to show that if the structure \mathbb{B} is allowed to belong to a slightly bigger class of structures, then Question (1) has a negative answer, too.

Theorem 1.5. *There exists a structure \mathbb{S} with the following properties.*

- (1) \mathbb{S} is an ω -categorical structure with less than double exponential orbit growth,
- (2) $\text{Pol}(\mathbb{S})$ has a minion homomorphism to \mathcal{P} ,
- (3) $\text{Pol}(\mathbb{S})$ has no uniformly continuous minion homomorphism to \mathcal{P} .

The structure \mathbb{S} is not a first-order reduct of a finitely bounded homogeneous structure (and therefore does not belong to the scope of Conjecture 1.2). However, its slow orbit growth shows that the techniques that were used to show that the two infinite-domain dichotomy conjectures are equivalent [BKO⁺17, BKO⁺] cannot be employed to remove the topological considerations from Conjecture 1.2.

2. NOTATION AND DEFINITIONS

We recall some basic notions from the algebraic approach to CSPs as well as from model theory. We refer to [BKW17] and [Hod97] for more detailed introductions to these topics.

2.1. Structures, polymorphisms

A *signature* is a list $\sigma = (R_i)_{i \in I}$ of symbols, where each symbol is associated with a natural number called its *arity*. A σ -structure \mathbb{A} is a tuple $(A, (R_i^{\mathbb{A}})_{i \in I})$ consisting of a set (the *domain*) together with a list $(R_i^{\mathbb{A}})_{i \in I}$ of relations on A , where for all $i \in I$ the relation $R_i^{\mathbb{A}}$ has the arity specified by σ .

A *graph* is a structure with a single binary symmetric relation; in particular, in this paper all graphs are undirected. We denote by \mathbb{K}_3 the complete graph on three vertices.

Let \mathbb{A}, \mathbb{B} be two structures with the same signature (e.g. two graphs). A map $h: A \rightarrow B$ is a *homomorphism* from \mathbb{A} to \mathbb{B} if it preserves all relations, i.e., for all $i \in I$,

$$(\clubsuit) \quad \text{if } (a_1, \dots, a_k) \in R_i^{\mathbb{A}}, \text{ then } (h(a_1), \dots, h(a_k)) \in R_i^{\mathbb{B}}.$$

Two structures \mathbb{A} and \mathbb{B} are *homomorphically equivalent* if there exist homomorphisms from \mathbb{A} to \mathbb{B} and from \mathbb{B} to \mathbb{A} . An *embedding* of \mathbb{A} into \mathbb{B} is an injective homomorphism from \mathbb{A} to \mathbb{B} such that the implication in (\clubsuit) is an equivalence.

For $n \geq 1$, we define the n -th power of a structure \mathbb{A} to be the structure \mathbb{A}^n with same signature, whose domain is A^n , and such that for all $i \in I$, a tuple $(\bar{a}^1, \dots, \bar{a}^k)$ of n -tuples is contained in $R_i^{\mathbb{A}^n}$ if, and only if, it is contained in $R_i^{\mathbb{A}}$ componentwise, i.e., $(a_j^1, \dots, a_j^k) \in R_i^{\mathbb{A}}$ for all $1 \leq j \leq n$. For graphs, this power is often called *tensor power* since the adjacency matrix of the power is a tensor power of the adjacency matrix of the original graph.

A *polymorphism* of a structure \mathbb{A} is a homomorphism from \mathbb{A}^n to \mathbb{A} , for some $n \geq 1$. We write $\text{Pol}(\mathbb{A})$ for the set of all polymorphisms of a structure \mathbb{A} . An *endomorphism* of \mathbb{A} is a homomorphism from \mathbb{A} to \mathbb{A} , i.e., a unary polymorphism of \mathbb{A} . An *automorphism* of \mathbb{A} is a bijective embedding of \mathbb{A} into \mathbb{A} .

2.2. Clones and height 1 conditions

Let A be a set. A *clone* is a set \mathcal{A} of finitary operations on A satisfying the following conditions:

- for all $n \geq 1$ and $1 \leq i \leq n$, the i -th n -ary projection $\text{pr}_i^n: (x_1, \dots, x_n) \mapsto x_i$ is a function in \mathcal{A} ;
- for all n -ary $f \in \mathcal{A}$ and all m -ary $g_1, \dots, g_n \in \mathcal{A}$, the composition $f \circ (g_1, \dots, g_n): (x_1, \dots, x_m) \mapsto f(g_1(x_1, \dots, x_m), \dots, g_n(x_1, \dots, x_m))$ is in \mathcal{A} .

For $n \geq 1$, we denote by $\mathcal{A}^{(n)}$ the set of n -ary functions in \mathcal{A} . We moreover write \mathcal{P} for the clone on the set $\{0, 1\}$ that consists only of projections. It is easy to see that for every structure \mathbb{A} the set $\text{Pol}(\mathbb{A})$ is a clone.

A *height 1 identity* is a statement of the form

$$\forall x_1, \dots, x_r. f(x_{\pi(1)}, \dots, x_{\pi(n)}) = g(x_{\rho(1)}, \dots, x_{\rho(m)})$$

where f, g are function symbols, and $\pi: \{1, \dots, n\} \rightarrow \{1, \dots, r\}, \rho: \{1, \dots, m\} \rightarrow \{1, \dots, r\}$ are any functions. We also write $f(x_{\pi(1)}, \dots, x_{\pi(n)}) \approx g(x_{\rho(1)}, \dots, x_{\rho(m)})$ for such an identity, omitting the universal quantification. An example of such an identity is the Siggers identity (\diamond).

A *height 1 condition* is a finite set Σ of height 1 identities (where several identities can use the same function symbol). Such a condition is said to be satisfied in a set of functions \mathcal{A} (e.g. a clone) if for each function symbol f appearing in Σ , there exists a function $f^{\mathcal{A}} \in \mathcal{A}$ of the corresponding arity such that every identity in Σ becomes a true statement when the symbols of Σ are instantiated by their counterparts in \mathcal{A} . In that case, we say that \mathcal{A} *satisfies* Σ .

The notion of satisfaction in a clone defines a natural quasi-order on height 1 conditions: if every clone satisfying Σ also satisfies Σ' , then we say that Σ *implies* Σ' (or that Σ' is *weaker* than Σ , or that Σ is *stronger* than Σ'). Two conditions Σ and Σ' are equivalent if a clone satisfies Σ if and only if it satisfies Σ' , i.e., if they belong to the same equivalence class of the quasi-order. The *strictly decreasing chain* in the statement of Theorem 1.4 is to be understood in this quasi-order, i.e., it is a sequence of conditions of strictly decreasing strength.

A height 1 condition is *trivial* if it is satisfied in every clone, or equivalently, if it holds in \mathcal{P} , or again equivalently, if it is implied by any other height 1 condition.

2.3. Minion homomorphisms

Let \mathcal{A}, \mathcal{B} be two clones. We say that a mapping $\xi: \mathcal{A} \rightarrow \mathcal{B}$ is a *minion homomorphism*² (introduced as *h1 clone homomorphism* in [BOP18]) if it preserves arities, and for all $f \in \mathcal{A}$ of arity n and all $\pi: \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ we have

$$\xi(f(x_{\pi(1)}, \dots, x_{\pi(n)})) = \xi(f)(x_{\pi(1)}, \dots, x_{\pi(n)}).$$

Note that a minion homomorphism preserves height 1 identities, and therefore also height 1 conditions, i.e., if there is a minion homomorphism from \mathcal{A} to \mathcal{B} and Σ is a height 1 condition such that \mathcal{A} satisfies Σ , then also \mathcal{B} satisfies Σ . In particular, if there exists a minion homomorphism from \mathcal{A} to the projection clone \mathcal{P} , then \mathcal{A} only satisfies trivial height 1 conditions. The converse of the

²A *minion* is an abstract algebraic structure and minion homomorphisms as introduced here correspond to the natural maps between such minions. The definition of a minion [BKO19, Definition 2.20] is irrelevant for our purposes so we omit it.

latter statement also holds as can be proved by a compactness argument (we give a proof in Lemma 3.5).

When \mathcal{A} and \mathcal{B} are clones, then a map $\xi: \mathcal{A} \rightarrow \mathcal{B}$ is called *uniformly continuous*³ if for every $n \geq 1$ there exists a finite set $S \subseteq A^n$ such that $f|_S = g|_S$ implies $\xi(f) = \xi(g)$ for all $f, g \in \mathcal{A}^{(n)}$. The non-existence of uniformly continuous minion homomorphisms from a clone \mathcal{A} to \mathcal{P} is equivalent to non-trivial height 1 conditions being satisfied in \mathcal{A} *locally*, in the following sense. Let Σ be a height 1 condition, and let $S \subseteq A$. We say that \mathcal{A} *satisfies* Σ *on* S if it satisfies Σ when the quantified variables in the identities of Σ only range over S (rather than A); i.e., the identities of Σ are replaced by formulas of the form

$$\forall x_1, \dots, x_r \in S. f(x_{\pi(1)}, \dots, x_{\pi(n)}) = g(x_{\rho(1)}, \dots, x_{\rho(m)}).$$

Then there is no uniformly continuous minion homomorphism from \mathcal{A} to \mathcal{P} if and only if for every finite subset S of A there exists a non-trivial height 1 condition that is satisfied by \mathcal{A} on S (again, this can be proved by a compactness argument).

