

ON GUARDED EXTENSIONS OF MMSNP

ALEXEY BARSUKOV ^a AND FLORENT R. MADELAINE ^b

^a Department of Algebra, Faculty of Mathematics and Physics, Charles University, Sokolovská 49/83,
186 00 Prague 8, Czechia
e-mail address: alexey.barsukov@matfyz.cuni.cz

^b Univ Paris Est Creteil, LACL, F-94010 Creteil, France
e-mail address: florent.madelaine@u-pec.fr

ABSTRACT. Feder and Vardi showed that the class *Monotone Monadic SNP without inequality* (MMSNP) has a P vs NP-complete dichotomy if and only if such a dichotomy holds for finite-domain *Constraint Satisfaction Problems* (CSPs). Moreover, they showed that none of the three classes obtained by removing one of the defining properties of MMSNP (monotonicity, monadicity, no inequality) has a dichotomy. The overall objective of this paper is to study the gaps between MMSNP and each of these three superclasses, where the existence of a dichotomy remains unknown. For the gap between MMSNP and Monotone SNP without inequality, we study the class *Guarded Monotone SNP without inequality* (GMSNP) introduced by Bienvenu, ten Cate, Lutz, and Wolter, and prove that GMSNP has a dichotomy if and only if a dichotomy holds for GMSNP problems over signatures consisting of a unique relation symbol. For the gap between MMSNP and MMSNP with inequality, we introduce a new class *MMSNP with guarded inequality*, that lies between MMSNP and MMSNP with inequality and that is strictly more expressive than the former and still has a dichotomy. For the gap between MMSNP and Monadic SNP without inequality, we introduce a logic that extends the class of Matrix Partitions in a similar way how MMSNP extends finite-domain CSP, and pose an open question about the existence of a dichotomy for this class. Finally, we revisit the theorem of Feder and Vardi, which claims that the class NP embeds into MMSNP with inequality. We give a detailed proof of this theorem as it ensures no dichotomy for the right-hand side class of each of the three gaps.

Key words and phrases: descriptive complexity and computational complexity and constraint satisfaction problem and dichotomy and mmsnp.

An extended abstract of this paper has appeared at CiE 2023.

This work was supported by ANR project DIFFERENCE (<https://anr.fr/Projet-ANR-20-CE48-0002>) and by the European Union (ERC, POCOCOP, 101071674). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

1. INTRODUCTION

1.1. The dichotomy question. A subclass of NP has a *dichotomy* if every problem in this class is either solvable in polynomial time or NP-complete. A consequence of Ladner's theorem [Lad75] is that if P is different from NP, then NP does not have a dichotomy. This gave motivation to study subclasses of NP with regard to the existence of a dichotomy. A large class of problems, which was believed to have this property was the class of finite-domain *Constraint Satisfaction Problems* (CSP). The existence of a dichotomy for finite-domain CSP was conjectured by Feder and Vardi [FV98]. Firstly, this conjecture was supported by several results confirming it for subclasses such as CSPs on a two-element domain [Sch78] and CSPs on undirected graphs [HN90]. Secondly, it generalised the conjecture of Bang-Jensen and Hell that a dichotomy holds for CSPs on directed graphs, where the template is a *smooth* digraph, i.e., the in- and out- degree of every vertex is greater than zero [BH90]. Today, both conjectures are known to be true: the weaker version for smooth digraphs was proved by Barto, Kozik, and Niven [BKN09], and the stronger version was proved independently by Bulatov [Bul17] and Zhuk [Zhu20].

Apart from conjecturing a dichotomy for CSPs, Feder and Vardi [FV98] built a link between NP and CSP through descriptive complexity. By Fagin's theorem [Fag74], the class NP is captured by *Existential Second Order logic* (ESO) which means that a class of finite structures is definable by a sentence in ESO if and only if it is decidable by a nondeterministic Turing machine running in polynomial time. Feder and Vardi studied subclasses of the logic *SNP* that itself is a fragment of ESO. They considered three syntactic properties that restrict the class SNP:

- *monotone* – all input relations have odd nested number of negations,
- *monadic* – all existentially quantified relations have arity at most one,
- *without inequality* – the “=” relation and its negation are not allowed.

Their main result states that, under these three restrictions, the resulting fragment of SNP, called *Monotone Monadic SNP without inequality* (MMSNP), is contained in CSP modulo polynomial-time equivalence. Moreover, they showed that it is necessary to have all the three properties in order to have a dichotomy, because removing any of these properties would give a class that does not have a dichotomy.

1.2. Three gaps above MMSNP. Feder and Vardi's result about MMSNP together with Bulatov and Zhuk's proofs of a dichotomy for CSP imply that the logic MMSNP also has a dichotomy. As removing any of the three properties from MMSNP gives a class that has no dichotomy, there appear three gaps, in which the existence of a dichotomy is unknown. Each gap consists of logical classes that contain MMSNP, where one of the three properties is relaxed.

The diagram in Figure 1 displays the current state of the art. There is an (undirected) edge between two classes if the class below is contained in the class above. Directed edges stand for the containment modulo polynomial-time equivalence, i.e., for every problem of the class at the tail there exists a polynomial-time equivalent problem in the class at the head. In Figure 1, the class obtained by removing the “monotone” property is called *monadic SNP*, removing “monadicity” gives the class *monotone SNP*, while allowing “inequality” results in the class $MMSNP_{\neq}$.

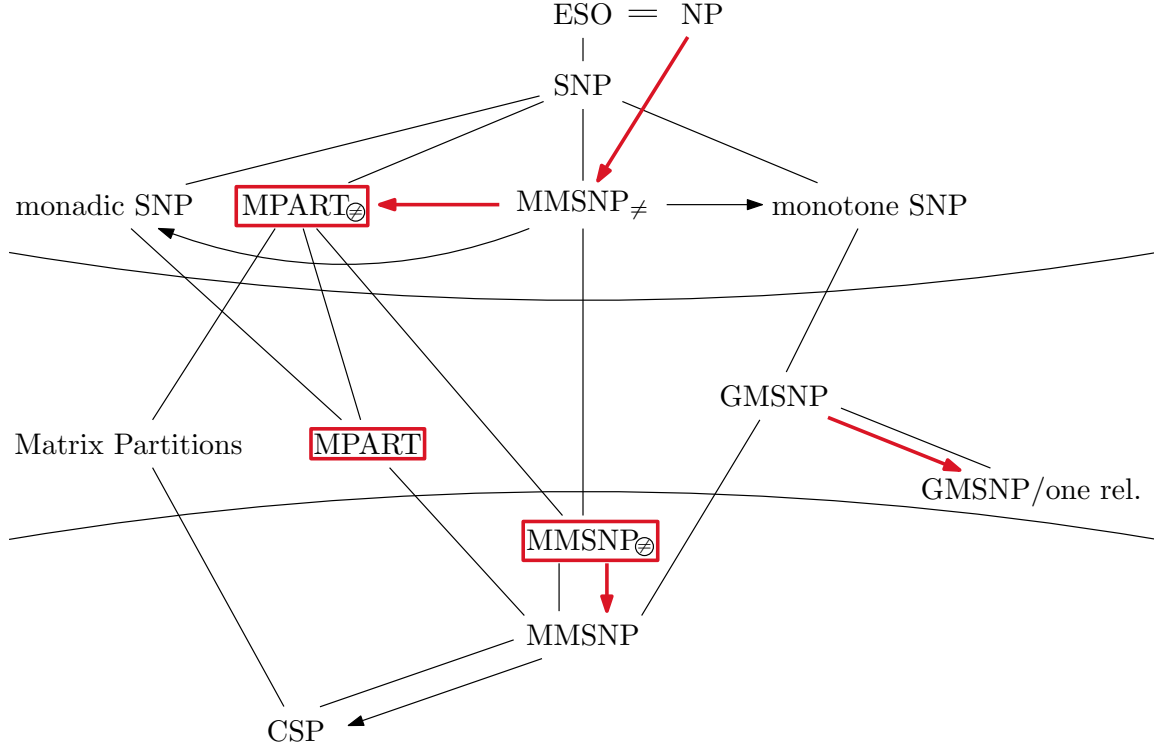


FIGURE 1. Classes at the bottom exhibit a PvsNP-complete dichotomy, while those at the top do not, and it remains open for classes in the middle. Undirected edges stand for straightforward inclusions and directed edges stand for inclusions under polynomial-time reductions. The classes and inclusions discussed in this paper are highlighted with red. The sign \oplus denotes classes extended with guarded inequality.

The contributions of the paper can be split in two parts. In one part, we argue that none of the three upper bound classes has a dichotomy. For this, we revisit Feder and Vardi’s paper [FV98] and give a detailed proof that every nondeterministic Turing machine is polynomial-time equivalent to a problem in monotone monadic SNP without inequality. In the other part, we focus on the three gaps and study them from the dichotomy question perspective. All the classes from the gaps are obtained from MMSN by making one of its three properties less strict:

- “monadicity” is relaxed by allowing existentially quantified relations of any arity as long as they are guarded by input relational tuples;
- “monotonicity” is relaxed by allowing negated input relational atoms as long as any two input atoms in the same negated conjunct are either both negated or both non-negated;
- “no inequality” is relaxed by allowing \neq -atoms as long as they are guarded by input relational tuples.

Each of these three relaxations is *guarded*, in some sense: either by input relational atoms or by negated conjuncts. This feature that all the classes in the gaps have in common, gave the title of this paper.

No dichotomy for MMSNP with inequality. When Feder and Vardi proved that there is no dichotomy for the three logics extending MMSNP, they first showed an inclusion of NP into MMSNP_{\neq} modulo polynomial-time equivalence, and then they showed such inclusions of MMSNP_{\neq} into the other two classes. In their paper, the proof of no dichotomy for MMSNP_{\neq} provides key ideas, in particular, the usage of *oblivious Turing machines*, but it is very brief and thus not transparent enough. In this paper we give a detailed proof of this statement, since it follows precisely from it that the right boundary of each interval has no dichotomy, which is a necessary criterion for studying the complexity of each of the three gaps. Having an explicit proof of this statement makes the gaps *well-defined*, i.e., the left-hand side class MMSNP has a dichotomy whereas all three right-hand side classes do not. The interested reader may consult [Bar22] for extensive proofs regarding the two other logics: monotone SNP and monadic SNP.

In the following three paragraphs, we describe which classes each gap contains and what our contributions are.

Relaxing monadicity. The gap between MMSNP and monotone SNP without inequality, among those three, is the most studied one. The reason is that the family of finite models of every sentence Φ in monotone SNP without inequality is closed under taking inverse homomorphisms, i.e., if Φ holds in \mathfrak{B} and if \mathfrak{A} maps homomorphically to \mathfrak{B} , then Φ also holds in \mathfrak{A} . For a family of finite structures over a finite relational signature, being closed under inverse homomorphisms and disjoint unions is equivalent to being the CSP of some (possibly infinite) structure (Lemma 1.1.8 in [Bod21]). CSPs with infinite domain templates contain the whole class NP modulo Turing equivalence (Corollary 13.2.3 in [Bod21]). Moreover, they can even be undecidable as, for example, solving Diophantine equations can be viewed as a CSP [Mat93].

The largest fragment of infinite CSP, that is wholly contained in NP, consists of those CSPs, where the template is a (first-order) reduct of a finitely bounded homogeneous structure. For this class, a dichotomy was conjectured by Bodirsky and Pinsker in 2011, see [BPP21]. The conjecture has been confirmed for many particular classes; for more information see the survey [Pin22].

One particularly interesting logical class that belongs to the gap between MMSNP and monotone SNP without inequality for which the dichotomy question is still open, is called *Guarded Monotone SNP* (GMSNP). It was introduced by Bienvenu, ten Cate, Lutz, and Wolter in [BtCLW14]. This class allows existentially quantified relations of arity greater than 1 but all such relations must be *guarded* by input relations, i.e., the image of every existential-relational tuple must be contained in the image of some input-relational tuple. In [BtCLW14], the authors showed, in particular, that GMSNP is strictly more expressive than MMSNP and that it is equally expressive as the class MMSNP_2 studied earlier by Madelaine [Mad09]. The latter generalises MMSNP by allowing existentially quantified unary relations whose domain is both set of input vertices and relational tuples. In [BKS20], it was shown that every problem in GMSNP is equivalent to a union (disjunction) of CSPs of first-order reducts of finitely bounded homogeneous structures. This means that the resolution of Bodirsky-Pinsker conjecture would yield a dichotomy for GMSNP. Conversely, showing a dichotomy for GMSNP would be a step towards proving the conjecture of Bodirsky and Pinsker.

In this paper, we show that it suffices to show a dichotomy only for those GMSNP problems, where the input relational signature consists of a single relation symbol.

Relaxing monotonicity. There has already been some research done in generalizing CSPs in terms of modifying the notion of the homomorphism. Typically, such problems are not closed under inverse homomorphisms. One intensively studied such class is called *Matrix Partitions*. Every such problem is described by a square matrix A of size $m \times m$ whose elements take values in the set $\{0, 1, *\}$. The input of the problem consists of graphs. For a given input graph \mathfrak{G} , the Matrix Partition problem parametrised by A asks: is there a partition of the set of vertices of \mathfrak{G} in sets P_1, \dots, P_m such that for all $i, j \in \{1, \dots, m\}$ and for every pair of *distinct* vertices $v_i \in P_i, v_j \in P_j$,

- $A_{ij} = 0$ implies that there must be no edge between v_i and v_j ,
- $A_{ij} = 1$ implies that there must be an edge between v_i and v_j ,
- $A_{ij} = *$ means that there is no restriction.

In particular, if the i th diagonal element A_{ii} is 0, then P_i induces an independent set, and if $A_{ii} = 1$, then P_i induces a clique.

Matrix Partitions contain all CSPs on graphs and the dichotomy question for them is open (Problem 2 in [Hel14]). This containment is strict because Matrix Partitions, in general, are not closed under inverse homomorphisms. For further reading, see the survey [Hel14].

Example 1.1. Consider the 3-colouring and the split graph recognition problems. They can be seen as Matrix Partition problems for the matrices A_3 and A_{split} , respectively, where

$$A_3 = \begin{pmatrix} 0 & * & * \\ * & 0 & * \\ * & * & 0 \end{pmatrix} \quad A_{\text{split}} = \begin{pmatrix} 0 & * \\ * & 1 \end{pmatrix}.$$

Another way to look at Matrix Partitions was given in [BK24]. There, it was proposed to see k -ary relations defined on a set A not as subsets of A^k but rather as mappings from A^k to some fixed poset. In this setting, a homomorphism between structures \mathfrak{A} and \mathfrak{B} was defined as a mapping between their domains A and B that sends every relational tuple to a tuple whose value is greater or equal (with respect to the partial order) than the value of the original tuple. If the poset has elements 0, 1, and the order is defined by the inequality $0 \preceq 1$, this is equivalent to usual CSPs. If the poset has elements 0, 1, *, and the order is defined by $0 \preceq *$ and $1 \preceq *$, then this is *almost* equivalent to Matrix Partitions.

These two definitions are not completely equivalent because, in Hell's definition, the partition constraints are applied to pairs of *distinct* vertices whereas the class studied in [BK24] does it for all pairs. As a result, trying to write down a problem from the former class as an SNP sentence would require the use of inequalities as input relations while the latter class does not have this issue. It remains unclear whether these two classes are polynomial-time equivalent. However, we feel that the variation from [BK24] is more amenable to the framework of descriptive complexity while retaining the key property of exhibiting non-monotonicity in a suitably restricted form that could prevent NP-intermediate complexity to arise.

