On the Complexity of MMSNP

Manuel Bodirsky^{*} Hubie Chen[†] Tomás Feder[‡]

Abstract

Monotone monadic strict NP (MMSNP) is a class of computational problems that is closely related to the class of constraint satisfaction problems for constraint languages over finite domains. It is known that one of those classes has a complexity dichotomy if and only if the other class has. Whereas the dichotomy conjecture has been verified for several subclasses of constraint satisfaction problems, little is known about the the computational complexity for subclasses of MMSNP.

In this paper we completely classify the complexity of MMSNP for the case where the obstructions are monochromatic and where loops in the input are forbidden. That is, we determine the computational complexity of natural partition problems of the following type. For fixed sets of finite structures S_1, \ldots, S_k , decide whether a given loopless structure can be vertex-partitioned into k parts such that for each $i \leq k$ none of the structures in S_i is homomorphic to the *i*-th part.

1 Introduction

An important topic in theoretical computer science is to classify computational problems with respect to their computational complexity, and to understand the border between problems that can be solved in polynomial time, and problems that are NP-hard. Occasionally, entire classes of computational problems have been studied concerning their computational complexity [13, 8, 6]. More generally, one might ask: for which classes of computational problems can we expect complete complexity classifications, and for which classes is this hopeless?

Feder and Vardi [7] showed that certain natural classes of computational problems will most likely never be completely classified with respect to their computational complexity. In fact, they showed that various very restricted classes of existential second-order logic (namely the three classes *monotone strict NP*, *monotone monadic strict NP with disequalities*, and *monadic strict NP*) are *rich* in the sense that every computational problem in NP has a polynomial-time equivalent problem in these classes.

^{*}École Polytechnique, Greater Paris, France. Email: bodirsky@lix.polytechnique.fr. Manuel Bodirsky has received funding from the *European Research Council* under the European Community's Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 257039).

[†]Universitat Pompeu Fabra, Barcelona, Spain. Email: hubie.chen@upf.edu. Hubie Chen is supported by the Spanish program "Ramon y Cajal" and MICINN grant TIN2010-20967-C04-02.

[‡]268 Waverley St., Palo Alto, CA 94301. E-mail: tomas@theory.stanford.edu.

For the class of monotone monadic strict NP (MMSNP) Feder and Vardi did not give such a result. They rather conjectured that the problems in MMSNP exhibit a *complexity dichotomy* in the sense that they are either in P or NP-complete. The conjecture is based on their result that every problem in MMSNP is polynomial-time equivalent [7, 9] to a constraint satisfaction problem (CSP) with a fixed finite constraint language over a finite domain. The dichotomy conjecture for MMSNP is thus implied by the dichotomy conjecture for CSPs.

Even though some progress has been made recently with the so-called universal-algebraic approach concerning CSP complexity, the dichotomy conjecture still appears to be wide open. The so-called algebraic *tractability conjecture* in constraint satisfaction says that every template that gives rise to an idempotent algebra that omits type 1 has a CSP in P [5]. Omitting type 1 is a very weak assumption in tame congruence theory; and so at present a proof that the corresponding CSPs are all tractable appears to be at far distance.

The reduction of Feder and Vardi that shows that every problem in MMSNP is computationally equivalent to a CSP usually generates templates over very large domains: the domain size is 2^m , where m is the number of unary predicates plus the number of articulation points in the obstructions (for formal definitions, see Section 2). And even if the mentioned tractability conjecture is true, verifying that an algebra which is given by the corresponding template omits type 1 is NP-hard. In fact, it is not difficult to come up with concrete examples of MMSNP sentences where it appears to be difficult to use the translation to CSPs to obtain complexity results. In this paper, we directly study the complexity of the computational problems in MMSNP, without the detour to CSPs. In contrast to the constraint satisfaction problems in the 'black region' where the universal algebraic approach has not yet provided answers, the problems in MMSNP that we study have a very concrete combinatorial flavour, and are very often natural graph-theoretic problems.

We completely classify the computational complexity of MMSNP for the special case where all obstructions are monochromatic, and where loops in the input are forbidden. Our result here is similar to a classification result by Achlioptas [1], who classified the complexity of *G*-free colorability, which is the problem to decide for a given graph H whether there is a vertex partition of H into two parts such that each part does not contain G as a (weak) subgraph. Our result is different from *G*-free coloring problems in the important aspect that we are studying coloring problems where a fixed set of structures is forbidden *homomorphically*, and not as subgraphs. It should be pointed out that our result

- applies not only to graphs, but to general relational structures;
- determines the complexity not only for partitions with two parts, but more generally for partitions into any finite number of parts; and
- explains the complexity of the problem not only for a single forbidden obstruction, but for any finite number of forbidden homomorphic obstructions.

We would also like to remark that every problem in MMSNP is a finite union of *connected* MMSNP sentences (see Section 3), and that connected MMSNP sentences describe constraint satisfaction problems with infinite ω -categorical templates [2]; hence, our result for monochromatic obstructions is also interesting in the context of obtaining systematic complexity results for CSPs with ω -categorical templates (as e.g. in [3, 4]).

2 MMSNP

A relational signature τ is a (in this paper always finite) set of relation symbols R_i , each of which has an associated finite arity k_i . A relational structure **A** over the signature τ (also called τ -structure) consists of a set A (the domain) together with a relation $R^{\mathbf{A}} \subseteq A^k$ for each relation symbol R of arity k from τ . We use bold-font letters **A**, **B**, **C** to denote relational structures with domains A, B, C, respectively. A relational structure **A** is loopless if no relation R in **A** contains a tuple where all entries are equal.

Let **A** be a relational τ -structure, and let **B** be a relational σ -structure with $\tau \subseteq \sigma$. If **A** and **B** have the same domain and $R^{\mathbf{A}} = R^{\mathbf{B}}$ for all $R \in \tau$, the **A** is called the τ -reduct (or simply reduct) of **B**, and **B** is called a σ -expansion (or simply expansion) of **A**.