2.4. Logic and model theory

The set of automorphisms of \mathbb{A} forms a group denoted by $\text{Aut}(\mathbb{A})$. For all $k \geq 1$, this group acts naturally on A^k by $\alpha \cdot (a_1, \dots, a_k) := (\alpha(a_1), \dots, \alpha(a_k))$. An *orbit* is a set of the form $\{\alpha \cdot \bar{a} \mid \alpha \in \text{Aut}(\mathbb{A})\}$ for some $\bar{a} \in A^k$. The number f_k of orbits of $\text{Aut}(\mathbb{A})$ on k -tuples is a non-decreasing function whose growth is an interesting measure of the combinatorial complexity of \mathbb{A} . If this number is finite for all $k \geq 1$, then we say that \mathbb{A} is *ω -categorical*. We say that \mathbb{A} has *less than double exponential orbit growth* if f_k is eventually dominated by 2^{2^k} , i.e., if $\lim_{k \rightarrow \infty} f_k / 2^{2^k} = 0$. A *stabilizer* of a group $\text{Aut}(\mathbb{A})$ (resp. a clone $\text{Pol}(\mathbb{A})$) is a group of the form $\text{Aut}(\mathbb{A}, a_1, \dots, a_k)$ (resp. a clone of the form $\text{Pol}(\mathbb{A}, a_1, \dots, a_k)$) where a_1, \dots, a_k are elements from A ; here, $(\mathbb{A}, a_1, \dots, a_k)$ denotes the expansion of \mathbb{A} by the unary relations $\{a_1\}, \dots, \{a_k\}$.

A first-order formula $\phi(x_1, \dots, x_n)$ is *primitive positive* (*pp*, for short) if it is of the form $\exists y_1, \dots, y_m. \bigwedge_i R_i(\bar{z}_i)$. A relation $R \subseteq A^n$ is *first-order definable* (resp. *pp-definable*) in \mathbb{A} if there exists a first-order formula ϕ (resp. pp-formula) such that $(a_1, \dots, a_n) \in R$ if and only if $\phi(a_1, \dots, a_n)$ holds in \mathbb{A} , for all $a_1, \dots, a_n \in A$. A structure \mathbb{B} is a *first-order reduct* of \mathbb{A} if \mathbb{B} and \mathbb{A} have the same domain and if every relation of \mathbb{B} is first-order definable in \mathbb{A} .

Uniformly continuous minion homomorphisms between clones have a counterpart for relational structures which we define next. Let \mathbb{A}, \mathbb{B} be relational structures. We say that \mathbb{B} is a *pp-power* of \mathbb{A} if it is isomorphic to a structure with domain A^n , where $n \geq 1$, whose relations are pp-definable from \mathbb{A} ; here, a k -ary relation on A^n is regarded as a $k \cdot n$ -ary relation on A . We say that \mathbb{B} is *pp-constructible* from \mathbb{A} if it is homomorphically equivalent to a pp-power of \mathbb{A} . The following theorem ties together the notions of pp-constructibility and minion homomorphisms.

³Clones can be naturally endowed with a uniform structure. The notion of uniform continuity corresponding to this uniform structure agrees with the definition given here; the interested reader will find details in [GP18, BOP18].

Theorem 2.1 (Theorem 1.8 in [BOP18]). *Let \mathbb{A} be an ω -categorical structure and let \mathbb{B} be a finite structure. Then \mathbb{B} is pp-constructible from \mathbb{A} if, and only if, there exists a uniformly continuous minion homomorphism from $\text{Pol}(\mathbb{A})$ to $\text{Pol}(\mathbb{B})$.*

We note, and are going to use, that the “only if” part of the statement above holds for arbitrary structures \mathbb{A} and \mathbb{B} .

If \mathbb{A} and \mathbb{B} are σ -structures such that $B \subseteq A$ and such that for every $R \in \sigma$ of arity k , $R^{\mathbb{A}} \cap B^k = R^{\mathbb{B}}$, then we say that \mathbb{B} is a *substructure* of \mathbb{A} . A structure \mathbb{A} is *homogeneous* if for every two finite substructures \mathbb{B}, \mathbb{C} and every isomorphism $f: \mathbb{B} \rightarrow \mathbb{C}$, there exists an automorphism α of \mathbb{A} such that $\alpha|_{\mathbb{B}} = f$. We note that if \mathbb{A} is homogeneous and its signature is finite, then \mathbb{A} has less than double exponential orbit growth.

A structure \mathbb{A} is *finitely bounded* if there exists a finite set \mathcal{F} of finite structures such that for every finite structure \mathbb{B} with the same signature as \mathbb{A} , \mathbb{B} embeds into \mathbb{A} if, and only if, no structure from \mathcal{F} embeds into \mathbb{B} . This is equivalent to saying that the class of finite substructures of \mathbb{A} is definable by a first-order universal sentence.

The structure constructed in Section 5 is finitely bounded but not necessarily homogeneous. However, it is *homogenizable* in the sense that by adding finitely many relations to the structure, it becomes homogeneous. In particular, the structure in Section 5 belongs to the class of *reducts of finitely bounded homogeneous structures with a finite signature*, which is the scope of the infinite-domain tractability conjecture (Conjecture 1.2).

An ω -categorical structure \mathbb{A} is a *model-complete core* if for every embedding $e: \mathbb{A} \rightarrow \mathbb{A}$ and every finite subset S of A , there exists an automorphism $\alpha \in \text{Aut}(\mathbb{A})$ such that $\alpha|_S = e|_S$.

Theorem 2.2 ([Bod07, BKO⁺17]). *Let \mathbb{A} be ω -categorical. There exists an ω -categorical model-complete core \mathbb{B} that is homomorphically equivalent to \mathbb{A} . Moreover, \mathbb{B} is unique up to isomorphism.*

The structure \mathbb{B} in the theorem above is referred to as *the model-complete core of \mathbb{A}* .

3. SIGGERS-LIKE CONDITIONS INDUCED BY GRAPHS

We show that for any non-trivial height 1 condition Σ , there is a non-trivial height 1 condition of a certain specific form, induced by a finite undirected graph, which is implied by Σ . Namely, from any finite undirected graph $\mathbb{G} = (V, E)$, one can construct a height 1 condition $\Sigma_{\mathbb{G}}$ in the following way: for each $v \in V$, one introduces a ternary function symbol f_v , and for each edge $(u, v) \in E$, one introduces a 6-ary symbol $g_{(u,v)}$, and adds to $\Sigma_{\mathbb{G}}$ the identities

$$\begin{aligned} f_u(x, y, z) &\approx g_{(u,v)}(x, y, x, z, y, z) \\ f_v(x, y, z) &\approx g_{(u,v)}(y, x, z, x, z, y). \end{aligned}$$

This corresponds to the condition $\Sigma(\mathbb{K}_3, \mathbb{G})$ constructed in [BKO19, Section 3.2]. To give a simple example, observe that if \mathbb{G} consist of a single vertex v with an edge (v, v) , then $\Sigma_{\mathbb{G}}$ is the Siggers condition (the function $g_{(v,v)}$ must satisfy the Siggers identity). We are now going to see that the Siggers condition is the strongest condition of this form; for clones over finite sets, it follows from [Sig10] that it is

also a weakest among all non-trivial height 1 conditions, and thus all non-trivial conditions of the form $\Sigma_{\mathbb{G}}$ are equivalent.

Lemma 3.1. *Let \mathbb{G} and \mathbb{H} be finite graphs. If \mathbb{G} maps homomorphically into \mathbb{H} , then $\Sigma_{\mathbb{H}}$ implies $\Sigma_{\mathbb{G}}$.*

Proof. Assume that $\Sigma_{\mathbb{H}}$ is satisfied in some clone \mathcal{C} , and fix functions $f_v \in \mathcal{C}$ for every vertex v of \mathbb{H} and functions $g_e \in \mathcal{C}$ for every edge e of \mathbb{H} witnessing this fact. Let $h: \mathbb{G} \rightarrow \mathbb{H}$ be a homomorphism. For every vertex v of \mathbb{G} we set $f'_v := f_{h(v)}$, and for every edge (u, v) of \mathbb{G} we set $g'_{(u,v)} = g_{(h(u), h(v))}$ (using the fact that h is a homomorphism). Then these functions witness the satisfaction of $\Sigma_{\mathbb{G}}$ in \mathcal{C} . \square

The condition $\Sigma_{\mathbb{G}}$ essentially forces the graph \mathbb{G} into any graph which is compatible with $\Sigma_{\mathbb{G}}$ and which contains \mathbb{K}_3 .