In the present article, we introduce a logic MPART that lies in the gap between MMSNP and monadic SNP without inequality. The class MPART extends MMSNP in a similar fashion as how Matrix Partitions extend graph colourings. However, using the inequality as an input relation symbol is not allowed (which is required to belong to the gap). Moreover, we show that, adding inequality to MPART in order to express Hell's Matrix Partitions, results in a logic, denoted by MPART_{\oplus} , that contains MMSNP_{\neq} modulo polynomial-time equivalence and, therefore, has no dichotomy.

Relaxing “no inequality”. To our knowledge, there are no classes from this gap that have been studied in the literature. We introduce a new logical class called *MMSNP with guarded inequality* (MMSNP_{\oplus}) that belongs to the gap. It generalises MMSNP by adding a restricted form of inequality, where each \neq -atomic formula is guarded by an atomic formula of some input relational symbol, similarly as atomic formulas of existential relations are guarded by input relational atomic formulas in GMSNP. We show that this class is strictly more expressive than MMSNP. However, it is still included in MMSNP modulo polynomial-time equivalence and, therefore, it has a dichotomy.

1.3. Outline. We give all the necessary definitions in Section 2. Then, in Section 3, we visit the three gaps one by one. In Subsection 3.1, we introduce MMSNP with guarded inequality (MMSNP_{\oplus}) into the gap between MMSNP and MMSNP with inequality and prove that it exhibits a dichotomy. In Subsection 3.2, we study the class GMSNP that is in the gap between MMSNP and monotone SNP without inequality. We establish that proving a dichotomy for signatures consisting of a single input relation would suffice. In Subsection 3.3, we give the definition for the logics MPART and MPART_{\oplus} and show that the latter does not have a dichotomy. Finally, in Section 4, we provide a complete proof of one crucial step of Feder and Vardi’s argument to establish that the three stronger siblings of MMSNP cannot have a dichotomy unless P equals NP, which makes the three gaps well-defined.

2. PRELIMINARIES

Denote by τ a finite relational signature. Let σ be a set of relation symbols $\{X_1, \dots, X_s\}$. Relation symbols of τ are called *input* and relation symbols of σ are called *existential*.

Every τ -sentence Φ in SNP can be written in the following form:

$$\exists X_1, \dots, X_s \forall \mathbf{x} \bigwedge_i \neg(\alpha_i \wedge \beta_i \wedge \epsilon_i), \quad (2.1)$$

where α_i is a conjunction of atomic or negated atomic τ -formulas, β_i is a conjunction of atomic or negated atomic σ -formulas, and ϵ_i is a conjunction of inequalities $x_i \neq x_j$ such that $x_i, x_j \in \mathbf{x}$. In this paper, an SNP sentence written in the form of Eq. (2.1) is said to be in *standard* form.

While SNP sentences are denoted by upper-case letters: Φ , their quantifier-free parts are denoted by lower-case letters: $\phi(\mathbf{x})$. Tuples of first-order variables or elements are denoted with bold lower-case letters: $\mathbf{x}, \mathbf{a}, \dots$

For the sake of brevity, we write “atom” instead of “atomic formula”. Every subformula $\neg(\alpha_i \wedge \beta_i \wedge \epsilon_i)$ of the quantifier-free part $\phi(\mathbf{x})$ is called a *negated conjunct*.

A $(\tau \cup \sigma)$ -structure \mathfrak{A}^σ is a σ -*expansion* of a τ -structure \mathfrak{A} if these two structures have the same domain and if, for every R in τ , the relations $R^{\mathfrak{A}}$ and $R^{\mathfrak{A}^\sigma}$ are the same. If \mathfrak{A}^σ is a σ -expansion of \mathfrak{A} , then \mathfrak{A} is called the τ -*reduct* of \mathfrak{A}^σ . For a τ -sentence Φ in standard form, we say that Φ *holds* in a τ -structure \mathfrak{A} , denoted $\mathfrak{A} \models \Phi$, if there exists a σ -expansion \mathfrak{A}^σ such that $\mathfrak{A}^\sigma \models \phi(\mathbf{a})$, for every assignment \mathbf{a} of elements of \mathfrak{A} to the variables of $\phi(\mathbf{x})$. When we want to highlight that a tuple \mathbf{a} of elements of a set A is assigned to the tuple \mathbf{x} of variables, we write $\mathbf{a} \in A^{\mathbf{x}}$.

In this paper, both the set of all finite τ -structures, where Φ holds, and the problem deciding the satisfiability of Φ , are denoted by $\text{fm}(\Phi)$.

An SNP-sentence Φ is called *monadic* if all existential symbols have arity one. It is called *monotone* if every α_i omits negated τ -atoms, and *without inequality* if every ϵ_i is empty.

Alternatively, one can think of being “monotone” and “without inequality” as of being closed under inverse homomorphisms. The latter means that a structure \mathfrak{A} satisfies Φ whenever \mathfrak{A} homomorphically maps to a structure $\mathfrak{B} \in \text{fm}(\Phi)$. Indeed, every sentence in monotone SNP without inequality is closed under inverse homomorphisms and the other direction is provided by the following.

Theorem 2.1 [FV03]. *For every sentence Φ in SNP, the class $\text{fm}(\Phi)$ is closed under inverse homomorphisms only if Φ is logically equivalent to a sentence in monotone SNP without inequality.*

Denote by MMSNP the fragment of SNP consisting of sentences that are monotone, monadic, and without inequality.

An SNP sentence Φ in standard form is in *Guarded Monotone Strict NP* (GMSNP) if (a) each α_i is a conjunction of non-negated τ -atoms, (b) each β_i is a conjunction of σ -atoms or negated σ -atoms, and (c) each ϵ_i is empty. Also, for every negated σ -atom $\neg X(\mathbf{x}_1)$ from every β_i there exists either (i) a τ -atom $R(\mathbf{x}_2)$ from α_i , or (ii) a non-negated σ -atom $X'(\mathbf{x}_2)$ from β_i such that $\mathbf{x}_1 \subseteq \mathbf{x}_2$. In this case, we say that $R(\mathbf{x}_2)$ (or $X'(\mathbf{x}_2)$) *guards* $X(\mathbf{x}_1)$.

Definition 2.2. An SNP sentence Φ in standard form is in *MMSNP with guarded inequality* (MMSNP_{\oplus}) if (a) each α_i is a conjunction of non-negated τ -atoms, (b) each β_i is a conjunction of unary σ -atoms or negated unary σ -atoms, and (c) for every inequality $x_j \neq x'_j$ of each ϵ_i there exists a τ -atom in α_i that contains both x_j and x'_j .

Remark 2.3. For every sentence that we consider, we assume that it is not trivial, i.e., it has some YES and some NO instances.

3. STUDYING THE THREE GAPS

3.1. MMSNP with guarded inequality. In this subsection, we prove the following theorem.

Theorem 3.1. *For every sentence Φ in MMSNP_{\oplus} there exists a sentence Φ' in MMSNP such that the problems $\text{fm}(\Phi)$ and $\text{fm}(\Phi')$ are polynomial-time equivalent.*

For the sake of brevity, we consider the case when $\tau = \{R\}$, where R has arity n . In the case of several relations, one can apply the same procedure for each relation independently.

The construction has two main steps. At first, we enrich every negated conjunct of Φ with inequalities while keeping the resulting sentence logically equivalent to the initial one. This transformation is the same as the one used in Lemma 4 in [FV03]. After that, we construct an equivalent MMSNP sentence Φ' by changing the input relational signature. Each new relation symbol is associated with an equivalence relation on an n -element set, where n is the arity of R .

Enriching the sentence with inequalities. For every negated conjunct $\neg\phi$ of Φ and for every two distinct variables x, y that appear in the same \mathbf{R} -atom of $\neg\phi$, replace this negated conjunct with two negated conjuncts: the first one is obtained from $\neg\phi$ by adding the inequality $x \neq y$ to the conjunction, the second one is obtained from $\neg\phi$ by replacing every occurrence of the variable y with the variable x . Repeat this procedure until every negated conjunct contains the inequality $x \neq y$ for every two distinct variables x, y of this conjunct. Call the resulting sentence Ψ .

Example 3.2. If Φ contains a negated conjunct $\neg(\mathbf{R}(x, y, z) \wedge X_1(x) \wedge X_2(y))$, then it is replaced by the following five negated conjuncts:

$$\begin{aligned} & \neg(\mathbf{R}(x, y, z) \wedge X_1(x) \wedge X_2(y) \wedge x \neq y \wedge x \neq z \wedge y \neq z) \\ & \wedge \neg(\mathbf{R}(x, x, y) \wedge X_1(x) \wedge X_2(x) \wedge x \neq y) \wedge \neg(\mathbf{R}(x, y, x) \wedge X_1(x) \wedge X_2(y) \wedge x \neq y) \\ & \wedge \neg(\mathbf{R}(x, y, y) \wedge X_1(x) \wedge X_2(y) \wedge x \neq y) \wedge \neg(\mathbf{R}(x, x, x) \wedge X_1(x) \wedge X_2(x)). \end{aligned}$$

The following lemma is a direct consequence of the construction.

Lemma 3.3. *Every MMSNP_{\oplus} sentence Φ is logically equivalent to a MMSNP_{\oplus} sentence Ψ such that, for every negated conjunct $\neg\psi_i$ of Ψ and for every two different variables x, y that appear within some \mathbf{R} -atom of ψ_i , this conjunct contains the inequality $x \neq y$.*

Proof. Let Φ and Ψ be as in the paragraph above Example 3.2, and let \mathfrak{A} be a τ -structure. Suppose that $\mathfrak{A} \not\models \Phi$; then, for every σ -expansion \mathfrak{A}^σ there exists a negated conjunct $\neg\phi_i(\mathbf{x})$ of Φ and a tuple $\mathbf{a} \in A^{\mathbf{x}}$ such that $\mathfrak{A}^\sigma \models \phi_i(\mathbf{a})$. Let $\neg\psi_i(\mathbf{x}')$ be the negated conjunct of Ψ that is obtained from $\neg\phi_i(\mathbf{x})$ by adding the inequality $x_i \neq x_j$, for every two distinct elements $a_i, a_j \in \mathbf{a}$, and by identifying x_i with x_j , for every two elements $a_i, a_j \in \mathbf{a}$ that are the same. By the construction of Ψ , such a conjunct exists, so it will not be satisfied in \mathfrak{A}^σ . If $\mathfrak{A} \not\models \Psi$, then, for every σ -expansion \mathfrak{A}^σ , there exists a negated conjunct $\neg\psi_i$ of Ψ that is false for some assignment. It is associated with some negated conjunct $\neg\phi_i$ of Φ , and, by construction, ϕ_i will not be satisfied in \mathfrak{A}^σ for the same assignment. We conclude that Φ and Ψ are logically equivalent. \square

Now, by Lemma 3.3, we may assume that, for every two variables x, y that appear within the same \mathbf{R} -atom in some negated conjunct of Φ , this negated conjunct contains $x \neq y$. For $n \in \mathbb{N}$, denote by B_n the number of equivalence relations on a set of n elements: $\sim_1, \dots, \sim_{B_n}$. For each \sim_k , denote by $n_k := |\{1, \dots, n\} / \sim_k|$ the number of equivalence classes of this relation. Every n -tuple (x_1, \dots, x_n) is associated with exactly one \sim_k such that $x_i = x_j$ if and only if $i \sim_k j$. If \sim_k is associated with \mathbf{x} , then we say that \mathbf{x} has *equivalence type* k . For every equivalence class $\{x_{c_1}, \dots, x_{c_l}\}$ of \sim_k , denote it by $[c]_{\sim_k}$, where $c = \min\{c_1, \dots, c_l\}$ – the smallest number of this set. Then, introduce a linear ordering \prec_k on the set $\{1, \dots, n\} / \sim_k$ by setting $[i]_{\sim_k} \prec_k [j]_{\sim_k}$ if $i < j$.

Definition 3.4. For every set X , define a function $\mathbf{p}: \biguplus_{n=1}^{\infty} X^n \rightarrow \biguplus_{n=1}^{\infty} X^n$ as follows. Let $\mathbf{x} = (x_1, \dots, x_n)$ in X^n be an n -tuple of equivalence type k . Let $[s_1]_{\sim_k}, \dots, [s_{n_k}]_{\sim_k}$ be the \sim_k -equivalence classes such that $[s_i]_{\sim_k} \prec_k [s_j]_{\sim_k}$ if and only if $i < j$. Then, put $\mathbf{p}(\mathbf{x}) := (x_{s_1}, \dots, x_{s_{n_k}})$. That is, the function \mathbf{p} removes an element from a tuple if it is not its first occurrence.

Example 3.5. Consider a 3-tuple $\mathbf{t} = (y, x, x)$, it is associated with the equivalence relation \sim_4 from Figure 2, the equivalence classes of \sim_4 on the set $\{1, 2, 3\}$ are $[1]_{\sim_4} = \{1\}$ and

$[2]_{\sim_4} = \{2, 3\}$. Then, $\mathbf{p}(\mathbf{t}) = (y, x)$ because y is on the first coordinate of \mathbf{t} , x is on the second, and $[1]_{\sim_4} \prec_4 [2]_{\sim_4}$.




FIGURE 2. All the 5 equivalence relations on a 3-element set.

Constructing the MMSN sentence. The MMSN sentence Φ' has the same existential signature σ as Φ , however the input signature τ' is not the same. Signature τ' has B_n -many relation symbols R_1, \dots, R_{B_n} . Each R_i is associated with one of the B_n equivalence relations \sim_k on the set $\{1, \dots, n\}$, the arity of R_k equals to the number n_k of equivalence classes of \sim_k . For example, if $n = 3$, then τ' contains $B_3 = 5$ relation symbols $R_1(\cdot, \cdot, \cdot), R_2(\cdot, \cdot, \cdot), R_3(\cdot, \cdot, \cdot), R_4(\cdot, \cdot, \cdot), R_5(\cdot, \cdot, \cdot)$, see Figure 2. We now describe how to construct the MMSN sentence Φ' from Φ .

- Firstly, remove all inequalities from Φ .
- Then, replace every τ -atom $R(\mathbf{x})$ with $R_k(\mathbf{p}(\mathbf{x}))$, where k is the equivalence type of \mathbf{x} .
- Finally, in order to simulate the inequalities that we have removed, we forbid the same variable to appear within τ' -atoms more than once. For every R_k in τ' and for every $i < j$ in $[n_k]$, add to Φ' the following:

$$\neg R_k(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_{j-1}, x, x_{j+1}, \dots, x_{n_k}). \quad (3.1)$$

Proof of Theorem 3.1. Let Φ' be obtained from Φ as above. Let \mathfrak{A} be a τ -structure. Firstly, we will construct a τ' -structure \mathfrak{A}' such that $\mathfrak{A} \models \Phi$ if and only if $\mathfrak{A}' \models \Phi'$. The structure \mathfrak{A}' has the same domain $A' = A$. For every R_k in τ' and for every tuple $\mathbf{a} = (a_1, \dots, a_n)$ of equivalence type k :

$$R_k^{\mathfrak{A}'}(\mathbf{p}(\mathbf{a})) \longleftrightarrow R^{\mathfrak{A}}(\mathbf{a}). \quad (3.2)$$

For every τ' -structure there exists a τ -structure \mathfrak{A} satisfying (3.2) if and only if all conditions of the type from (3.1) hold. Otherwise, such a structure does not satisfy Φ' , and we can output a fixed NO instance of Φ which always exists by Remark 2.3.