An SNP τ -sentence (SNP stands for strict NP [12, 7]) is an existential second-order sentence Φ of the form

$$\exists R_1, \ldots, R_k \, \forall x_1, \ldots, x_n. \phi$$

where ϕ is a quantifier-free first-order formula over the signature $\sigma = \tau \cup \{R_1, \ldots, R_k\}$, and $R_1, \ldots, R_k \notin \tau$. The signature τ is also called the *input signature* (of Φ). We refer to $\forall x_1, \ldots, x_n$. ϕ as the *first-order part of* Φ . Clearly, ϕ can always be rewritten in CNF, i.e., as a conjunction of disjunctions of literals, and we will always assume that ϕ is given in this form. If each literal in ϕ with a relation symbol from τ is negative, we say that the sentence is a *monotone SNP* sentence. If furthermore all existentially quantified relation symbols R_1, \ldots, R_k are *monadic* (i.e., unary), we say that the sentence is a *monotone monadic SNP sentence* or, shorter, *MMSNP sentence*.

An SNP sentence Φ describes a computational problem \mathcal{P} (viewed as a class of finite τ -structures) if a finite τ -structure **A** belongs to \mathcal{P} if and only if **A** satisfies Φ . We also denote by SNP (monotone SNP, MMSNP) the class of all computational problems on τ -structures that can be described by a (monotone, monotone monadic) SNP sentence.

Example 1 The problem No-Mono-Directed-Tri is the problem to decide whether a given finite graph can be partitioned into two graphs that do not contain directed triangles. This problem can be expressed by the monotone (and even monadic) SNP sentence

$$\exists M \,\forall x, y, z. \neg \big(M(x) \land M(y) \land M(z) \land E(x, y) \land E(y, z) \land E(z, x) \big) \\ \land \neg \big(\neg M(x) \land \neg M(y) \land \neg M(z) \land E(x, y) \land E(y, z) \land E(z, x) \big)$$

There is a close connection between MMSNP and so-called forbidden patterns problems [11] and lifts and shadows [10], which already becomes apparent in [7]. To describe this connection, we need the notion of homomorphisms. A homomorphism between two structures **A** and **B** of the same relational signature is a mapping $f : A \to B$ such that $(f(t_1), \ldots, f(t_k)) \in \mathbb{R}^{\mathbf{B}}$ whenever $(t_1, \ldots, t_k) \in \mathbb{R}^{\mathbf{A}}$. If **A** homomorphically maps to **B**, we write $\mathbf{A} \to \mathbf{B}$, and $\mathbf{A} \not\to \mathbf{B}$ otherwise.

Definition 1 Let \mathcal{A} be a class of τ -structures. A colored obstruction set for \mathcal{A} is a finite set \mathcal{F} of $(\tau \cup \rho)$ -structures, where ρ is a finite set of unary relation symbols such that $\mathbf{A} \in \mathcal{A}$ if and only if \mathbf{A} has a $(\tau \cup \rho)$ -expansion \mathbf{A}' where every element is contained in exactly one unary relation from ρ , and no structure $\mathbf{F} \in \mathcal{F}$ homomorphically maps to \mathbf{A}' .

If \mathcal{F} is a colored obstruction set for \mathcal{A} , we also write $\operatorname{Forb}(\mathcal{F})$ for \mathcal{A} (note that \mathcal{A} is uniquely described by \mathcal{F}). We can think of the unary relations in ρ as colors, and the expansion \mathbf{A}' of \mathbf{A} in Definition 1 can then be seen as a coloring of the vertices of \mathbf{A} : a vertex $v \in \mathcal{A}$ receives color C iff $v \in C^{\mathbf{A}'}$. The conditions for the expansion in Definition 1 make sure that every vertex receives exactly one color. The signature ρ is also called the *color signature* of \mathcal{F} , and $(\tau \cup \rho)$ -expansions of \mathbf{A} where every element is contained in exactly one unary relation from ρ are also called $(\rho$ -) colorings of \mathbf{A} . If Φ is an MMSNP sentence, and \mathcal{F} is a finite colored obstruction set for the set of τ -structures satisfying Φ , we say that \mathcal{F} is a *colored obstruction set of* Φ .

Example 2 Recall the example No-Mono-Directed-Tri, Example 1. The input signature τ is $\{E\}$ where E is a binary relation. To give a colored obstruction set for this problem, let $\rho = \{C_0, C_1\}$. Let \mathbf{T} be the τ -structure on domain $\{1, 2, 3\}$ where $E^{\mathbf{T}} = \{(1, 2), (2, 3), (3, 1)\}$. Let \mathbf{F}_1 be the $(\tau \cup \rho)$ -expansion of \mathbf{T} where $C_0 = \{1, 2, 3\}$, and $C_1 = \emptyset$, and let \mathbf{F}_2 be the $(\tau \cup \rho)$ -expansion of \mathbf{T} where $C_0 = \emptyset$, and $C_1 = \{1, 2, 3\}$. Then it is easy to see that $\mathcal{F} = \{\mathbf{F}_1, \mathbf{F}_2\}$ is a colored obstruction set for the No-Mono-Directed-Tri problem.

For long proofs of the following theorem of [7] in slightly different terminology, see [11] and [10].

Theorem 1 (of [7]) A set of τ -structures \mathcal{A} is in MMSNP if and only if \mathcal{A} has a colored obstruction set.

Let \mathbf{F} be the colored obstruction set of an MMSNP sentence Φ . A structure $\mathbf{F} \in \mathcal{F}$ is called *C*-chromatic if there is a $C \in \rho$ such that $C^{\mathbf{F}} = F$. If every structure $\mathbf{F} \in \mathcal{F}$ is *C*-chromatic for some $C \in \rho$, we say that \mathcal{F} is a monochromatic (colored) obstruction set of Φ . Our main result is the following.

Theorem 2 Let Φ be an MMSNP sentence that has a monochromatic obstruction set. Then Φ is in P or NP-complete.

