

A universal-algebraic proof of the complexity dichotomy for Monotone Monadic SNP

Manuel Bodirsky
Institut für Algebra, Technische
Universität Dresden
Dresden, Germany
manuel.bodirsky@tu-dresden.de

Florent Madelaine
LIMOS, Université Clermont Auvergne
Aubière, France
florent.madelaine@udamail.fr

Antoine Mottet
Institut für Algebra, Technische
Universität Dresden
Dresden, Germany
antoine.mottet@tu-dresden.de

Abstract

The logic MMSNP is a restricted fragment of existential second-order logic which allows to express many interesting queries in graph theory and finite model theory. The logic was introduced by Feder and Vardi who showed that every MMSNP sentence is computationally equivalent to a finite-domain constraint satisfaction problem (CSP); the involved probabilistic reductions were derandomized by Kun using explicit constructions of expander structures. We present a new proof of the reduction to finite-domain CSPs that does not rely on the results of Kun. This new proof allows us to obtain a stronger statement and to verify the Bodirsky-Pinsker dichotomy conjecture for CSPs in MMSNP. Our approach uses the fact that every MMSNP sentence describes a finite union of CSPs for countably infinite ω -categorical structures; moreover, by a recent result of Hubička and Nešetřil, these structures can be expanded to homogeneous structures with finite relational signature and the Ramsey property. This allows us to use the universal-algebraic approach to study the computational complexity of MMSNP.

ACM Reference Format:

Manuel Bodirsky, Florent Madelaine, and Antoine Mottet. 2018. A universal-algebraic proof of the complexity dichotomy for Monotone Monadic SNP. In *LICS '18: 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, July 9–12, 2018, Oxford, United Kingdom*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3209108.3209156>

1 Introduction

Monotone Monadic SNP (MMSNP) is a fragment of monadic existential second-order logic whose sentences describe problems of the form “given a structure \mathfrak{A} , is there a colouring of the elements of \mathfrak{A} that avoids some fixed family of forbidden patterns?” Examples of such problems are the classical k -colourability problem for graphs (where the forbidden patterns are edges whose endpoints have the same colour), or the problem of colouring the vertices of a graph so as to avoid monochromatic triangles (Figure 1).

Manuel Bodirsky has received funding from the ERC under the European Community’s Seventh Framework Programme (Grant Agreement no. 681988, CSP-Infinity), and the DFG-funded project “Homogene Strukturen, Bedingungserfüllungsprobleme, und topologische Klone” (Project number 622397).

Antoine Mottet is supported by DFG Gratuiertenkolleg 1763 (QuantLA).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LICS '18, July 9–12, 2018, Oxford, United Kingdom

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5583-4/18/07.

<https://doi.org/10.1145/3209108.3209156>

MMSNP has been introduced by Feder and Vardi [20], whose motivation was to find fragments of existential second-order logic that exhibit a complexity dichotomy between P and NP-complete. They proved that every problem described by an MMSNP sentence is equivalent under polynomial-time randomised reductions to a constraint satisfaction problem (CSP) over a finite domain, and conjectured that every finite-domain CSP is in P or NP-complete. Kun [22] later improved the result by derandomising the equivalence, thus showing that MMSNP exhibits a complexity dichotomy if and only if the Feder-Vardi dichotomy conjecture holds. Recently, Bulatov [14] and Zhuk [26] independently proved that the dichotomy conjecture indeed holds. Both authors establish a stronger form of the dichotomy, the so-called *tractability conjecture*, which gives a characterisation of the finite-domain CSPs that are solvable in polynomial time (assuming P is not NP). This characterisation is phrased in the language of universal algebra and is moreover decidable.

The universal algebraic approach can also be used to study CSPs over *infinite* domains, and there exists a generalisation of the tractability conjecture about the so-called *reducts of finitely bounded homogeneous structures*, see e.g. [1, 3, 4, 9]. Dalmau and Bodirsky [8] showed that every problem in MMSNP is a finite union of CSPs for ω -categorical structures. These structures can be expanded to finitely bounded homogeneous structures so that they fall into the scope of the mentioned infinite-domain tractability conjecture. This poses the question whether the complexity of MMSNP can be studied directly using the universal-algebraic approach, rather than via the detour to finite-domain CSPs which involves the technically involved reduction of Kun using expander structures. In particular, even though we now have a complexity dichotomy for MMSNP, it was hitherto unknown whether the CSPs in MMSNP satisfy the infinite-domain tractability conjecture.

The main result of this paper is the confirmation of the infinite-domain tractability conjecture for CSPs in MMSNP. As a by-product, we obtain a new proof of the complexity dichotomy for MMSNP that does not rely on the results of Kun. To the best of our knowledge, this is the first-time that the universal-algebraic approach for




Figure 1. The No-monochromatic-triangle problem: the input is a finite graph G , and the question is whether there exists a colouring of the vertices of G with two colours that avoids monochromatic triangles.

infinite-domain CSPs provides a classification for a class of computational problems that has been studied in the literature before¹, and which has been introduced without having the universal-algebraic approach in mind. We also solve an open question by Lutz and Wolter [23]. Informally, we prove that the existential second-order predicates of an MMSNP sentence can be added to the original (first-order) signature of the sentence without increasing the complexity of the corresponding problem; we refer the reader to Section 4 for a formal statement. Many proofs have been omitted because of space restrictions and we refer the interested reader to the long version of the present paper².

2 MMSNP and CSPs

2.1 MMSNP

Let τ be a relational signature (we also refer to τ as the *input signature*). *SNP* is a syntactically restricted fragment of existential second order logic. A sentence in SNP is of the form $\exists P_1, \dots, P_n. \phi$ where P_1, \dots, P_n are *predicates* (i.e., relation symbols) and ϕ is a *universal* first-order-sentence over the signature $\tau \cup \{P_1, \dots, P_n\}$. *Monotone Monadic SNP without inequality*, MMSNP, is the popular restriction thereof which consists of sentences Φ of the form

$$\exists P_1, \dots, P_n \forall \bar{x} \bigwedge_i \neg(\alpha_i \wedge \beta_i),$$

where P_1, \dots, P_n are *monadic* (i.e., unary) relation symbols not in τ , where \bar{x} is a tuple of first-order variables, and for every negated conjunct:

- α_i consists of a conjunction of atomic formulas involving relation symbols from τ and variables from \bar{x} ; and
- β_i consists of a conjunction of atomic formulas or negated atomic formulas involving relation symbols from P_1, \dots, P_n and variables from \bar{x} .

Notice that the equality symbol is not allowed in MMSNP sentences. In the following, τ will always denote the input signature of an MMSNP sentence and σ the corresponding set of monadic relation symbols.

Every MMSNP τ -sentence describes a computational problem: the input consists of a finite τ -structure \mathfrak{A} , and the question is whether $\mathfrak{A} \models \Phi$, i.e., whether the sentence Φ is true in \mathfrak{A} . We sometimes identify MMSNP with the class of all computational problems described by MMSNP sentences.

We say that a conjunction of atomic literals is *connected* if its conjuncts cannot be partitioned into two non-empty sets of conjuncts with disjoint sets of variables, and *disconnected* otherwise. A conjunction of atomic literals is called *biconnected* if its conjuncts cannot be partitioned into two non-empty sets of conjuncts that share at most one common variable. Note that formulas with only one variable might not be biconnected, e.g., the formula $R_1(x) \wedge R_2(x)$ is not biconnected. An MMSNP τ -sentence Φ is called *connected* (or *biconnected*) if for each conjunct $\neg(\alpha \wedge \beta)$ of Φ where α is a conjunction of τ -formulas and β is a conjunction of unary formulas, the formula α is connected (or *biconnected*, respectively).

2.2 Constraint Satisfaction Problems

Let \mathfrak{A} and \mathfrak{B} be two structures with the same relational signature τ . A *homomorphism* from \mathfrak{A} to \mathfrak{B} is a map from A (the domain of \mathfrak{A})

to B (the domain of \mathfrak{B}) that preserves all relations. An *embedding* is a homomorphism which is additionally injective and also preserves the complements of all relations. An *automorphism* of the structure \mathfrak{B} is a surjective embedding of \mathfrak{B} into itself. We write $\text{Aut}(\mathfrak{B})$ for the group of automorphisms of the structure \mathfrak{B} . Given a τ' -structure \mathfrak{B} and $\tau \subseteq \tau'$, we denote by \mathfrak{B}^τ the τ -reduct of \mathfrak{B} , that is, the τ -structure obtained by forgetting the relation symbols from $\tau' \setminus \tau$. For a relation symbol $R \in \tau$ and a τ -structure \mathfrak{B} , we denote by $R^\mathfrak{B}$ the interpretation of R in \mathfrak{B} . For a relational τ -structure \mathfrak{B} we write

- $\text{Age}(\mathfrak{B})$ for the class of all finite τ -structures that embed into \mathfrak{B} ;
- $\text{CSP}(\mathfrak{B})$ for the class of all finite τ -structures that homomorphically map into \mathfrak{B} .

For example, $\text{CSP}(K_3)$ is the 3-colouring problem: the signature $\tau := \{E\}$ is the signature of graphs, and $K_3 := (\{0, 1, 2\}; E)$ denotes the clique with three vertices, i.e., $E^\mathfrak{B} := \{0, 1, 2\}^3 \setminus \{(0, 0), (1, 1), (2, 2)\}$.

Let \mathcal{F} be a class of finite relational τ -structures. We write

- $\text{Forb}^{\text{ind}}(\mathcal{F})$ for the class of all finite τ -structures that do not embed a structure from \mathcal{F} ;
- $\text{Forb}^{\text{hom}}(\mathcal{F})$ stands for the class of all finite τ -structures \mathfrak{A} such that no structure in \mathcal{F} homomorphically maps to \mathfrak{A} .

A relational structure \mathfrak{B} is called *finitely bounded* if it has a finite signature τ and there exists a finite set of finite τ -structures \mathcal{F} (the *bounds*) such that $\text{Age}(\mathfrak{B}) = \text{Forb}^{\text{ind}}(\mathcal{F})$.

2.2.1 Logic perspective

We present the classical terminology to pass from structures to formulas and vice versa. Let \mathfrak{A} be a τ -structure. Then the *canonical query* of \mathfrak{A} is the formula whose variables are the elements of \mathfrak{A} , and which is a conjunction that contains for every $R \in \tau$ a conjunct $R(a_1, \dots, a_n)$ if and only if $(a_1, \dots, a_n) \in R^\mathfrak{A}$.