By construction, \mathfrak{A}' satisfies every negated conjunct from (3.1). Let \mathfrak{A}^σ and \mathfrak{A}'^σ be σ -expansions such that, for every X in σ and a in A , $X^{\mathfrak{A}^\sigma} = X^{\mathfrak{A}'^\sigma}$.

Suppose that \mathfrak{A}^σ does not satisfy the first-order part of Φ ; then there exist a negated conjunct $\neg\phi_i(\mathbf{x})$ of Φ and a tuple \mathbf{a} of A^* such that $\mathfrak{A}^\sigma \models \phi_i(\mathbf{a})$. For every σ -atom $X(a)$ of $\phi_i(\mathbf{a})$, we have $\mathfrak{A}^\sigma \models X(a)$, which implies that $\mathfrak{A}'^\sigma \models X(a)$. For every τ -atom $R(\mathbf{b})$ of $\phi_i(\mathbf{a})$, we have $\mathfrak{A}^\sigma \models R(\mathbf{b})$, which implies that $\mathfrak{A}'^\sigma \models R_k(\mathbf{p}(\mathbf{b}))$, by (3.2). Thus, $\mathfrak{A}'^\sigma \models \phi_i'(\mathbf{a})$. The other direction is similar. \square

In the following, we show that the class MMSN_{\oplus} strictly contains MMSN.

Proposition 3.6. *The class of τ -sentences in MMSN_{\oplus} is strictly more expressive than the class of τ -sentences in MMSN.*

Proof. Let $\tau = \{E(\cdot, \cdot)\}$ and let Φ be the sentence $\forall x, y \neg(E(x, y) \wedge x \neq y)$. This sentence is in MMSNP_{\oplus} , however it is not logically equivalent to any MMSNP sentence because $\text{fm}(\Phi)$ is not closed under inverse homomorphisms. \square

Remark 3.7. While adding guarded inequalities makes MMSNP strictly more expressive, it preserves GMSNP. Let Φ be a τ -sentence in GMSNP with guarded inequalities. Firstly, introduce a new existential binary relation Ineq . Then, for every $R \in \tau$ of arity k and for every $i \in [k]$, add to Φ the negated conjunct $\neg(\text{Ineq}(x_i, x_i) \wedge R(x_1, \dots, x_k))$. Finally, replace every inequality $x \neq y$ with $\text{Ineq}(x, y)$. The resulting sentence will be both in GMSNP and logically equivalent to Φ .

3.2. Guarded Monotone SNP over one-element signatures. In this section, we prove the following theorem.

Theorem 3.8. *For every finite relational signature τ there exists a relational signature τ_1 that consists of a single relation symbol such that the class of GMSNP τ -sentences has a dichotomy if and only if the class of τ_1 -sentences in GMSNP has a dichotomy.*

The single relation symbol is obtained by concatenating the original relation symbols, so its arity will be the sum of the arities of the original relations. The new relation P will hold on a tuple $(\mathbf{a}_1, \dots, \mathbf{a}_t)$ if and only if R_i holds on \mathbf{a}_i for all $i \in [t]$.

Why the obvious strategy fails. It is not possible to apply this transformation right away for two reasons:

- (a) There exists $R \in \tau$ such that the original sentence rejects all input structures that contain R -tuples, that is, the sentence contains the conjunct $\neg R(\mathbf{x})$. Applying the transformation to an input structure, where R is empty, returns a structure, where P is empty. So, the resulting input structure will be trivial while the original is not.
 - (b) There exists $R \in \tau$ and a forbidden conjunct of the original sentence Φ that does not contain R -atoms. We cannot apply the transformation to such a conjunct because the resulting conjunct would not contain any P -atoms, so it will become an empty constraint.
- For these two reasons, the proof consists of several stages. In the first stage, we check if Φ holds in a structure consisting of a single R -tuple, for R in τ . If not, then we remove this relation symbol because we need every relation of an input structure to be allowed to be non-empty. In the second stage, we enrich every negated conjunct with τ -atoms so that every relation symbol of τ is present in every conjunct. In the last stage, we construct a suitable sentence over a signature containing a single relation symbol.

A relational structure is called *connected* if it is not a disjoint union of two relational structures. An SNP τ -sentence Φ in the standard form is called *connected* if, for every negated conjunct $\neg\phi(\mathbf{x})$ of Φ , the $(\tau \cup \sigma)$ -structure \mathfrak{C}_ϕ is connected, where the domain of \mathfrak{C}_ϕ is the set of variables \mathbf{x} and, for every $R \in \tau \cup \sigma$, $R^{\mathfrak{C}_\phi} = \{\mathbf{y} \mid R(\mathbf{y}) \text{ is a non-negated atom of } \phi\}$.

The following lemma follows from the discussion after Proposition 1 in [BKS20].

Lemma 3.9. *The following statements are equivalent.*

- (1) *For every τ -sentence Φ in GMSNP, $\text{fm}(\Phi)$ is either in P or NP-complete.*
- (2) *For every connected τ -sentence Φ in GMSNP, $\text{fm}(\Phi)$ is either in P or NP-complete.*

The following lemma is Proposition 1.4.11 in [Bod21].

Lemma 3.10. *Let Φ be an SNP sentence. Then, $\text{fm}(\Phi)$ is closed under disjoint unions if and only if Φ is logically equivalent to a connected SNP sentence.*

By Lemma 3.9, we can consider only connected τ -sentences Φ in GMSNP and, by Lemma 3.10, $\text{fm}(\Phi)$ is closed under disjoint unions.

First stage: removing unnecessary relation symbols. For every k -ary R in τ , let \mathfrak{B}_R be the τ -structure that consists of one R -tuple $\mathbf{b}_R = (b_{R,1}, \dots, b_{R,k})$. That is, (a) \mathfrak{B}_R has the domain $B_R = \{b_{R,1}, \dots, b_{R,k}\}$, (b) $R^{\mathfrak{B}_R} = \{\mathbf{b}_R\}$, and (c) for every other $S \in \tau$, $S^{\mathfrak{B}_R} = \emptyset$. Denote by $\Phi^{\mathcal{R}}$ a connected $(\tau \setminus \{R\})$ -sentence that is obtained from Φ by removing all negated conjuncts that contain an R -atom.

Lemma 3.11. *If there exists R in τ such that a connected GMSNP τ -sentence Φ does not satisfy \mathfrak{B}_R , then $\text{fm}(\Phi)$ is polynomial-time equivalent to $\text{fm}(\Phi^{\mathcal{R}})$.*

Proof. Suppose that $\mathfrak{B}_R \not\models \Phi$. Let a τ -structure \mathfrak{A} be an input instance of $\text{fm}(\Phi)$. If $R^{\mathfrak{A}} \neq \emptyset$, then $\mathfrak{B}_R \rightarrow \mathfrak{A}$, so, by Theorem 2.1, $\mathfrak{A} \not\models \Phi$, and we reduce it to some NO instance of $\text{fm}(\Phi^{\mathcal{R}})$. Suppose that $R^{\mathfrak{A}} = \emptyset$; then we reduce \mathfrak{A} to its $(\tau \setminus \{R\})$ -reduct $\mathfrak{A}^{\mathcal{R}}$. For every σ -expansion of \mathfrak{A} , any negated conjunct with an R -atom is always satisfied, as $R^{\mathfrak{A}} = \emptyset$. Thus, $\mathfrak{A} \models \Phi$ if and only if $\mathfrak{A}^{\mathcal{R}} \models \Phi^{\mathcal{R}}$. For τ -structures \mathfrak{A} such that $R^{\mathfrak{A}} = \emptyset$, the correspondence $\mathfrak{A} \longleftrightarrow \mathfrak{A}^{\mathcal{R}}$ is one-to-one. This means that the two problems are polynomial-time equivalent. \square

A symbol $R \in \tau$ is called *redundant* for a connected GMSNP τ -sentence Φ if Φ rejects all input structures that contain R -tuples. Repeatedly applying Lemma 3.11, we remove all redundant symbols and prove that, for every connected GMSNP sentence with redundant symbols there exists a polynomial-time equivalent connected GMSNP sentence without redundant symbols.

Second stage: enriching negated conjuncts with τ -atoms. Call a GMSNP sentence Φ *enriched* if, for every negated conjunct $\neg\phi_i$ of Φ and every $R \in \tau$, $\neg\phi_i$ contains at least one positive R -atom. We now show, for every connected Φ without redundant symbols, how to obtain an enriched sentence Φ' such that the problem $\text{fm}(\Phi')$ is polynomial-time equivalent to $\text{fm}(\Phi)$. Replace every negated conjunct $\neg\phi_i$ of Φ with $\neg\phi'_i := \neg(\phi_i \wedge \bigwedge_{R \in \tau} R(\mathbf{x}_R))$, where each \mathbf{x}_R is a tuple of new variables. Observe that Φ' is not connected anymore. However, Φ is more restrictive than Φ' , as we enrich negated conjuncts of Φ with τ -atoms. That is, for every τ -structure \mathfrak{A} , $\mathfrak{A} \models \Phi$ implies $\mathfrak{A} \models \Phi'$. Below, we check that $\text{fm}(\Phi)$ and $\text{fm}(\Phi')$ are polynomial-time equivalent.

Lemma 3.12. *For every connected GMSNP τ -sentence Φ without redundant symbols there exists an enriched τ -sentence Φ' such that $\text{fm}(\Phi)$ is polynomial-time reducible to $\text{fm}(\Phi')$.*

Proof. For every τ -structure \mathfrak{A} , let \mathfrak{A}' be the disjoint union $\mathfrak{A} \uplus (\bigsqcup_{R \in \tau} \mathfrak{B}_R)$. Suppose that $\mathfrak{A} \models \Phi$. By Lemma 3.11, we assume that, for R in τ , $\mathfrak{B}_R \models \Phi$. Then, by Lemma 3.10, $\mathfrak{A}' \models \Phi$, as \mathfrak{A}' is the disjoint union of structures satisfying Φ . But then, as Φ' is less restrictive than Φ , we have $\mathfrak{A}' \models \Phi'$.

Suppose that $\mathfrak{A}' \not\models \Phi'$; then there exists a σ -expansion \mathfrak{A}'^σ that satisfies the first-order part ϕ' of Φ' . Define the σ -expansion \mathfrak{A}^σ as the substructure of \mathfrak{A}'^σ induced on A . Suppose that there exist a negated conjunct $\neg\phi_i(\mathbf{x}_i)$ of Φ and an assignment that makes it false in \mathfrak{A}^σ . But then we can take the negated conjunct $\neg(\phi_i \wedge \bigwedge_{R \in \tau} R(\mathbf{x}_R))$ of Φ' , and assign \mathbf{b}_R to

each new tuple of variables \mathbf{x}_R . Under this assignment the conjunct is false in \mathfrak{A}'^σ , this is a contradiction. \square

Lemma 3.13. *For every connected GMSNP τ -sentence Φ without redundant symbols there exists an enriched τ -sentence Φ' such that $\text{fm}(\Phi')$ is polynomial-time reducible to $\text{fm}(\Phi)$.*

Proof. Consider a τ -structure \mathfrak{A} that is an input instance of $\text{fm}(\Phi')$. Suppose that there exists R in τ such that $R^\mathfrak{A} = \emptyset$. Then, $\mathfrak{A} \models \Phi'$ because Φ' is enriched. We reduce it to some fixed YES instance of $\text{fm}(\Phi)$.

Suppose that, for every R in τ , $R^\mathfrak{A} \neq \emptyset$. Firstly, $\mathfrak{A} \models \Phi$ implies $\mathfrak{A} \models \Phi'$ because Φ is more restrictive than Φ' . Suppose that $\mathfrak{A} \not\models \Phi$; then, for every σ -expansion \mathfrak{A}^σ of \mathfrak{A} there exists a negated conjunct $\neg\phi_i(\mathbf{x}_i)$ of Φ and a tuple \mathbf{a}_i of elements of \mathfrak{A} such that $\mathfrak{A}^\sigma \models \phi_i(\mathbf{a}_i)$. But then, take the negated conjunct $\neg(\phi_i(\mathbf{x}_i) \wedge \bigwedge_{R \in \tau} R(\mathbf{x}_R))$, assign \mathbf{a}_i to \mathbf{x}_i and assign \mathbf{a}_R to \mathbf{x}_R such that, for every R in τ , $\mathbf{a}_R \in R^\mathfrak{A}$. At least one such \mathbf{a}_R exists as $R^\mathfrak{A} \neq \emptyset$. But then $\mathfrak{A}^\sigma \models (\phi_i(\mathbf{a}_i) \wedge \bigwedge_{R \in \tau} R(\mathbf{a}_R))$. This implies that $\mathfrak{A} \models \Phi'$. \square

Third stage: concatenation. In this part, we construct a τ_1 -sentence Φ^1 that is polynomial-time equivalent to Φ . By Lemmas 3.11, 3.12, and 3.13, we can assume that Φ is already enriched and has no redundant symbols. Set $\tau_1 := \{P\}$, where $\text{arity}(P) = \sum_{R \in \tau} \text{arity}(R)$. For two tuples $\mathbf{x} = (x_1, \dots, x_n)$, $\mathbf{y} = (y_1, \dots, y_m)$, set $(\mathbf{x}, \mathbf{y}) := (x_1, \dots, x_n, y_1, \dots, y_m)$. The tuple $(\mathbf{x}_1, \dots, \mathbf{x}_t)$ is defined similarly.

Let R_1, \dots, R_t be the relation symbols that constitute the signature τ . For every negated conjunct $\neg\phi_i$ of Φ , we transform it to a negated conjunct $\neg\phi_i^1$ of Φ^1 as follows. The σ -atoms of ϕ_i^1 are the same as in ϕ_i . For every τ -atom $R_j(\mathbf{x})$ of ϕ_i , add to ϕ_i^1 a τ_1 -atom $P(\mathbf{y}_1, \dots, \mathbf{y}_t)$, where $\mathbf{y}_j = \mathbf{x}$ and all other variables of this atomic formula are new and they are used only there. In the rest, we check the equivalence of the problems $\text{fm}(\Phi)$ and $\text{fm}(\Phi^1)$.

Lemma 3.14. *$\text{fm}(\Phi)$ is polynomial-time reducible to $\text{fm}(\Phi^1)$.*

Proof. For every τ -structure \mathfrak{A} , we will construct a τ_1 -structure \mathfrak{A}_1 such that $\mathfrak{A} \models \Phi$ if and only if $\mathfrak{A}_1 \models \Phi^1$. The structures have the same domain: $A_1 = A$. The relation $P^{\mathfrak{A}_1}$ is defined by the relations of τ as follows:

$$P^{\mathfrak{A}_1}(\mathbf{a}_1, \dots, \mathbf{a}_t) \longleftrightarrow R_1^{\mathfrak{A}}(\mathbf{a}_1) \wedge \dots \wedge R_t^{\mathfrak{A}}(\mathbf{a}_t). \quad (3.3)$$

Suppose that there exists R in τ such that $R^\mathfrak{A} = \emptyset$; then $P^{\mathfrak{A}_1} = \emptyset$ and $\mathfrak{A}_1 \models \Phi^1$. But also $\mathfrak{A} \models \Phi$, as Φ is enriched. So, we can now assume that $R^\mathfrak{A} \neq \emptyset$ for every R in τ .