3 Reduction to the Connected Case

The union $\mathbf{A} + \mathbf{B}$ of two τ -structures \mathbf{A} and \mathbf{B} is the structure \mathbf{C} with domain $A \cup B$ and relations $R^{\mathbf{C}} = R^{\mathbf{A}} \cup R^{\mathbf{B}}$ for all $R \in \tau$. The *disjoint union* of \mathbf{A} and \mathbf{B} is the union of isomorphic copies of \mathbf{A} and \mathbf{B} with disjoint domains. A structure is called *connected* if it is not the disjoint union of two non-empty structures, and *disconnected* otherwise. A colored obstruction set \mathbf{F} , and the MMSNP problem with colored obstruction set \mathbf{F} , are called *connected* if all structures in \mathbf{F} are connected.

It has been shown in [11] that every problem in MMSNP is in fact a *finite union* of connected MMSNP problems (defined below); we recall their proof below to observe that it preserves monochromaticity of the obstruction set.

Proposition 1 Let Φ be an MMSNP sentence. Then Φ is logically equivalent to $\Phi_1 \lor \cdots \lor \Phi_n$, where Φ_1, \ldots, Φ_n are MMSNP sentences with connected colored obstruction sets. When Φ has a monochromatic obstruction set, then Φ_1, \ldots, Φ_n can be chosen to be monochromatic. PROOF: Let \mathcal{F} be the obstruction set of Φ , and suppose that \mathcal{F} contains a disconnected structure $\mathbf{F} = \mathbf{F}_1 + \cdots + \mathbf{F}_k$. For $i \leq k$, let \mathcal{F}_i be $\mathcal{F} \setminus {\mathbf{F}_i} \cup {\mathbf{F}_i}$.

We claim that for every $(\tau \cup \sigma)$ -expansion \mathbf{A}' of \mathbf{A} there is a homomorphism from a structure $\mathbf{F}' \in \mathcal{F}$ to \mathbf{A}' if and only if for every $i \leq k$, there is a homomorphism from a structure in \mathcal{F}_i to \mathbf{A}' . If there is a structure $\mathbf{F}' \in \mathcal{F} \setminus {\mathbf{F}}$ that homomorphically maps to \mathbf{A}' , there is nothing to show since then $\mathbf{F}' \in \mathcal{F}_i$ for all $i \leq k$. Otherwise, if only $\mathbf{F}' = \mathbf{F}$ homomorphically maps to \mathbf{A}' , then the statement follows from the observation that \mathbf{F} homomorphically maps to \mathbf{A}' if and only if all its components $\mathbf{F}_1, \ldots, \mathbf{F}_k$ map to \mathbf{A}' .

By Theorem 1, for each $i \leq k$ there is an MMSNP sentence Φ_i having the colored obstruction set \mathcal{F}_i , and by the observation above a finite structure **A** satisfies Φ if and only if it satisfies $\Phi_1 \vee \cdots \vee \Phi_k$. Iterating this process for each Φ_i where \mathcal{F}_i is disconnected, we eventually arrive at a disjunction of *connected* MMSNP sentences (since Φ_i contains fewer disconnected structures than Φ). Finally, observe that when \mathcal{F} is monochromatic, then each of the \mathcal{F}_i constructed in the proof above is monochromatic as well.

We now show that we can reduce the complexity classification for MMSNP with monochromatic obstruction sets to the classification for *connected* MMSNP with monochromatic obstruction sets.

Proposition 2 Every MMSNP with monochromatic obstruction set is in P or NP-complete if and only if every connected MMSNP sentence with monochromatic obstruction set is in P or NP-complete.

PROOF: The forward direction of the statement holds trivially. For the backwards direction, assume that every connected MMSNP sentence with monochromatic obstruction set is either in P or NP-complete. Let Φ be an MMSNP sentence with monochromatic obstruction set, and let Φ_1, \ldots, Φ_k be the connected MMSNP sentences with monochromatic obstruction sets so that Φ is logically equivalent to the disjunction $\Phi_1 \vee \cdots \vee \Phi_k$ such that k is smallest possible. This is always possible by Proposition 1.

If each Φ_i can be decided in polynomial time by an algorithm A_i , then it is clear that Φ can be solved in polynomial time by running all of the algorithms A_1, \ldots, A_k on the input, and accepting if one of the algorithms accepts.

Otherwise, if one of the Φ_i describes an NP-complete problem, then Φ_i can be reduced to Φ as follows. Since k was chosen to be minimal, there exists a τ -structure **B** such that **B** satisfies Φ_i , but does not satisfy Φ_j for all $j \leq n$ that are distinct from i, since otherwise we could have removed Φ_i from the disjunction $\Phi_1 \vee \cdots \vee \Phi_k$ without affecting the equivalence of the disjunction to Φ .

To reduce Φ_i to Φ , we execute for a given finite τ -structure **A** the algorithm for Φ on **A** + **B**. We claim that **A** + **B** satisfies Φ if and only if **A** satisfies Φ_i . First suppose that **A** satisfies Φ_i . Since **B** also satisfies Φ_i by choice of **B**, and since Φ_i is closed under disjoint unions, we have that **A** + **B** satisfies Φ_i as well. The statement follows since Φ_i is a disjunct of Φ .

For the opposite direction, suppose that $\mathbf{A} + \mathbf{B}$ satisfies Φ . Since \mathbf{B} does not satisfy Φ_j for all j distinct from i, $\mathbf{A} + \mathbf{B}$ does not satisfy Φ_j as well, by monotonicity of Φ_j . Hence, $\mathbf{A} + \mathbf{B}$ must satisfy Φ_i . By monotonicity of Φ_i , it follows that \mathbf{A} satisfies Φ_i .

4 Reductions to CSPs

Our main technique to show that a problem in MMSNP is in P is the following translation into finite domain constraint satisfaction problems. The *constraint satisfaction problem* (CSP) for a structure Γ with a finite relational signature τ is the computational problem to decide for a given finite τ -structure **A** whether **A** homomorphically maps to Γ . We also write $CSP(\Gamma)$ for this problem; if Γ has a finite domain, we also say that the problem is a *finite domain CSP*. To reduce problems in MMSNP to finite domain CSPs, we need the following definitions.