A *primitive positive τ -formula* (also known as *conjunctive query* in database theory) is a formula that can be constructed from atomic formulas using conjunction \wedge and existential quantification \exists . We write $\phi(z_1, \dots, z_n)$ if the free variables of ϕ are contained in $\{z_1, \dots, z_n\}$. A formula without free variables is called a *sentence*.

Let ϕ be a primitive positive τ -formula without conjuncts of the form $y = y'$ and written in prenex normal form. Then the *canonical database* of ϕ is the τ -structure \mathfrak{A} whose elements are the variables of ϕ , and such that for every $R \in \tau$ we have $(a_1, \dots, a_n) \in R^\mathfrak{A}$ if and only if $R(a_1, \dots, a_n)$ is a conjunct of ϕ . We will apply the notion of canonical database also to primitive positive formulas in general, by first rewriting them into prenex form and then applying the definition above. Since the rewriting might require that some of the existentially quantified variables are renamed, the resulting canonical database is not uniquely defined; but since we usually consider structures up to isomorphism, this should not cause confusions. Also note that the information which variable is existentially quantified and which variable is free is lost in the passage from a primitive positive formula to the canonical database. The following is straightforward and well-known.

Proposition 1 ([17]). *Let \mathfrak{A} and \mathfrak{B} be two structures. The following are equivalent.*

- \mathfrak{A} has a homomorphism to \mathfrak{B} .
- $\mathfrak{B} \models \exists \bar{a}. \phi$ where ϕ is the canonical query for \mathfrak{A} and \bar{a} lists all the elements of \mathfrak{A} .

¹ https://complexityzoo.uwaterloo.ca/Complexity_Zoo:M#mmsnp.

² <http://arxiv.org/abs/1802.03255>

2.2.2 The finite-domain dichotomy theorem

We will use an important result from universal algebra, Theorem 2 below. A *polymorphism* of a structure \mathfrak{B} is a homomorphism from \mathfrak{B}^k (a finite direct power of \mathfrak{B}) to \mathfrak{B} . For every $i, j \in \mathbb{N}$, $i \leq k$, the projection $\pi_i^k: \mathfrak{B}^k \rightarrow \mathfrak{B}$ given by $\pi_i^k(x_1, \dots, x_k) := x_i$ is a polymorphism. The set of all polymorphisms of \mathfrak{B} is denoted by $\text{Pol}(\mathfrak{B})$; this set forms a *clone*, i.e., it is a set of operations on the set B that is closed under composition and contains the projections.

Let \mathcal{B} and \mathcal{C} be clones. A map $\xi: \mathcal{B} \rightarrow \mathcal{C}$ that preserves the arities is called

- a *clone homomorphism* if it maps every projection on B to the corresponding projection on C and it satisfies $\xi(f(g_1, \dots, g_n)) = \xi(f)(\xi(g_1), \dots, \xi(g_n))$ for all n -ary operations $f \in \mathcal{B}$ and all k -ary operations $g_1, \dots, g_n \in \mathcal{B}$.
- a *height 1 homomorphism* if it satisfies and $\xi(f(g_1, \dots, g_n)) = \xi(f)(g_1, \dots, g_n)$ for all n -ary operations $f \in \mathcal{B}$ and all m -ary projections g_1, \dots, g_n .

We write \mathcal{P} for the set of projections on the set $\{0, 1\}$.

Theorem 2 ([2, 3, 16]). *Let \mathfrak{B} be a finite structure. Then the following are equivalent.*

1. \mathfrak{B} has no polymorphism f of arity 6 which is Siggers, i.e., satisfies

$$\forall x, y, z. f(x, y, x, z, y, z) = f(y, x, z, x, z, y).$$

2. \mathfrak{B} has no polymorphism f of arity $k \geq 2$ which is cyclic, i.e., satisfies

$$\forall x_1, \dots, x_k. f(x_1, \dots, x_k) = f(x_2, \dots, x_k, x_1).$$

3. There exists a height 1 homomorphism from $\text{Pol}(\mathfrak{B})$ to \mathcal{P} .

It is known that if a finite structure \mathfrak{B} satisfies the equivalent items from Theorem 2, then $\text{CSP}(\mathfrak{B})$ is NP-hard [16]. Otherwise, we have the following recent result.

Theorem 3 (Finite-domain tractability theorem [14, 26]). *Let \mathfrak{B} be a finite structure with finite relational signature which does not satisfy the conditions from Theorem 2. Then $\text{CSP}(\mathfrak{B})$ is in P.*

2.2.3 Countable categoricity

Connected MMSNP sentences describe CSPs of countable structures that satisfy a strong property from model theory: *ω -categoricity*. A countably infinite structure \mathfrak{B} is called *ω -categorical* if all countable models of the first-order theory of \mathfrak{B} are isomorphic. A structure \mathfrak{B} is called *homogeneous* if every isomorphism between finite substructures of \mathfrak{B} can be extended to an automorphism of \mathfrak{B} . Homogeneous structures with finite relational signature are *ω -categorical*.

A finite or countably infinite ω -categorical structure \mathfrak{B} is called a *core* if all endomorphisms of \mathfrak{B} are embeddings, and it is called *model-complete* if all embeddings of \mathfrak{B} into \mathfrak{B} preserve all first-order formulas.

Theorem 4 ([5]). *Every ω -categorical structure \mathfrak{B} is homomorphically equivalent to a model-complete core, which is up to isomorphism unique, ω -categorical, and embeds into \mathfrak{B} .*

If \mathfrak{B} is an ω -categorical model-complete core, then adding a unary singleton relation to \mathfrak{B} does not change the computational complexity of $\text{CSP}(\mathfrak{B})$.

If B is a countable set, there is a natural metric on the set of all operations on B . For every $k \geq 1$, fix an enumeration $(\bar{b}_n)_{n \in \mathbb{N}}$ of

B^k . The distance between two functions $f, g: B^k \rightarrow B$ is defined to be 0 if $f = g$, and $\frac{1}{2^n}$ where n is the smallest index such that $f(\bar{b}_n) \neq g(\bar{b}_n)$ otherwise. This metric gives rise to the so-called *pointwise convergence topology*, and given a set \mathcal{B} of functions on B , we shall write $\bar{\mathcal{B}}$ for the topological closure of \mathcal{B} . With this metric, a map ξ from a clone \mathcal{B} on a set B to a clone \mathcal{C} on a set C is *uniformly continuous* if and only if for all $n \geq 1$ and all finite $C' \subseteq C$, there exists a finite $B' \subseteq B$ such that whenever two n -ary functions $f, g \in \mathcal{B}$ agree on B' , then $\xi(f)$ and $\xi(g)$ agree on C' .

Theorem 5 ([3]). *Let \mathfrak{B} be an ω -categorical structure. If $\text{Pol}(\mathfrak{B})$ has a uniformly continuous height 1 homomorphism to \mathcal{P} , then $\text{CSP}(\mathfrak{B})$ is NP-hard.*

2.2.4 The infinite-domain dichotomy conjecture

There are ω -categorical model-complete cores \mathfrak{B} (even homogeneous digraphs) that do not satisfy the conditions from Theorem 5 but $\text{CSP}(\mathfrak{B})$ is even undecidable [11]. So to generalise the finite-domain tractability theorem we consider a subclass of the class of all ω -categorical structures, namely structures that are homogeneous and finitely bounded. More generally, we also consider *first-order reducts* of such structures, i.e., structures \mathfrak{B} with the same domain as a homogeneous finitely bounded structure \mathfrak{C} such that all relations of \mathfrak{B} are first-order definable over \mathfrak{C} . For such structures, Bodirsky and Pinsker conjectured a pendant to the finite-domain tractability conjecture. We give here an equivalent statement to their original conjecture, reformulated using the results from [1].

Conjecture 6 (Infinite-domain tractability conjecture; see e.g. [13]). *Let \mathfrak{B} be a reduct of a finitely bounded homogeneous structure with finite relational signature. If the condition in Theorem 5 does not apply then $\text{CSP}(\mathfrak{B})$ is in P.*

2.3 Statement of the main result

The main result of this article is the proof of the infinite-domain tractability conjecture (Conjecture 6) for CSPs in MMSNP. We actually show a stronger formulation than the conjecture since we also provide a characterisation of the polynomial-time tractable cases using *pseudo-Siggers polymorphisms*. Given a set \mathcal{U} of unary operations on B , a function $f: B^6 \rightarrow B$ is called *pseudo-Siggers modulo \mathcal{U}* if there are $e_1, e_2 \in \mathcal{U}$ such that for all $x, y, z \in B$, the equation

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

is satisfied.

Theorem 7. *Let Φ be a connected MMSNP sentence. Then there exists an ω -categorical structure \mathfrak{B} such that Φ describes $\text{CSP}(\mathfrak{B})$ and such that exactly one of the following holds:*

1. $\text{Pol}(\mathfrak{B})$ has a uniformly continuous height 1 homomorphism to \mathcal{P} , and $\text{CSP}(\mathfrak{B})$ is NP-complete.
2. $\text{Pol}(\mathfrak{B})$ contains a pseudo-Siggers polymorphism modulo $\overline{\text{Aut}(\mathfrak{B})}$, and $\text{CSP}(\mathfrak{B})$ is in P.

In particular, Conjecture 6 holds for all CSPs in MMSNP.

Moreover, it is well-known that the complexity classification for MMSNP can be reduced to the complexity classification for *connected* MMSNP [20]. Together with our main result, we obtain the following corollary.

Corollary 8. *Every problem in MMSNP is in P or NP-complete.*

3 Normal Forms

We recall and adapt a normal form for MMSNP sentences that was initially proposed by Feder and Vardi [19, 20] and later extended in [25]. The normal form has been invented by Feder and Vardi to show that for every connected MMSNP sentence Φ there is a polynomial-time equivalent finite-domain CSP. In their proof, the reduction from an MMSNP sentence to the corresponding finite-domain CSP is straightforward, but the reduction from the finite-domain CSP to Φ is tricky: it uses the fact that hard finite-domain CSPs are already hard when restricted to high-girth instances. The fact that MMSNP sentences in normal form are *biconnected* is then the key to reduce high-girth instances to the problem described by Φ .

In our work, the purpose of the normal form is the reduction of the classification problem to MMSNP sentences that are *precoloured* in a sense that will be made precise in Section 4, which is later important to apply the universal-algebraic approach. Moreover, we describe a new *strong normal form* that is based on recolourings introduced by Madelaine [24]. Recolourings have been applied by Madelaine to study the computational problem whether one MMSNP sentence implies another. In our context, the importance of strong normal forms is that the templates that we construct for MMSNP sentences in strong normal form, expanded with the inequality relation \neq , are model-complete cores (Theorem 22). Let us mention that in order to get this result, the biconnectivity of the MMSNP sentences in normal form is essential (e.g. the proof of Theorem 22 uses Lemma 18, which crucially uses biconnectivity of Φ).