Lemma 3.2. *Let \mathbb{G} be a graph. Then \mathbb{G} maps homomorphically to any graph \mathbb{H} that contains \mathbb{K}_3 , and whose polymorphisms satisfy $\Sigma_{\mathbb{G}}$.*

Proof. Let v_1, v_2 and v_3 be vertices of some copy of \mathbb{K}_3 in \mathbb{H} , and assume that we have polymorphisms of \mathbb{H} satisfying the condition $\Sigma_{\mathbb{G}}$. Fix for every vertex v of \mathbb{H} a function f_v and for every edge e of \mathbb{H} a function g_e which witness this fact. We claim that the mapping $h: \mathbb{G} \rightarrow \mathbb{H}$ which sends every vertex v of \mathbb{G} to $f_v(v_1, v_2, v_3)$ is a homomorphism. Indeed, if (u, v) is an edge of \mathbb{G} then we get

$$\begin{aligned} f_u(v_1, v_2, v_3) &= g_{(u,v)}(v_1, v_2, v_1, v_3, v_2, v_3) \\ f_v(v_1, v_2, v_3) &= g_{(u,v)}(v_2, v_1, v_3, v_1, v_3, v_2) . \end{aligned}$$

Since $g_{(u,v)}$ is a polymorphism of \mathbb{H} , and since (v_i, v_j) is an edge in \mathbb{H} for all $i \neq j$, we get that $g_{(u,v)}(v_1, v_2, v_1, v_3, v_2, v_3)$ and $g_{(u,v)}(v_2, v_1, v_3, v_1, v_3, v_2)$ are related by an edge in \mathbb{H} . Hence, $(h(u), h(v)) = (f_u(v_1, v_2, v_3), f_v(v_1, v_2, v_3))$ is an edge of \mathbb{H} . \square

Finally, these tools allow us to provide a simple criterion for the triviality of conditions of the form $\Sigma_{\mathbb{G}}$. Even though the following lemma follows directly from [BKO19, Lemma 3.13], we include a proof for completeness.

Lemma 3.3. *(cf. [BKO19, Lemma 3.13]) For any finite graph \mathbb{G} , the condition $\Sigma_{\mathbb{G}}$ is trivial if and only if \mathbb{G} is 3-colorable.*

Proof. First, assume that \mathbb{G} is 3-colorable, i.e., it possesses a homomorphism to \mathbb{K}_3 . Then by the previous lemma, we have that $\Sigma_{\mathbb{G}}$ is implied by $\Sigma_{\mathbb{K}_3}$, and therefore it is enough to show that $\Sigma_{\mathbb{K}_3}$ is trivial. That is, we have to assign projections to the symbols of $\Sigma_{\mathbb{K}_3}$ in such a way that the identities are satisfied. Let 1, 2, 3 be the vertices of \mathbb{K}_3 , and define f_i to be the i -th ternary projection. Moreover, for $i \neq j$ assign to $g_{(i,j)}$ the unique 6-ary projection so that

$$\begin{aligned} f_i(x, y, z) &\approx g_{(i,j)}(x, y, x, z, y, z) \\ f_j(x, y, z) &\approx g_{(i,j)}(y, x, z, x, z, y) . \end{aligned}$$

are satisfied. By definition, this assignment satisfies $\Sigma_{\mathbb{K}_3}$.

If \mathbb{G} is not 3-colorable, then Lemma 3.2 implies that $\text{Pol}(\mathbb{K}_3)$ does not satisfy $\Sigma_{\mathbb{G}}$, and hence $\Sigma_{\mathbb{G}}$ is non-trivial. \square

Remark 3.4. The lemma implies that the problem of deciding the triviality of height 1 conditions is NP-hard, since it provides a reduction from the 3-coloring problem. The problem of deciding whether a given height 1 condition is trivial is known (in a different, but equivalent formulation) in computer science under the name *Label Cover* [ABSS97].

We now show that for each non-trivial height 1 condition Σ there is a non-3-colorable graph \mathbb{G} such that $\Sigma_{\mathbb{G}}$ is implied by Σ . We will use the folklore fact that $\text{Pol}(\mathbb{K}_3)$ does not satisfy any non-trivial height 1 condition since it only contains functions of the form $f(x_1, \dots, x_n) = \alpha(x_i)$ where $1 \leq i \leq n$ and $\alpha: \mathbb{K}_3 \rightarrow \mathbb{K}_3$ is a bijection. In particular, there exists a minion homomorphism from $\text{Pol}(\mathbb{K}_3)$ to \mathcal{P} .

Lemma 3.5. *Let \mathcal{A} be a clone that does not have a minion homomorphism to \mathcal{P} . Then there exists a finite graph \mathbb{G} which is not 3-colorable and such that \mathcal{A} satisfies $\Sigma_{\mathbb{G}}$.*

Proof. By [BKO19, Lemma 4.4], minion homomorphisms from \mathcal{A} to $\text{Pol}(\mathbb{K}_3)$ correspond precisely to 3-colorings of a certain graph $\mathbb{F} = (V, E)$, which we shall now describe (cf. [BKO19, Definition 4.1]). This (possibly infinite) graph will serve as a source of finite graphs \mathbb{G} such that \mathcal{A} satisfies $\Sigma_{\mathbb{G}}$.

We take $V := \mathcal{A}^{(3)}$, and define the edges of \mathbb{F} in the following way: $(f_1, f_2) \in E$ if and only if there exists $g \in \mathcal{A}^{(6)}$ such that

$$\begin{aligned} f_1(x, y, z) &\approx g(x, y, x, z, y, z) \\ f_2(x, y, z) &\approx g(y, x, z, x, z, y) \end{aligned}$$

holds in \mathcal{A} . Clearly, \mathcal{A} satisfies $\Sigma_{\mathbb{G}}$ for each finite subgraph \mathbb{G} of \mathbb{F} since the functions that correspond to the vertices of \mathbb{G} together with the witnesses for the edges of \mathbb{G} provide a solution to $\Sigma_{\mathbb{G}}$ (see also [BKO19, Lemma 4.3]).

Now [BKO19, Lemma 4.4] (applied to $\mathbf{A} = \mathbf{B} := \mathbb{K}_3$) states that the minion homomorphisms from \mathcal{A} to $\text{Pol}(\mathbb{K}_3)$ correspond precisely to the 3-colorings of \mathbb{F} . Since \mathcal{A} does not have any minion homomorphism to \mathcal{P} , it has none to $\text{Pol}(\mathbb{K}_3)$ either, and hence \mathbb{F} is not 3-colorable. By a standard compactness argument, there exists a finite subgraph \mathbb{G} of \mathbb{F} which is not 3-colorable. Since \mathcal{A} satisfies $\Sigma_{\mathbb{G}}$, the proof is complete. \square

Corollary 3.6. *For each non-trivial height 1 condition Σ there exists a graph \mathbb{G} that is not 3-colorable and such that $\Sigma_{\mathbb{G}}$ is implied by Σ .*

Proof. Let \mathcal{A} be the clone of term operations of the free countably generated algebra in the variety defined by Σ . Clearly, Σ witnesses that \mathcal{A} has no minion homomorphism to \mathcal{P} . Therefore, Lemma 3.5 provides a non 3-colorable graph \mathbb{G} such that \mathcal{A} satisfies $\Sigma_{\mathbb{G}}$. Since \mathcal{A} is free, we obtain that $\Sigma_{\mathbb{G}}$ is implied by Σ . \square

4. A DECREASING CHAIN OF HEIGHT 1 CONDITIONS

We now construct a sequence $(\Sigma_n)_{n \geq 1}$ of non-trivial height 1 conditions such that

- Σ_n implies Σ_{n+1} for all $n \geq 1$, and
- for every non-trivial height 1 condition Σ there exists $n \geq 1$ such that Σ implies Σ_n .




FIGURE 1. The gadget graph \mathbb{N} .

As we will see in the next section, these properties imply that for infinitely many $n \geq 1$, Σ_{n+1} does not imply Σ_n , so that by thinning out the sequence one could assume that the converse of the first property never holds. The key point of the construction is to find, for two non-trivial height 1 conditions Σ and Σ' , a non-trivial height 1 condition which is implied by both of them. Since by the results of the previous section any non-trivial height 1 condition implies a non-trivial condition of the form $\Sigma_{\mathbb{G}}$, it is sufficient to achieve this task for such conditions. Moreover, since whenever \mathbb{G} has a loop, then $\Sigma_{\mathbb{G}}$ implies all other non-trivial conditions $\Sigma_{\mathbb{G}'}$ (by Lemma 3.1), we can focus on loopless graphs.

Let us consider two loopless graphs \mathbb{G} and \mathbb{H} that are not 3-colorable, i.e., $\Sigma_{\mathbb{G}}$ and $\Sigma_{\mathbb{H}}$ are non-trivial. An edge e of a graph \mathbb{M} is called *critical* if the graph

$$\mathbb{M} - e$$

obtained from \mathbb{M} by removing e is 3-colorable. We first replace \mathbb{G} by a subgraph of \mathbb{G} that has a critical edge e . This can be done by repeatedly removing edges until we obtain a 3-colorable graph; the edge that we removed in the last step will be critical for the second-to-last graph of this procedure. Note that the height 1 condition induced by the subgraph obtained in this way is still non-trivial (since the subgraph is not 3-colorable) and implied by the height 1 condition of the original graph by Lemma 3.1. We modify \mathbb{H} in the same way as \mathbb{G} , and fix a critical edge f of \mathbb{H} .

Our next step is to glue together \mathbb{G} and \mathbb{H} at the critical edges e and f using a gadget graph \mathbb{N} , which is given in Fig. 1. The graph \mathbb{N} contains four special vertices that are labeled by x, x', y, y' , and a special edge labeled by d , and has the following properties:

- Every homomorphism $c: \mathbb{N} \rightarrow \mathbb{K}_3$ satisfies $c(x) \neq c(x')$ or $c(y) \neq c(y')$ but not both;
- Every mapping $c: \{x, x', y, y'\} \rightarrow \mathbb{K}_3$ that satisfies the property above extends to a homomorphism from \mathbb{N} to \mathbb{K}_3 .
- Every mapping $c: \{x, x', y, y'\} \rightarrow \mathbb{K}_3$ that satisfies $c(x) = c(x')$ and $c(y) = c(y')$ can be extended to a 3-coloring of $\mathbb{N} - d$.