Let ρ be a finite set of unary relation symbols; for the sake of notation we assume that ρ is of the form $\{C_i \mid i \in I\}$ for a finite index set I. For a τ -structure \mathbf{A} with domain $A = \{1, \ldots, n\}$, and elements $i_1, \ldots, i_n \in I$, write $\mathbf{A}[i_1, \ldots, i_n]$ for the $(\tau \cup \rho)$ -expansion of \mathbf{A} where $C_i^{\mathbf{A}} = \{j \in A \mid i_j = i\}$.

Definition 2 Let \mathcal{F} be a colored obstruction set for an MMSNP-sentence with input signature τ . For a τ -structure **G** with domain *G* of cardinality *k*, the relation induced by **G**, denoted by $\mathbb{R}_{\mathbf{G}}^{\mathcal{F}}$, is the *k*-ary relation over *I* defined as follows.

$$R_{\mathbf{G}}^{\mathcal{F}} := \{ (i_1, \dots, i_k) \in I^k \mid \mathbf{F} \not\to \mathbf{G}[i_1, \dots, i_k] \text{ for all } \mathbf{F} \in \mathcal{F} \}$$

Here we assume (wlog) that $G = \{1, \ldots, k\}$.

The induced constraint language of \mathcal{F} is a relational structure Γ with domain I that contains the relation $R_{\mathbf{G}}^{\mathcal{F}}$ for each τ -reduct \mathbf{G} of a structure from \mathcal{F} . We say that \mathbf{A} has an \mathcal{F} free ρ -coloring if there are $i_1, \ldots, i_n \in I$ such that no $\mathbf{F} \in \mathcal{F}$ homomorphically maps to $\mathbf{A}[i_1, \ldots, i_n]$.

Example 3 Let $\mathcal{F} = {\mathbf{F}_1, \mathbf{F}_2}$ be the colored obstruction set for the No-Mono-Directed-Tri Problem described in Example 2. Then the induced constraint language of \mathcal{F} contains one ternary relation R,

$$R = \{0, 1\}^3 \setminus \{(0, 0, 0), (1, 1, 1)\}$$
.

Lemma 1 Let \mathcal{F} be a colored obstruction set of an MMSNP sentence Φ , and Γ the induced constraint language of \mathcal{F} . Then there is a polynomial-time reduction from Φ to $CSP(\Gamma)$.

PROOF: Let $\tau \cup \rho$ be the signature of \mathcal{F} . Let \mathbf{A} be a τ -structure with domain $A = \{1, \ldots, n\}$. Again, assume without loss of generality that the domain of each structure \mathbf{F} in \mathcal{F} has domain $\{1, \ldots, k\}$, for some k. We create an instance \mathbf{B} of $\mathrm{CSP}(\Gamma)$ as follows. The domain B of \mathbf{B} equals the domain A of \mathbf{A} . For each homomorphism h from a τ -reduct \mathbf{G} of a structure from \mathcal{F} with |G| = k into \mathbf{A} we add the constraint $R(h(1), \ldots, h(k))$ to \mathbf{B} , where R is the relation induced by \mathbf{G} .

The structure **B** can clearly be constructed in polynomial time, since \mathcal{F} is finite and fixed. By Theorem 1, it suffices to show that **A** has an \mathcal{F} -free ρ -coloring if and only if **B** homomorphically maps to Γ . If f is a homomorphism from **B** to Γ , then the $(\tau \cup \rho)$ -expansion $\mathbf{A}[f(1), \ldots, f(n)]$ of **A** does not allow any homomorphism from a structure $\mathbf{F} \in \mathcal{F}$. The last implication is in fact an equivalence, and the statement follows.

5 Reductions from CSPs

In the previous section, we have seen how to reduce an MMSNP sentence Φ to a certain CSP. In many situations, we can go in the opposite direction, reducing the same CSP to Φ (Lemma 2). This will be our main tool to prove hardness of MMSNP sentences in the classification in Section 7.

A clique of size $k \ge 1$ (short, k-clique, also denoted by \mathbf{K}_k) over the signature τ is a τ structure \mathbf{A} with |A| = k where for each $R \in \tau$ of arity r the relation $R^{\mathbf{A}}$ contains all r-tuples
from A^r where not all entries are equal. It is clear that for l < k there is no homomorphism
from \mathbf{K}_k to \mathbf{K}_l , and that the structure induced by l elements in \mathbf{K}_k is isomorphic to \mathbf{K}_l .

A relational structure **A** is said to have girth greater than k if for any choice of $l \leq k$ tuples from relations R_i of arity r_i in **A**, the total number of elements from these l tuples is at least $1 + \sum_{i \leq l} (r_i - 1)$. In the proof of the main result of this section, we make use of the following theorem, which was shown in a randomized version in [7] and which was later derandomized in [9].

Theorem 3 (from [7] and [9]) Fix two integers k, d. Then for every τ -structure **A** on n elements there exists a τ -structure **B** on n^a elements (where a depends only on k and d) such that the girth of **B** is greater than k, there is a homomorphism from **B** to **A**, and for every τ -structure **C** on at most d elements there is a homomorphism from **B** to **C** if and only if there is a homomorphism from **A** to **C**. The structure **B** can be computed from **A** in polynomial time.

An element $a \in A$ of a relational structure **A** is called an *articulation point* if the structure induced by $A \setminus \{a\}$ in **A** is not connected. A structure is called *biconnected* if it does not contain articulation points. A subset of vertices of a structure **A** is called a *block* if it is a maximal biconnected induced substructure of **A**.

Lemma 2 Let \mathcal{F} be a connected monochromatic obstruction set of an MMSNP sentence Φ with color signature C_1, \ldots, C_c . Let Γ_l be $(\{1, \ldots, c\}; R_{K_l}^{\mathcal{F}})$. Then there is a polynomial-time reduction from $\text{CSP}(\Gamma_l)$ to Φ , for all $l \geq 3$.

PROOF: Let **A** be an instance of $\text{CSP}(\Gamma_l)$. We apply Theorem 3 to compute from **A** in polynomial time a structure **B** of girth greater than the maximal obstruction size of \mathcal{F} , such that **A** homomorphically maps to Γ_l iff **B** homomorphically maps to Γ_l . From **B**, we construct a τ -structure **C** with the same domain as **B** as follows. For any tuple (x_1, \ldots, x_l) of the relation $R^{\mathbf{B}}$, we add tuples to the relations from τ such that x_1, \ldots, x_l induces \mathbf{K}_l in **C**.