3.1 The normal form for MMSNP

Every connected MMSNP sentence can be rewritten to a connected MMSNP sentence of a very particular shape, and this shape will be crucial for the results that we prove in the following sections.

Definition 9 (originates from [20]; also see [25]). Let Φ be an MMSNP sentence where M_1, \dots, M_n , for $n \geq 1$, are the existentially quantified predicates (also called the *colours* in the following). Then Φ is said to be in *normal form* if it is connected and

1. (Every vertex has a colour) the first conjunct of Φ is

$$\neg(\neg M_1(x) \wedge \dots \wedge \neg M_n(x));$$
2. (Every vertex has at most one colour) Φ contains the conjunct

$$\neg(M_i(x) \wedge M_j(x))$$
 for all distinct $i, j \in \{1, \dots, n\}$;
3. (Clauses are fully coloured) for each conjunct $\neg\phi$ of Φ except the first, and for each variable x that appears in ϕ , there is an $i \leq n$ such that ϕ has a literal of the form $M_i(x)$;
4. (Clauses are biconnected) if a conjunct $\neg\phi$ of Φ is not of the form as described in item 1 and 2, the formula ϕ is biconnected;
5. (Small clauses are explicit) any $(\tau \cup \{M_1, \dots, M_n\})$ -structure \mathfrak{A} with at most k elements satisfies the first-order part of Φ if \mathfrak{A} satisfies all conjuncts of Φ with at most k variables.

Note that when Φ is in normal form, then in all conjuncts $\neg\phi$ of Φ except for the first we can drop conjuncts where predicates appear negatively in ϕ ; hence, we assume henceforth that ϕ is a conjunction of atomic formulas.

Lemma 10. *Every connected MMSNP sentence Φ is equivalent to an MMSNP sentence Ψ in normal form, and Ψ can be computed from Φ .*

The following lemma states a key property that we have achieved with our normal form (in particular, we use the biconnectivity assumption).

Lemma 11. *Let ϕ be the first-order part of an MMSNP τ -sentence in normal form with colour set σ and let $\psi_1(x, \bar{y})$ and $\psi_2(x, \bar{z})$ be two conjunctions of atomic $(\tau \cup \sigma)$ -formulas such that*

- \bar{y} and \bar{z} are vectors of disjoint sets of variables;
- the canonical databases of ψ_1 and of ψ_2 satisfy ϕ ;
- for every $M \in \sigma$, ψ_1 contains the literal $M(x)$ if, and only if, ψ_2 does.

Then the canonical database \mathfrak{A} of $\psi_1(x, \bar{y}) \wedge \psi_2(x, \bar{z})$ satisfies ϕ .

3.2 Templates for sentences in normal form

Let Φ be an MMSNP τ -sentence in normal form. Let σ be the set of colours of Φ . We will now construct an ω -categorical $(\tau \cup \sigma)$ -structure \mathbb{C}_Φ for Φ ; this structure will have several important properties:

1. a structure \mathfrak{A} satisfies Φ if and only if \mathfrak{A} homomorphically maps to \mathbb{C}_Φ^τ , the τ -reduct of \mathbb{C}_Φ ;
2. for every colour M of Φ , $M^{\mathbb{C}_\Phi}$ is an orbit of elements under $\text{Aut}(\mathbb{C}_\Phi)$; moreover, every orbit is of this form.
3. (\mathbb{C}_Φ, \neq) is a model-complete core;
4. if Φ is furthermore in *strong normal form* (to be introduced in Section 3.3) then even $(\mathbb{C}_\Phi^\tau, \neq)$ is a model-complete core.

In order to do this, we follow a strategy similar to the one used in [8]: we associate with Φ a set of finite connected structures \mathcal{F} , and use a theorem by Cherlin, Shelah, and Shi [18] to obtain an ω -categorical structure $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$ whose age is equal to $\text{Forb}^{\text{hom}}(\mathcal{F})$. The structure \mathbb{C}_Φ is then obtained by considering a substructure of $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$.

Theorem 12 (Theorem 4 in [18]). *Let \mathcal{F} be a finite set of finite connected τ -structures. Then there exists a countable model-complete τ -structure $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$ such that $\text{Age}(\mathfrak{B}_{\mathcal{F}}^{\text{ind}}) = \text{Forb}^{\text{hom}}(\mathcal{F})$. The structure $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$ is up to isomorphism unique and ω -categorical.*

Let $\mathfrak{B}_{\mathcal{F}}^{\text{hom}}$ be so that $(\mathfrak{B}_{\mathcal{F}}^{\text{hom}}, \neq)$ is the model-complete core of $(\mathfrak{B}_{\mathcal{F}}^{\text{ind}}, \neq)$. By Theorem 4, we can see $\mathfrak{B}_{\mathcal{F}}^{\text{hom}}$ as a substructure of $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$.

Theorem 13. *Let \mathcal{F} be a finite set of finite connected τ -structures. Then for every finite τ -structure \mathfrak{A} , \mathfrak{A} homomorphically and injectively maps to $\mathfrak{B}_{\mathcal{F}}^{\text{hom}}$ if and only if $\mathfrak{A} \in \text{Forb}^{\text{hom}}(\mathcal{F})$.*

We apply the previous theorem to the following set of finite structures.

Definition 14. Let Φ be an MMSNP τ -sentence in normal form with colours σ . The *coloured obstruction set* for Φ is the set \mathcal{F} of all canonical databases (in the signature $\tau \cup \sigma$) for formulas ϕ such that $\neg\phi$ is a conjunct of Φ , except for the first conjunct.

The structure $\mathfrak{B}_{\mathcal{F}}^{\text{hom}}$ might contain some elements that do not belong to the interpretation of any symbol $M \in \sigma$. Since we are interested in a template for Φ , this naturally motivates the following definition.

Definition 15. Let Φ be an MMSNP τ -sentence in normal form and let \mathcal{F} be the coloured obstruction set of Φ . Then \mathbb{C}_Φ denotes the substructure of $\mathfrak{B}_{\mathcal{F}}^{\text{hom}}$ induced by the coloured elements of $\mathfrak{B}_{\mathcal{F}}^{\text{hom}}$.

The τ -reduct \mathbb{C}_Φ^τ of the structure \mathbb{C}_Φ that we constructed for an MMSNP sentence Φ in normal form is indeed a template for the CSP described by Φ .

Lemma 16. *Let Φ be an MMSNP τ -sentence in normal form and let \mathfrak{A} be a τ -structure. Then the following are equivalent.*

1. $\mathfrak{A} \models \Phi$;
2. \mathfrak{A} homomorphically and injectively maps to \mathbb{C}_Φ^τ ;
3. \mathfrak{A} homomorphically maps to \mathbb{C}_Φ^τ .

Proof. Let σ be the colour set and let \mathcal{F} be the coloured obstruction set of Φ .

(1) \Rightarrow (2). If \mathfrak{A} satisfies Φ it has a $(\tau \cup \sigma)$ -expansion \mathfrak{A}' such that no structure in \mathcal{F} homomorphically maps to \mathfrak{A}' . So \mathfrak{A}' homomorphically and injectively maps to $\mathfrak{B}_{\mathcal{F}}^{\text{hom}}$ by Theorem 13. Moreover, every element of \mathfrak{A}' is contained in one predicate from σ (because of the first conjunct of Φ) and hence the image of the embedding must lie in \mathbb{C}_Φ .

(2) \Rightarrow (3) is trivial.

(3) \Rightarrow (1). Let h be the homomorphism from \mathfrak{A} to \mathbb{C}_Φ^τ . Expand \mathfrak{A} to a $(\tau \cup \sigma)$ -structure \mathfrak{A}' by colouring each element $a \in A$ by the colour of $h(a)$ in \mathbb{C}_Φ ; then there is no homomorphism from a structure $\mathfrak{F} \in \mathcal{F}$ to \mathfrak{A}' , since the composition of such a homomorphism with h would give a homomorphism from \mathfrak{F} to $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$, a contradiction. The expansion \mathfrak{A}' also satisfies the first conjunct of Φ , and hence $\mathfrak{A} \models \Phi$. \square

Example 17. Consider the set \mathcal{F} of obstructions given in Figure 1. The structure $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$ is an infinite graph whose vertices are partitioned into three parts: uncoloured vertices, magenta (round) vertices, and blue (square) vertices. The graph induced by the uncoloured vertices is isomorphic to the random graph with random loops, while the parts induced by the magenta and blue vertices are isomorphic to the countable homogeneous triangle-free graph. The edges with endpoints in two different parts form a graph that is isomorphic to the countable homogeneous tripartite graph. The corresponding structure \mathbb{C}_Φ does not contain the uncoloured vertices; moreover, the graph formed by the edges with endpoints in two different parts form a complete bipartite graph.

We now describe some effects that the normal form has on our template. The orbit of a tuple $(a_1, \dots, a_k) \in (\mathbb{C}_\Phi)^k$ under $\text{Aut}(\mathbb{C}_\Phi)$ is the set of tuples $(\alpha(a_1), \dots, \alpha(a_k))$, where $\alpha \in \text{Aut}(\mathbb{C}_\Phi)$.

Lemma 18. *Let Φ be an MMSNP τ -sentence in normal form with colours σ . Then the orbits of elements of \mathbb{C}_Φ under $\text{Aut}(\mathbb{C}_\Phi)$ are precisely the sets $M^{\mathbb{C}_\Phi}$ for $M \in \sigma$.*

Lemma 19. *Let Φ be an MMSNP τ -sentence in normal form. Then (\mathbb{C}_Φ, \neq) is a model-complete core.*

3.3 The strong normal form

Let Φ_1 and Φ_2 be two MMSNP τ -sentences in normal form with colour sets σ_1 and σ_2 , respectively. For $r: \sigma_1 \rightarrow \sigma_2$ and a $(\tau \cup \sigma_1)$ -structure \mathfrak{A} we write $r(\mathfrak{A})$ for the structure obtained from \mathfrak{A} by renaming each predicate $P \in \sigma_1$ to $r(P) \in \sigma_2$.

