On the Computational Power of Extensional ESO*

Manuel Bodirsky[†1] and Santiago Guzmán-Pro[‡1]

[1]Institut für Algebra, TU Dresden, Germany

Abstract

Extensional ESO is a fragment of existential second-order logic (ESO) that captures the following family of problems. Given a fixed ESO sentence Ψ and an input structure $\mathbb{A}$ the task if to decide whether there is an *extension* $\mathbb{B}$ of $\mathbb{A}$ that satisfies the first-order part of Ψ, i.e., a structure $\mathbb{B}$ such that $R^{\mathbb{A}} \subseteq R^{\mathbb{B}}$ for every existentially quantified predicate R of Ψ, and $R^{\mathbb{A}} = R^{\mathbb{B}}$ for every non-quantified predicate R of Ψ. In particular, extensional ESO describes all pre-coloured finite-domain constraint satisfaction problems (CSPs).

In this paper we study the computational power of extensional ESO; we ask, *for which problems in* NP *is there a polynomial-time equivalent problem in extensional ESO?*. One of our main results states that extensional ESO has the same computational power as *hereditary first-order logic*. We also characterize the computational power of the fragment of extensional ESO with monotone universal first-order part in terms of finitely bounded CSPs. These results suggest a rich computational power of this logic, and we conjecture that extensional ESO captures NP-intermediate problems. We further support this conjecture by showing that extensional ESO can express current candidate NP-intermediate problems such as Graph Isomorphism, and Monotone Dualization (up to polynomial-time equivalence). On the other hand, another main result proves that extensional ESO does not have the full computational power of NP: there are problems in NP that are not polynomial-time equivalent to a problem in extensional ESP (unless E=NE).

Contents

1 Introduction **1**

2 Preliminaries **5**
2.1 (First-order) structures 5
2.2 SNP and its fragments 6
2.3 Polynomial-time reductions and subclasses of NP 6

3 Extensional ESO **7**
3.1 Surjective CSPs 7
3.2 Graph Sandwich Problems 8
3.3 Orientation Completion Problems 8
3.4 Constraint Topological Sorting 9

*This project has received funding from the European Union (Project POCOCOP, ERC Synergy Grant 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.

†manuel.bodirsky@tu-dresden.de

‡santiago.guzman_pro@tu-dresden.de

4 Extensional ESO and HerFO **9**
4.1 HerFO captures extensional ESO 10

5 Monotone extensional SNP **14**
5.1 Finitely bounded structures and full homomorphisms 14
5.2 Monotone extensional SNP and CSPs 16
5.3 CSPs of finitely bounded structures 18

6 Candidate NP-intermediate problems in extensional ESO **19**

7 Limitations of the computational power **22**
7.1 Disjunctive self-reducible sets 22
7.2 NP-intermediate disjunctive self-reducible sets 23
7.3 Search reduces to decision 23
7.4 Promise problems 25
7.5 Failure of NP-richness 26

8 Conclusion and Open Problems **26**

9 References **27**

A A remark concerning the definition of extensional ESO **31**

B Monotone connected extensional SNP **31**

1 Introduction

First-order logic (FO) has good computational properties: given a fixed first-order sentence ϕ it can be verified in deterministic logarithmic space whether an input structure $\mathbb{A}$ satisfies ϕ. However, it has been realized long ago that FO has a very limited expressive power when it comes to describing classes of finite structures. For instance, there is no first-order sentence which holds on a finite graph if and only if the graph does not contain cycles (i.e., is a forest). In contrast, second-order logic is very powerful and captures the entire polynomial-time hierarchy [37]. Even restricted fragments of second-order logic can be very expressive. For instance, Fagin [28] showed that *existential second-order logic (ESO)* captures the complexity class NP. Taking complements, *universal second-order logic (USO)* captures coNP.

Further restrictions of ESO might no longer have the full expressive power, but still the full *computational* power of NP, in the sense that every problem in NP is polynomial-time equivalent to a problem that can be expressed in the restricted logic; we will call such classes *NP-rich*.[1] A prominent example of an NP-rich ESO fragment is *SNP* (for *strict NP*), which is defined as the restriction of ESO to sentences with a universal first-order part. SNP is quite expressive, and contains for instance all of Datalog [29], but clearly it cannot express every class of finite structures in NP, because SNP can only describe *hereditary classes*, i.e., classes $\mathcal{C}$ closed under taking induced substructures. However, Feder and Vardi [29] proved that SNP is NP-rich (see [8] for some of the missing details of the proof given in [29]). In particular, it follows from Ladner's theorem [41]

[1] If we consider disjunctive, conjunctive, or Turing reductions, we call such classes *conjunctive, disjunctive,* or *Turing NP-rich*, respectively. These reductions are defined in Section 2.

that whenever a fragment $\mathcal{L}$ of ESO is NP-rich, then $\mathcal{L}$ does not exhibit a P versus NP-complete dichotomy.

On the dichotomy side, Feder and Vardi showed that a fragment of ESO called *monotone monadic SNP (MMSNP)* has the same computational power as finite-domain CSPs. Hence, it follows from the finite-domain CSP dichotomy [20, 48] that MMSNP exhibits a P vs. NP-complete dichotomy. Bienvenu, ten Cate, Lutz, and Wolter [11] asked whether the dichotomy for MMSNP extends to a broader fragment called *guarded monotone SNP (GMSNP)*. In this direction, Bodirsky, Knäuer, and Starke [16] showed that GMSNP should indeed exhibit a dichotomy according to the Bodirsky–Pinsker conjecture [12, Conjecture 3.7.1] — yet, the GMSNP (conjectured) dichotomy seems elusive. It is worth mentioning that so far, the main technique to show that a fragment of ESO does not exhibit a P versus NP-complete dichotomy is by showing that it is NP-rich [39].

The story. In this paper we initiate the study of the computational power of fragments of ESO beyond complexity dichotomies and NP-richness. We propose to find a *natural* and *decidable* syntactic fragments of ESO that is neither NP-rich nor exhibits a P vs. NP-complete dichotomy. To this end, we introduce *extensional ESO*, and we exhibit several families of problems that can be expressed in extensional ESO, ranging from the world of CSPs to graph theory and to language theory. We then address the question "what is the computational power of extensional ESO?". We show that extensional ESO has the same computational power as the recently introduced formalism *hereditary first-order logic*. We also show that extensional ESO is not NP-rich, unless E = NE. However, we conjecture that extensional ESO captures NP-intermediate problems. We support this conjecture by showing that some examples of candidate NP-intermediate problems such as Graph Isomorphism and Monotone Dualization are expressible in extensional ESO. We also show that a very large class of CSPs, namely the class of CSPs for *finitely bounded structures*, can be expressed in extensional ESO; we believe that this class contains NP-intermediate problems.

In Figure 1 we present an illustration of this story, and we now present a brief description of *emphasized characters* (all missing details will be carefully introduced later in the paper).