We claim that **C** satisfies Φ if and only if **A** homomorphically maps to Γ_l . Suppose first that there is an \mathcal{F} -free ρ -coloring **C'** of **C**. We show that the ρ -coloring gives rise to a homomorphism h from **B** to Γ_l , which suffices by Theorem 3. The mapping h is defined as follows: if v is colored by C_i in **C'**, then h(v) := i. If there is a tuple (x_1, \ldots, x_l) in $\mathbb{R}^{\mathbf{B}}$ such that $(h(x_1), \ldots, h(x_l)) \notin \mathbb{R}^{\Gamma_l}$, then by definition of \mathbb{R}^{Γ_l} there must be an $\mathbf{F} \in \mathcal{F}$ such that $\mathbf{F} \to \mathbf{K}_l[h(x_1), \ldots, h(x_l)]$, in contradiction to the assumption that h corresponds to a \mathcal{F} -free coloring of **C**. Hence, h is a homomorphism from **B** to Γ_l . For the opposite direction of the claim, assume that there is a homomorphism from \mathbf{A} to Γ_l . Then there is also a homomorphism h from \mathbf{B} to Γ_l by Theorem 3. We now show that if we color the elements of \mathbf{C} according to h, we obtain an \mathcal{F} -free ρ -coloring \mathbf{C}' of \mathbf{C} . Suppose for contradiction that there is an $\mathbf{F} \in \mathcal{F}$ and a homomorphism f from \mathbf{F} to \mathbf{C}' .

First consider the case that the image of \mathbf{F} under f is fully contained in a substructure of \mathbf{C}' with elements x_1, \ldots, x_l isomorphic to $K_l[i_1, \ldots, i_l]$. Note that $h(x_j) = i_j$ for all $1 \leq j \leq l$. Since \mathbf{F} homomorphically maps to $K_l[i_1, \ldots, i_l]$, we have that $(i_1, \ldots, i_l) \notin R_{\mathbf{K}_l}^{\mathcal{F}}$, by definition of $R_{\mathbf{K}_l}^{\mathcal{F}}$. This contradicts the assumption that $(h(x_1), \ldots, h(x_l)) \in R^{\Gamma_l}$.

The other case is that the image of \mathbf{F} under f is not contained in a clique of \mathbf{C} . Due to the high girth of \mathbf{B} , the image of \mathbf{F} induces an acyclic structure in \mathbf{B} . Since all edges in \mathbf{C} have been introduced via cliques that replace edges of \mathbf{B} , and since \mathbf{F} is connected, the image of \mathbf{F} under f must in this case induce in \mathbf{C}' a structure \mathbf{H} that contains articulation points. Let \mathbf{G} be a block of \mathbf{H} of maximal cardinality. The elements of \mathbf{G} induce a clique in \mathbf{C} . Since \mathbf{F} and \mathbf{G} are monochromatic, this clique will be monochromatic in \mathbf{C}' . By the choice of \mathbf{G} , there is also a homomorphism f' from \mathbf{F} to \mathbf{C}' that maps all of \mathbf{F} to this clique. As in the previous case, this contradicts the definition of $R_{\mathbf{K}_l}^{\mathcal{F}}$ and the assumption that h is a homomorphism.

6 The two-chromatic case

With the result of the previous section we can deduce the NP-hardness of an MMSNP sentences Φ with colored obstruction sets \mathcal{F} and color signature C_1, \ldots, C_c from the NP-hardness of $\text{CSP}((\{1, \ldots, c\}; R_{K_l}^{\mathcal{F}})))$. But for some NP-hard sentences Φ , this CSP is *not* NP-hard. This section is devoted to hardness proofs in those situations.

The chromatic number $\chi(\mathbf{A})$ of a loopless τ -structure \mathbf{A} is the least $k \geq 1$ such that there is a homomorphism from \mathbf{A} to the k-clique over τ . Note that if $\chi(\mathbf{A}) = 1$ then all relations of \mathbf{A} must be empty. If \mathcal{F} is a monochromatic obstruction set and $C \in \rho$, then $\chi(\mathcal{F}, C)$ is the least $k \geq 1$ such that $\chi(\mathbf{F}) = k$ for some C-chromatic $\mathbf{F} \in \mathcal{F}$; if there is no C-chromatic $\mathbf{F} \in \mathcal{F}$, we set $\chi(\mathbf{F}) = -\infty$.

A core of a finite structure \mathbf{A} is the smallest substructure \mathbf{B} of \mathbf{A} such that $\mathbf{A} \to \mathbf{B}$ (all cores of \mathbf{A} are isomorphic, and we therefore speak of *the* core of \mathbf{A}). A finite structure \mathbf{A} is a *core* if the core of \mathbf{A} is \mathbf{A} .

We call a colored obstruction set \mathcal{F} reduced if there are no two structures \mathbf{F}_1 and \mathbf{F}_2 in \mathcal{F} such that \mathbf{F}_1 homomorphically maps to \mathbf{F}_2 , and all structures \mathbf{F} in \mathcal{F} are loopless. Note that if there are structures \mathbf{F}_1 and \mathbf{F}_2 in \mathcal{F} such that \mathbf{F}_1 homomorphically maps to \mathbf{F}_2 , then we can remove \mathbf{F}_2 from \mathcal{F} without affecting the set $\operatorname{Forb}(\mathcal{F})$. Also recall that we study the complexity of Φ on loopless input structures. If there is a structure $\mathbf{F} \in \mathcal{F}$ that is not loopless, that is, with a relation R that contains a tuple of the form (a, \ldots, a) , then we can remove \mathbf{F} from \mathcal{F} and obtain a colored obstruction set that is equivalent to the previous one on loopless inputs. Hence it is clear that for every obstruction set \mathcal{F} there is a reduced obstruction set $\mathcal{F}' \subseteq \mathcal{F}$ such that $\operatorname{Forb}(\mathcal{F}) = \operatorname{Forb}(\mathcal{F}')$. So, if \mathcal{F} is connected, then \mathcal{F}' is reduced and connected. **Lemma 3** Let \mathcal{F} be a reduced monochromatic obstruction set of an MMSNP sentence Φ with colors C_0 and C_1 and $\chi(\mathcal{F}, C_0) = \chi(\mathcal{F}, C_1) = 2$. If \mathcal{F} contains a connected obstruction **F** of size at least three then Φ is NP-hard on loopless input structures.