Definition 20. A *recolouring* (from Φ_1 to Φ_2) is given by a function $r: \sigma_1 \rightarrow \sigma_2$ such that for every $(\tau \cup \sigma_1)$ -structure \mathfrak{A} , if a coloured obstruction of Φ_2 homomorphically maps to $r(\mathfrak{A})$, then a coloured obstruction of Φ_1 homomorphically maps to \mathfrak{A} . A recolouring $r: \sigma_1 \rightarrow \sigma_2$ is said to be *proper* if r is non-injective.

An MMSNP sentence Φ is defined to be in *strong normal form* if it is in normal form and there is no proper recolouring from Φ to Φ . In other words, a formula is in strong normal form if it cannot be simplified by removing colours without changing the set of structures described by the formula.

Theorem 21. *For every connected MMSNP sentence Φ there exists an equivalent connected MMSNP Ψ in strong normal form, and Ψ can be effectively computed from Φ .*

The first consequence of the strong normal form for our template is that already the τ -reduct \mathbb{C}_Φ^τ of \mathbb{C}_Φ , when expanded with the inequality relation, is a model-complete core.

Theorem 22. *Let Φ be an MMSNP sentence in strong normal form and with input signature τ . Then $(\mathbb{C}_\Phi^\tau, \neq)$ is a model-complete core.*

The second consequence of the strong normal form is described in the next section.

4 Precoloured MMSNP

An MMSNP τ -sentence Φ in normal form is called *precoloured* if, informally, for each colour of Φ there is a corresponding unary relation symbol in τ that forces elements to have this colour. In this section we show that every MMSNP sentence is polynomial-time equivalent to a precoloured MMSNP sentence; this answers a question posed in [23]. We first formally introduce precoloured MMSNP and state some basic properties in Section 4.1. We then prove a stronger result than the complexity statement above: we show that the Bodirsky-Pinsker tractability conjecture is true for CSPs in MMSNP if and only if it is true for CSPs in precoloured MMSNP (Corollary 28). The main results are stated in Section 4.3. In Section 4.4 we complete the proofs of the results in this section.

4.1 Basic properties of precoloured MMSNP

Formally, an MMSNP τ -sentence Φ is *precoloured* if it is in normal form and for every colour M of Φ there exists a unary symbol $P_M \in \tau$ such that for every colour M' of Φ which is distinct from M the formula Φ contains the conjunct $\neg(P_M(x) \wedge M'(x))$.

Lemma 23. *Every precoloured MMSNP sentence is in strong normal form.*

Proof. Let Φ be a precoloured MMSNP sentence with colour set σ . We will show that every recolouring $r: \sigma \rightarrow \sigma$ of Φ must be the identity. Let $M \in \sigma$, and let A be the canonical database of $P_M(x) \wedge M(x)$. Note that A does not contain any obstruction of Φ as homomorphic image. But if $M' := r(M) \neq M$, then $r(A)$ contains the canonical database of $P_{M'}(x) \wedge M'(x)$, in contradiction to the assumption that r is a recolouring. Hence, $r(M) = M$ for all $M \in \sigma$. \square

Lemma 24. *Let Φ be a precoloured MMSNP sentence. Then for each colour M , the symbol P_M and M have the same interpretation in \mathbb{C}_Φ .*

Proof. By Lemma 19 the structure (\mathbb{C}_Φ, \neq) is a model-complete core. Note that the ω -categorical structures $(\mathbb{C}_\Phi, \neq, M)$ and $(\mathbb{C}_\Phi, \neq, P_M)$

have the same CSP, and hence are *homomorphically equivalent*, i.e., there are homomorphisms between the two structures in both directions. Thus, the statement follows from the fact that ω -categorical model-complete cores are up to isomorphism unique [5]. \square

4.2 Adding inequality

We first show that adding the inequality relation to \mathbb{C}_Φ^τ does not increase the complexity of its CSP. Since $(\mathbb{C}_\Phi^\tau, \neq)$ is a model-complete core, this then allows us to add constants to the language of the structure, also without increasing the complexity of the CSP.

Proposition 25. *CSP(\mathbb{C}_Φ^τ) and CSP($\mathbb{C}_\Phi^\tau, \neq$) are polynomial-time equivalent.*

Proof. If a given instance of CSP($\mathbb{C}_\Phi^\tau, \neq$), viewed as a primitive positive sentence, contains conjuncts of the form $x \neq x$, then the instance is unsatisfiable. Otherwise, we only consider the constraints using relations from τ , and let \mathfrak{A} be the canonical database of those constraints. If \mathfrak{A} has no homomorphism to \mathbb{C}_Φ^τ then the instance is unsatisfiable. Otherwise, by Lemma 16 there is an injective homomorphism from \mathfrak{A} to \mathbb{C}_Φ^τ . The injectivity implies that the homomorphism also satisfies all the inequality constraints, so we have a polynomial-time reduction from CSP($\mathbb{C}_\Phi^\tau, \neq$) to CSP(\mathbb{C}_Φ^τ). \square

Together with the previous proposition, the following lemma shows that \mathbb{C}_Φ^τ satisfies the Bodirsky-Pinsker conjecture if, and only if, $(\mathbb{C}_\Phi^\tau, \neq)$ does.

Lemma 26. *There is a uniformly continuous height 1 homomorphism $\text{Pol}(\mathbb{C}_\Phi^\tau) \rightarrow \mathcal{P}$ if, and only if, there is a uniformly continuous height 1 homomorphism $\text{Pol}(\mathbb{C}_\Phi^\tau, \neq) \rightarrow \mathcal{P}$.*

4.3 The standard precolouration

Let Φ be an MMSNP sentence in strong normal form with colour set σ , and let Ψ be the following precoloured MMSNP sentence: we obtain Ψ from Φ by adding for each $M \in \sigma$ a new input predicate P_M and adding the conjunct $\neg(P_M(x) \wedge M'(x))$ for each colour $M' \in \sigma \setminus \{M\}$. We call this sentence the *standard precolouration* of Φ . The main result of this section is the following.

Theorem 27. *Let Φ be an MMSNP sentence in strong normal form with input signature τ . Let Ψ be the standard precolouration of Φ , and let ρ be the input signature of Ψ . Then CSP(\mathbb{C}_Φ^τ) and CSP(\mathbb{C}_Ψ^ρ) are equivalent under polynomial-time reductions. Moreover, there exists a uniformly continuous height 1 homomorphism $\text{Pol}(\mathbb{C}_\Phi^\tau) \rightarrow \mathcal{P}$ if, and only if, there exists a uniformly continuous height 1 homomorphism $\text{Pol}(\mathbb{C}_\Psi^\rho) \rightarrow \mathcal{P}$.*

The proof of this theorem will be given in Section 4.4. We first point out an immediate consequence.

Corollary 28. *Conjecture 6 holds for CSPs in MMSNP if, and only if, it holds for CSPs described by a precoloured MMSNP sentence.*

4.4 Proof of the precolouring theorem

Let \mathfrak{A} be a properly coloured $(\tau \cup \sigma)$ -structure, i.e., every element appears in the interpretation of precisely one symbol from σ . For an element $a \in A$, denote by $\mathfrak{A}[a \mapsto *]$ the structure obtained by uncolouring a . For $M \in \sigma$ and a tuple \bar{a} of elements \mathfrak{A} , denote by $\mathfrak{A}[\bar{a} \mapsto M]$ the structure obtained by uncolouring the elements of \bar{a} , and giving them the colour M . Let $C(\mathfrak{A}, a)$ be the subset of \mathbb{C}_Φ

containing all elements c such that there exists a homomorphism

$$h: \mathfrak{A}[a \mapsto *] \rightarrow \mathbb{C}_\Phi$$

that satisfies $h(a) = c$. Note that $C(\mathfrak{A}, a)$ is, by Lemma 18, a union of colours. So we can also see $C(\mathfrak{A}, a)$ as the union of $M^{\mathbb{C}_\Phi}$ for $M \in \sigma$ such that $\mathfrak{A}[a \mapsto M]$ is \mathcal{F} -free.

Lemma 29. *Suppose that Φ is in strong normal form, and let M be a colour of Φ . Then $M^{\mathbb{C}_\Phi} = \bigcap C(\mathfrak{A}, a)$ where the intersection ranges over all $\mathfrak{A} \in \mathcal{F}$ and $a \in F$ such that $M^{\mathbb{C}_\Phi} \subseteq C(\mathfrak{A}, a)$.*

Proof. The left-to-right inclusion is clear. We prove the other inclusion. To do this, it suffices to show that for every $M' \in \sigma \setminus \{M\}$, there exists $\mathfrak{G} \in \mathcal{F}$ and $b \in G$ such that $M^{\mathbb{C}_\Phi} \subseteq C(\mathfrak{G}, b)$ but $(M')^{\mathbb{C}_\Phi} \not\subseteq C(\mathfrak{G}, b)$. Let $r: \sigma \rightarrow \sigma$ be defined by $r(M) = M'$ and $r(N) = N$ for all $N \in \sigma \setminus \{M\}$. Since Φ is in strong normal form and r is not surjective, it cannot be a recolouring of Φ . This means that there exists a \mathcal{F} -free structure \mathfrak{A} and $\mathfrak{F} \in \mathcal{F}$ such that there exists a homomorphism $h: \mathfrak{A} \rightarrow r(\mathfrak{A})$. Let a_1, \dots, a_k be the elements of F that are mapped to $M^{\mathfrak{A}}$ by h . In $r(\mathfrak{A})$, these elements are in M' , so since h is a homomorphism and \mathfrak{F} is completely coloured, we have that $a_1, \dots, a_k \in (M')^{\mathfrak{F}}$. Moreover, since \mathfrak{A} is \mathcal{F} -free, the structure $\mathfrak{F}[a_1, \dots, a_k \mapsto M]$ is \mathcal{F} -free. Let $0 \leq j \leq k$ be minimal such that $\mathfrak{F}[a_1, \dots, a_j \mapsto M]$ is \mathcal{F} -free. Since $\mathfrak{F} \in \mathcal{F}$, we have $j \geq 1$. Let now $\mathfrak{G} \in \mathcal{F}$ be such that there exists $g: \mathfrak{G} \rightarrow \mathfrak{F}[a_1, \dots, a_{j-1} \mapsto M]$, which exists by minimality of j . Note that a_j is in the image of g , otherwise g would be a homomorphism $g: \mathfrak{G} \rightarrow \mathfrak{F}[a_1, \dots, a_j \mapsto M]$, in contradiction to the choice of j . Thus, let $b \in G$ be such that $g(b) = a_j$, and note that $b \in (M')^{\mathfrak{G}}$, so that $(M')^{\mathbb{C}_\Phi} \not\subseteq C(\mathfrak{G}, b)$. Since g is a homomorphism $\mathfrak{G}[b \mapsto M] \rightarrow \mathfrak{F}[a_1, \dots, a_j \mapsto M]$, the structure $\mathfrak{G}[b \mapsto M]$ is \mathcal{F} -free. This implies that $M^{\mathbb{C}_\Phi} \subseteq C(\mathfrak{G}, b)$. We therefore found a $\mathfrak{G} \in \mathcal{F}$ and $b \in G$ such that $M^{\mathbb{C}_\Phi} \subseteq C(\mathfrak{G}, b)$ but $(M')^{\mathbb{C}_\Phi} \not\subseteq C(\mathfrak{G}, b)$. \square

If the sets of the form $C(\mathfrak{A}, a)$ were pp-definable (i.e., definable by a primitive positive formula) in an expansion of (\mathbb{C}_Φ, \neq) by finitely many constants, we would be done for the proof of Theorem 27 since the intersection in Lemma 29 is finite. We show how to approximate these sets by pp-definable subsets.