In our glueing construction, we will only need these three properties of \mathbb{N} , i.e., any other graph with the same properties would work as well. We construct a new graph, denoted by $(\mathbb{G}, e) \oplus (\mathbb{H}, f)$, in the following way:

- (1) We first glue together \mathbb{N} and \mathbb{G} by replacing the edge e by the pair (x, x') of \mathbb{N} (the pair (x, x') remaining a non-edge), and leaving the other vertices disjoint, and then
- (2) we add the graph \mathbb{H} to the construction by replacing the edge f by the pair (y, y') of \mathbb{N} (the pair (y, y') remaining a non-edge).

Lemma 4.1. *Let $\mathbb{W} := (\mathbb{G}, e) \oplus (\mathbb{H}, f)$ be the graph as constructed above. Then:*

- (1) \mathbb{W} is not 3-colorable;
- (2) The edge d is a critical edge of \mathbb{W} ;
- (3) $\Sigma_{\mathbb{W}}$ is implied by both $\Sigma_{\mathbb{G}}$ and $\Sigma_{\mathbb{H}}$.

Proof. To prove (1), let us assume that there is a homomorphism $c: \mathbb{W} \rightarrow \mathbb{K}_3$. Then neither of its restrictions to the vertices of \mathbb{G} and the vertices of \mathbb{H} , respectively, is a 3-coloring of \mathbb{G} or \mathbb{H} , since these graphs are not 3-colorable. Since all the edges of \mathbb{G} except e are included in \mathbb{W} , these facts are witnessed on (x, x') and (y, y') , i.e., we have that $c(x) = c(x')$ and $c(y) = c(y')$. This implies that the restriction of c to \mathbb{N} is not a homomorphism (by the properties of \mathbb{N} above), a contradiction.

For (2), we have to show that removing the edge d from \mathbb{W} we obtain a 3-colorable graph. To find such a coloring, we first pick 3-colorings of $\mathbb{G} - e$ and of $\mathbb{H} - f$, and let c be the union of the two. Then c extends to a 3-coloring of \mathbb{W} , since $c(x) = c(x')$ and $c(y) = c(y')$, and by the properties of \mathbb{N} above.

We now prove (3). Due to the symmetry of the statement it is enough to prove that $\Sigma_{\mathbb{G}}$ implies $\Sigma_{\mathbb{W}}$. Let us assume that \mathcal{A} is a clone which satisfies $\Sigma_{\mathbb{G}}$, i.e., there are functions $f_v^{\mathcal{A}}$ and $g_{(u,v)}^{\mathcal{A}}$ for all vertices v of \mathbb{G} and all edges (u, v) of \mathbb{G} which witness the satisfaction of $\Sigma_{\mathbb{G}}$. We extend this family of functions to a solution of $\Sigma_{\mathbb{W}}$. Before we do that let us fix a 3-coloring c of the subgraph of \mathbb{W} induced by the vertices of \mathbb{N} and \mathbb{H} such that $c(x) = 1$ and $c(x') = 2$. Such a coloring exists by the properties of $\mathbb{H} - f$, of \mathbb{N} , and the construction of \mathbb{W} . Now, define

$$(\spadesuit.1) \quad f_1^{\mathcal{A}}(x, y, z) := g_e^{\mathcal{A}}(x, y, x, z, y, z)$$

$$(\spadesuit.2) \quad f_2^{\mathcal{A}}(x, y, z) := g_e^{\mathcal{A}}(y, x, z, x, z, y)$$

$$(\spadesuit.3) \quad f_3^{\mathcal{A}}(x, y, z) := g_e^{\mathcal{A}}(z, z, y, y, x, x) .$$

Note that $f_x^{\mathcal{A}} = f_1^{\mathcal{A}}$ and $f_{x'}^{\mathcal{A}} = f_2^{\mathcal{A}}$. For any vertex v of \mathbb{W} which is not a vertex of \mathbb{G} , we put $f_v^{\mathcal{A}} := f_{c(v)}^{\mathcal{A}}$; for any edge (u, v) of \mathbb{W} which is not an edge of \mathbb{G} , we define

$$g_{(u,v)}(x_1, \dots, x_6) := g_e(x_{\sigma(1)}, \dots, x_{\sigma(6)}) ,$$

where σ is a permutation of $\{1, \dots, 6\}$ such that the identities

$$f_u^{\mathcal{A}}(x, y, z) \approx g_{(u,v)}^{\mathcal{A}}(x, y, x, z, y, z)$$

$$f_v^{\mathcal{A}}(x, y, z) \approx g_{(u,v)}^{\mathcal{A}}(y, x, z, x, z, y)$$

hold. This is always possible since when considering any two rows of (\spadesuit) , the columns of the right-hand side contain all combinations of pairs of different variables. It is clear that these functions are defined so that they satisfy all identities of $\Sigma_{\mathbb{W}}$, which concludes the proof. \square

Let us conclude with a recursive construction of the promised sequence of height 1 conditions.

Proof of Theorem 1.4 given Theorem 1.3. We will first construct a chain of non-trivial height 1 conditions of the form $\Sigma_{\mathbb{G}}$, and then prove that for any non-trivial height 1 condition there is one in the chain that is weaker. We fix an enumeration

$$(\mathbb{G}_1, e_1), (\mathbb{G}_2, e_2), \dots$$

of all pairs where each \mathbb{G}_n is a finite loopless graph that is not 3-colorable, and e_n is a critical edge of \mathbb{G}_n . From this, we construct inductively a sequence of loopless graphs $\mathbb{H}_1, \mathbb{H}_2, \dots$. None of the graphs will be 3-colorable and all of them will have a critical edge. We start by setting $\mathbb{H}_1 := \mathbb{G}_1$ which clearly satisfies these two requirements. Assume now that we have constructed \mathbb{H}_n , and let f_n be a critical edge of it. We define

$$\mathbb{H}_{n+1} := (\mathbb{H}_n, f_n) \oplus (\mathbb{G}_{n+1}, e_{n+1}).$$

We know from Lemma 4.1 that \mathbb{H}_{n+1} is not 3-colorable and loopless, that it contains a critical edge f_{n+1} , and moreover that $\Sigma_{\mathbb{H}_n}$ and $\Sigma_{\mathbb{G}_{n+1}}$ both imply $\Sigma_{\mathbb{H}_{n+1}}$. By Lemma 3.3, $\Sigma_{\mathbb{H}_n}$ is non-trivial for all $n \geq 1$.

It remains to be verified that every non-trivial height 1 condition Σ implies $\Sigma_{\mathbb{H}_n}$ for some $n \geq 1$. Starting with a non-trivial height 1 condition Σ , we obtain by Corollary 3.6 a graph \mathbb{G} that is not 3-colorable such that Σ implies $\Sigma_{\mathbb{G}}$. It is thus sufficient to show that $\Sigma_{\mathbb{G}}$ implies $\Sigma_{\mathbb{H}_n}$ for some $n \geq 1$. We distinguish two cases: (1) \mathbb{G} contains a loop, and (2) \mathbb{G} is loopless.

(1) If \mathbb{G} contains a loop, then any graph has a homomorphism to it, and hence $\Sigma_{\mathbb{G}}$ implies $\Sigma_{\mathbb{H}_n}$ for all $n \geq 1$ (see Lemma 3.1). (2) If \mathbb{G} does not contain a loop, then we keep removing edges of \mathbb{G} until we obtain a graph that is 3-colorable, and we let \mathbb{G}' be the graph in the second-to-last step; \mathbb{G}' is not 3-colorable but contains a critical edge e' . Since \mathbb{G}' maps homomorphically to \mathbb{G} , we have that $\Sigma_{\mathbb{G}}$ implies $\Sigma_{\mathbb{G}'}$ (Lemma 3.1). Finally, there exists $n \geq 1$ such that $(\mathbb{G}', e') = (\mathbb{G}_n, e_n)$. By our construction, $\Sigma_{\mathbb{G}'}$ implies $\Sigma_{\mathbb{H}_n}$ (Lemma 4.1). Thus, $\Sigma_{\mathbb{G}}$ also implies $\Sigma_{\mathbb{H}_n}$.

We have established that $(\Sigma_{\mathbb{H}_1}, \Sigma_{\mathbb{H}_2}, \dots)$ is a sequence of non-trivial height 1 conditions of decreasing strength with the property that any non-trivial height 1 condition implies one of the conditions of the sequence. It follows from Theorem 1.3, which we are going to prove in the next section, that any sequence with these properties must strictly decrease infinitely often. \square

5. THERE IS NO WEAKEST HEIGHT 1 CONDITION

We now show that there is no weakest height 1 condition, even when restricted to reducts of finitely bounded homogeneous structures in a finite relational language. More precisely, for each non-trivial height 1 condition Σ there is a structure \mathbb{A} such that

- \mathbb{A} has a finite relational language, and is a reduct of a finitely bounded homogeneous structure;
- $\text{CSP}(\mathbb{A})$ is in P;
- $\text{Pol}(\mathbb{A})$ satisfies some non-trivial height 1 condition (or equivalently, does not possess a minion homomorphism to \mathcal{P});
- $\text{Pol}(\mathbb{A})$ does not satisfy Σ .