Let \mathfrak{A}^σ and \mathfrak{A}_1^σ be two σ -expansions of \mathfrak{A} and \mathfrak{A}_1 such that, for every X in σ , $X^{\mathfrak{A}^\sigma} = X^{\mathfrak{A}_1^\sigma}$. It suffices to show that $\mathfrak{A}^\sigma \models \forall \mathbf{x} \phi(\mathbf{x})$ if and only if $\mathfrak{A}_1^\sigma \models \forall \mathbf{x}, \mathbf{y} \phi^1(\mathbf{x}, \mathbf{y})$, where \mathbf{y} represents variables that are added during the construction of Φ^1 .

Suppose that $\mathfrak{A}^\sigma \not\models \forall \mathbf{x} \phi(\mathbf{x})$; then there exist a negated conjunct $\neg\phi_i(\mathbf{x}_i)$ of Φ and a tuple $\mathbf{a}_i \in A^{\mathbf{x}_i}$ such that $\mathfrak{A}^\sigma \models \phi_i(\mathbf{a}_i)$. Take the i th negated conjunct $\neg\phi_i^1(\mathbf{x}_i, \mathbf{y}_i)$ of Φ^1 and assign to \mathbf{x}_i exactly the same tuple $\mathbf{a}_i \in A^{\mathbf{x}_i}$. As Φ is enriched, the conjunct $\phi_i(\mathbf{x}_i)$ contains an atom $R_j(\mathbf{z}_j)$, for every $R_j \in \tau$. For every $j \in [t]$, let $\mathbf{b}_j \in A^{\mathbf{z}_j}$ be the restriction of \mathbf{a}_i on \mathbf{z}_j . Let $P(\mathbf{y}_1, \dots, \mathbf{y}_t)$ be a τ_1 -atom that contains a tuple of variables \mathbf{y}_j that is not contained in \mathbf{x}_i and is thus unassigned. By construction, all the variables of \mathbf{y}_j appear in ϕ^1 just once. Therefore, we can assign the tuple \mathbf{b}_j to every such tuple \mathbf{y}_j , for every $j \in [t]$. Equation (3.3) implies that $\mathfrak{A}_1^\sigma \not\models \forall \mathbf{x}, \mathbf{y} \phi^1(\mathbf{x}, \mathbf{y})$.

For the other direction, suppose that $\mathfrak{A}_1^\sigma \not\models \forall \mathbf{x}, \mathbf{y} \phi^1(\mathbf{x}, \mathbf{y})$; then there exists a negated conjunct $\neg\phi_i^1(\mathbf{x}_i, \mathbf{y}_i)$ and tuples $\mathbf{a}_i, \mathbf{b}_i$ of elements of A such that $\mathfrak{A}_1^\sigma \models \phi_i^1(\mathbf{a}_i, \mathbf{b}_i)$. Take the

negated conjunct $\neg\phi_i(\mathbf{x}_i)$ of Φ associated with $\neg\phi_i^1$ and assign \mathbf{a}_i to \mathbf{x}_i . We have $\mathfrak{A}^\sigma \models \phi_i(\mathbf{a}_i)$, by (3.3). \square

Lemma 3.15. $\text{fm}(\Phi^1)$ is polynomial-time reducible to $\text{fm}(\Phi)$.

Proof. For every τ_1 -structure \mathfrak{B}_1 , we construct a τ -structure \mathfrak{B} . They have the same domain B , and every τ -relation $R_j^\mathfrak{B}$ is defined as follows:

$$R_j^\mathfrak{B}(\mathbf{x}_j) \longleftrightarrow \exists \mathbf{x}_1, \dots, \mathbf{x}_{j-1}, \mathbf{x}_{j+1}, \dots, \mathbf{x}_t \ P^{\mathfrak{B}_1}(\mathbf{x}_1, \dots, \mathbf{x}_t). \quad (3.4)$$

Similarly as above, we consider two σ -expansions \mathfrak{B}^σ and \mathfrak{B}_1^σ such that, for all X in σ , $X^{\mathfrak{B}^\sigma} = X^{\mathfrak{B}_1^\sigma}$.

Suppose that $\mathfrak{B}_1^\sigma \not\models \forall \mathbf{x}, \mathbf{y} \ \phi^1(\mathbf{x}, \mathbf{y})$; then there exists a negated conjunct $\neg\phi_i^1(\mathbf{x}_i, \mathbf{y}_i)$ and assignments $\mathbf{a}_i \in B^{\mathbf{x}_i}$ and $\mathbf{b}_i \in B^{\mathbf{y}_i}$ such that $\mathfrak{B}_1^\sigma \models \phi_i^1(\mathbf{a}_i, \mathbf{b}_i)$. Take the i th negated conjunct $\neg\phi_i(\mathbf{x}_i)$ of Φ and assign \mathbf{a}_i to \mathbf{x}_i . As ϕ_i and ϕ_i^1 have the same σ -part, \mathfrak{B}^σ satisfies every σ -atom and negated σ -atom of ϕ_i . Any τ -atom of ϕ_i is satisfied, by (3.4). So $\mathfrak{B}^\sigma \models \phi_i(\mathbf{a}_i)$.

For the other direction, suppose that there exist a negated conjunct $\neg\phi_i(\mathbf{x}_i)$ of Φ and a tuple $\mathbf{a}_i \in B^{\mathbf{x}_i}$ such that $\mathfrak{B}^\sigma \models \phi_i(\mathbf{a}_i)$. For every τ -atom $R_j(\mathbf{z})$ of ϕ_i , the conjunct ϕ_i^1 contains a τ_1 -atom $P(\mathbf{y}_1, \dots, \mathbf{y}_t)$, where $\mathbf{y}_j = \mathbf{z}$. Let $\mathbf{c} \in B^{\mathbf{z}}$ be the restriction of \mathbf{a}_i taken on \mathbf{z} . In particular, we have that $\mathfrak{B}^\sigma \models R_j(\mathbf{c})$. By (3.4), for every $k \in [t] \setminus j$, there exists $\mathbf{b}_k \in B^{\mathbf{y}_k}$ such that $\mathfrak{B}_1^\sigma \models P(\mathbf{b}_1, \dots, \mathbf{b}_t)$, where $\mathbf{b}_j = \mathbf{c}$. So, for all $\mathbf{y}_1, \dots, \mathbf{y}_t$ except for $\mathbf{y}_j = \mathbf{z}$, we assign $\mathbf{b}_1, \dots, \mathbf{b}_t$ to them. After doing this procedure for every τ -atom of the conjunct ϕ_i^1 , all τ -atoms of ϕ_i^1 will be satisfied. All σ -atoms and negated σ -atoms of ϕ_i are also satisfied because the σ -relations of \mathfrak{B}_1^σ are the same as in \mathfrak{B}^σ . This implies that $\mathfrak{B}_1^\sigma \not\models \forall \mathbf{x}, \mathbf{y} \ \phi^1(\mathbf{x}, \mathbf{y})$ and we are done. \square

Proof of Theorem 3.8. By Lemma 3.9, we can assume that a given GMSNP τ -sentence Φ is connected. By Lemma 3.11, we can assume that Φ has no redundant symbols. Let Φ' be constructed from Φ as in the *second stage*. By Lemmas 3.12 and 3.13, $\text{fm}(\Phi)$ and $\text{fm}(\Phi')$ are polynomial-time equivalent. Let Φ^1 be constructed from Φ' as in the *third stage*. By Lemmas 3.14 and 3.15, $\text{fm}(\Phi')$ and $\text{fm}(\Phi^1)$ are polynomial-time equivalent. \square

3.3. Matrix Partitions. In this subsection, we define the classes MPART_\oplus and MPART and prove Theorem 3.17.

Definition 3.16. An SNP sentence Φ in standard form is in MPART_\oplus if (a) each α_i is either a conjunction only of non-negated τ -atoms or a conjunction only of negated τ -atoms, (b) each β_i is a conjunction of unary σ -atoms or negated unary σ -atoms, and (c) for every inequality $x_j \neq x'_j$ of each ϵ_i there exists a τ -atom or a negated τ -atom in α_i that contains both x_j and x'_j .

The class MPART_\oplus contains both Matrix Partition problems and MMSNP_\oplus from Subsection 3.1 but it is not contained in monadic SNP without inequality as MPART_\oplus allows inequalities. Denote by MPART the fragment of MPART_\oplus where each ϵ_i is empty. Therefore, it will be contained in the gap between MMSNP and monadic SNP without inequality, see Figure 3.

Theorem 3.17. *The class MMSNP is strictly contained in MPART , and MPART is strictly contained in MPART_\oplus . Moreover, for every problem in MMSNP_\neq there exists a polynomial-time equivalent problem in MPART_\oplus .*

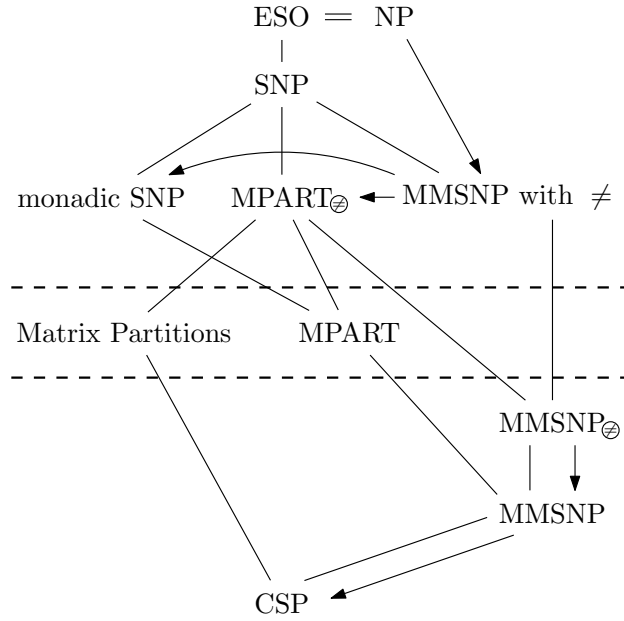


FIGURE 3. The classes introduced in Subsection 3.3 and their relation to other classes from this article. Dashed lines divide the figure in 3 parts: “dichotomy”, “unknown”, “no dichotomy”. Undirected edges stand for inclusions and directed edges stand for inclusions under polynomial-time reductions.

Proof. The containment in both cases follows from the definition. Below, we show, that the containment is strict.

Let $\Phi \in \text{MPART}$ be a sentence that accepts only complete directed graphs with loops:

$$\forall x, y \neg(\neg E(x, y)).$$

As $\text{fm}(\Phi)$ is not closed under inverse homomorphisms, Φ is not expressible in MMSNPN, by Theorem 2.1.

Let $\Phi \in \text{MPART}_\otimes$ be a sentence that accepts only an isolated vertex or an isolated vertex with a loop:

$$\forall x, y \neg(E(x, y) \wedge x \neq y) \wedge \neg(\neg E(x, y) \wedge x \neq y)$$

Suppose that a graph \mathfrak{G} (with loops allowed) satisfies a sentence Ψ in MPART. Take any vertex v of \mathfrak{G} and add to \mathfrak{G} a twin v' of v , i.e., the sets of vertices adjacent to v and to v' are the same. The resulting graph \mathfrak{G}' still satisfies Ψ as one can choose the same σ -relations for v' as were chosen for v in the σ -expansion \mathfrak{G}^σ witnessing $\mathfrak{G} \models \Psi$. Let $(\mathfrak{G}')^\sigma$ be this σ -expansion. If $(\mathfrak{G}')^\sigma$ does not satisfy the first-order part of Ψ , then, for some negated conjunct $\neg\psi(\mathbf{x})$ of Ψ and for some $\mathbf{g} \in (G')^x$, we have $(\mathfrak{G}')^\sigma \models \psi(\mathbf{g})$. Then, as v and v' are twins, we have $\mathfrak{G}^\sigma \models \psi(\mathbf{g}_{v' \rightarrow v})$, where the tuple $\mathbf{g}_{v' \rightarrow v}$ is obtained from \mathbf{g} by replacing each appearance of v' with v . In particular, if Ψ holds in the graph consisting of an isolated vertex, it will hold in the graph consisting of two isolated vertices. Therefore, Φ is not expressible in MPART.

Now, we prove that every problem in MMSNPN_\neq is polynomial-time equivalent to a problem in MPART_\otimes which implies no dichotomy for the latter class. Let Φ be of some signature τ . Add to τ a new binary relation symbol \mathbf{N} . The $\tau \cup \{\mathbf{N}\}$ -sentence Ψ is constructed

as follows. (a) Replace every atomic formula of Φ of the form $x \neq y$ with $N(x, y)$. (b) Add to the first-order part the two following negated conjuncts:

$$\forall x, y \neg N(x, x) \wedge \neg(\neg N(x, y) \wedge x \neq y).$$

Let \mathfrak{A} be a τ -structure and \mathfrak{A}' be its $\{N\}$ -expansion, where N is interpreted as the inequality relation. Then, associating \mathfrak{A}' with \mathfrak{A} we reduce Φ to Ψ . For the other direction, if N is not interpreted as the inequality in the input $\tau \cup \{N\}$ -structure, then Ψ rejects it, so we reduce it to some fixed NO instance of Φ . Otherwise, we reduce it to the τ -reduct of the input instance. \square

Although MPART_{\oplus} is related to MPART in the same way as MMSNP_{\oplus} is related to MMSNP , the idea of the proof of Theorem 3.1 do not transfer to them because now negated conjuncts may contain negated τ -atoms. This gives motivation to ask the following question.

Question 3.18. Does the class MPART have a dichotomy?

4. NP IS POLYNOMIAL-TIME EQUIVALENT TO MMSNP WITH INEQUALITY

In this section, we provide a detailed proof of the following theorem.

Theorem 4.1 [FV98]. *For every problem Ω in NP there exists a polynomial-time equivalent sentence Φ in MMSNP_{\neq} .*

The signature of the sentence Φ is denoted with τ . We will construct two mappings: (i) r_{easy} from the set of input strings of Ω to the set of τ -structures and (ii) r_{hard} in the opposite direction, such that

- for every input string s of Ω , s is accepted by Ω if and only if $r_{\text{easy}}(s) \models \Phi$, and
- for every τ -structure \mathfrak{A} , we have $\mathfrak{A} \models \Phi$ if and only if $r_{\text{hard}}(\mathfrak{A})$ is accepted by Ω .

First, without loss of generality we assume that Ω is decided by an *oblivious* Turing machine M . This property means that the head movement depends only on the size of the input string. We build a sentence Φ in MMSNP_{\neq} that shall operate on an input \mathfrak{A} which ought to describe the space-time diagram of an input string \mathbf{x} of M . We use two binary input predicates `succ` and `next` for space and time respectively, together with additional predicates to obtain a grid like structure.

Obliviousness allows us to check the computation with monadic predicates only, and it is not hard to prove that there exists a polynomial-time reduction r_{easy} from Ω to $\text{fm}(\Phi)$. The difficult part is the converse direction, where one has to deal appropriately with structures that are degenerate in the sense that they are not in the image of the reduction r_{easy} . It is precisely this part of the proof that is fully absent in [FV98]. Our construction of Φ ensures the following.