For $M \in \sigma$, let $P(M)$ be the set of pairs (\mathfrak{A}, a) such that $M^{\mathbb{C}_\Phi} \subseteq C(\mathfrak{A}, a)$. Let $(\mathfrak{A}, a) \in P(M)$. Let a_1, \dots, a_k be the elements of \mathfrak{A} that are distinct from a . Let $\phi_{\mathfrak{A}}(a, a_1, \dots, a_k)$ be the canonical query of \mathfrak{A} . Let M_1, \dots, M_k be the colours of these elements in \mathfrak{A} . Fix $\psi_{\mathfrak{A}, a}(x, U_1, \dots, U_k)$ to be the formula

$$\exists y_1, \dots, y_k \left(\phi_{\mathfrak{A}}(x, y_1, \dots, y_k) \wedge \bigwedge_{i \in \{1, \dots, k\}} U_i(y_i) \right),$$

in the language $\tau \cup \{U_1, \dots, U_k\}$. Let $\chi_M^{(0)}$ be $M(x)$. We define $\chi_M^{(n)}$ inductively. For $n \geq 0$, let

$$\chi_M^{(n+1)}(x) := \bigwedge_{(\mathfrak{A}, a) \in P(M)} \psi_{\mathfrak{A}, a}(x, \chi_{M_1}^{(n)}, \dots, \chi_{M_k}^{(n)}).$$

Lemma 30. *For any $n \in \mathbb{N}$ and $M \in \sigma$ the formula $\chi_M^{(n)}(x)$ defines $M^{\mathbb{C}_\Phi}$ over \mathbb{C}_Φ .*

Proof. We prove the result by induction, the case $n = 0$ being trivial. Suppose that the result is proved for some $n \geq 0$. From Lemma 29 and the induction hypothesis follows that $\chi_M^{(n+1)}(x)$ defines a subset of $M^{\mathbb{C}_\Phi}$, so we just have to prove that the formula is satisfiable (then




Figure 2. Illustration of the formula $\chi_M^{(2)}(x)$, for the MMSNP sentence of Example 31. All the variables except for x are existentially quantified.

by Lemma 18, we get that $\chi_M^{(n+1)}$ defines $M^{\mathfrak{C}_\Phi}$. By Lemma 11, if $\chi_M^{(n+1)}$ is not satisfiable then there must exist $(\mathfrak{F}, a) \in P(M)$ such that $\psi_{\mathfrak{F}, a}(x, \chi_{M_1}^{(n)}, \dots, \chi_{M_k}^{(n)})$ is not satisfiable, i.e.,

$$\phi_{\mathfrak{F}}(x, y_1, \dots, y_k) \wedge \bigwedge_{i \in \{1, \dots, k\}} \chi_{M_i}^{(n)}(y_i)$$

is not satisfiable, where M_1, \dots, M_k are the colours in \mathfrak{F} of the elements other than a . By Lemma 11 again, and since $\phi_{\mathfrak{F}}(x, y_1, \dots, y_k)$ is clearly satisfiable, there must exist $i \in \{1, \dots, k\}$ such that $\chi_{M_i}^{(n)}(y_i)$ is not satisfiable, in contradiction to our induction hypothesis. Therefore, $\chi_M^{(n+1)}$ is satisfiable. \square

Example 31. We show in Figure 2 the construction of the formula $\chi_M^{(2)}$ in the case of the MMSNP sentence given by the obstructions in Figure 1, where M is represented in magenta (filled round vertices). Note that if \mathfrak{F} is the triangle made of blue (square) vertices and a is a vertex of this triangle then $C(\mathfrak{F}, a) = M^{\mathfrak{C}_\Phi}$. Note that each y_i must be coloured blue (otherwise the triangle in magenta would appear), so that x necessarily belongs to $M^{\mathfrak{C}_\Phi}$. This shows that $\chi_M^{(2)}(x)$ defines a subset of $M^{\mathfrak{C}_\Phi}$.

Let $n > |\Phi|$. It is a consequence of Lemma 30 that for each $M \in \sigma$, the formula $\chi_M^{(n)}(x)$ is satisfiable in \mathfrak{C}_Φ . Let \mathfrak{A} be the canonical query of $\chi_M^{(n)}(x)$, where we additionally colour the elements of \mathfrak{A} according to an arbitrary satisfying assignment for $\chi_M^{(n)}$. Then \mathfrak{A} homomorphically maps to \mathfrak{C}_Φ^τ , so by Lemma 16 it also injectively maps to \mathfrak{C}_Φ^τ . We deduce from this that $\chi_M^{(n)}$ is satisfiable by an injective assignment h . For every $M' \in \sigma$, replace in $\chi_M^{(n)}$ each literal $M'(y)$ (the vertices at the bottom level, in Figure 2) by the literal $y = h(y)$. The resulting formula, $\tilde{\chi}_M(x)$, is then a primitive positive formula in an expansion of \mathfrak{C}_Φ^τ by finitely many constants \bar{c} .

Lemma 32. *The formula $\tilde{\chi}_M(x)$ defines a subset of $M^{\mathfrak{C}_\Phi}$ in $(\mathfrak{C}_\Phi^\tau, \bar{c})$.*

Proof. Immediate from Lemma 30 and the definition of $\tilde{\chi}_M$. \square

We claim that the formulas $\tilde{\chi}$ define a universal substructure of \mathfrak{C}_Φ , in the sense that any structure \mathfrak{A} that has a homomorphism to \mathfrak{C}_Φ has a homomorphism h to \mathfrak{C}_Φ such that $\mathfrak{C}_\Phi \models \tilde{\chi}_M(h(a))$ for every $a \in M^{\mathfrak{A}}$.

Proposition 33. *Let \mathfrak{A} be a finite structure that has a homomorphism to \mathfrak{C}_Φ , and let $\phi_{\mathfrak{A}}(a_1, \dots, a_k)$ be the canonical query of \mathfrak{A} . Let*

M_i be the colour of a_i in \mathfrak{A} . Let $n > |\Phi|$. Then the formula

$$\phi_{\mathfrak{A}}(x_1, \dots, x_k) \wedge \bigwedge_{1 \leq i \leq k} \tilde{\chi}_{M_i}(x_i)$$

is satisfiable in $(\mathfrak{C}_\Phi, \bar{c})$.

Finally, we need the following concepts for the proof of Theorem 27. A *pp-power* of \mathfrak{B} is a structure with domain B^d , for $d \in \mathbb{N}$, whose k -ary relations are primitive positive definable when viewed as dk -ary relations over \mathfrak{B} . A structure \mathfrak{C} is said to have a *pp-construction* over \mathfrak{B} if it is homomorphically equivalent to a pp-power of \mathfrak{B} . It is known [3] that the expansion of a model-complete core \mathfrak{C} by finitely many constants is pp-constructible in \mathfrak{C} . Moreover, we have the following.

Lemma 34 ([3]). *Let \mathfrak{B} and \mathfrak{C} be two relational structures with finite signature. If \mathfrak{C} has a pp-construction in \mathfrak{B} , then:*

- *CSP(\mathfrak{C}) reduces to CSP(\mathfrak{B}) in polynomial time, and*
- *there is a uniformly continuous height 1 homomorphism from $\text{Pol}(\mathfrak{B})$ to $\text{Pol}(\mathfrak{C})$.*

Proof of Theorem 27. We first show that \mathfrak{C}_Ψ^ρ is pp-constructible in $(\mathfrak{C}_\Phi^\tau, \neq)$. By Lemma 34 and Proposition 25, this shows that $\text{CSP}(\mathfrak{C}_\Psi^\rho)$ reduces in polynomial time to $\text{CSP}(\mathfrak{C}_\Phi^\tau)$. Let \mathfrak{D} be the expansion with signature ρ of the structure \mathfrak{C}_Φ^τ such that for every colour $M \in \sigma$ of Φ the symbol $P_M \in \rho$ denotes the relation defined by the formula $\tilde{\chi}_M$ from Lemma 32. Since $(\mathfrak{C}_\Phi^\tau, \neq)$ is a model-complete core and \mathfrak{D} is pp-definable in \mathfrak{C}_Φ^τ after having added finitely many constants, we obtain that \mathfrak{D} is pp-constructible from $(\mathfrak{C}_\Phi^\tau, \neq)$. Hence, it suffices to show that \mathfrak{D} and \mathfrak{C}_Ψ^ρ are homomorphically equivalent. We first show that \mathfrak{D} satisfies Ψ . Consider the expansion of \mathfrak{D} where $M \in \sigma$ denotes $M^{\mathfrak{C}_\Phi}$. This expansion satisfies for distinct $M, M' \in \sigma$ the clause $\forall x. \neg(P_M(x) \wedge M'(x))$ of Ψ as a consequence of Lemma 32. The expansion clearly satisfies all other conjuncts of Ψ . Therefore, \mathfrak{D} satisfies Ψ and we obtain a homomorphism $\mathfrak{D} \rightarrow \mathfrak{C}_\Psi^\rho$. Conversely, Proposition 33 gives that every finite substructure of \mathfrak{C}_Ψ^ρ has a homomorphism to \mathfrak{D} . By the ω -categoricity of \mathfrak{D} , we get a homomorphism from \mathfrak{C}_Ψ^ρ to \mathfrak{D} .

Secondly, we prove that \mathfrak{C}_Φ^τ is pp-constructible in \mathfrak{C}_Ψ^ρ . For this it suffices to note that the structures \mathfrak{C}_Φ^τ and \mathfrak{C}_Ψ^τ are homomorphically equivalent, and that \mathfrak{C}_Ψ^τ is obtained from \mathfrak{C}_Ψ^ρ by dropping the relations from $\rho \setminus \tau$, and is in particular a pp-power of \mathfrak{C}_Ψ^ρ .