It follows that in the dichotomy conjecture for reducts of finitely bounded homogeneous structures, which currently characterizes tractability of a CSP by the

existence of a pseudo-Siggers polymorphism, the latter cannot be replaced by any height 1 condition. This is in contrast with the CSP dichotomy for finite structures (Theorem 1.1), which does draw the borderline between tractability and hardness by such a condition, for example, a Siggers polymorphism.

Our structures will be obtained as universal structures for graphs with forbidden homomorphic images, first constructed by Cherlin, Shelah, and Shi [CSS99] and later refined by Hubička and Nešetřil [HN16].

Definition 5.1. For a family of σ -structures \mathcal{G} , we set $\text{Forb}(\mathcal{G})$ to be the class of all σ -structures which do not contain a homomorphic image of any member of \mathcal{G} . A countable structure is *universal* for $\text{Forb}(\mathcal{G})$ if it embeds precisely those countable structures which are elements of $\text{Forb}(\mathcal{G})$.

In the following, a *cut* of a relational structure \mathbb{G} is defined to be a set of elements of \mathbb{G} whose removal disconnects the Gaifman graph of \mathbb{G} (the graph with same domain as \mathbb{G} and where there is an edge $\{x, y\}$ iff x and y appear together in some tuple of some relation of \mathbb{G}). The structure \mathbb{G} is *connected* if its Gaifman graph is. A structure \mathbb{A} has *no algebraicity* if for all $k \geq 0$ and $a_1, \dots, a_k \in A$, the finite orbits of the stabilizer $\text{Aut}(\mathbb{A}, a_1, \dots, a_k)$ are $\{a_1\}, \dots, \{a_k\}$.

Theorem 5.2 ([CSS99], Corollary of [HN16, Theorem 3.3]). *Let \mathcal{G} be a finite family of finite connected structures. There exists a countable ω -categorical structure $\text{CSS}(\mathcal{G})$ with the following properties:*

- $\text{CSS}(\mathcal{G})$ is universal for $\text{Forb}(\mathcal{G})$,
- $\text{CSS}(\mathcal{G})$ has no algebraicity,
- there exists a homogeneous expansion of $\text{CSS}(\mathcal{G})$ by finitely many pp-definable relations whose arities are the size of the minimal cuts of the structures in \mathcal{G} . Moreover, this expansion is finitely bounded.

We simply write $\text{CSS}(\mathbb{G})$ when \mathcal{G} consists of a single structure \mathbb{G} . Note that a consequence of the third item in Theorem 5.2 is that $\text{CSS}(\mathbb{G})$ resides within the scope of the infinite-domain CSP dichotomy conjecture.

Definition 5.3. We say that a function $f: A^n \rightarrow A$ is a *quasi near unanimity operation* if it satisfies the identities

$$f(y, x, \dots, x) \approx f(x, y, x, \dots, x) \approx \dots \approx f(x, \dots, x, y) \approx f(x, \dots, x),$$

i.e., if it takes the same value on all tuples that consist of a single value $x \in A$ with at most one exception.

Note that every near unanimity operation is an idempotent quasi near unanimity operation. Also note that the identities in Definition 5.3 constitute a non-trivial height 1 condition.

Lemma 5.4. *Let \mathbb{G} be a finite connected graph which is not 3-colorable, and let \mathbb{H} be universal for $\text{Forb}(\mathbb{G})$. Then:*

- $\text{Pol}(\mathbb{H})$ does not satisfy $\Sigma_{\mathbb{G}}$;
- $\text{Pol}(\mathbb{H})$ has quasi near unanimity polymorphisms of all arities larger than the number of edges of \mathbb{G} .

Proof. First, we prove that $\text{Pol}(\mathbb{H})$ does not satisfy $\Sigma_{\mathbb{G}}$. Observe that \mathbb{H} contains an isomorphic copy of \mathbb{K}_3 ; on the other hand, it does not contain a homomorphic image of \mathbb{G} . The latter is clear from the definition, and the former follows from

the assumption that \mathbb{G} is not 3-colorable, which implies that there is no homomorphism from \mathbb{G} to \mathbb{K}_3 , hence \mathbb{K}_3 embeds into \mathbb{H} by universality. The claim then follows from Lemma 3.2.

For the second claim, let n be larger than the number of edges of \mathbb{G} . To show that $\text{Pol}(\mathbb{H})$ contains a quasi near unanimity operation of arity n , we use the *indicator structure* for this condition. It is obtained by factoring the n -th Cartesian power \mathbb{H}^n of \mathbb{H} by the equivalence relation \sim which identifies all sets of tuples of the form

$$\{(x, \dots, x, y), \dots, (y, x, \dots, x), (x, \dots, x)\} .$$

There is an edge in \mathbb{H}^n/\sim between two equivalence classes A and B if and only if there exist $(u_1, \dots, u_n) \in A$ and $(v_1, \dots, v_n) \in B$ such that (u_i, v_i) is an edge in \mathbb{H} for all $1 \leq i \leq n$. We now argue that the graph \mathbb{H}^n/\sim thus obtained is an element of $\text{Forb}(\mathbb{G})$, since if that is the case, then it embeds into \mathbb{H} by universality. This embedding provides the requested quasi near unanimity polymorphism of \mathbb{H} by composing it with the factor map from \mathbb{H}^n to \mathbb{H}^n/\sim .

Assume for contradiction that there exists a homomorphism $h: \mathbb{G} \rightarrow \mathbb{H}^n/\sim$. Let us call n -tuples which are constant except for at most one value *almost constant*. These are precisely the tuples whose equivalence class with respect to \sim consists of more than one element, or equivalently, contains a constant tuple. When u is a vertex of \mathbb{G} , then we write (u_1, \dots, u_n) for the representative of the equivalence class $h(u)$ which is constant, when $h(u)$ contains such a representative, and which is the only representative of its class otherwise. Observe that if (u, v) is an edge of \mathbb{G} and

- $h(u)$ and $h(v)$ are both not almost constant, then (u_i, v_i) is an edge of \mathbb{H} for all $1 \leq i \leq n$;
- $h(u)$ is not almost constant and $h(v)$ is, then (u_i, v_i) is an edge of \mathbb{H} for all but at most one $1 \leq i \leq n$;
- $h(u)$ and $h(v)$ are both almost constant, then (u_i, v_i) is an edge of \mathbb{H} for all $1 \leq i \leq n$.

Therefore, there exists $1 \leq i \leq n$ such that (u_i, v_i) is an edge of \mathbb{H} for all edges (u, v) of \mathbb{G} . But then the mapping which sends every $u \in \mathbb{G}$ to u_i is a homomorphism from \mathbb{G} into \mathbb{H} , a contradiction. \square

Lemma 5.5. *Let \mathbb{G} be a finite graph, and let \mathbb{H} be universal for $\text{Forb}(\mathbb{G})$. Then $\text{CSP}(\mathbb{H})$ is solvable in polynomial time.*

Proof. This is obvious, as $\text{CSP}(\mathbb{H})$ corresponds to the problem of determining whether there exists a homomorphism from \mathbb{G} (which is fixed) to an input graph \mathbb{H}' , and there are at most $|\mathbb{H}'|^{|\mathbb{G}|}$ such homomorphisms. \square

We finally prove Theorem 1.3 stated in the introduction.

Proof. Let Σ be a non-trivial height 1 condition. By Corollary 3.6 there exists a non 3-colorable graph \mathbb{G} such that $\Sigma_{\mathbb{G}}$ is weaker than Σ . If \mathbb{G} is not connected, then one of its connected components \mathbb{C} is non 3-colorable. The height 1 condition $\Sigma_{\mathbb{C}}$ is non-trivial by Lemma 3.3, and clearly weaker than $\Sigma_{\mathbb{G}}$. By Lemma 5.4, $\text{Pol}(\text{CSS}(\mathbb{C}))$ does not satisfy $\Sigma_{\mathbb{C}}$ but has a quasi near unanimity operation of sufficiently large arity. In particular, $\text{Pol}(\text{CSS}(\mathbb{C}))$ does not satisfy Σ , but satisfies some non-trivial height 1 condition. \square

6. TOPOLOGY IS RELEVANT

The original CSP dichotomy conjecture for reducts of finitely bounded homogeneous structures due to Bodirsky and Pinsker (see [BPP]) claims the following borderline between tractability and hardness.

Conjecture 6.1. *Let \mathbb{A} be a reduct of a finitely bounded homogeneous structure. Exactly one of the following holds:*

- (1) *some stabilizer of the polymorphism clone of its model-complete core possesses a continuous clone homomorphism to \mathcal{P} , and $\text{CSP}(\mathbb{A})$ is NP-complete,*
- (2) *no stabilizer of the polymorphism clone of its model-complete core possesses a continuous clone homomorphism to \mathcal{P} , and $\text{CSP}(\mathbb{A})$ is in P.*

Barto and Pinsker showed in [BP16, BP18] that topology was irrelevant in this conjectured borderline, since the word ‘continuous’ can simply be dropped without changing the conjecture. More precisely, when \mathbb{A} is any ω -categorical structure with model-complete core \mathbb{B} , then some stabilizer of $\text{Pol}(\mathbb{B})$ possesses a clone homomorphism to \mathcal{P} if and only if some stabilizer of $\text{Pol}(\mathbb{B})$ possesses a continuous such homomorphism, and this is witnessed by the non-satisfaction of the pseudo-Siggers identity in $\text{Pol}(\mathbb{B})$.