The main character. The central role of this paper is played by *extensional ESO*. An ESO τ-sentence Ψ is called *extensional* if for every existentially quantified predicate R there is a distinct symbol $R' \in \tau$ of the same arity such that Ψ contains a conjunct $\forall \overline{x}\big(R'(\bar{x}) \Rightarrow R(\bar{x})\big)$, and R' does not appear anywhere else in Ψ. In this case we say that R *extends* R'. Clearly, every problem described by an ESO sentence Φ is in NP. To present two simple examples, we consider *digraph acyclicity*, and *precoloured* 3*-coloring problem* from graph theory; further examples can be found in Section 3.

Example 1 (Acyclic digraphs)**.** The class of strict linear orders can be easily axiomatized in universal first-order logic with a binary relation symbol $<$. Let ϕ be such a sentence, and notice that a digraph $\mathbb{D}$ is a acyclic if and only if is a spanning graph of a transitive tournament, in other words, of a strict linear order. Hence, the class of acyclic digraphs is described by the extensional ESO sentence

$$\exists{<}\ \forall x, y\big((E(x,y) \Rightarrow\ <(x,y))\ \wedge\ \phi\big).$$

Example 2 (Pre-coloured 3-colourability)**.** Given an input graph G with a partial colouring $c\colon U \subseteq V \to \{R, G, B\}$, the task is to decide whether there is a 3-colouring of G such that its restriction to U coincides with c. This problem can be expressed by the extensional ESO sentence

$$\begin{aligned}\exists R, G, B\, \forall x, y \big(& (R'(x) \Rightarrow R(x)) \wedge (G'(x) \Rightarrow G(x)) \wedge (B'(x) \Rightarrow B(x)) \\ & \wedge (E(x,y) \Rightarrow \neg(R(x) \wedge R(y)) \wedge \neg(G(x) \wedge G(y)) \wedge \neg(B(x) \wedge B(y))) \\ & \wedge (R(x) \vee G(x) \vee B(x))\big).\end{aligned}$$

Here, R, G, and B denote the color classes of the graph, and must denote extensions of input predicates R', G', and B' which can be used to enforce that some variables are colored by a specific color.

HerFO and the tractability problem. As supporting character we have *hereditary first-order logic*. This is a recently introduced logic that can naturally model several computational problems such as digraph acyclicity, graph chordality, and several infinite-domain CSPs [14] — parametrized variants of this problem were previously studied, e.g., in [7, 30]. A class $\mathcal{C}$ of structures is in hereditary first-order logic (HerFO) if there exists a first-order sentence ϕ such that a structure $\mathbb{A}$ is in $\mathcal{C}$ if and only if all (non-empty) substructures of $\mathbb{A}$ satisfy ϕ. HerFO can be naturally defined as a fragment of universal monadic second-order logic [14, Observation 4.1].

The *tractability problem* for second-order logic asks whether a given SO sentence Φ describes a polynomial-time solvable problem. The tractability problem is already undecidable for ESO (and thus, for USO as well); see, e.g., [12]. As a rule of thumb, we expect that if a fragment $\mathcal{L}$ of ESO (or of USO) has a rich computational power, e.g, is NP-rich, then the tractability problem is undecidable for $\mathcal{L}$. It is still open whether HerFO captures coNP-intermediate problems, nonetheless, it is known that the tractability problem is already undecidable for HerFO [14].

Finitely bounded CSPs. Our lead secondary character if the family of CSPs of finitely bounded structures. A structure $\mathbb{B}$ is called *finitely bounded* if there exists a finite set $\mathcal{F}$ of finite structures such that a finite structure $\mathbb{A}$ embeds into $\mathbb{B}$ if and only if no structure from $\mathcal{F}$ embeds into $\mathbb{A}$. In particular, every finite structure is finitely bounded, but also the countable homogeneous linear order $(\mathbb{Q}, <)$ is finitely bounded, as well as the generic partial order and the generic semilinear order [42]. Examples from graph theory include the universal homogeneous K_k-free graphs [40] and the generic circular triangle-free graph [15].

CSPs of finitely bounded structures are closely related to the previously introduced logic SNP. In particular, there is a syntactic fragment of SNP that captures CSPs of *reducts* of finitely bounded structures [12, Corollary 1.4.12] — a reduct of a structure $\mathbb{A}$ is a structure obtained $\mathbb{A}'$ by forgetting some relations in $\mathbb{A}$. Actually, this family of CSPs (and the corresponding fragment of SNP) has the full computational power of NP [12, Proposition 13.2.2]. This has lead some people (including us) to believe that CSPs of finitely bounded structures also capture NP-intermediate problems. It would be an amazing surprise if CSPs of finitely bounded structures exhibit a P vs. NP-complete dichotomy, but as soon as one considered their reducts, we obtain the full computational power of NP!

Our contributions

Our first main result is a characterization of the computational power of extensional ESO in terms of hereditary first-order logic (Theorem 4.6): we show that for every sentence Ψ in extensional ESO there exists a first-order formula ϕ such that $\mathrm{Her}(\phi)$ is polynomial-time (in fact, log-space) equivalent to the problem described by the negation of Ψ, and vice versa. As a byproduct of this connection, we get that the tractability problem for extensional ESO is undecidable (Corollary 4.4).