By Lemma 34, these pp-constructions give uniformly continuous height 1 homomorphisms $\text{Pol}(\mathfrak{C}_\Psi^\rho) \rightarrow \text{Pol}(\mathfrak{C}_\Phi^\tau)$ and $\text{Pol}(\mathfrak{C}_\Phi^\tau, \neq) \rightarrow \text{Pol}(\mathfrak{C}_\Psi^\rho)$. From the former homomorphism we get that if there is a uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{C}_\Phi^\tau) \rightarrow \mathcal{P}$, there is also one $\text{Pol}(\mathfrak{C}_\Psi^\rho) \rightarrow \mathcal{P}$. The latter homomorphism gives us that if there exists a uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{C}_\Psi^\rho) \rightarrow \mathcal{P}$, there is one $\text{Pol}(\mathfrak{C}_\Phi^\tau, \neq) \rightarrow \mathcal{P}$. We conclude by Lemma 26. \square

5 An Algebraic Dichotomy for MMSNP

We prove in this section that MMSNP exhibits a complexity dichotomy, that is, that every problem in MMSNP is in P or NP-complete. Moreover, we show that the tractability border can be described in terms of homomorphisms to \mathcal{P} , thus confirming the general conjecture of Bodirsky and Pinsker for the class of constraint satisfaction problems in MMSNP.

Theorem 35. *Let \mathfrak{B} be an ω -categorical structure such that $\text{CSP}(\mathfrak{B})$ is in MMSNP. Then exactly one of the following holds:*

- (i) *there is a uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{B}) \rightarrow \mathcal{P}$ and $\text{CSP}(\mathfrak{B})$ is NP-complete, or*
- (ii) *$\text{CSP}(\mathfrak{B})$ is solvable in polynomial time.*

We briefly describe the road to proving Theorem 35. In virtue of Corollary 28, it suffices to focus on the case that $\text{CSP}(\mathfrak{B})$ is described by a precoloured MMSNP sentence. For each precoloured sentence Φ , we consider the structure \mathfrak{C}_Φ^τ whose CSP is described by Φ . We prove that the complexity of $\text{CSP}(\mathfrak{C}_\Phi^\tau)$ and the existence of a homomorphism $\text{Pol}(\mathfrak{C}_\Phi^\tau) \rightarrow \mathcal{P}$ are determined by the existence of a clone homomorphism $\mathcal{C} \rightarrow \mathcal{P}$, where \mathcal{C} is the subset of $\text{Pol}(\mathfrak{C}_\Phi^\tau)$ that contains the functions that are *canonical* with respect to $(\mathfrak{C}_\Phi, <)$ (see Section 5.1 for the definition).

5.1 Canonical functions

Let \mathfrak{B} be a relational structure. A function $f: B^k \rightarrow B$ is said to be *canonical with respect to \mathfrak{B}* if for all $m \geq 1$ and all $\bar{t}_1, \dots, \bar{t}_k \in B^m$ the orbit of $f(\bar{t}_1, \dots, \bar{t}_k) \in B^m$ only depends on the orbits of $\bar{t}_1, \dots, \bar{t}_k$ with respect to the componentwise action of $\text{Aut}(\mathfrak{B})$ on B^m . The following theorem states, informally, that every function behaves like a function that is canonical with respect to \mathfrak{B} on some infinite subset of \mathfrak{B} .

Theorem 36 ([12]). *Let \mathfrak{B} be a countable ω -categorical Ramsey structure. Then for any map $h: B^k \rightarrow B$ there exists a function in*

$$\overline{\{\beta \circ h \circ (\alpha_1, \dots, \alpha_k) \mid \alpha_1, \dots, \alpha_k, \beta \in \text{Aut}(\mathfrak{B})\}}$$

that is canonical with respect to \mathfrak{B} .

We write $\overline{\text{Aut}(\mathfrak{B})f \text{Aut}(\mathfrak{B})}$ for the set in the statement of Theorem 36.

The assumption that \mathfrak{B} is a Ramsey structure is very strong; however, it was proven recently that this property holds for our templates of interest.

Theorem 37 (implied by Theorem 2.1 in [21]). *For all finite sets of finite connected τ -structures \mathcal{F} there exists a linear order $<$ on $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$ such that $(\mathfrak{B}_{\mathcal{F}}^{\text{ind}}, <)$ is Ramsey.*

Moreover, it is known that the Ramsey property transfers to model-complete cores [7], so we obtain the following:

Corollary 38 (implied by Corollary 3.8 in [21]). *For every MMSNP sentence Φ in normal form, there exists a linear order $<$ on \mathfrak{C}_Φ such that $(\mathfrak{C}_\Phi, <)$ is Ramsey.*

Let \mathcal{C} be a clone that consists of canonical functions with respect to structure \mathfrak{B} . Every $f \in \mathcal{C}$ induces a natural operation on orbits of elements under $\text{Aut}(\mathfrak{B})$, due to the fact that it is canonical with respect to \mathfrak{B} . We denote this operation by $\xi_1^{\text{typ}}(f)$ or simply f^{typ} . Moreover, we write $\mathcal{C}_1^{\text{typ}}$ for the clone of functions of the form $\xi_1^{\text{typ}}(f)$, with $f \in \mathcal{C}$.

We finish this section by stating a consequence of assuming that Φ is precoloured and in normal form on the clone $\mathcal{C}_1^{\text{typ}}$. An operation f is called *idempotent* if it satisfies $f(x, \dots, x) = x$, for all x in the domain of f .

Proposition 39. *Let Φ be a precoloured MMSNP sentence in normal form. Let \mathcal{C} be the set of polymorphisms of \mathfrak{C}_Φ^τ that are canonical with respect to $(\mathfrak{C}_\Phi, <)$. Then all functions in $\mathcal{C}_1^{\text{typ}}$ are idempotent.*

5.2 The tractable case

In this section, we prove that $\text{CSP}(\mathfrak{C}_\Phi^\tau)$ is polynomial-time tractable, under the assumption that \mathfrak{C}_Φ^τ has a polymorphism that is canonical with respect to $(\mathfrak{C}_\Phi, <)$ and whose behaviour on orbits of elements is Siggers. For that we use the infinite-to-finite reduction from [9] and the recent solutions to the Feder-Vardi conjecture [14, 26].

Proposition 40. *Let \mathcal{C} be the clone of functions in $\text{Pol}(\mathfrak{C}_\Phi^\tau)$ that are canonical with respect to $(\mathfrak{C}_\Phi, <)$. Suppose that $\mathcal{C}_1^{\text{typ}}$ does not have a homomorphism to \mathcal{P} . Then $\overline{\text{Pol}(\mathfrak{C}_\Phi)}$ contains an operation that is pseudo-Siggers modulo $\text{Aut}(\mathfrak{C}_\Phi, <)$ and canonical with respect to $(\mathfrak{C}_\Phi, <)$.*

Theorem 41 (Corollary 15 in [10]). *Let \mathfrak{A} be a finite-signature reduct of a finitely bounded homogeneous structure \mathfrak{B} . If \mathfrak{A} has a pseudo-Siggers polymorphism modulo $\text{Aut}(\mathfrak{B})$ that is canonical with respect to \mathfrak{B} , then $\text{CSP}(\mathfrak{A})$ is in P.*

In order to use Theorem 41, it remains to prove that \mathfrak{C}_Φ^τ is a reduct of a finitely bounded homogeneous structure.

Proposition 42. *The structure \mathfrak{C}_Φ has a homogeneous expansion by finitely many primitive positive definable relations. Moreover, the expansion is finitely bounded.*

Theorem 43. *If there is no clone homomorphism $\mathcal{C}_1^{\text{typ}} \rightarrow \mathcal{P}$, then $\text{CSP}(\mathfrak{C}_\Phi^\tau)$ is in P.*

5.3 The hard case

Let Φ be a precoloured MMSNP sentence in normal form, and let \mathcal{C} be the clone of polymorphisms of \mathfrak{C}_Φ^τ that are canonical with respect to $(\mathfrak{C}_\Phi, <)$. In this section, we deal with the case that there exists a clone homomorphism $\xi: \mathcal{C}_1^{\text{typ}} \rightarrow \mathcal{P}$, and prove that there exists a uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{C}_\Phi^\tau) \rightarrow \mathcal{P}$.

There is a natural candidate for a height 1 homomorphism $\text{Pol}(\mathfrak{C}_\Phi^\tau) \rightarrow \mathcal{P}$, which we describe now. By Theorem 36, for every $f \in \text{Pol}(\mathfrak{C}_\Phi)$, the set $\overline{\text{Aut}(\mathfrak{C}_\Phi, <)f \text{Aut}(\mathfrak{C}_\Phi, <)}$ has a non-empty intersection with \mathcal{C} . Thus, a natural definition of a uniformly continuous height 1 homomorphism $\phi: \text{Pol}(\mathfrak{C}_\Phi) \rightarrow \mathcal{P}$ is given by

$$\phi(f) := \xi(g^{\text{typ}}) \text{ where } g \in \mathcal{C} \cap \overline{\text{Aut}(\mathfrak{C}_\Phi, <)f \text{Aut}(\mathfrak{C}_\Phi, <)}.$$

This map is well-defined only if for every choice of g, h in $\mathcal{C} \cap \overline{\text{Aut}(\mathfrak{C}_\Phi, <)f \text{Aut}(\mathfrak{C}_\Phi, <)}$ we have $\xi(g^{\text{typ}}) = \xi(h^{\text{typ}})$. We focus on proving that this map (potentially after replacing ξ with another clone homomorphism from $\mathcal{C}_1^{\text{typ}} \rightarrow \mathcal{P}$) is well-defined in the following series of propositions. The proof that the then-defined map is a uniformly continuous height 1 homomorphism is due to Bodirsky and Mottet [9, Theorem 17].

Let ρ be a subset of σ such that ρ is preserved by $\mathcal{C}_1^{\text{typ}}$. Let Θ be an equivalence relation on ρ that is preserved by $\mathcal{C}_1^{\text{typ}}$ and with two equivalence classes $S, T \subseteq \rho$. We call $\{S, T\}$ a *subfactor* of $\mathcal{C}_1^{\text{typ}}$. The clone $\mathcal{C}_1^{\text{typ}}$ naturally induces a clone on the two-element set $\{S, T\}$. If this clone is (isomorphic to) the projection clone \mathcal{P} , then we call $\{S, T\}$ a *trivial subfactor*. The theory of finite idempotent algebras implies that $\mathcal{C}_1^{\text{typ}}$ has a homomorphism to \mathcal{P} if, and only if, $\mathcal{C}_1^{\text{typ}}$ has a trivial subfactor $\{S, T\}$ (see [15, Proposition 4.14], for example).