- If \mathfrak{A} is *over-complete* (i.e., contains too many relational tuples), then $\mathfrak{A} \not\models \Phi$. Here, \mathfrak{A} does not satisfy some first-order properties that are defined using “inequality”. This is the case, where the expressive power of “inequality” is used.
- If \mathfrak{A} is *incomplete* either locally – it is like a grid with holes – or globally – too small to simulate a complete run of M , then $\mathfrak{A} \models \Phi$. This is achieved by a monadic marking scheme of \mathfrak{A} by the sentence Φ and a relativisation of the verification of M ’s computation to the marked part.
- If \mathfrak{A} is in an *appropriate form* and can simulate a run of M , then $\mathfrak{A} \models \Phi$ if and only if M accepts.

Convention on Turing machines. Every Turing machine that we further consider is nondeterministic. We follow the convention of [Pap94] and assume that the input tape is one-way infinite. The alphabet is a set $\Sigma \cup \{\triangleright, \sqcup\}$, where \triangleright is the *first* symbol marking the left end of the tape, \sqcup is the *blank* symbol, and Σ contains all other symbols and may vary for different machines. The set of states of the machine is denoted by \mathcal{Q} . The machine input is a string of the form $\triangleright(\Sigma)^n\sqcup^\infty$ for $n \in \mathbb{N}$, i.e., the first symbol is followed by a string of symbols from Σ of length n , and the rest consists of blank symbols. The number n of elements from Σ in the input string is called the *size* of the input. The cell that contains \triangleright is called the *first cell*. If the head arrives in the first cell during the execution, then it will write \triangleright and move to the right. Before the execution starts, the head of the machine is at the first cell. We assume that every machine M halts in polynomial time in the size of the input. That is, for every machine M there exists $k > 0$ and a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that $f(n) \in O(n^k)$ and, for every input of size n , the machine M halts in at most $f(n)$ steps. Finally, we assume that every machine M that we consider is not trivial, i.e., there exists at least one input string which is accepted and at least one which is rejected. We fix a string s_Y which is accepted by M and a string s_N which is rejected.

4.1. Obliviousness. A Turing machine M is called *oblivious* if there exists a function $f: \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ such that, for every input \mathbf{x} of size $|\mathbf{x}| = n$, at any moment of time $t \in \mathbb{N}$, the head of the machine is at the $f(n, t)$ th cell. Oblivious Turing machines are introduced by Pippenger and Fischer in [PF79].

In order to ensure that an input instance \mathfrak{A} of a sentence in MMSNP_\neq can simulate the whole execution of M , we need to know (a) the position of its head at any moment of time and (b) the precise time when the machine will halt.

In the MMSNP_\neq sentence which is supposed to simulate the oblivious Turing machine, the head movement will be described by an existentially defined relation. In order to make the reductions r_{easy} and r_{hard} universal, we assume that the head movement in every oblivious Turing machine is the same and similar to the one from Theorem 3 in [PF79]. That is, for every input of size n , the head first does $n + 1$ steps to the right until it reaches a cell with the blank symbol, then returns back to the first cell, then does $n + 2$ steps to the right, returns to the first cell again, and so on. When the machine halts, the head is at the first cell. Call an oblivious machine *back-and-forth* if its head movement is as described above. Say that the head of a back-and-forth machine does a *round trip* if it starts at the first cell, goes right until it scans a blank symbol, and returns back to the first cell.

The ideas from the proof of Theorem 3 in [PF79] imply the following statement.

Lemma 4.2. *For every machine M there exists an equivalent machine M' which is back-and-forth.*

Proof. For every transition of the machine M , the back-and-forth machine M' spends one round trip. Those transitions of M where the head goes from the i th cell to the $(i + 1)$ th cell are simulated when the head of M' is going from left to right. The transitions of M where the head goes from the i th cell to the $(i - 1)$ th cell are simulated when the head of M' is going in the opposite direction. \square

Say that a back-and-forth machine is *precise* if there exist $k \in \mathbb{N}$ and $g(n) \in O(n^k)$ such that, for every input string of length n , the head does precisely $g(n)$ round trips as above before it halts. If the value $g(n)$ is known, then one can obtain the exact halting time $f(n)$

as follows. To do the first round trip on an input of size n , the head will need to make $(n+1) + (n+1)$ steps. For each next round trip, the number of steps will increase by 2, so the last round trip is done in $(n + g(n)) + (n + g(n))$ steps. Therefore, the total execution time would be

$$f(n) := g(n)(g(n) + 2n + 1)$$

The preciseness condition is necessary for us because we need to know whether a given input space-time diagram is large enough to simulate the execution of the machine. The goal of Subsection 4.1 is to prove the following.

Theorem 4.3. *For every machine M there exists a polynomial-time equivalent machine M' which is precise.*

Proof. By Lemma 4.2, we may assume that M is back-and-forth. The goal is to construct a precise machine M' that simulates M .

Recall that, by convention, there exist $k > 0$ and $f \in O(n^k)$ such that, for every input of size n , the machine M halts in at most $f(n)$ steps. Therefore, there exists $g \in O(n^k)$ such that M does at most $g(n)$ round trips. We may assume without loss of generality that $g(n) = cn^k$, where $c, k \in \mathbb{N}$. It suffices to modify M so that it does precisely $g'(n)$ round trips, where $g'(n) \in O(n^k)$. This can be done by adding some procedure that has running time a little bit greater than cn^k and that requires the same amount of time for all input strings of the same size.

Recall that the input string of M is of the form $\triangleright(\Sigma)^n \sqcup^\infty$. We want it to be of the form $\triangleright 0^m(\Sigma)^n \sqcup^\infty$. That is, we want to add a string consisting of 0s of length m between the first cell and the input string so that the machine M' would write on this string all numbers from 0 to 2^m in binary encoding. We want m to be at least $\lceil \log_2 cn^k \rceil$ and that $2^m \in O(n^k)$. We call this process the *binary enumeration*.

The alphabet of the new machine M' is $\{\triangleright, \sqcup, 0, 1\} \cup (\Sigma \times \{0, 1\})$, where $\{\triangleright, \sqcup\} \cup \Sigma$ is the alphabet of M . At the first round trip, the machine makes the second coordinate equal to the parity of the cell, i.e., for every $i \in [n]$, if a cell contains a symbol (s, t) , $s \in \Sigma$, $t \in \{0, 1\}$, then it writes $(s, i \bmod 2)$ instead. After the first round trip, the string of the second coordinate symbols has the form $(01)^{n/2}$. Then, at each next round trip, the machine M' decides whether it preserves or changes the second-coordinate symbol depending on one of the two states: Q_+ , Q_- , which are defined as follows. The machine M' starts a round trip having the state Q_+ . If, right after reading “0”, it reads “1”, M' changes to the state Q_- until it reads “01” for the next time. If M' has the state Q_- , then it rewrites “0” with “1” and vice versa, otherwise, it preserves the second coordinate symbol. Therefore, after the second round trip, the second coordinate of the string will have the form $(0011)^{n/4}$, after the third, it will look like $(00001111)^{n/8}$, and so on. The procedure runs until the second coordinate string has the form 0^n . It requires $\lceil \log_2 n \rceil + 1$ round trips. Moreover, at each round trip, the head must move the input string one position to the right. This can be done during the movement from left to right: the machine reads the symbol (s_i, t_i) from the i th cell, remembers it, writes the previously scanned symbol (s_{i-1}, t_{i-1}) , moves to the right, scans (s_{i+1}, t_{i+1}) and writes (s_i, t_i) instead, and so on.

If we repeat this procedure k times, then the machine M' will make $k \lceil \log_2 n \rceil + k$ round trips. There exists some constant $c' \in \mathbb{N}$ such that c' depends only on k and c and that $k \lceil \log_2 n \rceil + k + c' \geq \lceil \log_2 cn^k \rceil$. Notice that $2^{k \lceil \log_2 n \rceil + k + c'} \in O(n^k)$, so we can put m to be equal to $k \lceil \log_2 n \rceil + k + c'$. As c' depends only on k and c , we can move the first coordinate string to the right by c' cells.

Now we can assume that the input string looks like $\triangleright 0^m(\Sigma)^n \sqcup^\infty$. The machine M' can simulate the movement of the head of M and do the binary enumeration during the same round trip. By our assumption, the simulation of M finishes earlier than the binary enumeration. Therefore, we require the head of M' to repeat round trips until the binary enumeration is finished in order to keep it back-and-forth. The binary enumeration is done in $O(n^k)$ steps. The number of steps depends only on the size of the input, so the resulting machine is precise. \square

4.2. Construction of the space-time diagram. By Theorem 4.3, we can assume that the machine M is precise. Therefore, we can fix the functions $f(n)$ and $g(n)$ such that M halts precisely in time $f(n)$ and does precisely $g(n)$ round trips for every input of size n , where

$$f(n) := g(n)(g(n) + 2n + 1)$$

Let $\mathbf{s} \in \Sigma^n$ be an input string of M . In this section, we describe the relational structure \mathfrak{A} such that $\mathbf{r}_{\text{easy}}(\mathbf{s}) = \mathfrak{A}$. The structure \mathfrak{A} will have a form of a space-time diagram, where the horizontal axis represents the string of symbols, i.e., *space*, and the vertical axis represents *time*.

Domain of \mathfrak{A} . The domain A of \mathfrak{A} is the following disjoint union:

$$A = A_0^{\leftarrow} \uplus \biguplus_{i=1}^{g(n)} (A_i^{\rightarrow} \uplus A_i^{\leftarrow}) \uplus A_{g(n)+1}^{\rightarrow} \quad (4.1)$$

Here, the set A_0^{\leftarrow} describes the machine tape before the execution starts. For every $i \in [g(n)]$, the set A_i^{\rightarrow} represents the machine tape during the time when the head goes from the first cell to the end of the input string while the set A_i^{\leftarrow} does the same for the movement in the opposite direction.

We put $A_0^{\leftarrow} = [n+1] \times [1]$. For every $i \in [g(n)]$, we put $A_i^{\rightarrow} = [n+i] \times [n+i-2]$ and $A_i^{\leftarrow} = [n+i+1] \times [n+i-2]$, where the first coordinate denotes the number of columns and the second coordinate denotes the number of rows. Finally, $A_{g(n)+1}^{\rightarrow} = [1] \times [1]$ is a one-element set. See Figure 8 and Figure 4a for an illustration.

Notice, that the total number of rows in this grid is equal to

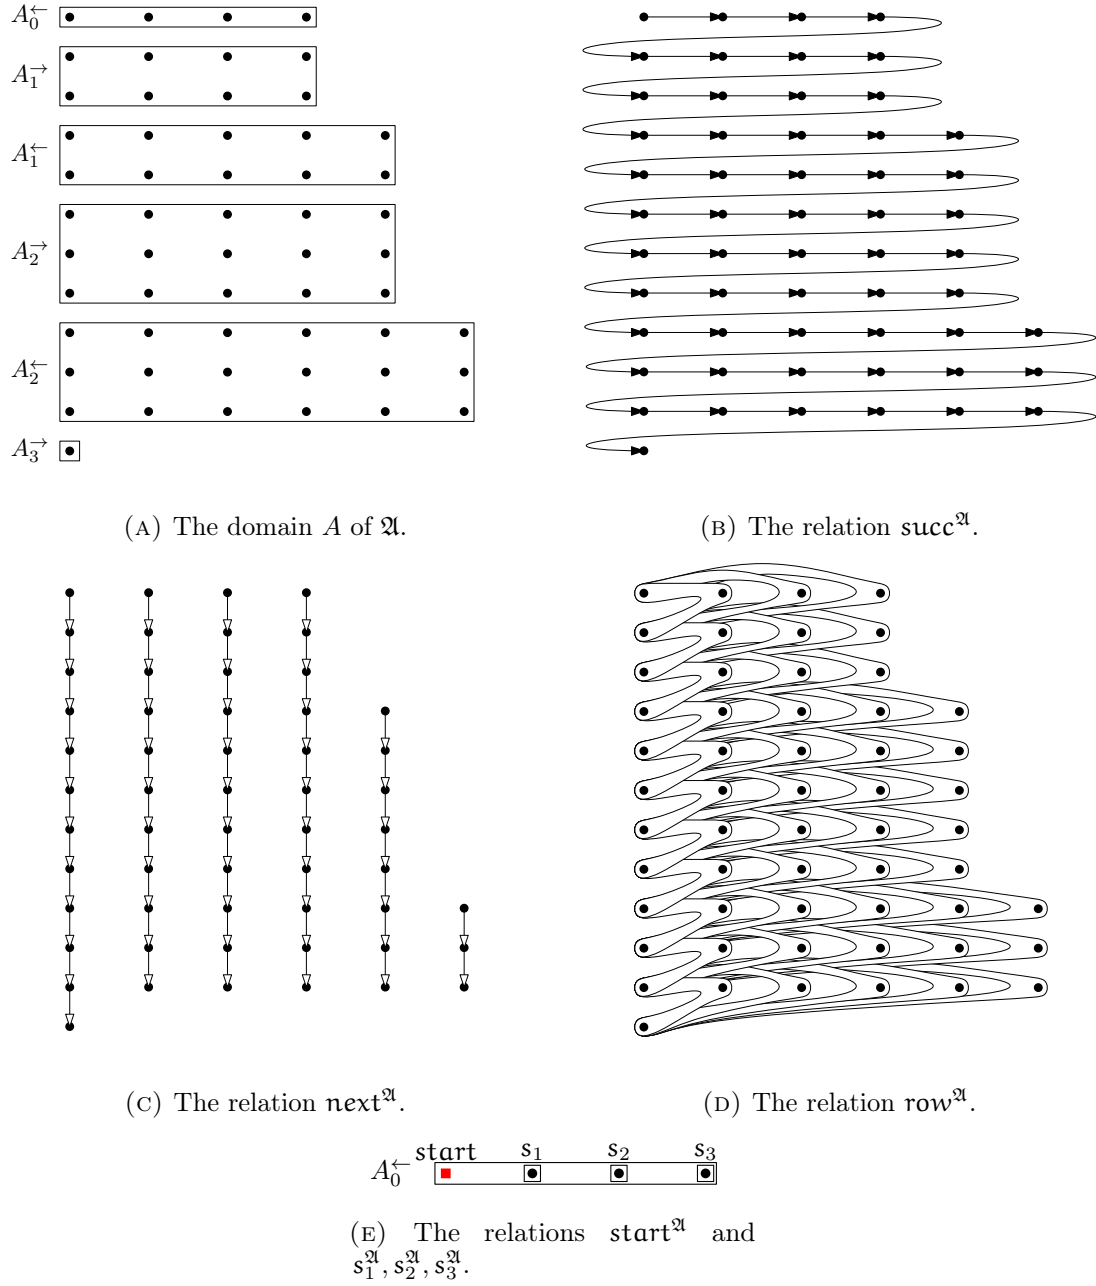
$$1 + 1 + \sum_{i=1}^{g(n)} 2(n+i-2) = (g(n))^2 + (2n-3)g(n) + 2$$

For every set of $\{A_0^{\leftarrow}, A_1^{\rightarrow}, A_1^{\leftarrow}, \dots, A_{g(n)+1}^{\rightarrow}\}$, let \mathbf{l} be the function that outputs the maximal value of the first component, that is,

$$\mathbf{l}(A_0^{\leftarrow}) = n+1 \text{ and, for } i \in [g(n)], \mathbf{l}(A_i^{\rightarrow}) = n+i, \mathbf{l}(A_i^{\leftarrow}) = n+i+1 \text{ and } \mathbf{l}(A_{g(n)+1}^{\rightarrow}) = 1$$

Similarly, define the function \mathbf{h} that outputs the maximal value of the second component:

$$\mathbf{h}(A_0^{\leftarrow}) = 1 \text{ and, for } i \in [g(n)], \mathbf{h}(A_i^{\rightarrow}) = \mathbf{h}(A_i^{\leftarrow}) = n+i-2 \text{ and } \mathbf{h}(A_{g(n)+1}^{\rightarrow}) = 1$$

FIGURE 4. The domain and the relations for a structure \mathfrak{A} from Example 4.4.