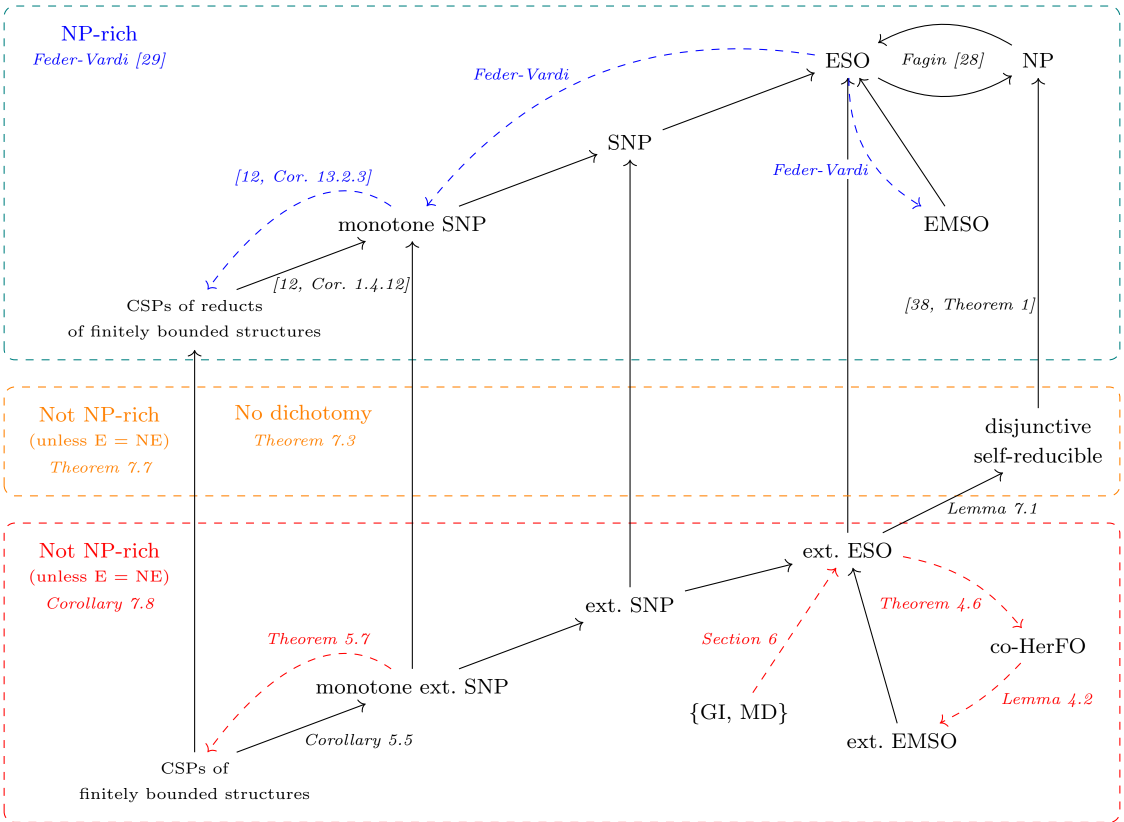

Figure 1: Classes of computational problems studied in this paper. Solid arcs indicate inclusions, and a dashed arc from A to B indicates that every problem in A is polynomial-time equivalent to a problem in B — for a cleaner picture we omit the solid arc from extensional EMSO to EMSO. GI stands for the Graph Isomorphism problem, and MD for the Monotone Dualization problem.

We then look at fragments of extensional ESO defined by restrictions on the quantification and on the polarity of atoms of the first-order part. In particular, our second main result characterizes the computational power of *monotone extensional SNP* (introduced below) in terms of CSPs of finitely bounded structures (Theorem 5.7). We also show that extensional ESO captures surjective CSPs, as well as several graph sandwich problems, orientation completion problems, and constraint topological sorting problems. We further present examples of candidate NP-intermediate problems expressible in extensional ESO; namely, Graph Isomorphism and Monotone Dualization (Section 6).

These results point towards a rich computational power of extensional ESO, and we thus conjecture that extensional ESO does not exhibit a P vs. NP-complete dichotomy (Conjecture 6.1). However, our last main result shows that extensional ESO cannot be NP-rich (Corollary 7.8): there are problems in NP that are not polynomial-time equivalent to a problem in extensional ESO, unless E = NE (the class E is the class of problems computable in exponential time with linear exponent; NE is its non-deterministic pendant). The assumption that E $\neq$ NE is stronger than the usual assumption that P $\neq$ NP, but still a common assumption in complexity theory [10, 34, 46]. To prove this result, we show that every problem in extensional ESO is *disjunctively self-reducible*, and that this class is not NP-rich, unless E = NE (Theorem 7.7); we complement this result by showing that disjunctive self-reducible sets do not exhibit a P versus NP-complete complexity dichotomy (Theorem 7.3). It also follows from this result that CSPs of finitely bounded structures do not capture the full computational power of NP, and that HerFO cannot be coNP-rich.

Outline. We begin by recalling some basic concepts from finite model theory (Section 2) and continue to present several families of examples captured by extensional ESO in Section 3. In Section 4 we show that extensional ESO and HerFO have the same computational power. An important restriction of extensional ESO is monotone extensional SNP (Section 5), which we show has the same computational power as CSPs of finitely bounded structures (Theorem 5.7). In Section 6 we prove that Graph Isomorphism and Monotone Dualization are captured by extensional ESO (up-to polynomial-time equivalence). We discuss the limitations of the computational power of extensional ESO in Section 7, and in particular we prove that extensional ESO is not NP-rich, unless E = NE (Corollary 7.8). We close with a series of problems that are left open (Section 8).

2 Preliminaries

We assume basic familiarity with first-order logic and we follow standard notation from model theory, as, e.g., in [36]. We also use standard notions from complexity theory.

2.1 (First-order) structures

Given a relational signature τ and a τ-structure $\mathbb{A}$, we denote by $R^{\mathbb{A}}$ the interpretation in $\mathbb{A}$ of a relation symbol $R \in \tau$. Also, we denote relational structures with letters $\mathbb{A}, \mathbb{B}, \mathbb{C}, \dots$, and their domains by $A, B, C, \dots$. In this article, structures have non-empty domains.

If $\mathbb{A}$ and $\mathbb{B}$ are τ-structures, then a *homomorphism* from $\mathbb{A}$ to $\mathbb{B}$ is a map $f\colon A \to B$ such that for all $a_1, \dots, a_k \in A$ and $R \in \tau$ of arity k, if $(a_1, \dots, a_k) \in R^{\mathbb{A}}$, then $R(f(a_1), \dots, f(a_k)) \in R^{\mathbb{B}}$ We write $\mathbb{A} \to \mathbb{B}$ if there exists a homomorphism from $\mathbb{A}$ to $\mathbb{B}$, and we denote by CSP($\mathbb{B}$) the class of finite structures $\mathbb{A}$ such that $\mathbb{A} \to \mathbb{B}$. Let $\mathcal{C}$ be a class of finite τ-structures. We say that $\mathcal{C}$ is

- *closed under homomorphisms* if for every $\mathbb{B} \in \mathcal{C}$, if $\mathbb{B} \to \mathbb{A}$, then $\mathbb{A} \in \mathcal{C}$ as well.
- *closed under inverse homomorphisms* if for every $\mathbb{B} \in \mathcal{C}$, if $\mathbb{A} \to \mathbb{B}$, then $\mathbb{A} \in \mathcal{C}$ as well (i.e., the complement of $\mathcal{C}$ in the class of all finite τ-structures is closed under homomorphisms).

Note that CSP($\mathbb{B}$) is closed under inverse homomorphisms.

If $\mathbb{A}$ and $\mathbb{B}$ are τ-structures with disjoint domains A and B, respectively, then the *disjoint union* $\mathbb{A} \uplus \mathbb{B}$ is the τ-structure $\mathbb{C}$ with domain $A \cup B$ and the relation $R^{\mathbb{C}} = R^{\mathbb{A}} \cup R^{\mathbb{B}}$ for every $R \in \tau$. Note that CSP($\mathbb{B}$) is *closed under disjoint unions*, i.e., if $\mathbb{A} \in \mathcal{C}$ and $\mathbb{B} \in \mathcal{C}$, then $\mathbb{A} \uplus \mathbb{B} \in \mathcal{C}$.

For examples we will often use graphs and digraphs. Since this paper is deeply involved with logic and finite model theory, we will think of digraphs as binary structures with signature $\{E\}$, following and adapting standard notions from graph theory [17] to this setting. In particular, given a digraph $\mathbb{D}$, we call $E^{\mathbb{D}}$ the *edge set* of $\mathbb{D}$, and its elements we call *edges* or *arcs*. Also, a *graph* is a digraph whose edge set is a symmetric relation, and an *oriented graph* is a digraph whose edge set is an anti-symmetric relation. Given a positive integer n, we denote by K_n the complete graph on n vertices, i.e., the graph with vertex set $[n]$ and edge set (i, j) where $i \neq j$. A *tournament* is an oriented graph whose symmetric closure is a complete graph, and an *oriented path* is an oriented graph whose symmetric closure is a path. We denote by $\vec{P_n}$ the *directed path* on n vertices, i.e., the oriented path with vertex set $[n]$ and edges $(i, i+1)$ for $i \in [n-1]$.