Following the discovery of the importance of minion homomorphisms for the complexity of CSPs in [BOP18], it was then shown that whenever \mathbb{A} is any ω -categorical structure with less than double exponential orbit growth (a condition satisfied in particular by all structures in the range of the conjecture), then the above hardness criterion is equivalent to the existence of a uniformly continuous minor preserving map from $\text{Pol}(\mathbb{A})$ to \mathcal{P} [BKO⁺, BKO⁺17].

Naturally, the question of whether topology was irrelevant also for minion homomorphisms was raised in this context [BP16, BOP18, BP18], in particular for ω -categorical structures with less than double exponential orbit growth.

Question 6.2. *Let \mathbb{A} be an ω -categorical structure. If there exists a minion homomorphism $\text{Pol}(\mathbb{A}) \rightarrow \mathcal{P}$, does there exist a uniformly continuous one?*

While a positive answer was obtained in some special cases [BKO⁺, BKO⁺17], we are going to provide a negative answer to the question in general. The remainder of this section will be devoted to the construction of the structure \mathbb{S} and the verification of the properties claimed in Theorem 1.5.

6.1. Encoding graphs in higher arities

Our first step will be a standard construction which allows us to encode graphs as structures on n -tuples, for arbitrary $n \geq 1$.

Lemma 6.3. *Let \mathbb{G} be a finite connected loopless graph and $n \geq 1$. Then there exists a structure $\mathbb{S}(\mathbb{G}, n)$ with a single relation R of arity $2n$ such that*

- (1) *The expansion $(\mathbb{S}(\mathbb{G}, n), \neq)$ of $\mathbb{S}(\mathbb{G}, n)$ by the inequality relation is an ω -categorical model-complete core without algebraicity;*
- (2) *$\mathbb{S}(\mathbb{G}, n)$ pp-constructs the Cherlin-Shelah-Shi structure $\text{CSS}(\mathbb{G})$;*
- (3) *The relation R of $\mathbb{S}(\mathbb{G}, n)$ only contains tuples with pairwise distinct entries;*
- (4) *$\text{Aut}(\mathbb{S}(\mathbb{G}, n))$ has for every $k \geq 2$ at most $3^{k^{n|\mathbb{G}|}}$ orbits of k -tuples.*

Proof. The structure $\mathbb{S}(\mathbb{G}, n)$ is itself obtained via a CSS structure for a finite family \mathcal{G} of structures. Let \mathbb{G}' be the structure obtained from \mathbb{G} by replacing each vertex x of \mathbb{G} by an n -tuple \bar{x} of new distinct elements, and requiring, for vertices x, y of \mathbb{G} , the $2n$ -ary relation $R(\bar{x}, \bar{y})$ to hold if (x, y) is an edge in \mathbb{G} . Note that \mathbb{G}' is connected because \mathbb{G} is connected.

Let \mathcal{G} contain \mathbb{G}' as well as every connected structure on $< 2n$ elements and containing a single R -tuple (such a structure is called *loop-like* in the following). Let \mathbb{F}' be the CSS structure for \mathcal{G} . Then \mathbb{F}' is ω -categorical and has no algebraicity. We set $(\mathbb{S}(\mathbb{G}, n), \neq)$ to be the model-complete core of the structure (\mathbb{F}', \neq) ; it is also ω -categorical [Bod07], and it follows from its construction that it has no algebraicity either (cf. the proof of Theorem 27 in [BMM18]). Hence, item (1) is satisfied.

We now show item (2). Let \mathbb{T} be the graph whose vertices are the n -tuples of $\mathbb{S}(\mathbb{G}, n)$, and where two tuples (x_1, \dots, x_n) and (y_1, \dots, y_n) are related if and only if $R(x_1, \dots, x_n, y_1, \dots, y_n)$ holds in $\mathbb{S}(\mathbb{G}, n)$. Clearly, \mathbb{T} is a pp-power of $\mathbb{S}(\mathbb{G}, n)$. We claim that \mathbb{T} and $\text{CSS}(\mathbb{G})$ are homomorphically equivalent; this implies that $\text{CSS}(\mathbb{G})$ is pp-constructible from $\mathbb{S}(\mathbb{G}, n)$, as required.

To prove the claim, by a standard compactness argument it is sufficient to show that a finite graph homomorphically maps into \mathbb{T} if and only if it homomorphically maps into $\text{CSS}(\mathbb{G})$; in other words, a finite graph homomorphically maps into \mathbb{T} if and only if it does not contain a homomorphic image of \mathbb{G} . Suppose first that there existed a homomorphism from \mathbb{G} into \mathbb{T} . Then \mathbb{G}' constructed as above would have a homomorphism into $\mathbb{S}(\mathbb{G}, n)$, and hence also into \mathbb{F}' , a contradiction. Conversely, if \mathbb{H} is a graph which does not contain a homomorphic image of \mathbb{G} , then \mathbb{G}' does not homomorphically map into \mathbb{H} , and therefore \mathbb{H} embeds into \mathbb{F}' , and hence homomorphically maps into $\mathbb{S}(\mathbb{G}, n)$. But this implies that \mathbb{H} homomorphically maps to \mathbb{T} .

Item (3) of the lemma holds since we have included loop-like obstructions in the definition of \mathbb{F}' , and since \mathbb{F}' and $\mathbb{S}(\mathbb{G}, n)$ are homomorphically equivalent.

To see item (4), note that the orbit-growth of a homogeneous structure with relations of arity at most r is bounded by 3^{kr} for large enough k . By Theorem 5.2, $\text{CSS}(\mathbb{G}')$ has a homogeneous expansion by relations with arity at most $|\mathbb{G}'| = n|\mathbb{G}|$. Thus, $\text{Aut}(\mathbb{F}')$ has for large k at most $3^{k^{n|\mathbb{G}|}}$ orbits of k -tuples. Whence, the same holds for the model-complete core $\mathbb{S}(\mathbb{G}, n)$, which has at most the number of orbits of the original structure (see [BKO⁺, BKO⁺17]). \square

6.2. Superposition of the encodings

It is well-known (see [Cam90, Section 2.7]) that if two ω -categorical structures \mathbb{A} and \mathbb{B} in disjoint signatures σ and τ have no algebraicity, then there exists a generic superposition $\mathbb{A} \odot \mathbb{B}$ of the two in the signature $\sigma \cup \tau$ and which is unique up to isomorphism. This generic superposition is again ω -categorical and without algebraicity. It is obtained as follows:

- (1) Expand \mathbb{A} by all relations that have a first-order definition in \mathbb{A} , and similarly for \mathbb{B} . Call \mathbb{A}' and \mathbb{B}' the resulting structures and let σ' and τ' be their signatures (that we take to be disjoint without loss of generality).
- (2) Since \mathbb{A} and \mathbb{B} are without algebraicity, so are \mathbb{A}' and \mathbb{B}' (expanding by first-order definable relations does not change the automorphism groups

of the structures). Thus, the class of finite substructures of \mathbb{A}' and \mathbb{B}' have the strong amalgamation property (see Proposition 2.15 in [Cam90]).

- (3) The class of finite $(\sigma' \cup \tau')$ -structures whose σ' - and τ' -reducts embed into \mathbb{A}' and \mathbb{B}' , respectively, has the strong amalgamation property, and we call $\mathbb{A}' \odot \mathbb{B}'$ its Fraïssé limit. The $(\sigma \cup \tau)$ -reduct of $\mathbb{A}' \odot \mathbb{B}'$ is then our structure $\mathbb{A} \odot \mathbb{B}$.

As an example, take \mathbb{A} to be $(\mathbb{Q}, <)$ and \mathbb{B} to be the random graph (i.e., the graph $\text{CSS}(\mathbb{L})$ where \mathbb{L} is a graph on a single vertex with a loop). Then $\mathbb{A} \odot \mathbb{B}$ is the *random ordered graph*, i.e., the Fraïssé limit of the class of finite simple graphs with a total ordering on the vertices.

The same construction works for generic superpositions of infinitely many ω -categorical structures without algebraicity. The generic superposition will have an infinite signature, but will be ω -categorical if the Fraïssé class which yields the superposition has finitely many inequivalent atomic formulas of each arity.

In our construction of the structure of Theorem 1.5, we would like to superpose the graphs from the proof of Theorem 1.4; this superposition would however not be ω -categorical as there would be infinitely many orbits of pairs of vertices. This is why we superpose encodings of these graphs on tuples of increasing arity instead.