Relations of \mathfrak{A} . The structure \mathfrak{A} has the following relational signature:

$$\tau = \{s(\cdot) \mid s \in \Sigma \setminus \{\sqcup\}\} \uplus \{\text{start}(\cdot), \text{succ}(\cdot, \cdot), \text{next}(\cdot, \cdot), \text{row}(\cdot, \cdot, \cdot)\}$$

- The binary relation $\text{succ}^{\mathfrak{A}}$ denotes horizontal arcs. It links two neighbours in the same row and also it connects the last element of a row to the first element of the next row. Everywhere in this section, it is denoted by black-headed arcs. Formally, a

pair $(\mathbf{x}, \mathbf{y}) \in A^2$ is added to $\text{succ}^{\mathfrak{A}}$ if and only if there exists $i \in \{0, \dots, g(n)\}$ such that one of the following holds:

- $\mathbf{x}, \mathbf{y} \in A_i^{\rightarrow}$ and $x_1 + 1 = y_1$ and $x_2 = y_2$;
 - $\mathbf{x}, \mathbf{y} \in A_i^{\leftarrow}$ and $x_1 + 1 = y_1$ and $x_2 = y_2$;
 - $\mathbf{x}, \mathbf{y} \in A_i^{\rightarrow}$ and $x_1 = l(A_i^{\rightarrow})$ and $y_1 = 1$ and $x_2 + 1 = y_2$;
 - $\mathbf{x}, \mathbf{y} \in A_i^{\leftarrow}$ and $x_1 = l(A_i^{\leftarrow})$ and $y_1 = 1$ and $x_2 + 1 = y_2$;
 - $\mathbf{x} \in A_i^{\rightarrow}$ and $\mathbf{y} \in A_i^{\leftarrow}$ and $\mathbf{x} = (l(A_i^{\rightarrow}), h(A_i^{\rightarrow}))$ and $\mathbf{y} = (1, 1)$;
 - $\mathbf{x} \in A_i^{\leftarrow}$ and $\mathbf{y} \in A_{i+1}^{\rightarrow}$ and $\mathbf{x} = (l(A_i^{\leftarrow}), h(A_i^{\leftarrow}))$ and $\mathbf{y} = (1, 1)$.
- The binary relation $\text{next}^{\mathfrak{A}}$ denotes vertical arcs. It links two neighbours in the same column and is denoted by white-headed arcs elsewhere. Formally, a pair $(\mathbf{x}, \mathbf{y}) \in A^2$ is added to $\text{next}^{\mathfrak{A}}$ if and only if there exists $i \in \{0, \dots, g(n)\}$ such that one of the following holds:
 - $\mathbf{x}, \mathbf{y} \in A_i^{\rightarrow}$ and $x_1 = y_1$ and $x_2 + 1 = y_2$;
 - $\mathbf{x}, \mathbf{y} \in A_i^{\leftarrow}$ and $x_1 = y_1$ and $x_2 + 1 = y_2$;
 - $\mathbf{x} \in A_i^{\rightarrow}$ and $\mathbf{y} \in A_i^{\leftarrow}$ and $x_1 = y_1$ and $x_2 = h(A_i^{\rightarrow})$ and $y_2 = 1$;
 - $\mathbf{x} \in A_i^{\leftarrow}$ and $\mathbf{y} \in A_{i+1}^{\rightarrow}$ and $x_1 = y_1$ and $x_2 = h(A_i^{\leftarrow})$ and $y_2 = 1$.
 - The relation $\text{row}^{\mathfrak{A}}$ is ternary. It describes all the triples, where the first and the third coordinates are the leftmost elements of two consecutive rows while the second coordinate belongs to the same row as the first one. Formally, a triple $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \in A^3$ is added to $\text{row}^{\mathfrak{A}}$ if and only if the following holds:

$$x_1 = 1 \text{ and } \text{next}(\mathbf{x}, \mathbf{z}) \text{ and } y_1 > 1 \text{ and } x_2 = y_2.$$

- The unary relation $\text{start}^{\mathfrak{A}}$ represents the point of reference of the diagram. there exists a unique element in $\text{start}^{\mathfrak{A}}$ – the first element of the first row. It is denoted by a red square everywhere in this section. Formally, an element $\mathbf{x} \in A$ belongs to $\text{start}^{\mathfrak{A}}$ if and only if $\mathbf{x} \in A_0^{\leftarrow}$ and $\mathbf{x} = (1, 1)$.
- For every symbol $s \in \Sigma$, the unary relation $s^{\mathfrak{A}}$ captures all the cells of the input string \mathbf{s} which contain the symbol s . Formally, an element $\mathbf{x} \in A$ belongs to $s^{\mathfrak{A}}$ if and only if $\mathbf{x} \in A_0^{\leftarrow}$ and the $(x_1 - 1)$ th cell of \mathbf{s} contains the symbol s .

Example 4.4. Suppose that an input string $\mathbf{s} = (s_1 s_2 s_3)$ of the machine M has size 3 and that $g(3) = 2$. This means, that the machine M has to make two round trips and then halt. Then the domain of the τ -structure $\mathfrak{A} = \mathbf{r}_{\text{easy}}(\mathbf{s})$ is displayed on Figure 4a and is the following disjoint union

$$([4] \times [1]) \uplus ([4] \times [2]) \uplus ([5] \times [2]) \uplus ([5] \times [3]) \uplus ([6] \times [3]) \uplus ([1] \times [1])$$

The relations $\text{succ}^{\mathfrak{A}}$, $\text{next}^{\mathfrak{A}}$, and $\text{row}^{\mathfrak{A}}$ are exactly as on Figures 4b to 4d respectively. The Figure 4e displays the relation $\text{start}^{\mathfrak{A}}$ and the relations $s_1^{\mathfrak{A}}$, $s_2^{\mathfrak{A}}$, $s_3^{\mathfrak{A}}$ which represent the symbols of the string \mathbf{s} .

4.3. Construction of the sentence. The existential relations of Φ in MMSNP_{\neq} are as follows:

$$\sigma = \{I(\cdot), M(\cdot), H(\cdot)\} \uplus \{S(\cdot) \mid s \in \Sigma\} \uplus \{Q(\cdot) \mid q \in \mathcal{Q}\}$$

- I (*initial*) contains the elements of the first row, it is used in order to generate all M elements and to assign the initial symbols: $s(x) \rightarrow S(x)$.
- M (*marked*) determines the region, where the head movement is mandatory.
- H (*head*) contains all the positions of M 's head at every moment of time.




FIGURE 5. The rule (left) initiating the relations I and H (denoted by \square and by a circle, respectively) and the rule (right) that forces the relation I to spread. The I- and H-atoms in the head of each rule are coloured in green.

- Relations S describe the string symbols at any moment of time.
- Relations Q represent the current state of the machine.

Section 4.3 is divided into several paragraphs. Each of them describes a certain subroutine of the simulation of M .

Initiating. Here, we explain how to find the part of the input from which we later construct the input string of M and how the simulation of M is supposed to start.

The relations I and H are initiated by the **succ**-neighbour of the **start** if the structure induced on the **start**-element and its **succ**- and **next**-neighbours repeats the grid-structure from Section 4.2. Then, the relation I is forced to spread along the **succ**-arcs either until the end of the first row or until an element that is not contained in a relation s for $s \in \Sigma$. See (4.2), (4.3), and Figure 5 for an illustration. Although, in Eq. (4.2), the head contains a conjunction of two atomic formulas which is not allowed by MMSNP_{\neq} , this clause can be transformed to a conjunction of allowed clauses as follows:

$$A \rightarrow (B \wedge C) \equiv (A \rightarrow B) \wedge (A \rightarrow C) \equiv \neg(A \wedge \neg B) \wedge \neg(A \wedge \neg C)$$

$$\bigwedge_{s \in \Sigma} (\text{start}(x) \wedge \text{next}(x, z) \wedge \text{succ}(x, y) \wedge \text{row}(x, y, z) \wedge s(y) \rightarrow I(y) \wedge H(y)) \quad (4.2)$$

$$\bigwedge_{s \in \Sigma} (\text{start}(w) \wedge \text{next}(w, z) \wedge I(x) \wedge \text{succ}(x, y) \wedge \text{row}(w, y, z) \wedge s(y) \rightarrow I(y)) \quad (4.3)$$

If the machine M is at the state $q_0 \in Q$ at the start of the execution, then every element in I must be also in Q_0 :

$$I(x) \rightarrow Q_0(x) \quad (4.4)$$

If an element is forced to be in I, then it is contained in some input relation s . In this case, it is also forced to be in the existential relation S associated with s :

$$\bigwedge_{s \in \Sigma} (I(x) \wedge s(x) \rightarrow S(x)) \quad (4.5)$$

Determining the substructure where the simulation of M happens. During the reduction r_{hard} , we determine in polynomial-time, whether the input structure is sufficiently large to simulate the execution of M . For this reason, we have the relation M which describes the subset of input elements, where the relation H that simulates the head of the machine is forced to spread.

The relation M is initiated by the I -elements:

$$I(x) \rightarrow M(x) \quad (4.6)$$

Then M must spread along the vertical **next**-arcs from top to bottom according to the four rules displayed on Figure 6. The **next**-neighbour must be in M only if its **next**-predecessor is in M and if the structure locally looks like a grid from Section 4.2. Every rule of Figure 6 is also given as a logical formula

$$\mu_A(x, x_1, x_2, y, y_1, y_2, z) \rightarrow M(y_1), \quad \mu_B(x, x_1, x_2, x_3, y, y_1, y_2, y_3, z) \rightarrow M(y_2), \quad (4.7)$$

$$\mu_C(x, x_1, x_2, y, y_1, y_2, z) \rightarrow M(y_2), \quad \mu_D(x, x_1, x_2, y, y_1, y_2, y_3, z) \rightarrow M(y_2) \wedge M(y_3), \quad (4.8)$$

where

$$\begin{aligned} \mu_A(x, x_1, x_2, y, y_1, y_2, z) &\equiv \text{succ}(x, x_1) \wedge \text{succ}(x_1, x_2) \wedge \text{succ}(y, y_1) \wedge \text{succ}(y_1, y_2) \\ &\wedge \text{next}(x, y) \wedge \text{next}(y, z) \wedge \text{next}(x_1, y_1) \wedge \text{next}(x_2, y_2) \\ &\wedge \text{row}(x, x_1, y) \wedge \text{row}(x, x_2, y) \wedge \text{row}(y, y_1, z) \wedge \text{row}(y, y_2, z) \\ &\wedge M(x_1) \wedge M(x_2) \end{aligned} \quad (4.9)$$

$$\begin{aligned} \mu_B(x, x_1, x_2, x_3, y, y_1, y_2, y_3, z) &\equiv \text{succ}(x_1, x_2) \wedge \text{succ}(x_2, x_3) \wedge \text{succ}(y_1, y_2) \wedge \text{succ}(y_2, y_3) \\ &\wedge \text{next}(x, y) \wedge \text{next}(y, z) \wedge \text{next}(x_1, y_1) \wedge \text{next}(x_2, y_2) \wedge \text{next}(x_3, y_3) \\ &\wedge \text{row}(x, x_1, y) \wedge \text{row}(x, x_2, y) \wedge \text{row}(x, x_3, y) \\ &\wedge \text{row}(y, y_1, z) \wedge \text{row}(y, y_2, z) \wedge \text{row}(y, y_3, z) \\ &\wedge M(x_1) \wedge M(x_2) \wedge M(x_3) \wedge M(y_1) \end{aligned} \quad (4.10)$$

$$\begin{aligned} \mu_C(x, x_1, x_2, y, y_1, y_2, z) &\equiv \text{succ}(x_1, x_2) \wedge \text{succ}(x_2, y) \wedge \text{succ}(y_1, y_2) \wedge \text{succ}(y_2, z) \\ &\wedge \text{next}(x, y) \wedge \text{next}(y, z) \wedge \text{next}(x_1, y_1) \wedge \text{next}(x_2, y_2) \\ &\wedge \text{row}(x, x_1, y) \wedge \text{row}(x, x_2, y) \wedge \text{row}(y, y_1, z) \wedge \text{row}(y, y_2, z) \\ &\wedge \neg H(x_2) \wedge M(x_1) \wedge M(x_2) \wedge M(y_1) \end{aligned} \quad (4.11)$$

$$\begin{aligned} \mu_D(x, x_1, x_2, y, y_1, y_2, y_3, z) &\equiv \\ &\equiv \text{succ}(x_1, x_2) \wedge \text{succ}(x_2, y) \wedge \text{succ}(y_1, y_2) \wedge \text{succ}(y_2, y_3) \wedge \text{succ}(y_3, z) \\ &\wedge \text{next}(x, y) \wedge \text{next}(y, z) \wedge \text{next}(x_1, y_1) \wedge \text{next}(x_2, y_2) \\ &\wedge \text{row}(x, x_1, y) \wedge \text{row}(x, x_2, y) \wedge \text{row}(y, y_1, z) \wedge \text{row}(y, y_2, z) \wedge \text{row}(y, y_3, z) \\ &\wedge H(x_2) \wedge M(x_1) \wedge M(x_2) \wedge M(y_1) \end{aligned} \quad (4.12)$$

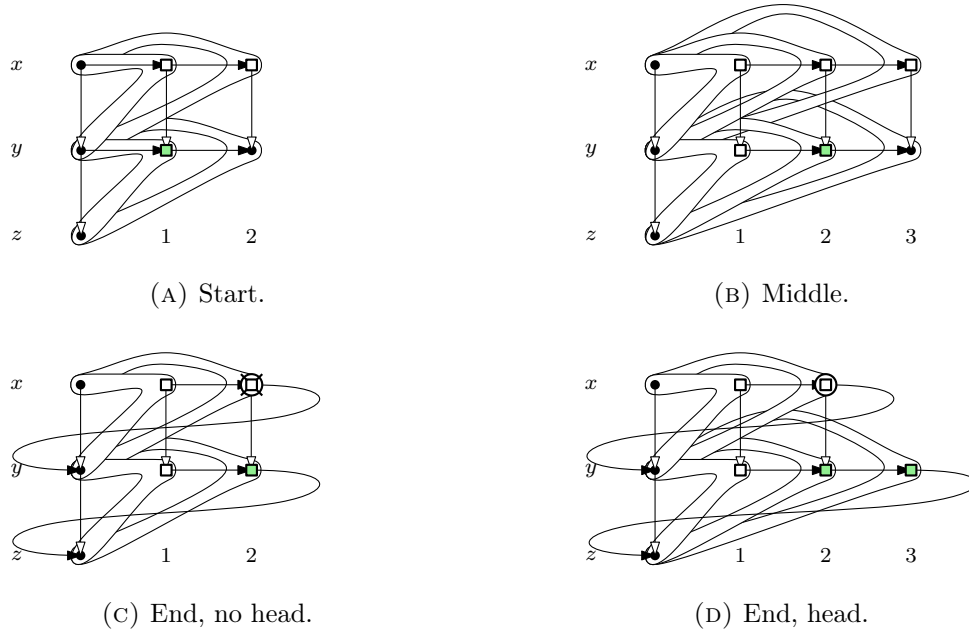


FIGURE 6. All the four cases when M (denoted by \square) is forced to spread down. Crossed white circle at Figure 6c stands for a negated H-atom. White circle at Figure 6d stands for a non-negated H-atom. The M -atoms in the head of each implication are coloured in green.