2.2 SNP and its fragments

Let τ be a finite set of relation symbols. An *SNP τ-sentence* (short for *strict non-deterministic polynomial-time*) is a sentence of the form

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

where ψ is a quantifier free $\tau \cup \{R_1, \ldots, R_k\}$-formula. If the structure $\mathbb{A}$ satisfies the sentence Ψ, we write $\mathbb{A} \models \Psi$. We say that a class of finite τ-structures $\mathcal{C}$ is in *SNP* if there exists an SNP τ-sentence Φ such that $\mathbb{A} \models \Phi$ if and only if $\mathbb{A} \in \mathcal{C}$. Note that the class of (finite) models of an SNP sentence is closed under substructures.

A first-order τ-formula ϕ is called *positive* if it does not use the negation symbol $\neg$ (so it just uses the logical symbols $\forall, \exists, \wedge, \vee, =$, variables and symbols from τ). The negation of a positive formula is called *negative*; note that every negative formula is equivalent to a formula in prenex conjunctive normal form where every atomic formula appears in negated form, and all negation symbols are in front of atomic formulas; such formulas will also be called negative.

Consider a first-order sentence ϕ in prenex conjunctive normal form without (positive) literals of the form $x = y$. Let ψ be a clause of ϕ, and let $\mathbb{C}$ be the structure whose vertex set is the set of variables appearing in ψ, and the interpretation of $R \in \tau$ consists of those tuples $\overline{x}$ such that ψ contains the negative literal $\neg R(\overline{x})$ (this structure is sometimes called the *canonical database* of ψ). We say ψ is *connected* if $\mathbb{C}$ is connected, and that ϕ is *connected* if every clause of ϕ is connected — notice that in particular, the notion of connectedness is well-defined for negative first-order logic. Given an SNP τ-sentence Φ with first-order part ϕ in conjunctive normal form (CNF), we say that Φ is

- *connected* if ϕ is connected (so ϕ does not contain (positive) literals of the form $x = y$),
- *monotone* if every symbol from $\tau \cup \{=\}$ is negative in ϕ,
- *monadic* if every existentially quantified symbol is monadic (i.e., unary), and
- *without equalities* if ϕ does not use the equality symbol (so it just uses the logical symbols $\forall, \exists, \wedge, \vee, \neg$, variables and symbols from τ).

Notice that for every SNP sentence Φ we may assume that it does not contain literals of the form $\neg(x = y)$; otherwise, we obtain an equivalent formula by deleting every such literal and replacing each appearance of the variable y in the same clause by the variable x. Hence, in this paper we assume without loss of generality that if Φ is in monotone SNP, then it is in SNP without equalities.[2] Connected monotone SNP is closely related to CSPs in the following sense.

Theorem 2.1 (Corollary 1.4.12 in [12]). *An* SNP *sentence Φ describes a problem of the form* CSP($\mathbb{A}$) *for some (infinite) structure $\mathbb{A}$ if and only if Φ is equivalent to a connected monotone* SNP *sentence.*

2.3 Polynomial-time reductions and subclasses of NP

We consider three kinds of polynomial-time reductions which we specify now. We say that a set L *reduces in polynomial time* to a set L' if there is a polynomial-time computable function f such that $x \in L$ if and only if $f(x) \in L'$; we also say that f is a *polynomial-time many-one reduction*.

[2] We deviate from the notation from [29]: they call this fragment SNP without inequalities since they assume negation normal form and we do not.

More generally, we say that there is a *disjunctive polynomial-time reduction* from L to L' if there is a polynomial-time computable function g which computes a set of queries $g(x)$ such that $x \in L$ if and only if $y \in L'$ for some $y \in g(x)$. Similarly, if g satisfies that $x \in L$ if and only if $y \in L'$ for all $y \in g(x)$, we say that there is a *conjunctive polynomial-time reduction* from L to L'. Finally, we say that there is a *polynomial-time Turing-reduction* from L to L', if there is a polynomial-time deterministic oracle machine M that accepts L using L' as an oracle. Clearly, if there exists a polynomial-time reduction from L to L', then there are also disjunctive and conjunctive polynomial-time reductions, and, in turn, if there exists a disjunctive or a conjunctive polynomial-time reduction from L to L', then there is also a Turing-reduction from L to L'. We say that L and L' are *polynomial-time equivalent* (*disjunctive-*, *conjunctive-*, and *Turing-equivalent*, respectively) if there are polynomial-time reductions (disjunctive-, conjunctive-, and Turing-reductions, respectively) from L to L' and from L' to L. All of these notions of reducibility also exist in the variant with *log-space* instead of *polynomial-time*; in this case we require the existence of a Turing machine which has access to a (read-only) input tape, a (write-only) output tape, and a work tape whose size is logarithmic in the size of the input. Note that any function which is log-space computable in this sense is also polynomial-time computable.

3 Extensional ESO

The aim of this section is to present several families of computational problems considered in the literature that are captured by extensional ESO. First, we recall the definition of extensional ESO. An ESO τ-sentence Ψ is called *extensional* if for every existentially quantified predicate R there is a distinct symbol $R' \in \tau$ of the same arity such that Ψ contains a conjunct $\forall \bar{x}\big(R'(\bar{x}) \Rightarrow R(\bar{x})\big)$, and R' does not appear anywhere else in Ψ. In this case we say that R *extends* R'.

Actually, one can reinterpret extensional ESO as follows. Consider a signature σ which is the disjoint union of two signatures τ_1 and τ_2, and let ϕ be a first-order σ-sentence. On an input σ-structure, the task is to decide whether there is an extension $\mathbb{A}'$ of $\mathbb{A}$ where we are not allowed to modify the interpretation of symbols $R \in \tau_1$, i.e., $A' = A$, $R(\mathbb{A}') = R(\mathbb{A})$ for each $R \in \tau_1$, and $R(\mathbb{A}) \subseteq R(\mathbb{A}')$ for each $R \in \tau_2$. Clearly, these problems can be expressed with extensional ESO sentences.

3.1 Surjective CSPs

Arguably, the most natural ESO sentence describing k-colourability has k unary quantified predicates $U_1, \dots, U_k$, and the first-order part states that every vertex belongs to exactly one U_i, and that each U_i is an independent set. Such a sentence is not in extensional ESO, but k-colourability can be expressed in extensional ESO as follows.

Example 3 (k-colouring problem)**.** Let E_1 be a binary relational symbol. For every positive integer k, denote by ϕ_k the universal $\{E_1\}$-formula stating that $(V; E_1)$ is a loopless symmetric graph with no induced $K_1 + K_2$ (the three-element graph which is formed as the disjoint union of an isolated vertex and an edge) and no K_{k+1}. Hence, $(V; E_1) \models \phi_k$ if and only if $(V; E_1)$ is a *complete k-partite* graph, i.e., there is partition P of V with at most k parts such that there is an edge $uv \in E_1$ if and only if u and v lie in different parts of P. Now, consider the extensional ESO sentence

$$\Phi := \exists E_1 \left(\phi_k \wedge \forall x, y \left(E(x,y) \Rightarrow E_1(x,y)\right)\right)$$

Then a graph $(V; E)$ satisfies Φ if and only if there is an edge set $E_1 \supseteq E$ such that $(V; E_1)$ is a complete k-partite graph. It thus follows that $(V; E) \models \Phi$ if and only if $(V; E)$ is k-colourable.