Construction 6.4. Let $\mathbb{H}_1, \mathbb{H}_2, \dots$ be an enumeration of the graphs as in the proof of Theorem 1.4 such that $\Sigma_{\mathbb{H}_1}, \Sigma_{\mathbb{H}_2}, \dots$ is a decreasing chain of height 1 conditions. Let $\alpha: \mathbb{N} \rightarrow \mathbb{N}$ be an increasing function to be determined later. Let \mathbb{S} be the generic superposition of all of the structures $\mathbb{S}(\mathbb{H}_n, \alpha(n))$, for $n \geq 1$:

$$\mathbb{S} := \bigodot_{n \geq 1} \mathbb{S}(\mathbb{H}_n, \alpha(n)) .$$

We note that by Theorem 5.2, each $\mathbb{S}(\mathbb{H}_n, \alpha(n))$ has an expansion by finitely many relations which is homogeneous. The structure obtained by expanding \mathbb{S} by this infinite set of relations is itself homogeneous. In the proof below, we call this expansion ‘the’ homogenization of \mathbb{S} , even though it is not unique.

Lemma 6.5. *The structure (\mathbb{S}, \neq) is an ω -categorical model-complete core without algebraicity.*

Proof. The generic superposition of structures without algebraicity always has no algebraicity, and expanding a structure by \neq does not introduce algebraicity since \mathbb{S} and (\mathbb{S}, \neq) have the same orbits.

We prove that \mathbb{S} is ω -categorical (which implies that (\mathbb{S}, \neq) is ω -categorical, by the sentence above). First, we prove that every atomic formula $\phi(x_1, \dots, x_r)$ over $\mathbb{S}(\mathbb{H}_n, \alpha(n))$ is either equivalent to “false” or has at least $\alpha(n)$ different variables. Suppose that ϕ is the relation symbol R_n (and thus $r = 2\alpha(n)$). By construction of $\mathbb{S}(\mathbb{H}_n, \alpha(n))$, since all loop-like structures have been forbidden, we have that either all the variables are distinct, or $\phi(x_1, \dots, x_r)$ is not satisfiable in $\mathbb{S}(\mathbb{H}_n, \alpha(n))$ and is equivalent to false. Suppose now that ϕ is a relation symbol added for the homogenization of $\mathbb{S}(\mathbb{H}_n, \alpha(n))$. Let \mathbb{H}'_n be the structure obtained from \mathbb{H}_n as in the proof of Lemma 6.3. Note that the cuts of \mathbb{H}'_n have size at least $\alpha(n)$, so we know from Theorem 5.2 that $r \geq \alpha(n)$. Moreover, Theorem 5.2 gives that ϕ is equivalent to a pp-formula over $\mathbb{S}(\mathbb{H}_n, \alpha(n))$. Then at least $\alpha(n)$ of the variables of ϕ are different, for otherwise a clause in ϕ would be of the form $R_n(y_1, \dots, y_{2\alpha(n)})$ with

fewer than $2\alpha(n)$ distinct variables, and ϕ would be equivalent to “false”. In conclusion, we obtain that all non-trivial atomic formulas over the homogenization of $\mathbb{S}(\mathbb{H}_n, \alpha(n))$ have arity at least $\alpha(n)$, and there are only finitely many of them since this homogenization has a finite signature. Thus, since α is an increasing function, the homogenization of \mathbb{S} has only finitely many atomic formulas of each arity. It follows that the homogenization of \mathbb{S} has finitely many orbits of each arity, so that this homogenization is ω -categorical, and thus \mathbb{S} itself is ω -categorical.

To see that (\mathbb{S}, \neq) is a model-complete core, let e be an endomorphism of (\mathbb{S}, \neq) , and let F be a finite subset of its domain. Then e is also, in particular, an endomorphism of $(\mathbb{S}(\mathbb{H}_n, \alpha(n)), \neq)$ for all $n \geq 1$, and since the latter structures are model-complete cores, the restriction of e to F has an expansion to an automorphism of $\mathbb{S}(\mathbb{H}_n, \alpha(n))$ for each $n \geq 1$. It then follows that this restriction is a partial isomorphism of the Fraïssé structure \mathbb{S}' of which \mathbb{S} is the reduct. By homogeneity, it extends to an automorphism of \mathbb{S}' , which is also an automorphism of \mathbb{S} and of (\mathbb{S}, \neq) . \square

We show in the next lemma that the orbit growth of \mathbb{S} can be controlled by picking a suitable α in the construction.

Lemma 6.6. *For every increasing $f: \mathbb{N} \rightarrow \mathbb{N}$ that dominates every polynomial, there exists an $\alpha: \mathbb{N} \rightarrow \mathbb{N}$ such that the number of orbits of k -tuples of \mathbb{S} is not asymptotically larger than $3^{f(k)}$. In particular, there exists an α such that \mathbb{S} has less than double exponential orbit growth.*

Proof. We construct α by induction, first setting $\alpha(1) = 1$. Suppose now that $\alpha(1), \dots, \alpha(n)$ are defined. Since f dominates every polynomial, there exists a $k_n > \alpha(n)$ such that $\sum_{i=1}^n k_n^{\alpha(i)|\mathbb{H}_i|} < f(k_n)$ for all $k \geq k_n$. Let $\alpha(n+1) := k_n + 1$.

Let now $n \geq 1$. Orbits of k_n -tuples in \mathbb{S} are uniquely determined by orbits of k_n -tuples in $\mathbb{S}(\mathbb{H}_m, \alpha(m))$ for $m \leq n$; this follows from the fact that $k_n < \alpha(m)$ for $m > n$ and that the orbits of k -tuples of $\mathbb{S}(\mathbb{H}_m, \alpha(m))$ are that of the empty structure if $k < \alpha(m)$. By Lemma 6.3, the number of orbits of k_n -tuples in \mathbb{S} is then at most $3^{k_n^{\alpha(1)|\mathbb{H}_1|}} \dots 3^{k_n^{\alpha(n)|\mathbb{H}_n|}} = 3^{\sum k_n^{\alpha(i)|\mathbb{H}_i|}} < 3^{f(k_n)}$. Therefore, the number of orbits of \mathbb{S} is bounded above by $3^{f(k)}$ infinitely often.

To prove the final remark, simply note that if the orbit growth of \mathbb{S} is doubly exponential then in particular it is asymptotically larger than $3^{3^{\sqrt{k}}}$. \square

6.3. Identities in \mathbb{S}

Lemma 6.7. *$\text{Pol}(\mathbb{S})$ does not satisfy any non-trivial height 1 condition.*

Proof. For each non-trivial height 1 condition Σ there exists an $n \geq 1$ such that $\Sigma_{\mathbb{H}_n}$ is weaker than Σ . Since $\text{CSS}(\mathbb{H}_n)$ does not satisfy $\Sigma_{\mathbb{H}_n}$, and $\text{CSS}(\mathbb{H}_n)$ is pp-constructible from $\mathbb{S}(\mathbb{H}_n, \alpha(n))$, Theorem 2.1 implies that the latter does not satisfy $\Sigma_{\mathbb{H}_n}$ either, and in particular does not satisfy Σ . Therefore, $\text{Pol}(\mathbb{S})$ does not satisfy Σ . \square

Since $\text{Pol}(\mathbb{S}, \neq) \subseteq \text{Pol}(\mathbb{S})$, we obtain in particular that $\text{Pol}(\mathbb{S}, \neq)$ does not satisfy any non-trivial height 1 condition either.

We are not going to show directly that (\mathbb{S}, \neq) satisfies non-trivial height 1 conditions locally, but will expose other (not height 1) conditions it satisfies, and then use its slow orbit growth to deduce the satisfaction of local height 1 conditions.

A polymorphism f of a structure is called a *pseudo-Siggers operation* if there are endomorphisms e_1, e_2 of the structure such that for all x, y, z of the domain

$$e_1 \circ f(x, y, x, z, y, z) = e_2 \circ f(y, x, z, x, z, y)$$

holds.

Lemma 6.8. *The structure (\mathbb{S}, \neq) has a pseudo-Siggers polymorphism.*

Proof. For each $n \geq 1$, let \mathbb{S}_n be the generic superposition

$$\mathbb{S}(\mathbb{H}_1, \alpha(1)) \odot \cdots \odot \mathbb{S}(\mathbb{H}_n, \alpha(n)) .$$

Then $\text{Pol}(\mathbb{S}_n)$ satisfies a quasi near unanimity identity of some sufficiently large arity. To see this, note that there exists $\ell \geq 1$ such that $\text{CSS}(\mathbb{H}_1), \dots, \text{CSS}(\mathbb{H}_n)$ all have a quasi near unanimity polymorphism of arity ℓ , by Lemma 5.4. Similarly, such an ℓ exists for the CSS-structures on tuples constructed in the proof of Lemma 6.3. One also sees that in Lemma 5.4, taking ℓ large enough ensures that the constructed polymorphism is also a polymorphism of $(\text{CSS}(\mathbb{H}_i), \neq)$. Thus, the model-complete cores of these structures, i.e., the structures $(\mathbb{S}(\mathbb{H}_1, \alpha(1)), \neq), \dots, (\mathbb{S}(\mathbb{H}_n, \alpha(n)), \neq)$, also have a quasi near unanimity polymorphism. Moreover, these quasi near unanimity polymorphisms have the property that they do not identify any tuples other than those required to be identified by the quasi near unanimity identities. Hence, since the superposition \mathbb{S}_n is generic, (\mathbb{S}_n, \neq) has a quasi near unanimity polymorphism of arity ℓ as well.