Let X be a pp-definable subset of \mathfrak{C}_Φ . A binary symmetric relation $N \subseteq X^2$ defines an undirected graph on σ : there is an edge between

M and M' if, and only if, there exist $x \in M^{\mathbb{C}_\Phi}$ and $y \in M'^{\mathbb{C}_\Phi}$ such that $(x, y) \in N$. If N is pp-definable in \mathbb{C}_Φ , we call the resulting graph on σ a *definable colour graph* over X . The following proposition ties the existence of a trivial subfactor to the existence of a particular definable colour graph.

Proposition 44. *Let Φ be a precoloured MMSNP sentence in strong normal form and let \mathcal{C} be the clone of polymorphisms of \mathbb{C}_Φ that are canonical with respect to $(\mathbb{C}_\Phi, <)$. Let $\{S, T\}$ be a trivial subfactor of $\mathcal{C}_1^{\text{typ}}$. Then there exists a pp-definable subset X of \mathbb{C}_Φ and a pp-definable binary symmetric relation $N \subseteq X^2$ that defines a colour graph with an edge from S to T and whose every connected component is either included in S , included in T , or is a bipartite graph whose bipartition is induced by S and T .*

Theorem 45. *Let Φ be a precoloured MMSNP sentence in strong normal form. Let \mathcal{C} be the clone of polymorphisms of \mathbb{C}_Φ that are canonical with respect to $(\mathbb{C}_\Phi, <)$. If there is a clone homomorphism $\mathcal{C}_1^{\text{typ}} \rightarrow \mathcal{P}$, then there exists a uniformly continuous height 1 homomorphism from $\text{Pol}(\mathbb{C}_\Phi)$ to \mathcal{P} that is invariant under left-composition by $\text{Aut}(\mathbb{C}_\Phi)$, and $\text{CSP}(\mathbb{C}_\Phi^{\mathcal{P}})$ is NP-hard.*

Proof. As we have mentioned before, if the finite idempotent algebra $\mathcal{C}_1^{\text{typ}}$ has a homomorphism to \mathcal{P} , then $\mathcal{C}_1^{\text{typ}}$ has a trivial subfactor $\{S, T\}$ (see [15, Proposition 4.14]).

Let $\xi: \mathcal{C}_1^{\text{typ}} \rightarrow \mathcal{P}$ be the clone homomorphism defined as follows. Let $R \in S$ and $B \in T$ be arbitrary. For a k -ary $f \in \mathcal{C}_1^{\text{typ}}$, let $i \in \{1, \dots, k\}$ be the unique index such that $f(B, \dots, B, R, B, \dots, B) \in S$, where the argument R is in the i th position. Such an i exists because of the assumption that $\{S, T\}$ is a trivial subfactor of $\mathcal{C}_1^{\text{typ}}$. Define $\xi(f)$ to be the i th projection. Note that the definition of ξ does not depend on the choice of R and B , by the fact that the equivalence relation on $S \cup T$ whose equivalence classes are S and T is assumed to be preserved by the operations in $\mathcal{C}_1^{\text{typ}}$. It is straightforward to check that the map ξ thus defined is a clone homomorphism.

Let $X \subseteq \mathbb{C}_\Phi$ and $N \subseteq X^2$ be the pp-definable relations given by Proposition 44. Fix $f \in \text{Pol}(\mathbb{C}_\Phi)$ a k -ary operation and g, h two operations in $\mathcal{C} \cap \text{Aut}(\mathbb{C}_\Phi, <)f \text{Aut}(\mathbb{C}_\Phi, <)$.

As explained in the beginning of this section, it suffices to prove that $\xi(g^{\text{typ}}) = \xi(h^{\text{typ}})$. For ease of notation, assume that $\xi(g^{\text{typ}})$ is the first projection, the general case being treated in the same way. Since ξ is the clone homomorphism induced by $\{S, T\}$, this means that for all $R \in S$ and $B \in T$, we have $g^{\text{typ}}(R, B, \dots, B) \in S$. In order to prove that $\xi(h^{\text{typ}})$ is also the first projection, it suffices to prove that there exist $R \in S$ and $B \in T$ such that $h^{\text{typ}}(R, B, \dots, B) \in S$. Let $R \in S$ and $B \in T$ be adjacent colours in the colour graph defined by N . Let (a_1, \dots, a_k) be any tuple in $R^{\mathbb{C}_\Phi} \times B^{\mathbb{C}_\Phi} \times \dots \times B^{\mathbb{C}_\Phi}$. Since f interpolates g and h modulo $\text{Aut}(\mathbb{C}_\Phi, <)$, there are automorphisms $\alpha, \beta_1, \dots, \beta_k$ such that

$$g(a_1, \dots, a_k) = \alpha f(\beta_1 a_1, \dots, \beta_k a_k)$$

and automorphisms $\gamma, \delta_1, \dots, \delta_k$ such that

$$h(a_1, \dots, a_k) = \gamma f(\delta_1 a_1, \dots, \delta_k a_k).$$

Let \mathfrak{S} be the substructure of \mathbb{C}_Φ induced by the elements of the form $\beta_i a_i$ and $\delta_i a_i$, for $i \in \{1, \dots, k\}$. Since (\mathbb{C}_Φ, \neq) is a model-complete core (Lemma 19), it follows from Theorem 3.6.11 in [6] that the orbit of the tuple

$$(\beta_1 a_1, \dots, \beta_k a_k, \delta_1 a_1, \dots, \delta_k a_k)$$

in \mathbb{C}_Φ is defined by a formula $\theta(x_1, \dots, x_k, y_1, \dots, y_k)$ that is primitive positive in the language of (\mathbb{C}_Φ, \neq) . Let θ^* be θ where the atomic formulas \neq have been removed. Let $\phi_N(x, y)$ be a primitive positive formula defining the relation $N \subseteq (\mathbb{C}_\Phi)^2$ in \mathbb{C}_Φ . Fix an integer ℓ such that $2\ell > |\Phi|$. For every $i \in \{1, \dots, k\}$, let $z_1^i, \dots, z_{2\ell-1}^i$ be fresh variables. In the following, we also write z_0^i for x_i and $z_{2\ell}^i$ for y_i . Let $\psi(x_1, \dots, x_k, y_1, \dots, y_k)$ be the primitive positive formula whose conjuncts are:

- $\theta^*(x_1, \dots, x_k, y_1, \dots, y_k)$,
- $\phi_N(z_j^i, z_{j+1}^i)$, for every $i \in \{1, \dots, k\}$ and $j \in \{0, \dots, 2\ell - 1\}$,
- $R(z_j^i)$ for even $j \in \{1, \dots, 2\ell - 1\}$ and $B(z_j^i)$ for odd $j \in \{1, \dots, 2\ell - 1\}$,
- for $i \in \{2, \dots, k\}$, the conjunct $B(z_j^i)$ for even $j \in \{1, \dots, 2\ell - 1\}$ and $R(z_j^i)$ for odd $j \in \{1, \dots, 2\ell - 1\}$.

We claim that ψ is satisfiable in \mathbb{C}_Φ . We first prove that it is satisfiable in $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$, where \mathcal{F} is the coloured obstruction set of Φ . Let \mathfrak{S}' be the canonical database of ψ (see Figure 3). By Lemma 11, ψ is satisfiable if and only if all the biconnected components of \mathfrak{S}' are \mathcal{F} -free. Suppose that there exists an obstruction $\mathfrak{F} \in \mathcal{F}$ of Φ and a homomorphism $e: \mathfrak{F} \rightarrow \mathfrak{S}'$ to one of the biconnected components of \mathfrak{S}' . By the choice of ℓ we have that $|\mathfrak{F}| < 2\ell$. It follows that either the image of e is included in \mathfrak{S} , or it is included in the subset induced by the canonical database of some $N(z_j^i, z_{j+1}^i)$ for some $i \in \{1, \dots, k\}$ and $j \in \{0, \dots, 2\ell - 1\}$. But the assumption on N is that there is $(a, b) \in N$ such that $a \in R^{\mathbb{C}_\Phi}$ and $b \in B^{\mathbb{C}_\Phi}$. Therefore, the conjunct $\phi_N(z_j^i, z_{j+1}^i)$ is satisfiable by an assignment that maps z_j^i and z_{j+1}^i to the appropriate colours. We conclude that there exists an embedding e of \mathfrak{S}' into $\mathfrak{B}_{\mathcal{F}}^{\text{ind}}$.

Let $d: \mathfrak{B}_{\mathcal{F}}^{\text{ind}} \rightarrow \mathfrak{B}_{\mathcal{F}}^{\text{hom}}$ be an injective homomorphism (whose existence follows from Theorem 13). Note that the image of the restriction of d to $\mathbb{C}_\Phi \subseteq \mathfrak{B}_{\mathcal{F}}^{\text{ind}}$ is in \mathbb{C}_Φ since d must preserve the colours. Since $d \circ e$ is injective, the tuple $(d \circ e)(x_1, \dots, x_k, y_1, \dots, y_k)$ satisfies θ . This means that $d \circ e: \mathfrak{S}' \rightarrow \mathbb{C}_\Phi$ is a satisfying assignment that maps $(x_1, \dots, x_k, y_1, \dots, y_k)$ to a tuple that is in the same orbit as $(\beta_1 a_1, \dots, \beta_k a_k, \delta_1 a_1, \dots, \delta_k a_k)$. By composing with an automorphism of \mathbb{C}_Φ , we can suppose that the image of $(x_1, \dots, x_k, y_1, \dots, y_k)$ is exactly this tuple.

It must therefore be the case that the elements $f(\beta_1 a_1, \dots, \beta_k a_k)$ and $f(\delta_1 a_1, \dots, \delta_k a_k)$ of \mathbb{C}_Φ are connected by an N -path of even length, that is, there are $b_1, \dots, b_{2\ell-1} \in \mathbb{C}_\Phi$ such that $(b_j, b_{j+1}) \in N$ for all $j \in \{1, \dots, 2\ell\}$ and $(f(\beta_1 a_1, \dots, \beta_k a_k), b_1) \in N$ and $(b_{2\ell-1}, f(\delta_1 a_1, \dots, \delta_k a_k)) \in N$. This means that $f(\beta_1 a_1, \dots, \beta_k a_k)$ and $f(\delta_1 a_1, \dots, \delta_k a_k)$ are in the same component in the colour graph defined by N . If this connected component is included in S , then there is $Y \in S$ such that $f(\delta_1 a_1, \dots, \delta_k a_k) \in Y$, i.e., $h(a_1, \dots, a_k) \in Y$. Otherwise, the connected component of these elements is bipartite, and since there is a path of even length between the two elements, it must be the case that there is $Y \in S$ such that $f(\delta_1 a_1, \dots, \delta_k a_k)$ belongs to Y . In both cases, we obtain that $h(R, B, \dots, B) \in S$. \square

5.4 The dichotomy: conclusion

Summing up the results of the previous two sections, we obtain the following dichotomy for precoloured MMSNP sentences.