Let $\mathbb{A}$ and $\mathbb{B}$ be relational τ-structures. A *full-homomorphism* is a homomorphism $f\colon \mathbb{A} \to \mathbb{B}$ such that for every r-ary relation symbol R and every r-tuple $\bar{a}$ of A it is the case that $\bar{a} \in R^{\mathbb{A}}$ if and only if $f(\bar{a}) \in R^{\mathbb{B}}$. In the following, will assume that τ is finite. We denote by $\mathrm{CSP}_{\mathrm{full}}(\mathbb{A})$ the class of finite structure $\mathbb{A}$ that admit a full homomorphism to $\mathbb{B}$.

It is straightforward to observe that $\mathbb{A}$ maps homomorphically to $\mathbb{B}$ if and only if there is an extension $\mathbb{A}'$ of $\mathbb{A}$ that admits a full-homomorphism to $\mathbb{B}$. It was proved in [5] that for every finite τ-structure $\mathbb{B}$ there is a finite set $\mathcal{F}$ of finite τ-structures such that $\mathrm{CSP}_{\mathrm{full}}(\mathbb{B})$ is the class of $\mathcal{F}$-free finite structures. In particular, there is a first-order sentence ϕ such that $\mathrm{CSP}_{\mathrm{full}}(\mathbb{B})$ coincides with the finite models of ϕ. Hence, if $\tau = \{R_1, \ldots, R_k\}$ is the signature of $\mathbb{B}$, and ϕ' is obtained from ϕ by replacing each symbol R_i by R_i', the CSP of $\mathbb{B}$ is described by the extensional ESO sentence

$$\exists R_1', \ldots, R_k' \big(\bigwedge_{R \in \tau} \forall \bar{x}\ (R(\bar{x}) \Rightarrow R'(\bar{x})) \wedge \phi' \big).$$

Surjective CSPs are the variant of CSPs where the task is to decide whether there exists a surjective homomorphism $f\colon \mathbb{A} \to \mathbb{B}$ from a given input structure $\mathbb{A}$ to a fixed template structure $\mathbb{B}$. Similarly as for CSPs one can show that for every finite structure $\mathbb{B}$ the surjective CSP for $\mathbb{B}$ is in extensional ESO as follows. For a finite τ-structure $\mathbb{B}$, let $\delta_{\mathbb{B}}$ denote the *diagram* of $\mathbb{B}$, that is, the conjunctive quantifier-free formula whose variables x_b are indexed by elements $b \in B$ and such that for each symbol $R \in \tau$ of arity r, the formula $\delta_{\mathbb{B}}$ contains the conjunct $R(x_{b_1}, \ldots, x_{b_r})$ if $(b_1, \ldots, b_r) \in R(\mathbb{B})$, and contains the conjunct $\neg R(x_{b_1}, \ldots, x_{b_r})$ otherwise. We denote by $\delta_{\mathbb{B}}'$ the existential sentence obtained from the diagram $\delta_{\mathbb{B}}$ by existentially quantifying each variable and replacing each occurrence of $R \in \tau$ by R'. Now, using the arguments above, it follows that there is a surjective homomorphism form $\mathbb{A}$ to $\mathbb{B}$ if and only if $\mathbb{A}$ satisfies the extensional ESO sentence

$$\exists R_1', \ldots, R_k' \big(\bigwedge_{R \in \tau} \forall \bar{x} (R(\bar{x}) \Rightarrow R'(\bar{x})) \wedge \phi' \wedge \delta_{\mathbb{B}}' \big).$$

3.2 Graph Sandwich Problems

Graph Sandwich Problems were introduced by Golumbic, Kaplan, and Shamir in [32]. The *sandwich problem* (SP) for a graph class $\mathcal{C}$ is the following computational problem. The input is a pair of graphs $((V, E_1), (V, E_2))$ where $E_1 \subseteq E_2$, and the task is to decide whether there is an edge set E such that $E_1 \subseteq E \subseteq E_2$ and the graph (V, E) belongs to $\mathcal{C}$. Equivalently, the SP for $\mathcal{C}$ can receives as an input a triple (V, E, N) where V is a set of vertices, E is a set of edges, and N is a set of non-edges. The task is to determine whether there is a graph $(V, E') \in \mathcal{C}$ such that $E \subseteq E'$ and $E' \cap N = \varnothing$.

Considerable attention has been placed on the SP for $\mathcal{F}$-free graphs for finite sets of graphs $\mathcal{F}$ [1, 22, 23, 24]. In particular, for every such finite set $\mathcal{F}$, there is a universal sentence ϕ such that a $\{E\}$-structure $\mathbb{G}$ satisfies ϕ if and only if $\mathbb{G}$ is an $\mathcal{F}$-free graph. Consider the binary signature $\{E, N\}$, and let ϕ' be the universal sentence obtained by replacing each occurrence of a E by E', and a every occurrence of $\neg E$ by N'. It is now straightforward to verify that the SP for $\mathcal{F}$-free graphs is described by the extensional ESO sentence

$$\exists E', N'\ \forall x, y \big((E(x,y) \Rightarrow E'(x,y)) \wedge (N(x,y) \Rightarrow N'(x,y)) \wedge (N'(x,y) \Leftrightarrow \neg E'(x,y)) \wedge \phi' \big).$$

3.3 Orientation Completion Problems

An *oriented graph* is a digraph $\mathbb{D} = (V, A)$ with no symmetric pair of edges, i.e., there are no $x, y \in V$ such that $(x, y) \in A$ and $(y, x) \in A$. Orientation completion problems were introduced by Bang-Jensen, Huang, and Zhu [6]. Given a finite set of oriented graphs $\mathcal{F}$, the orientation completion to $\mathcal{F}$-free graphs is the following. The input is a partially oriented graph $(\mathbb{G}, A')$, i.e., a graph $\mathbb{G}$ together with a set of edges $A' \subseteq E(\mathbb{G})$ such that $(V(\mathbb{G}), A')$ is an oriented graph. The task is to decide whether there exists an *$\mathcal{F}$-free orientation completion of* $(\mathbb{G}, A)$, i.e., an edge set A such that

- $A' \subseteq A$,
- (G, A) is an oriented graph that does not embed any $\mathbb{F} \in \mathcal{F}$, and
- the symmetric closure of A equals $E(\mathbb{G})$.