By [BP16, BP18], it follows that $\text{Pol}(\mathbb{S}_n, \neq)$ has a pseudo-Siggers operation for all $n \geq 1$. Fix, for each $n \geq 1$, a pseudo-Siggers operation $p_n \in \text{Pol}(\mathbb{S}_n, \neq)$. We can write

$$\text{Pol}(\mathbb{S}, \neq) = \bigcap_{n \geq 1} \text{Pol}(\mathbb{S}_n, \neq) .$$

By a standard compactness argument, there exist $\alpha_n \in \text{Aut}(\mathbb{S})$ for all $n \geq 1$ such that the sequence $(\alpha_n \circ p_n)_{n \geq 1}$ converges pointwise to a function p . Clearly, $p \in \text{Pol}(\mathbb{S}, \neq)$.

We finish the proof by showing that p is a pseudo-Siggers polymorphism of (\mathbb{S}, \neq) . Let F be a finite subset of the domain of \mathbb{S} . Then on F we have $p = \alpha_n \circ p_n$ for almost all $n \geq 1$. By the same argument as for (\mathbb{S}, \neq) , one sees that each (\mathbb{S}_n, \neq) is a model-complete core. Hence, since p_n is a pseudo-Siggers polymorphism of (\mathbb{S}_n, \neq) , there exists $\beta_n \in \text{Aut}(\mathbb{S}_n)$ such that $p_n(x, y, x, z, y, z) = \beta_n \circ p_n(y, x, z, x, z, y)$ for all $x, y, z \in F$. Altogether, we get that for almost all $n \geq 1$ we have that for all $x, y, z \in F$

$$\begin{aligned} p(x, y, x, z, y, z) &= \alpha_n \circ p_n(x, y, x, z, y, z) \\ &= \alpha_n \circ \beta_n \circ p_n(y, x, z, x, z, y) \\ &= \alpha_n \circ \beta_n \circ (\alpha_n)^{-1} \circ p(y, x, z, x, z, y) . \end{aligned}$$

This means that for almost all $n \geq 1$, there exists an automorphism of \mathbb{S}_n such that $p(x, y, x, z, y, z)$ can be composed with that automorphism from the outside to obtain $p(y, x, z, x, z, y)$ on F . By a standard compactness argument, there exists an automorphism of \mathbb{S} with this property. Again by a standard compactness argument, there exist endomorphisms of \mathbb{S} witnessing that p is a pseudo-Siggers polymorphism of (\mathbb{S}, \neq) . \square

We can therefore apply the following result from [BKO⁺].

Theorem 6.9. *Let \mathcal{C} be the polymorphism clone of an ω -categorical model-complete core. Suppose that*

- \mathcal{C} satisfies a non-trivial height 1 identity modulo outer unary functions, and
- \mathcal{C} has a uniformly continuous minion homomorphism to \mathcal{P} .

Then \mathcal{C} has at least double exponential orbit growth.

Lemma 6.10. *There is no uniformly continuous minion homomorphism from $\text{Pol}(\mathbb{S}, \neq)$ to \mathcal{P} .*

Proof. By Lemma 6.8, $\text{Pol}(\mathbb{S}, \neq)$ contains a pseudo-Siggers operation, and by Lemmas 6.5 and 6.6, (\mathbb{S}, \neq) is an ω -categorical model-complete core with less than double exponential orbit growth. Hence, Theorem 6.9 implies that $\text{Pol}(\mathbb{S}, \neq)$ has no uniformly continuous minion homomorphism to \mathcal{P} . \square

Proof of Theorem 1.5. The structure (\mathbb{S}, \neq) of Construction 6.4 is an ω -categorical model-complete core without algebraicity and with less than double exponential orbit growth by Lemmas 6.5 and 6.6. Moreover, $\text{Pol}(\mathbb{S}, \neq)$ has a minion homomorphism to \mathcal{P} by Lemma 6.7, but no uniformly continuous such map by Lemma 6.10. \square

References

- [ABSS97] Sanjeev Arora, László Babai, Jacques Stern, and Z Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *J. Comput. Syst. Sci.*, 54(2):317–331, April 1997.
- [BG08] Manuel Bodirsky and Martin Grohe. Non-dichotomies in constraint satisfaction complexity. In *International Colloquium on Automata, Languages, and Programming*, pages 184–196. Springer, 2008.
- [BJ17] Manuel Bodirsky and Peter Jonsson. A model-theoretic view on qualitative constraint reasoning. *Journal of Artificial Intelligence Research*, 58:339–385, 2017.
- [BJK05] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.
- [BJP17] 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 conference STACS 2016.
- [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).
- [BK14] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *Journal of the ACM*, 61(1):3:1–3:19, 2014.
- [BKO⁺] Libor Barto, Michael Kompatscher, Miroslav Olšák, Trung Van Pham, and Michael Pinsker. Equations in oligomorphic clones and the constraint satisfaction problem for ω -categorical structures. *Journal of Mathematical Logic*. To appear. Preprint arXiv:1612.07551.
- [BKO⁺17] Libor Barto, Michael Kompatscher, Miroslav Olšák, Trung Van Pham, and Michael Pinsker. 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.
- [BKO19] Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. Preprint arXiv:1811.00970v2, a conference version is to appear in the Proceedings of STOC 2019., April 2019.
- [BKW17] Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them. In Andrei Krokhin and Stanislav Živný, editors, *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 1–44. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017.

- [BMM18] Manuel Bodirsky, Florent Madelaine, and Antoine Mottet. A universal-algebraic proof of the complexity dichotomy for Monotone Monadic SNP. In *Proceedings of the Symposium on Logic in Computer Science – LICS’18*, 2018. Preprint available under ArXiv:1802.03255.
- [BMPP16] Manuel Bodirsky, Barnaby Martin, Michael Pinsker, and András Pongrácz. Constraint satisfaction problems for reducts of homogeneous graphs. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, pages 119:1–119:14, 2016.
- [Bod07] Manuel Bodirsky. Cores of countably categorical structures. *Logical Methods in Computer Science*, 3(1):1–16, 2007.
- [BOP18] Libor Barto, Jakub Opršal, and Michael Pinsker. The wonderland of reflections. *Israel Journal of Mathematics*, 223(1):363–398, 2018.
- [BP15] 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.
- [BP16] 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. Preprint arXiv:1602.04353.
- [BP18] Libor Barto and Michael Pinsker. Topology is irrelevant (in the infinite domain dichotomy conjecture for constraint satisfaction problems). Preprint, 2018.
- [BPP] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Projective clone homomorphisms. *Journal of Symbolic Logic*. To appear. Preprint arXiv:1409.4601.
- [Bul17] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330, Oct 2017.
- [Cam90] Peter J. Cameron. *Oligomorphic permutation groups*. Cambridge University Press, 1990.
- [CSS99] Gregory Cherlin, Saharon Shelah, and Niandong Shi. Universal graphs with forbidden subgraphs and algebraic closure. *Advances in Applied Mathematics*, 22(4):454–491, 1999.
- [FV98] 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(1):57–104, 1998.
- [GP18] Mai Gehrke and Michael Pinsker. Uniform Birkhoff. *Journal of Pure and Applied Algebra*, 222(5):1242–1250, 2018.
- [HN16] Jan Hubička and Jaroslav Nešetřil. Homomorphism and embedding universal structures for restricted classes. *Journal of Multiple-Valued Logic and Soft Computing*, 27:229–253, 2016.
- [Hod97] Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, 1997.
- [Jón67] Bjarni Jónsson. Algebras whose congruence lattices are distributive. *Mathematica Scandinavica*, 21:110–121, 1967.
- [KP17] Michael Kompatscher and Trung Van Pham. A Complexity Dichotomy for Poset Constraint Satisfaction. In *34th Symposium on Theoretical Aspects of Computer Science (STACS 2017)*, volume 66 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47:1–47:12, 2017.
- [Olš17] Miroslav Olšák. The weakest nontrivial idempotent equations. *Bulletin of the London Mathematical Society*, 49(6):1028–1047, 2017.
- [Sig10] Mark H. Siggers. A strong Mal’cev condition for locally finite varieties omitting the unary type. *Algebra universalis*, 64(1):15–20, 2010.
- [Tay88] Walter Taylor. Some very weak identities. *Algebra Universalis*, 25(1):27–35, 1988.
- [Zhu17] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 331–342, Oct 2017.

INSTITUTE OF ALGEBRA, TECHNISCHE UNIVERSITÄT DRESDEN, DRESDEN, GERMANY
Email address: manuel.bodirsky@tu-dresden.de

DEPARTMENT OF ALGEBRA, CHARLES UNIVERSITY IN PRAGUE, CZECH REPUBLIC
Email address: mottet@karlin.mff.cuni.cz

DEPARTMENT OF ALGEBRA, CHARLES UNIVERSITY IN PRAGUE, CZECH REPUBLIC
Email address: mirek@olsak.net

DEPARTMENT OF COMPUTER SCIENCE, DURHAM UNIVERSITY, DURHAM, UK
Email address: jakub.oprsal@durham.ac.uk

INSTITUTE OF DISCRETE MATHEMATICS AND GEOMETRY, TECHNISCHE UNIVERSITÄT WIEN, AUSTRIA, AND DEPARTMENT OF ALGEBRA, CHARLES UNIVERSITY IN PRAGUE, CZECH REPUBLIC
Email address: marula@gmx.at

DEPARTMENT OF PURE MATHEMATICS, UNIVERSITY OF WATERLOO, WATERLOO, ON, CANADA
Email address: ross.willard@uwaterloo.ca