PROOF: Let τ be the input signature of Φ . Let Γ be the structure $(\{0, 1\}; R, E)$ where R is the Boolean relation $R_{\mathbf{G}}$ induced by the τ -reduct \mathbf{G} of \mathbf{F} , and E is the Boolean relation $R_{\mathbf{K}_2}^{\mathcal{F}}$ induced by the 2-clique \mathbf{K}_2 over signature τ . We first show that $\mathrm{CSP}(\Gamma)$ is NP-hard, and then reduce $\mathrm{CSP}(\Gamma)$ to Φ .

We claim that the relation $E = R_{\mathbf{K}_2}^{\mathcal{F}}$ is $\{(0,1), (1,0)\}$. Because $\chi(\mathcal{F}, C_0) = 2$, there exists an $\mathbf{H} \in \mathcal{F}$ whose τ -reduct \mathbf{H}' maps homomorphically to \mathbf{K}_2 and hence \mathbf{H} maps homomorphically to $\mathbf{K}_2[0,0]$. This shows that (0,0) is not in E, and similarly we can also exclude that the tuple (1,1) is in E. On the other hand, suppose that (0,1) is not in E. Then there exists $\mathbf{H} \in \mathcal{F}$ that homomorphically maps to $K_2[0,1]$. Since \mathbf{H} is monochromatic, \mathbf{H} homomorphically maps to $K_1[0]$, in contradiction to $\chi(\mathcal{F}, C_0) > 1$; or, \mathbf{H} homomorphically maps to $K_1[1]$, in contradiction to $\chi(\mathcal{F}, C_1) > 1$. Similarly, one can show that $(1,0) \in E$, which shows the claim.

We now discuss properties of the relation $R = R_{\mathbf{G}}$. Assume without loss of generality that \mathbf{F} is C_0 -chromatic. Then the relation R does not contain the tuple $(0, \ldots, 0)$. We claim that it contains the tuples $t_1 := (1, 0, \ldots, 0), t_2 := (0, 1, \ldots, 0), \ldots, t_n := (0, \ldots, 0, 1)$. Suppose otherwise that there is $\mathbf{H} \in \mathcal{F}$ with $\mathbf{H} \to \mathbf{G}[t_i]$ for some $i \leq n$. The structure \mathbf{H} cannot be C_0 -chromatic, since then $\mathbf{H} \to \mathbf{F}$, in contradiction to the assumption that no two structures in \mathcal{F} are homomorphically related. Since \mathcal{F} is monochromatic, this shows that \mathbf{H} must be C_1 -chromatic. Because only one entry of t_i is 1, the homomorphism from \mathbf{H} to $\mathbf{G}[t_i]$ must be constant. If there is some relation $R \in \tau$ such that $R^{\mathbf{H}}$ is non-empty, then $R^{\mathbf{G}}$ must therefore contain a tuple of the form (a, \ldots, a) , which we excluded. But then \mathbf{H} maps to a 1-clique, in contradiction to $\chi(\mathcal{F}, C_1) > 1$.

To verify that $\text{CSP}(\Gamma)$ is NP-complete, we use the well-known classification of Boolean CSPs by Schaefer [13]. Since Γ contains the relation E it is neither 0-valid nor 1-valid. Therefore, it suffices to verify that Γ is not closed under the ternary majority (the bijunctive case), the ternary minority (the affine case), the binary 'and' operation (the Horn case), and the binary 'or' operation (the dual Horn case).

The relation $R_{\mathbf{G}}$ is not preserved by the majority operation, since majority $(t_1, t_2, t_3) = (0, \ldots, 0)$ is not contained in $R_{\mathbf{G}}$. Since Γ contains the relation $E = \{(0, 1), (1, 0)\}$, it cannot be closed under the binary operation 'and' and the binary operation 'or'. Finally, assume for contradiction that $R_{\mathbf{G}}$ is closed under the minority operation $(x, y, z) \mapsto x \oplus y \oplus z$. Then the tuple $s := (1, 1, 0, \ldots, 0)$ cannot be in $R_{\mathbf{G}}$ since otherwise minority $(t_1, t_2, s) = (0, \ldots, 0)$ would be in $R_{\mathbf{G}}$. Therefore, there must be $\mathbf{H} \in \mathcal{F}$ such that $\mathbf{H} \to \mathbf{G}[s]$. From $\mathbf{H} \not\to \mathbf{F}$ and the fact that \mathbf{H} is monochromatic we conclude that \mathbf{H} is C_1 -chromatic. Hence, all tuples from $\{0, 1\}^k$ where the first two arguments are 1 are not contained in $R_{\mathbf{G}}$ as well. But minority $(t_1, t_2, t_3) = (1, 1, 1, 0, \ldots, 0) \in R_{\mathbf{G}}$, a contradiction.

We now show that $\text{CSP}(\Gamma)$ can be reduced to Φ . From a given instance \mathbf{A} of $\text{CSP}(\Gamma)$ with variables x_1, \ldots, x_n and tuples t_0, \ldots, t_{m-1} from $R^{\mathbf{A}}$ and $E^{\mathbf{A}}$, we create an instance \mathbf{B} of Φ (i.e., a τ -structure) as follows. The vertices of \mathbf{B} are $v_{1,1}, v_{1,2}, \ldots, v_{1,2m-2}, v_{2,1}, \ldots, v_{n,2m-2}$. For each tuple $t_j = (x_{i_1}, \ldots, x_{i_k}) \in R^{\mathbf{A}}$ we add tuples to the relations $S^{\mathbf{B}}$ for $S \in \tau$ such that $(v_{i_1,2j}, \ldots, v_{i_k,2j})$ induces \mathbf{G} in \mathbf{B} . Similarly, for each tuple $t_j = (x_{i_1}, x_{i_2}) \in E^{\mathbf{A}}$ we add tuples to the relations $S^{\mathbf{B}}$ for $S \in \tau$ such that $(v_{i_1,2j}, v_{i_2,2j})$ induces \mathbf{K}_2 in \mathbf{B} . Moreover, for all $i \leq n, 0 \leq j < 2m-2$ we add tuples to \mathbf{B} such that the structure induced by $\{v_{i,j}, v_{i,j+1}\}$ is isomorphic to the 2-clique over signature τ . The idea here is that the different occurrences of the variables in constraints in \mathbf{A} are linked in \mathbf{B} by paths of 2-cliques of even length. Clearly, \mathbf{B} can be constructed in polynomial time from \mathbf{B} . We claim that the \mathbf{B} satisfies Φ if and only if \mathbf{A} homomorphically maps to Γ .