Theorem 46. *Let Φ be a precoloured MMSNP sentence in normal form. Then exactly one of the following statements hold:*




Figure 3. Depiction of the structure \mathfrak{C}' , in the case that $k = 2$.

1. there is a uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{C}'_{\Phi}) \rightarrow \mathcal{P}$ that is invariant under left-composition by $\overline{\text{Aut}(\mathfrak{C}_{\Phi})}$, and $\text{CSP}(\mathfrak{C}'_{\Phi})$ is NP-complete,
2. $\text{Pol}(\mathfrak{C}'_{\Phi})$ contains a pseudo-Siggers operation modulo $\overline{\text{Aut}(\mathfrak{C}_{\Phi})}$, and $\text{CSP}(\mathfrak{C}'_{\Phi})$ is in P.

Proof. It is clear that at most one of the two statements hold. Let \mathcal{C} be the clone of polymorphisms of \mathfrak{C}'_{Φ} that are canonical with respect to $(\mathfrak{C}_{\Phi}, <)$. If the first item does not hold then by Theorem 45 there is no clone homomorphism $\mathcal{C}_{1}^{\text{typ}} \rightarrow \mathcal{P}$. By Theorem 43, the CSP of \mathfrak{C}'_{Φ} is in P and by Proposition 40, \mathfrak{C}'_{Φ} has a pseudo-Siggers polymorphism modulo $\overline{\text{Aut}(\mathfrak{C}_{\Phi}, <)}$, and in particular this polymorphism is pseudo-Siggers modulo $\overline{\text{Aut}(\mathfrak{C}_{\Phi})}$. \square

We can finally prove Theorem 7 from Section 2.3.

Proof. Let Φ be a connected MMSNP sentence with input signature τ . By Theorem 21, we can assume that Φ is in strong normal form. Let $\mathfrak{B} := \mathfrak{C}'_{\Phi}$. Let Ψ be the standard precolouration of Φ with input signature ρ . Assume that there is no uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{B}) \rightarrow \mathcal{P}$. By Theorem 27, there is no uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{C}'_{\Psi}) \rightarrow \mathcal{P}$. By Theorem 46, this means that \mathfrak{C}'_{Ψ} has a pseudo-Siggers polymorphism and $\text{CSP}(\mathfrak{C}'_{\Psi})$ is in P. Since \mathfrak{B} is isomorphic to the τ -reduct of \mathfrak{C}'_{Ψ} and the two structures have the same automorphism group, this means that \mathfrak{B} has a pseudo-Siggers polymorphism. Moreover, again by Theorem 27, we have that $\text{CSP}(\mathfrak{B})$ is in P.

It remains to prove that the two cases are mutually exclusive. Assume that there is a uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{B}) \rightarrow \mathcal{P}$. Then by Lemma 26, there is a uniformly continuous height 1 homomorphism $\text{Pol}(\mathfrak{B}, \neq) \rightarrow \mathcal{P}$. Since (\mathfrak{B}, \neq) is a model-complete core, Corollary I.8 in [1] implies the existence of finitely many constants c_1, \dots, c_n and a clone homomorphism $\text{Pol}(\mathfrak{B}, \neq, c_1, \dots, c_n) \rightarrow \mathcal{P}$. In particular, $\text{Pol}(\mathfrak{B}, \neq, c_1, \dots, c_n)$ cannot contain a pseudo-Siggers operation. A standard argument using the fact that (\mathfrak{B}, \neq) is a model-complete core shows that $\text{Pol}(\mathfrak{B}, \neq)$ cannot contain a pseudo-Siggers operation. Finally, showing that $\text{Pol}(\mathfrak{B})$ cannot contain a pseudo-Siggers operation can be done similarly as in Lemma 26.

Finally, we show that Conjecture 6 holds for all CSPs in MMSNP. Let \mathfrak{C} be an ω -categorical structure such that $\text{CSP}(\mathfrak{C})$ is described by an MMSNP connected sentence Φ , and suppose that there is no uniformly continuous clonoid homomorphism from $\text{Pol}(\mathfrak{C})$ to \mathcal{P} . Let \mathfrak{B} be the structure in the statement of Theorem 7. Since $\text{CSP}(\mathfrak{C})$ and $\text{CSP}(\mathfrak{B})$ are equal and both \mathfrak{C} and \mathfrak{B} are ω -categorical,

we obtain that \mathfrak{C} and \mathfrak{B} are homomorphically equivalent. It follows that there is a uniformly continuous clonoid homomorphism from $\text{Pol}(\mathfrak{C})$ to $\text{Pol}(\mathfrak{B})$. Hence, there cannot be any uniformly continuous clonoid homomorphism from $\text{Pol}(\mathfrak{B})$ to \mathcal{P} , so that by Theorem 7, $\text{CSP}(\mathfrak{B})$ is tractable in polynomial time. Thus, $\text{CSP}(\mathfrak{C})$ is in P. \square

References

- [1] Libor Barto, Michael Kompatscher, Miroslav Olšák, Michael Pinsker, and Trung Van Pham. 2017. The Two Dichotomy Conjectures for Infinite Domain Constraint Satisfaction Problems Are Equivalent. In *Proceedings of the 32nd Annual IEEE Symposium on Logic in Computer Science – LICS'17*.
- [2] Libor Barto and Marcin Kozik. 2012. Absorbing subalgebras, cyclic terms and the constraint satisfaction problem. *Logical Methods in Computer Science* 8/1, 07 (2012), 1–26.
- [3] Libor Barto, Jakub Opršal, and Michael Pinsker. 2017. The wonderland of reflections. *Israel Journal of Mathematics* (2017). To appear. Preprint arXiv:1510.04521.
- [4] Libor Barto and Michael Pinsker. 2016. 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*. 615–622.
- [5] Manuel Bodirsky. 2007. Cores of countably categorical structures. *Logical Methods in Computer Science* 3, 1 (2007), 1–16.
- [6] Manuel Bodirsky. 2012. Complexity Classification in Infinite-Domain Constraint Satisfaction. (2012). Mémoire d'habilitation à diriger des recherches, Université Diderot – Paris 7. Available at arXiv:1201.0856v8.
- [7] Manuel Bodirsky. 2015. Ramsey Classes: Examples and Constructions. In *Surveys in Combinatorics. London Mathematical Society Lecture Note Series 424*. Cambridge University Press. Invited survey article for the British Combinatorial Conference.
- [8] Manuel Bodirsky and Victor Dalmau. 2013. Datalog and Constraint Satisfaction with Infinite Templates. *Journal on Computer and System Sciences* 79 (2013), 79–100. A preliminary version appeared in the proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS'05).
- [9] Manuel Bodirsky and Antoine Mottet. 2016. Reducts of finitely bounded homogeneous structures, and lifting tractability from finite-domain constraint satisfaction. In *Proceedings of the 31th Annual IEEE Symposium on Logic in Computer Science – LICS'16*. 623–632. Preprint available at ArXiv:1601.04520.
- [10] Manuel Bodirsky and Antoine Mottet. 2018. A Dichotomy for First-Order Reducts of Unary Structures. *Logical Methods of Computer Science* (2018). To appear.
- [11] Manuel Bodirsky and Jaroslav Nešetřil. 2006. Constraint Satisfaction with Countable Homogeneous Templates. *Journal of Logic and Computation* 16, 3 (2006), 359–373.
- [12] Manuel Bodirsky and Michael Pinsker. 2016. Canonical Functions: a Proof via Topological Dynamics. (2016). Preprint available under <http://arxiv.org/abs/1610.09660>.
- [13] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. 2014. Projective clone homomorphisms. (2014). Preprint arXiv:1409.4601.
- [14] Andrei A. Bulatov. 2017. A dichotomy theorem for nonuniform CSPs. In *Proceedings of FOCS 2017*.
- [15] Andrei A. Bulatov and Peter Jeavons. 2001. Algebraic Structures in Combinatorial Problems. Technical report MATH-AL-4-2001, Technische Universität Dresden. (2001).
- [16] Andrei A. Bulatov, Andrei A. Krokhin, and Peter G. Jeavons. 2005. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.* 34 (2005), 720–742.
- [17] Ashok K. Chandra and Philip M. Merlin. 1977. Optimal Implementation of Conjunctive Queries in Relational Data Bases. In *Proceedings of the Symposium on Theory of Computing (STOC)*. 77–90.
- [18] Gregory Cherlin, Saharon Shelah, and Niangdong Shi. 1999. Universal graphs with forbidden subgraphs and algebraic closure. *Advances in Applied Mathematics* 22 (1999), 454–491.
- [19] Tomás Feder and Moshe Y. Vardi. 1993. Monotone monadic SNP and constraint satisfaction. In *Proceedings of the Symposium on Theory of Computing (STOC)*. 612–622.
- [20] Tomás Feder and Moshe Y. Vardi. 1999. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM J. Comput.* 28 (1999), 57–104.
- [21] J. Hubička and J. Nešetřil. 2016. All those Ramsey classes (Ramsey classes with closures and forbidden homomorphisms). (2016). arXiv:1606.07979v2.
- [22] Gábor Kun. 2013. Constraints, MMSNP, and Expander Relational Structures. *Combinatorica* 33, 3 (2013), 335–347.
- [23] Carsten Lutz and Frank Wolter. 2015. On the Relationship between Consistent Query Answering and Constraint Satisfaction Problems. In *Proceedings of the 18th International Conference on Database Theory (ICDT15)*.
- [24] Florent Madelaine. 2003. Constraint satisfaction problems and related logic. PhD-thesis, University of Leicester. (2003).
- [25] Florent Madelaine and Iain A. Stewart. 2007. Constraint Satisfaction, Logic and Forbidden Patterns. *SIAM J. Comput.* 37, 1 (2007), 132–163.
- [26] Dmitry Zhuk. 2017. The Proof of CSP Dichotomy Conjecture. In *Proceedings of FOCS 2017*.