Clearly, for every finite set of oriented graphs $\mathcal{F}$ there is a universal $\{A\}$-sentence ϕ that describes the class of $\mathcal{F}$-free oriented graphs. Hence, the $\mathcal{F}$-free orientation problem is captured by

$$\Phi := \exists A\ \forall x, y\big((A'(x,y) \Rightarrow A(x,y))\ \wedge\ ((A(x,y) \vee A(y,x)) \Leftrightarrow E(x,y))\ \wedge\ \phi\big).$$

3.4 Constraint Topological Sorting

A *topological sort* of a directed acyclic graph $\mathbb{D}$ is a linear ordering $<$ of D such that if $(x,y) \in E(\mathbb{D})$, then $x < y$. *Constraint topological sorting problems (CTSs)* arise in the context of automata theory and regular languages [2]. These problems are described by a regular language L; the input is a directed acyclic graph $\mathbb{D}$ whose vertices are labeled with letters of some alphabet Σ and the task is to decide whether there exists a topological sort of $\mathbb{D}$ where the word obtained by reading the label of the vertices according to the order $<$ belongs to the language L. Amarilli and Paperman conjecture that CTS for regular languages exhibit an NL vs. NP-complete dichotomy [2, Conjecture 2.3]. By considering a slight modification on the problem, they show that such a dichotomy holds for aperiodic regular languages [2, Theorem 5.2]. Aperiodic regular languages correspond to regular languages expressible in first-order logic [47] where the signature consists of a unary relation symbol for each letter of the alphabet Σ and a binary relation symbol s encoding the successor relation. Clearly, the successor relation can be first-order defined from the order relation, e.g.,

$$s(x,y) := (x < y) \wedge (\forall z.\ (x < z < y \Rightarrow (z = x \vee z = y)).$$

Hence, for every aperiodic regular language L there is a first-order $\Sigma \cup \{<\}$-sentence ϕ that describes the language L. Therefore, the CTS for L is expressible in extensional ESO by

$$\exists <\ \big(\forall x, y(E(x,y) \Rightarrow x < y) \wedge \phi\big).$$

4 Extensional ESO and HerFO

Every problem in HerFO is polynomial-time equivalent to the complement of a problem described by monadic extensional ESO. This follows from the following observation [14].

Observation 4.1. *Consider a first-order τ-formula $\phi := Q_1 x_1 \dots Q_n x_n.\psi(x_1, \dots, x_n)$. If $\psi(x)$ is a first-order τ-formula, then there is a τ-formula ϕ_ψ such that a τ-structure $\mathbb{A}$ satisfies ϕ_ψ if and only if the substructure of $\mathbb{A}$ with domain $\{a \in A \colon \mathbb{A} \models \psi(a)\}$ satisfies ϕ. Namely,*

$$\phi_\psi :=\ Q_1 x_1 \dots Q_n x_n \left(\bigwedge_{i \in U} \psi(x_i)\ \Longrightarrow\ \bigwedge_{i \in E} \psi(x_i) \wedge \phi(x_1, \dots, x_n) \right),$$

where U (respectively, E) is the set of indices $i \in \{1, \dots, n\}$ such that Q_i is a universal (respectively, existential) quantifier.

We will use the formula ϕ_ψ from Observation 4.1 in the following lemma and in the proof of Theorem 4.5.

Lemma 4.2. *Consider a first-order τ-formula $\phi := Q_1x_1 \dots Q_nx_n.\psi(x_1,\dots,x_n)$. Then,* $\mathrm{Her}(\phi)$ *and the complement of the problem described by the monadic extensional $(\tau \cup \{S\})$-sentence Ψ are log-space equivalent, where*

$$\Psi := \exists \bar{S}.\, Q_1x_1 \dots Q_nx_n \left(\forall y (S(y) \Rightarrow \bar{S}(y)) \wedge \bigwedge_{i\in E} \neg \bar{S}(x_i) \wedge \big(\bigwedge_{i \in U} \neg \bar{S}(x_i) \Rightarrow \psi(x_1,\dots,x_n)\big)\right).$$

Proof. First notice that if an expansion $\mathbb{A}' := (\mathbb{A}, \bar{S})$ satisfies the first-order part of Φ and $\psi(x) := \neg\bar{S}(x)$, then $(\mathbb{A}, \bar{S}) \models \phi_\psi$ (see Observation 4.1 for the definition of ϕ_ψ). Moreover, $(\mathbb{A}, \bar{S})$ satisfies the first-order part of Φ if and only if $S^{\mathbb{A}} \subseteq \bar{S}^{\mathbb{A}'}$, $A \not\subseteq \bar{S}^{\mathbb{A}'}$, and $\mathbb{A}' \models \phi_\psi$. Hence, the log-space reduction from $\mathrm{Her}(\phi)$ to $\neg\Psi$ is straightforward: every τ-structure $\mathbb{A}$ is a $(\tau \cup \{S\})$-structure, and it thus follows via Observation 4.1 and the definition of ϕ_ψ that $\mathbb{A} \models \neg\Psi$ if and only if every non-empty substructure of $\mathbb{A}$ satisfies ϕ, i.e., if and only if $\mathbb{A} \in \mathrm{Her}(\phi)$. The converse reduction follows with similar arguments. Indeed, if $\mathbb{A}$ is a $(\tau \cup \{S\})$-structure, then $\mathbb{A} \models \neg\Psi$ if and only if $S^{\mathbb{A}} \neq A$ and the substructure $\mathbb{B}$ of $\mathbb{A}$ with domain $A \setminus S^{\mathbb{A}}$ models $\neg\Psi$. Similarly as above, $\mathbb{B} \models \neg\Psi$ if and only if $\mathbb{B}$ hereditarily satisfies ϕ. Clearly, $\mathbb{B}$ can be computed from $\mathbb{A}$ in polynomial-time, and thus $\neg\Psi$ reduces in polynomial time (and in logarithmic space) to $\mathrm{Her}(\phi)$. □

Combining this lemma and Corollary 7.8, we get the following application of our results.

Corollary 4.3. HerFO *is not* coNP*-rich, unless E = NE.*

Now we use the fact that the tractability problem for HerFO is undecidable [14] to conclude that the tractability problem for extensional ESO is also undecidable (via Lemma 4.2).

Corollary 4.4. *The tractability problem for (monadic) extensional ESO is undecidable.*

4.1 HerFO captures extensional ESO

Now we see that for every problem in extensional ESO, there is a polynomial-time equivalent problem in HerFO. That is, for every extensional ESO τ-sentence Ψ we propose a first-order σ-formula ϕ such that $\neg\Psi$ is polynomial-time equivalent to $\mathrm{Her}(\phi)$. First note that $\neg\Psi$ can be viewed as a universal second-order sentence. The idea for one reduction is that given a τ-structure $\mathbb{A}$, we construct a σ-structure $\mathbb{B}$ such that every expansion of $\mathbb{A}$ is encoded by some substructure of $\mathbb{B}$, and every substructure of $\mathbb{B}$ either encodes an expansion of $\mathbb{A}$, or it trivially satisfies ϕ. We illustrate a construction of $\mathbb{B}$ in Figure 2. Naturally, we will define ϕ in such a way that a substructure $\mathbb{B}'$ of $\mathbb{B}$ encoding an expansion $\mathbb{A}'$ of $\mathbb{A}$ satisfies ϕ if and only if $\mathbb{A}'$ satisfies the first-order part of $\neg\Psi$. Hence, we obtain a polynomial-time reduction from $\neg\Psi$ to $\mathrm{Her}(\phi)$. For the remaining reduction we need to know how to handle σ-structures that are not in the image of the previous construction. This will also be taken care of by ϕ: we will include some clauses in ϕ that will allow us to find a substructure $\mathbb{B}_1$ of $\mathbb{B}$ such that $\mathbb{B}$ hereditarily satisfies ϕ if and only if $\mathbb{B}_1$ hereditarily satisfies ϕ, and $\mathbb{B}_1$ is constructed (as before) from some τ-structure $\mathbb{A}$. This will yield a polynomial-time reduction from $\mathrm{Her}(\phi)$ to $\neg\Psi$.