First suppose that there is a homomorphism h from \mathbf{A} to Γ . We claim that the following ρ -expansion \mathbf{B}' of \mathbf{B} is \mathcal{F} -free: for all $1 \leq i \leq n, 1 \leq j \leq m-1$, the vertex $v_{i,2j}$ receives color $C_{h(x_i)}$. For all $1 \leq j < m-1$, the vertex $v_{i,2j+1}$ receives color $C_{1-h(x_i)}$. Suppose for contradiction that there is $\mathbf{F} \in \mathcal{F}$ with a homomorphism f to \mathbf{B}' . Since for all pairs of elements that induce a 2-clique in \mathbf{B}' one element is colored by C_0 , and one is colored by C_1 , and since $\chi(\mathcal{F}) > \infty$, and because \mathbf{F} is connected, the image of \mathbf{F} under f must be contained in $\{v_{i_1,2j},\ldots,v_{i_k,2j}\}$ for some tuple (x_{i_1},\ldots,x_{i_k}) from $R^{\mathbf{A}}$. However, this contradicts the definition of R and the construction of \mathbf{B}' from the homomorphism h.

Now suppose that **B** satisfies Φ . That is, there is a \mathcal{F} -free ρ -coloring of **B**. Then for all $1 \leq i \leq n, 1 \leq j < 2m-2$ the vertex $v_{i,j}$ has color C_0 if and only if the vertex $v_{i,j+1}$ has the color C_1 . Therefore, for all $1 \leq j \leq 2m-4$, the vertex $v_{i,j}$ has the same color as the vertex $v_{i,j+2}$, and consequently all vertices from $\{v_{i,2j} \mid 1 \leq j < m-1\}$ have the same color C_{c_i} . Then the assignment that maps x_i to c_i is a homomorphism from **A** to Γ .

7 Classification

In this section we classify the computational complexity of MMSNP on loopless input structures for all MMSNP sentences Φ with a monochromatic obstruction set. By Corollary 2, we can assume that Φ is connected.

For a colored obstruction set \mathcal{F} , define $||\mathcal{F}||$ to be $\sum_{\mathbf{F}\in\mathcal{F}}|\mathbf{F}|$.

Theorem 4 Let Φ be a connected MMSNP sentence that has a monochromatic obstruction set. On loopless input structures, Φ is in P or NP-complete.

PROOF: Let \mathcal{F} be a reduced connected monochromatic colored obstruction set for Φ . We prove the statement by induction on the lexicographic ordering of the pairs $(|\rho|, ||\mathcal{F}||)$.

If $\chi(\mathcal{F}, C) = -\infty$ for some $C \in \rho$, then a given **A** can be trivially expanded to a \mathcal{F} -free $(\tau \cup \rho)$ -structure by coloring all elements of **A** by C, and Φ is in P.

If $\chi(\mathcal{F}, C) = 1$ for some $C \in \rho$, then there is some *C*-chromatic $\mathbf{F} \in \mathcal{F}$ such that all relations of \mathbf{F} except for *C* are empty. Hence, every expansion \mathbf{A}' of a given τ -structure \mathbf{A} where $C^{\mathbf{A}'}$ is non-empty admits a homomorphism from \mathbf{F} . Let \mathcal{F}' be the set of all $((\rho \cup \tau) \setminus \{C\})$ -reducts of structures \mathbf{F} in \mathcal{F} where $C^{\mathbf{F}}$ is empty. Then \mathcal{F}' is a colored obstruction set for a sentence that is polynomial-time equivalent to Φ , but has a smaller signature, and so can be handled by induction. Thus, we can assume that $\chi(\mathcal{F}, C) \geq 2$ for all $C \in \rho$.

If ρ contains just one element, then Φ can be solved in polynomial time by checking for a given τ -structure whether the τ -reduct of some structure in \mathcal{F} homomorphically maps to **A**; this is the case if and only if the structure **A** does not satisfy Φ . I. The two color case. Suppose next that ρ has two elements, say C_0 and C_1 . If every structure in \mathcal{F} has two elements, then the induced constraint language Γ of Φ is bijunctive, and hence $\text{CSP}(\Gamma)$ can be solved in polynomial time [13]. Lemma 1 implies that also Φ can be solved in polynomial time. So in the following we assume that \mathcal{F} contains a structure \mathbf{F} with at least three elements.

I.1. Chromatic number two. Suppose $\chi(\mathcal{F}, C_0) = \chi(\mathcal{F}, C_1) = 2$. Therefore, the structure $\mathbf{F} \in \mathcal{F}$ with $k \geq 3$ elements satisfies $\chi(\mathbf{F}) = 2$. By assumption, \mathbf{F} is connected. In this case we can apply Lemma 3 to show that Φ is NP-hard.