Simulating the transitions of the machine. Both families of relations $S(\cdot)$, for $s \in \Sigma$, and $Q(\cdot)$, for $q \in \mathcal{Q}$, partition the elements of the input.

$$\neg \left(\bigwedge_{s \in \Sigma} \neg S(x) \wedge M(x) \right) \wedge \bigwedge_{s, s' \in \Sigma} \neg (S(x) \wedge S'(x) \wedge M(x)) \quad (4.13)$$

$$\neg \left(\bigwedge_{q \in \mathcal{Q}} \neg Q(x) \wedge M(x) \right) \wedge \bigwedge_{q, q' \in \mathcal{Q}} \neg (Q(x) \wedge Q'(x) \wedge M(x)) \quad (4.14)$$

Moreover, no two elements of the same row can belong to distinct Q and Q' :

$$\bigwedge_{q, q' \in \mathcal{Q}} \neg (\text{row}(a, x, a') \wedge \text{row}(a, y, a') \wedge Q(x) \wedge Q'(y) \wedge M(x) \wedge M(y)). \quad (4.15)$$

After being initiated, the relation H spreads down according to the movement of the head of M . As the machine M is back-and-forth, the states of M that are not halting states can be divided in two groups: $\mathcal{Q}^{\rightarrow}$ and \mathcal{Q}^{\leftarrow} . If the head of the machine M is in a state from $\mathcal{Q}^{\rightarrow}$ and if it is not at the right end of the tape, then it keeps moving to the right. Similarly, if it is in a state from \mathcal{Q}^{\leftarrow} and if it is not at the left end of the tape, then it keeps moving to the left. Otherwise, the head is at the left or at the right end of the tape and it has to start moving in the opposite direction. Notice that, when the head of the machine M is at the first symbol \triangleright , it is already determined what it has to write in the cell and move to the right, so only the state can change. Similarly, when the head is at the blank symbol \sqcup , it will move to the left. Therefore, the transition function Δ of M can be described by the four functions

- (a) $\Delta_{\triangleright}: \mathcal{Q}^{\leftarrow} \rightarrow 2^{\mathcal{Q}^{\rightarrow}}$, when the head is at the first cell,
- (b) $\Delta_{\rightarrow}: \Sigma \times \mathcal{Q}^{\rightarrow} \rightarrow 2^{\Sigma \times \mathcal{Q}^{\rightarrow}}$, when the head is moving right,
- (c) $\Delta_{\sqcup}: \mathcal{Q}^{\rightarrow} \rightarrow 2^{\Sigma \times \mathcal{Q}^{\leftarrow}}$, when the head is at the right end of the tape,
- (d) $\Delta_{\leftarrow}: \Sigma \times \mathcal{Q}^{\leftarrow} \rightarrow 2^{\Sigma \times \mathcal{Q}^{\leftarrow}}$, when the head is moving left.

Let $\Delta_{\leftarrow}: \Sigma \times \mathcal{Q}^{\rightarrow} \rightarrow 2^{\Sigma \times \Sigma \times \mathcal{Q}^{\leftarrow}}$ return, for a pair $(s_0, q_0) \in \Sigma \times \mathcal{Q}^{\rightarrow}$, the set of all triples $(s_2, s_3, q_3) \in \Sigma \times \Sigma \times \mathcal{Q}^{\leftarrow}$ such that

$$\begin{aligned} \exists s_1 \in \Sigma, q_1 \in \mathcal{Q}^{\rightarrow}, q_2 \in \mathcal{Q}^{\leftarrow} \\ (s_1, q_1) \in \Delta_{\rightarrow}(s_0, q_0) \wedge (s_2, q_2) \in \Delta_{\sqcup}(q_1) \wedge (s_3, q_3) \in \Delta_{\leftarrow}(s_1, q_2) \end{aligned} \quad (4.16)$$

Similarly, let $\Delta_{\rightarrow}: \Sigma \times \mathcal{Q}^{\leftarrow} \rightarrow 2^{\Sigma \times \mathcal{Q}^{\rightarrow}}$ return, for a pair $(s_0, q_0) \in \Sigma \times \mathcal{Q}^{\leftarrow}$, the set of all pairs $(s_2, q_3) \in \Sigma \times \mathcal{Q}^{\rightarrow}$ such that

$$\begin{aligned} \exists s_1 \in \Sigma, q_1 \in \mathcal{Q}^{\leftarrow}, q_2 \in \mathcal{Q}^{\rightarrow} \\ (s_1, q_1) \in \Delta_{\leftarrow}(s_0, q_0) \wedge q_2 \in \Delta_{\triangleright}(q_1) \wedge (s_2, q_3) \in \Delta_{\rightarrow}(s_1, q_2) \end{aligned} \quad (4.17)$$

With Δ_{\leftarrow} and Δ_{\rightarrow} , the relation H does not have to simulate the moments when the head is at the left or at the right end of the tape, it changes its moving direction right before it comes to \triangleright or to \sqcup respectively. The relation H is forced to simulate the movement of the head of the machine only if the underlying element is in M and if locally the structure looks like a grid from Section 4.2. Let μ_A, μ_B, μ_D be as in Eqs. (4.9), (4.10) and (4.12). Then, the execution of M is simulated according to Eqs. (4.18) to (4.21), see Figure 7. Each of Eqs. (4.18) to (4.21) contains a disjunction of conjunctions in the head of the implication. It is not allowed by MMSN_{\neq} but each such clause can be transformed to an allowed form as follows:

$$\begin{aligned} A \rightarrow ((B \wedge C) \vee (D \wedge E)) &\equiv \neg(A \wedge (\neg B \vee \neg C) \wedge (\neg D \vee \neg E)) \equiv \\ &\equiv \neg(A \wedge \neg B \wedge \neg D) \wedge \neg(A \wedge \neg B \wedge \neg E) \wedge \neg(A \wedge \neg C \wedge \neg D) \wedge \neg(A \wedge \neg C \wedge \neg E) \end{aligned}$$

$$\begin{aligned} \bigwedge_{(s,q) \in \Sigma \times \mathcal{Q}^{\leftarrow}} \left(\left(\mu_A(x, x_1, x_2, y, y_1, y_2, z) \wedge H(x_1) \wedge S(x_1) \wedge Q(x_1) \right) \right. \\ \left. \rightarrow \left(H(y_2) \wedge \bigvee_{(s',q') \in \Delta_{\rightarrow}(s,q)} S'(y_1) \wedge Q'(y_2) \right) \right) \end{aligned} \quad (4.18)$$

$$\begin{aligned} \bigwedge_{(s,q) \in \Sigma \times \mathcal{Q}^{\rightarrow}} \left(\left(\mu_B(x, x_1, x_2, x_3, y, y_1, y_2, y_3, z) \wedge H(x_2) \wedge S(x_2) \wedge Q(x_2) \right) \right. \\ \left. \rightarrow \left(H(y_3) \wedge \bigvee_{(s',q') \in \Delta_{\rightarrow}(s,q)} S'(y_2) \wedge Q'(y_3) \right) \right) \end{aligned} \quad (4.19)$$

$$\begin{aligned} \bigwedge_{(s,q) \in \Sigma \times \mathcal{Q}^{\rightarrow}} \left(\left(\mu_D(x, x_1, x_2, y, y_1, y_2, y_3, z) \wedge H(x_2) \wedge S(x_2) \wedge Q(x_2) \right) \right. \\ \left. \rightarrow \left(H(y_1) \wedge \bigvee_{(s'_{\sqcup}, s', q') \in \Delta_{\leftarrow}(s,q)} S'(y_2) \wedge S'_{\sqcup}(y_3) \wedge Q'(y_1) \right) \right) \end{aligned} \quad (4.20)$$

$$\bigwedge_{(s,q) \in \Sigma \times \mathcal{Q}^{\leftarrow}} \left(\left(\mu_B(x, x_1, x_2, x_3, y, y_1, y_2, y_3, z) \wedge H(x_2) \wedge S(x_2) \wedge Q(x_2) \right) \right. \\ \left. \rightarrow \left(H(y_1) \wedge \bigvee_{(s',q') \in \Delta \leftarrow (s,q)} S'(y_2) \wedge Q'(y_1) \right) \right) \quad (4.21)$$

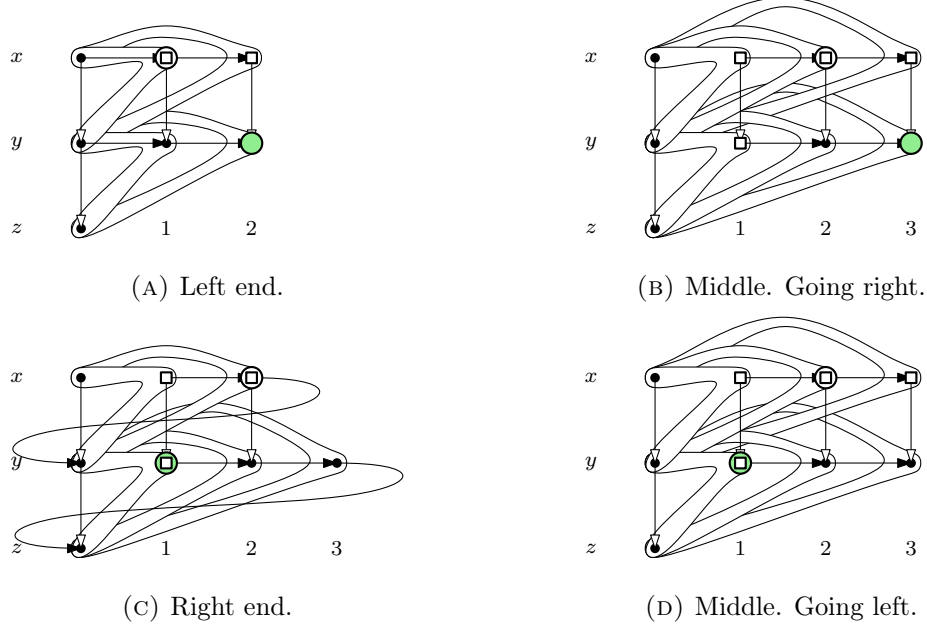


FIGURE 7. All the four cases when H (denoted by a circle) is forced to spread. The H -atoms in the head of each implication are coloured in green.

If an element x does not belong to H , then its **next**-successor must have the same symbol relation as x has:

$$\bigwedge_{s \in \Sigma} (M(x) \wedge \neg H(x) \wedge S(x) \wedge \text{next}(x, x') \rightarrow S(x')) \quad (4.22)$$

The head cannot visit two distinct cells at the same moment of time, so we need the following rule:

$$\text{row}(a, x, a') \wedge \text{row}(a, y, a') \wedge M(x) \wedge M(y) \wedge H(x) \wedge H(y) \rightarrow x = y \quad (4.23)$$

Let Q_r be a relation associated with a rejecting state q_r in \mathcal{Q} . Then, for every such Q_r , we add the following rule that forbids a head to be in this state during the execution:

$$\neg (M(x) \wedge H(x) \wedge Q_r(x)) \quad (4.24)$$

First-order restrictions. The expressiveness of MMSN_{\neq} permits Φ to require the following properties that reject *over-complete* input structures. Notice that the following negated conjuncts have no σ -atoms (existentially quantified). Therefore, we will call them *first-order restrictions*.

(a) There is at most one **start** element: $\text{start}(x) \wedge \text{start}(y) \rightarrow x = y$.




FIGURE 8. The unique choices for the relations I (grey), M (grey and white), and H (white circles) for the structure from Example 4.4.

(b) The in- and out- degrees of succ and next are at most 1:

$$(\text{succ}(x, y) \wedge \text{succ}(x, y') \rightarrow y = y') \wedge (\text{succ}(x, y) \wedge \text{succ}(x', y) \rightarrow x = x')$$

$$\wedge (\text{next}(x, y) \wedge \text{next}(x, y') \rightarrow y = y') \wedge (\text{next}(x, y) \wedge \text{next}(x', y) \rightarrow x = x').$$

(c) An element y cannot be in different triples $\text{row}(x, y, z)$ and $\text{row}(x', y, z')$:

$$(\text{row}(x, y, z) \wedge \text{row}(x', y, z') \rightarrow x = x') \wedge (\text{row}(x, y, z) \wedge \text{row}(x', y, z') \rightarrow z = z').$$

(d) Loops $\text{succ}(x, x)$, $\text{next}(x, x)$ are forbidden as well as $\text{row}(x, y, x)$.

(e) An element x is contained in at most one symbol relation $s(x)$:

$$\bigwedge_{\text{distinct } s, s' \in \Sigma} \neg(s(x) \wedge s'(x)).$$

4.4. Reduction from M to Φ . Let M be a precise Turing machine and let Φ be a MMSNP_{\neq} sentence constructed from M according to Subsection 4.3. Let r_{easy} be a mapping from the set of finite strings Σ^* to the set of finite τ -structures defined in Subsection 4.2.

Lemma 4.5. *The mapping r_{easy} describes a polynomial-time reduction from M to Φ .*

Proof. Let $s \in \Sigma^*$ be an input string of M and let $\mathfrak{A} = r_{\text{easy}}(s)$.

By the construction of Φ , there are unique interpretations of M , H , I that satisfy the rules of Φ and the relation H simulates exactly $g(n)$ round trips of the head of M , see Figure 8 for example. By the construction of Φ , there exists a one-to-one correspondence between transitions of the head of the machine and the set of rules of Φ that describe these transitions. So, if the head of the machine is in a rejecting state at the end, then, for every possible choice of existentially quantified relations, the rule (4.24) will be violated. Otherwise, if there exists a sequence of transitions that finishes in an accepting state, then the construction of Φ allows to repeat exactly the same sequence, so the rule (4.24) will not be violated. \square

4.5. Reduction from Φ to M . It remains to describe the mapping r_{hard} that takes as input a finite τ -structure \mathfrak{A} and outputs in time polynomial in $|A|$ a string $\mathbf{s} \in \Sigma^*$ such that the precise machine M accepts \mathbf{s} if and only if $\mathfrak{A} \models \Phi$.

Rejecting over-complete input. One can check whether the input structure \mathfrak{A} violates one of the first-order restrictions in polynomial time in the size of \mathfrak{A} . If so, then the mapping r_{hard} assigns to \mathfrak{A} the fixed NO instance \mathbf{s}_N .