Theorem 4.5. *For every sentence Ψ in extensional* ESO *there exists a first-order formula ϕ such that the problem described by $\neg\Psi$ is log-space equivalent to* $\mathrm{Her}(\phi)$*.*

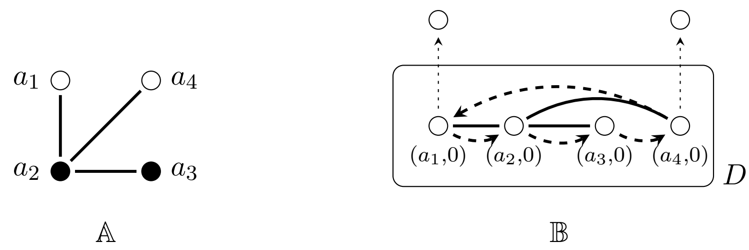

Figure 2: Let R be a binary and U' be a unary relation symbol. Consider an extensional ESO $\{R, U'\}$-formula Φ of the form $\exists U \forall x((U'(x) \Rightarrow U(x)) \wedge \phi)$ where ϕ is a first-order $\{R, U\}$-formula. On the left, we depict an $\{R, U'\}$-structure where the interpretation of U' is indicated by black vertices, and the interpretation of R is indicated by undirected edges. On the right, we depict the structure $\mathbb{B}$ with signature $\{R, D, E, <, S_U\}$, where $D^{\mathbb{B}}$ is indicated by the vertices inside the rectangle, $E^{\mathbb{B}}$ is indicated by bended dashed directed edges, $<^{\mathbb{B}}$ is the strict left-to-right linear order on $A \times \{0\}$, and the interpretation of S_U is indicated by straight dotted arcs. Notice that every substructure $\mathbb{B}'$ of $\mathbb{B}$ either does not contain a directed E-cycle, or it encodes an expansion of $\mathbb{A}$ by a unary predicate U where $U(x)$ if and only if $\mathbb{B}' \models \neg\exists y. S_U(x, y)$. Moreover, this expansion satisfies the clause $\forall x(U'(x) \Rightarrow U(x))$. We use these facts in the proof of Theorem 4.5.

Proof. Let τ be the input signature of Ψ. We may assume that Ψ is of the form

$$\exists R_1, \dots, R_n \, Q_1 x_1 \dots Q_m x_m . \psi$$

for some quantifier-free prefix $Q = (Q_1 \dots Q_m)$ and quantifier-free formula ψ in conjunctive normal form. For each $i \in [n]$ let R'_i be an r_i-ary predicate in τ such that ψ contains the clause $R'_i(y_1, \dots, y_{r_i}) \Rightarrow R_i(y_1, \dots, y_{r_i})$ where each y_i is a universally quantified variable, and the symbol R' does not appear in any other clause of ψ. We denote by τ' the subset of τ consisting of R'_i for every $i \in [n]$.

We first consider the trivial cases where either all finite structures satisfy Ψ, or no finite structure satisfies Ψ. In the former case, we may pick ϕ to be the formula that is always false, and in the latter, we pick the formula that is always true. Then we have a trivial reduction from $\mathrm{Her}(\phi)$ to $\neg\Psi$ and vice versa.

Let ρ be the signature that contains τ, a unary symbol D, two binary symbols E and $<$, and for each $i \in [n]$ a symbol S_i of arity $r_i + 1$ (as we will see below, these symbols S_i will encode the existentially quantified symbol R_i). Let ϕ be the formula $(\phi_1 \vee \phi_2) \wedge \phi_3 \wedge \phi_4 \wedge \phi_5 \wedge \phi_6$ defined as follows.

1. We define ϕ_1 as follows: first, let $\chi := \neg(Q_1 x_1 \dots Q_m x_m . \psi)$ and $\delta(x) := D(x)$; secondly, let $\chi' := \chi_\delta$ (see Observation 4.1); finally, let ϕ_1 be the formula obtained from χ' by replacing occurrences of $R_i(y_1, \dots, y_r)$ by $\forall z. \neg S_i(y_1, \dots, y_r, z)$. It will also be convenient to denote by $\phi'_1(x_1, \dots, x_m)$ the formula where the bounded variables $x_1, \dots, x_m$ are now free variables (but the universally quantified variables z remain bounded).

2. Let ϕ_2 be the formula stating that there is some y such that for all $x_1, \dots, x_m$ the formulas $\neg E(x_1, y)$ and $\neg S_i(x_1, \dots, x_{r_i}, y)$ hold for all $i \in [n]$.

3. Let ϕ_3 be the formula that asserts that if $S_i(\overline{x}, y)$ holds, then y does not appear in any other tuple of the structure.

4. Let ϕ_4 be the formula stating that if $E(x, y)$ holds, then $D(x)$ and $D(y)$ hold.

5. Let ϕ_5 be the formula stating that the in-degree and out-degree of every vertex with respect to the relation E is at most 1.

6. Let ϕ_6 be the formula stating that $<$ is a strict linear order on the subset D such that there is at most one edge (of E) *going backwards with respect to* $<$, i.e.,

$$\forall x, y, z, w \in D\big((x < y \wedge z < w \wedge E(y, x) \wedge E(w, z)) \Rightarrow (x = z \wedge y = w)\big).$$

We claim that $\neg\Psi$ and $\mathrm{Her}(\phi)$ are polynomial-time equivalent.

For the polynomial-time reduction from $\neg\Psi$ to $\mathrm{Her}(\phi)$ consider a finite τ-structure $\mathbb{A}$ with domain $A = \{a_0, \dots, a_{\ell-1}\}$. We first construct an auxiliary structure $\mathbb{F}$ with domain

$$A \times \{0\} \cup A^{r_1} \times \{1\} \cup \cdots \cup A^{r_n} \times \{n\}$$

such that

- the map $a \mapsto (a, 0)$ is an isomorphism between $\mathbb{A}$ and the $(\tau \setminus \tau')$-reduct of the substructure of $\mathbb{F}$ with domain $A \times \{0\}$,
- $D(x)$ holds for every $x \in A \times \{0\}$,
- $E((a_{i-1}, 0), (a_i, 0))$ holds for every $i \in [\ell]$ (where indices are modulo ℓ),
- $(a_i, 0) < (a_j, 0)$ if and only if $i < j$, for every $i, j \in \{0, \dots, \ell - 1\}$,
- $S_i((a_{j_1}, 0), \dots, (a_{j_{r_i}}, 0), (a_{j_1}, \dots, a_{j_{r_i}}, i))$ holds for all $i \in [n]$ and all $a_{j_1}, \dots, a_{j_{r_i}} \in A$, and
- no other relations hold in $\mathbb{F}$.