I.2. Larger chromatic number. In the remaining case, $3 \leq \chi(\mathcal{F}, C_0)$ and $2 \leq \chi(\mathcal{F}, C_1) \leq \chi(\mathcal{F}, C_0)$; or $3 \leq \chi(\mathcal{F}, C_1)$ and $2 \leq \chi(\mathcal{F}, C_0) \leq \chi(\mathcal{F}, C_1)$. Suppose the former case applies; the latter is analogous. We also write l for $\chi(\mathcal{F}, C_0)$. Then the relation $R_{\mathbf{K}_l}^{\mathcal{F}}$ doesn't contain the tuple $(0, \ldots, 0)$ and the tuple $(1, \ldots, 1)$, and any tuple in $R_{\mathbf{K}_l}^{\mathcal{F}}$ must have t entries with value 1 and c - t entries with value 0, in any order, with $1 \leq t < \chi(\mathcal{F}, C_1)$. Let Γ be $(\{0, 1\}; R)$ where R denotes $R_{\mathbf{K}_l}^{\mathcal{F}}$. It follows easily from Schaefer's theorem that $\mathrm{CSP}(\Gamma)$ is NP-hard. NP-hardness of Φ follows from Lemma 2.

II. More than two colors. Finally, suppose that ρ has more than two elements, say C_0, \ldots, C_c . Let k be $\sum_{C \in \rho} (\chi(\mathcal{F}, C) - 1)$. Then every tuple from $R := R_{\mathbf{K}_k}^{\mathcal{F}}$ must have less than $\chi(\mathcal{F}, C_i)$ entries with value i, for every $i \leq c$. Moreover, all tuples with this property belong to R, so R contains all tuples with exactly $\chi(\mathcal{F}, C_i) - 1$ entries with value i. Let Γ be the structure (I; R). We claim that $\mathrm{CSP}(\Gamma)$ is NP-hard. This implies NP-hardness of Φ by Lemma 2.

To prove the claim, first observe that the finite structure Γ is a core. It is well known (see Corollary 4.8 in [5]) that $\text{CSP}(\Gamma)$ is polynomial-time equivalent to $\text{CSP}(\Delta)$ where Δ is an expansion of Γ by all relations R_a^{Δ} of the form $\{a\}$ for $a \in I$. We prove NP-hardness of $\text{CSP}(\Delta)$ by reduction from |I|-colorability (recall that |I| > 2).

Let **G** be a finite graph. We create an instance **A** of $\text{CSP}(\Delta)$ as follows. The vertices of **A** are the vertices of **G** and some additional vertices. For each edge $\{u, v\}$ from **G**, we introduce new variables w_1, \ldots, w_{k-2} , and add the tuple $(u, v, w_1, \ldots, w_{k-2})$ to $R^{\mathbf{A}}$. Moreover, for each $i \in I$, we place $\chi(\mathcal{F}, i) - 2$ elements from w_1, \ldots, w_{k-1} in $R_a^{\mathbf{A}}$ such that no element from w_1, \ldots, w_{k-1} is contained in $R_a^{\mathbf{A}}$ and $R_b^{\mathbf{A}}$ for $a \neq b$. It is clear that this structure **A** can be computed in polynomial time from **G**. It is straightforward to verify that **A** homomorphically maps to Δ if and only if **G** is |I|-colorable.

Theorem 2 follows from Theorem 4 and Corollary 2. For MMSNP formulas that describe problems over graphs or directed graphs, or more generally problems for structures with a single relation, we can drop the assumption in Theorem 4 that the instances are loopless.

Corollary 1 Let Φ be an MMSNP sentence that has a monochromatic obstruction set \mathcal{F} , and suppose that the input signature τ contains a single relation symbol R. Then Φ is in P or NP-complete.

PROOF: We claim that the problem described by Φ is polynomial-time equivalent to the problem restricted to loopless input; and on loopless input structures the problem is in P or NP-complete, by Theorem 4.

It suffices to show that the problem described by Φ can be reduced to the problem for loopless input. We can assume that for every color $C \in \rho$ there is a *C*-chromatic obstruction $\mathbf{F} \in \mathcal{F}$ (otherwise, the problem is trivial). But then any ρ -expansion of a τ -structure **A** that is not loopless does not satisfy Φ . To see this, let (a, \ldots, a) be the tuple in $\mathbb{R}^{\mathbf{A}}$ with only equal entries. If *a* is colored by *C* in the expansion, then there is a homomorphism from a *C*-colored $\mathbf{F} \in \mathcal{F}$ to **A** that maps all elements of **F** to *a*.

Acknowledgement. This work arose from the meeting "Applications of universal algebra and logic to the constraint satisfaction problem" sponsered by and held at the American Institute of Mathematics (AIM). The authors would like to thank the anonymous referee for his very useful remarks, and Florent Madelaine for valuable discussions around MMSNP.

References

- D. Achlioptas. The complexity of G-free colourability. Discrete Mathematics, 165:21–30, 1997.
- [2] Manuel Bodirsky and Víctor Dalmau. Datalog and constraint satisfaction with infinite templates. An extended abstract appeared in the proceedings of STACS'06. The full version is available online at arXiv:0809.2386 [cs.LO], 2008.
- [3] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2), 2009. An extended abstract appeared in the proceedings of STOC'08.
- [4] Manuel Bodirsky and Michael Pinsker. Schaefer's theorem for graphs. In Proceedings of STOC, pages 655–664, 2011. Preprint of the long version available at arxiv.org/abs/1011.2894.
- [5] Andrei Bulatov, Andrei Krokhin, and Peter G. Jeavons. Classifying the complexity of constraints using finite algebras. SIAM Journal on Computing, 34:720–742, 2005.
- [6] Nadia Creignou, Sanjeev Khanna, and Madhu Sudan. Complexity Classifications of Boolean Constraint Satisfaction Problems. SIAM Monographs on Discrete Mathematics and Applications 7, 2001.
- [7] Tomás Feder and Moshe Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. SIAM Journal on Computing, 28:57–104, 1999.
- [8] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. Journal of Combinatorial Theory, Series B, 48:92–110, 1990.

- [9] Gabor Kun. Constraints, MMSNP, and expander relational structures. Available at arXiv:0706.1701, 2007.
- [10] Gábor Kun and Jaroslav Nešetřil. NP by means of lifts and shadows. In MFCS, pages 171–181, 2007.
- [11] Florent Madelaine and Iain A. Stewart. Constraint satisfaction, logic and forbidden patterns. SIAM Journal on Computing, 37(1):132–163, 2007.
- [12] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43:425–440, 1991.
- [13] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of STOC*, pages 216–226, 1978.