Unique choices for the relations M and H . Now we may assume that \mathfrak{A} satisfies the first-order restrictions. We describe the minimal possible sets M_{\min} and H_{\min} consisting of elements of \mathfrak{A} that ought to be in the relations M and H in order to satisfy Φ . We will show in Lemma 4.6 that without loss of generality we can consider exclusively the elements from these sets and assume that the relations M and H are interpreted as M_{\min} and H_{\min} respectively.

Denote by a_0 the element of \mathfrak{A} such that $\mathfrak{A} \models \text{start}(a_0)$. If such an element does not exist, then no element is forced to be in M , so \mathfrak{A} is immediately accepted by Φ and r_{hard} can assign to \mathfrak{A} the string \mathbf{s}_Y . Otherwise, the sets M_{\min} and H_{\min} are defined by induction.

Firstly, we add to M_{\min} the right *succ*-neighbour of a_0 according to Eq. (4.2). Then we add to M_{\min} the neighbour of the neighbour of a_0 , and so on, by using the rule from Eq. (4.3). Eventually, we add all the elements that are forced to be in I . By Eq. (4.6), every element in I must also be in M . Then, we add to M_{\min} every element that is forced to be in M according to Eqs. (4.7) and (4.8), and add to H_{\min} every element that is forced to be in H according to Eqs. (4.18) to (4.21). Repeat this until we do not have to add any new element to M_{\min} and H_{\min} .

As \mathfrak{A} satisfies the first-order restrictions, the in- and out-degrees of *succ*- and *next*-arcs of \mathfrak{A} cannot be greater than 1. So, we can assign a pair of natural numbers to every element of M_{\min} . Assign $(1, 0)$ to the right *succ*-neighbour of a_0 . If (i, j) is assigned to x and if $\mathfrak{A} \models \text{succ}(x, y)$, then we assign $(i + 1, j)$ to y . If (i, j) is assigned to x and if $\mathfrak{A} \models \text{next}(x, y)$, then we assign $(i, j + 1)$ to y . It follows from the definition of M_{\min} and from the assumption that \mathfrak{A} satisfies the first-order restrictions, that this assignment is well-defined.

Let S be the set of *next*-successors of a_0 . That is, a_0 belongs to S and, if $\mathfrak{A} \models \text{next}(x, y)$ and x is in S , then y belongs to S . Let \mathfrak{A}_{\min} be the substructure of \mathfrak{A} induced on $A_{\min} := S \cup M_{\min}$.

Lemma 4.6. *Let \mathfrak{A} satisfy the first-order restrictions of Φ . Then, Φ holds in \mathfrak{A} if and only if it holds in \mathfrak{A}_{\min} .*

Proof. It is sufficient to show only that $\mathfrak{A}_{\min} \models \Phi$ implies $\mathfrak{A} \models \Phi$. The other direction is clear because SNP problems are closed under taking substructures.

Let $\mathfrak{A}_{\min}^\sigma$ be a σ -expansion of \mathfrak{A}_{\min} that satisfies the first-order part of Φ . An example of such a structure $\mathfrak{A}_{\min}^\sigma$ is given on Figure 9. By the construction of M_{\min} and H_{\min} , we know that $M_{\min} \subseteq M_{\min}^{\mathfrak{A}_{\min}^\sigma}$ and that $H_{\min} \subseteq H_{\min}^{\mathfrak{A}_{\min}^\sigma}$. We can assume without loss of generality that $M_{\min}^{\mathfrak{A}_{\min}^\sigma} = M_{\min}$ because we can just remove every element of $(A_{\min} \setminus M_{\min})$ from $M_{\min}^{\mathfrak{A}_{\min}^\sigma}$ without violating a rule of Φ .

Let \mathfrak{A}^σ be any σ -expansion of \mathfrak{A} that contains $\mathfrak{A}_{\min}^\sigma$ as a substructure and such that $M^{\mathfrak{A}^\sigma} = M_{\min}^{\mathfrak{A}_{\min}^\sigma} = M_{\min}$ and $H^{\mathfrak{A}^\sigma} = H_{\min}^{\mathfrak{A}_{\min}^\sigma}$. By the definition of M_{\min} , \mathfrak{A}^σ satisfies all the rules from Eqs. (4.2), (4.3) and (4.6) to (4.8). In all other conjuncts except for the variables a and a' of (4.15) and (4.23), all variables are assumed to be in M , so such conjuncts will




FIGURE 9. A σ -expansion of the substructure of \mathfrak{A} induced on M_{\min} that satisfies the first-order part of Φ . Grey zone highlight the set S . The red square is **start**, horizontal arcs are **succ**, vertical are **next**. Circles are the elements of H . The red line highlights the moment of time when M halts.

not be violated in \mathfrak{A}^σ . Suppose that one of the negated conjuncts from (4.15) or (4.23) is violated. Let a, a', x , and y witness it. As x and y are in M , they must be in M_{\min} . So, at least one of a or a' is not in A_{\min} . But, by the definition of M_{\min} , if x and y are in M_{\min} , then both a and a' are in S . So, all of them belong to A_{\min} , this contradicts $\mathfrak{A}_{\min} \models \Phi$. \square

The reduction r_{hard} . We are ready to construct the mapping r_{hard} that witnesses a polynomial-time reduction from Φ to M .

If a τ -structure \mathfrak{A} does not satisfy the first-order restrictions, then $r_{\text{hard}}(\mathfrak{A}) = s_N$, hence, it is also rejected by M . Otherwise, by Lemma 4.6, we can assume that $\mathfrak{A} = \mathfrak{A}_{\min}$. Suppose that the first row of M_{\min} consists of the elements $(1, 0), \dots, (n, 0)$. The construction of the sets M_{\min} and H_{\min} is deterministic and requires polynomial time in the size of \mathfrak{A} .


Recall from Section 4.2 that, in order to simulate every execution of M , the set M_{\min} must have at least $h(n) = (g(n))^2 + (2n - 3)g(n) + 1$ rows. If the set H_{\min} does not contain the element $(1, h(n))$, then we know that \mathfrak{A} cannot simulate the whole execution of M , so the head will never be in a rejecting state, i.e., the rule (4.24) will not be violated. See Figure 10a for an illustration. Other rules can be satisfied for the σ -expansion \mathfrak{A}^σ , where $M^{\mathfrak{A}^\sigma} = M_{\min}$, $H^{\mathfrak{A}^\sigma} = H_{\min}$, and the relations S (for every $s \in \Sigma$) and Q (for every $q \in \mathcal{Q}$) are chosen in such a way that no rule among Eqs. (4.4), (4.5), (4.13) to (4.15) and (4.18) to (4.22) is violated. In this case, $r_{\text{hard}}(\mathfrak{A}) = s_Y$.

If the set H_{\min} contains the element $(1, h(n))$, then we know that \mathfrak{A} allows to simulate the whole execution of M on the input of size n . Define $r_{\text{hard}}(\mathfrak{A})$ to be a string $s = s_1 \dots s_n$ such that, for every $i \in [n]$, the i th element s_i is equal to some s in Σ if and only if the element $(i, 0)$ is in $s^{\mathfrak{A}}$. See Figure 10b for an illustration. As \mathfrak{A} satisfies the first-order restrictions, no element of \mathfrak{A} belongs to two distinct “symbol” relations $s_1^{\mathfrak{A}}$ and $s_2^{\mathfrak{A}}$. As each

element of the first row is in M_{\min} , it must belong to at least one “symbol” relation $s^{\mathfrak{A}}$. So, the mapping r_{hard} is well-defined. The following lemma directly follows from the definition of r_{hard} .

Lemma 4.7. *For every τ -structure \mathfrak{A} , the machine M accepts $r_{\text{hard}}(\mathfrak{A})$ if and only if $\mathfrak{A} \models \Phi$.*

Lemma 4.7 together with Lemma 4.5 imply Theorem 4.1.



5. CONCLUSION

In this paper, we have explored some fragments that sit above Feder and Vardi’s logic MMSNP and below the three logics obtained by relaxing one of the three syntactic restrictions: namely, (non monotone) monadic SNP, MMSNP with \neq and (non monadic) monotone SNP. These three logics cannot exhibit a dichotomy unless P is equal to NP by Ladner’s theorem. In order to extend a dichotomy to some superclasses of MMSNP, we have highlighted three more general fragments thereof: namely, MPART, MMSNP_{\oplus} and GMSNP. Apart from MMSNP_{\oplus} which we embed up to polynomial-time equivalence in MMSNP to prove dichotomy, the question remains open for its two siblings.

A weaker yet interesting result would consist in relating directly the dichotomy status among two classes for which the dichotomy question is open. For example, is it the case that for every problem in MPART there exists a polynomial-time equivalent problem in GMSNP?

Another puzzling question concerns the the class MPART that contains the variation of Matrix Partitions (MP) defined in [BK24]. In particular is it the case that one can find for each sentence of MPART a polynomial-time equivalent problem in MP? The analogous question for CSP and MMSNP was settled by Feder and Vardi using Erdős’ lemma which allows to construct in randomised polynomial-time a structure of high girth and fixed chromatic number. The technique breaks down in the context of MPART. Indeed, the main difficulty when working in the context of MP and MPART is that it is a genuine challenge to combine structures as is common when working in the monotone world of CSP and MMSNP.

In the latter, we are free to take the disjoint union of accepted instances to produce a larger accepted instance or to combine two accepted instances sharing some common feature (*e.g.* free amalgamation). In the former, absence of an edge between vertices does not usually mean that we may freely map these vertices. While it is possibly the lack of freedom of this feature that means that the problems could not show NP-intermediate complexity, it poses a real challenge when working with these non monotone problems.

If we could somehow show that each problem in GMSNP enjoys a polynomial-time equivalent problem in MMSNP, then we would obtain a dichotomy for GMSNP. The challenge in this endeavour is that some input structures in the target problem of MMSNP are not in the image of the reduction from GMSNP. This is precisely the problem that was solved by Feder and Vardi using Erdős' Lemma in the context of MMSNP and CSP.

The above question seems hard and a perhaps weaker question that could serve as a first step would consist in showing that a problem in GMSNP over a signature with symbols of arity at most r enjoys a polynomial-time equivalent problem in GMSNP over a signature with symbols of arity at most $r - 1$. In particular, the case $r = 3$ would be an interesting first step.

ACKNOWLEDGMENT

The authors sincerely appreciate the anonymous referees whose valuable comments helped to improve the presentation and writing of this article.

REFERENCES

- [Bar22] Alexey Barsukov. *On dichotomy above Feder and Vardi's logic*. PhD thesis, Université Clermont Auvergne, 2022.
- [BH90] Jørgen Bang-Jensen and Pavol Hell. The effect of two cycles on the complexity of colourings by directed graphs. *Discret. Appl. Math.*, 26(1):1–23, 1990. doi:10.1016/0166-218X(90)90017-7.
- [BK24] Alexey Barsukov and Mamadou Moustapha Kanté. Generalisations of matrix partitions: Complexity and obstructions. *Theor. Comput. Sci.*, 1007:114652, 2024. URL: <https://doi.org/10.1016/j.tcs.2024.114652>, doi:10.1016/J.TCS.2024.114652.
- [BKN09] Libor Barto, Marcin Kozik, and Todd Niven. The CSP dichotomy holds for digraphs with no sources and no sinks (A positive answer to a conjecture of Bang-Jensen and Hell). *SIAM J. Comput.*, 38(5):1782–1802, 2009. doi:10.1137/070708093.
- [BKS20] Manuel Bodirsky, Simon Knäuer, and Florian Starke. ASNP: A tame fragment of existential second-order logic. In Marcella Anselmo, Gianluca Della Vedova, Florin Manea, and Arno Pauly, editors, *CiE 2020, Proceedings*, volume 12098 of *Lecture Notes in Computer Science*, pages 149–162. Springer, 2020. URL: https://doi.org/10.1007/978-3-030-51466-2_13, doi:10.1007/978-3-030-51466-2_13.
- [Bod21] Manuel Bodirsky. *Complexity of Infinite-Domain Constraint Satisfaction*. Lecture Notes in Logic. Cambridge University Press, 2021. Postprint available online: <https://wwwpub.zih.tu-dresden.de/~bodirsky/Book.pdf>. doi:10.1017/9781107337534.
- [BPP21] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Projective clone homomorphisms. *J. Symb. Log.*, 86(1):148–161, 2021. URL: <https://doi.org/10.1017/jsl.2019.23>, doi:10.1017/JSL.2019.23.
- [BtCLW14] Meghyn Bienvenu, Balder ten Cate, Carsten Lutz, and Frank Wolter. Ontology-based data access: A study through disjunctive datalog, CSP, and MMSNP. *ACM Trans. Database Syst.*, 39(4):33:1–33:44, 2014. doi:10.1145/2661643.
- [Bul17] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 319–330. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.37.

- [Fag74] Ronald Fagin. Generalized first-order spectra and polynomial-time recognizable sets. In *Complexity of computation (Proc. SIAM-AMS Sympos., New York, 1973)*, volume Vol. VII of *SIAM-AMS Proc.*, pages 43–73. Amer. Math. Soc., Providence, RI, 1974.
- [FV98] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1998. doi:10.1137/S0097539794266766.
- [FV03] Tomás Feder and Moshe Y. Vardi. Homomorphism closed vs. existential positive. In *18th IEEE Symposium on Logic in Computer Science (LICS 2003)*, pages 311–320. IEEE Computer Society, 2003. doi:10.1109/LICS.2003.1210071.
- [Hel14] Pavol Hell. Graph partitions with prescribed patterns. *Eur. J. Comb.*, 35:335–353, 2014. doi:10.1016/j.ejc.2013.06.043.
- [HN90] Pavol Hell and Jaroslav Nešetřil. On the complexity of H -coloring. *J. Comb. Theory, Ser. B*, 48(1):92–110, 1990. doi:10.1016/0095-8956(90)90132-J.
- [Lad75] Richard E. Ladner. On the structure of polynomial time reducibility. *J. ACM*, 22(1):155–171, 1975. doi:10.1145/321864.321877.
- [Mad09] Florent R. Madelaine. Universal structures and the logic of forbidden patterns. *Log. Methods Comput. Sci.*, 5(2), 2009. doi:10.2168/LMCS-5(2:13)2009.
- [Mat93] Y. V. Matiyasevich. *Hilbert’s Tenth Problem*. The MIT Press, 1993.
- [Pap94] Christos H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.
- [PF79] Nicholas Pippenger and Michael J. Fischer. Relations among complexity measures. *J. ACM*, 26(2):361–381, 1979. doi:10.1145/322123.322138.
- [Pin22] Michael Pinsker. Current challenges in infinite-domain constraint satisfaction: Dilemmas of the infinite sheep. In *52nd IEEE International Symposium on Multiple-Valued Logic, ISMVL 2022, Dallas, TX, USA, May 18-20, 2022*, pages 80–87. IEEE, 2022. doi:10.1109/ISMVL52857.2022.00019.
- [Sch78] Thomas J. Schaefer. The complexity of satisfiability problems. In Richard J. Lipton, Walter A. Burkhard, Walter J. Savitch, Emily P. Friedman, and Alfred V. Aho, editors, *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, May 1-3, 1978, San Diego, California, USA*, pages 216–226. ACM, 1978. doi:10.1145/800133.804350.
- [Zhu20] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. doi:10.1145/3402029.