Now, we consider the substructure $\mathbb{B}$ of $\mathbb{F}$ with domain

$$B := \{(a, 0) \mid a \in A\} \cup \{(\bar{a}, i) \mid \mathbb{A} \models \neg R_i'(\bar{a})\}.$$

Clearly, $\mathbb{B}$ can be computed from $\mathbb{A}$ in polynomial time and logarithmic space — we provide an illustration of this construction in Figure 2. We have to argue that $\mathbb{A} \models \neg\Psi$ if and only if all substructures of $\mathbb{B}$ satisfy ϕ. Note that $\mathbb{B}$ satisfies ϕ_3, ϕ_4, ϕ_5, and ϕ_6 by construction. Since $\phi_3 \wedge \phi_4 \wedge \phi_5 \wedge \phi_6$ is a universal sentence, all substructures of $\mathbb{B}$ satisfy $\phi_3 \wedge \phi_4 \wedge \phi_5 \wedge \phi_6$. Now, suppose that a substructure $\mathbb{B}'$ of $\mathbb{B}$ does not contain $(a_i, 0)$ for some $i \in \{0, \dots, \ell - 1\}$. First consider the case that $(a_j, 0) \in B'$ for some $j \in \{0, \dots, \ell - 1\}$. It follows that $\mathbb{B}'$ satisfies ϕ_2, because the vertex $(a_k, 0)$ where k is the first index i' from $i, i + 1, \dots, j$ (where indices are modulo ℓ) such that $(a_{i'}, 0)$ belongs to the domain of $\mathbb{B}'$ provides a witness for the existentially quantified variable y in ϕ_2 (the conjuncts $\neg S_i(x_1, \dots, x_{r_i}, y)$ are satisfied by $(a_k, 0)$ because no element from $A \times \{0\}$ is the last coordinate of some tuple in $S_i^{\mathbb{B}}$). If $B' \cap (A \times \{0\}) = \varnothing$, then any element $b \in B'$ is a witness for the existentially quantified variable y in ϕ_2. This proves that if $A \times \{0\} \not\subseteq B'$, then $\mathbb{B}'$ satisfies ϕ_2. Conversely, if $A \times \{0\} \subseteq B'$, then $\mathbb{B}'$ does not satisfy ϕ_2. Therefore, we have to argue that $\mathbb{A} \models \neg\Psi$ if and only if all substructures $\mathbb{B}'$ of $\mathbb{B}$ with $A \times \{0\} \subseteq B'$ satisfy ϕ_1.

To see this, consider the following mapping e from substructures $\mathbb{B}'$ of $\mathbb{B}$ with $A \times \{0\} \subseteq B'$ to expansions of $\mathbb{A}$. For such a $\mathbb{B}'$, let $\mathbb{C} := e(\mathbb{B}')$ be the expansion of $\mathbb{A}$ where $R_i(x_1, \dots, x_{r_i})$ holds if $(x_1, \dots, x_{r_i}, i) \notin B'$, equivalently, if $\neg S_i((x_1, 0), \dots, (x_{r_i}, 0), y)$ holds in $\mathbb{B}'$ for all $y \in B'$. It is straightforward to verify that e is injective. Since $\mathbb{B}'$ is a substructure of $\mathbb{B}$, then $\mathbb{C}$ satisfies the clause $R'(\bar{x}) \Rightarrow R(\bar{x})$ for every $R \in \tau$. Moreover, the image of e is exactly the set of expansions of $\mathbb{A}$

that satisfy $R'(\bar{x}) \Rightarrow R(\bar{x})$ for every $R \in \tau$. Hence, we regard e as a bijection between substructures of $\mathbb{B}$ containing $A \times \{0\}$ and the expansions of $\mathbb{A}$ satisfying $R'(\bar{x}) \Rightarrow R(\bar{x})$ for every $R \in \tau$.

It follows from the definition of ϕ_1 and of $\mathbb{C}$ that for every $a_{j_1}, \dots, a_{j_m} \in A$

$$\mathbb{C} \models \neg\psi(a_{j_1}, \dots, a_{j_m}) \text{ if and only if } \mathbb{B}' \models \phi_1'((a_{j_1}, 0), \dots, (a_{j_m}, 0)).$$

Then, since $A \times \{0\} \subseteq B'$, it follows that $\mathbb{B}' \models \phi_1$ if and only if $\mathbb{C} \models \neg(Q_1 x_1 \dots Q_m x_m.\psi)$. Therefore, if $\mathbb{A} \models \neg\Psi$ and $\mathbb{B}'$ is a substructure of $\mathbb{B}$ with $A \times \{0\} \subseteq B'$, then $\mathbb{C} \models \neg(Q_1 x_1 \dots Q_m x_m.\psi)$ and thus $\mathbb{B}' \models \neg\phi_1$. This implies that if $\mathbb{A} \models \neg\Psi$, then $\mathbb{B} \in \mathrm{Her}(\phi)$. Conversely, suppose $\mathbb{A} \models \Psi$, and let $\mathbb{C}$ be an $\{R_1, \dots, R_n\}$-expansion of $\mathbb{A}$ such that $\mathbb{C} \models Q_1 x_1 \dots Q_m x_m.\psi$. In particular, $\mathbb{C}$ satisfies $R'(\bar{x}) \Rightarrow R(\bar{x})$. Hence, there is a substructure $\mathbb{B}'$ of $\mathbb{B}$ such that $e(\mathbb{B}') = \mathbb{C}$. By the previous arguments we conclude that $\mathbb{B}' \models \neg\phi_1$, and since $A \times \{0\} \subseteq B'$, we see that $\mathbb{B}' \models \neg\phi_2$, and so $\mathbb{B}' \models \neg\phi$. Therefore, $\mathbb{A} \models \neg\Psi$ if and only if $\mathbb{B} \in \mathrm{Her}(\phi)$, and since $\mathbb{B}$ can be constructed in logarithmic space from $\mathbb{A}$, we have a log-space reduction from $\neg\Psi$ to $\mathrm{Her}(\phi)$.

We now present a log-space reduction from $\mathrm{Her}(\phi)$ to $\neg\Psi$. Recall that there are structures $\mathbb{D}$ and $\mathbb{F}$ such that $\mathbb{D} \models \Psi$ and $\mathbb{F} \models \neg\Psi$ — we fix $\mathbb{D}$ and $\mathbb{F}$ for this reduction. Let $\mathbb{B}$ be a finite ρ-structure. If $\mathbb{B}$ does not satisfy $\phi_3 \wedge \phi_4 \wedge \phi_5 \wedge \phi_6$, then the reduction returns $\mathbb{D}$. Otherwise, let $X \subseteq B$ be the set of elements x for which there is no $\overline{w}$ such that $S_i(\overline{w}, x)$ holds for some $i \in [n]$. In particular, for every element b in $B \setminus X$, there is some tuple $\overline{w}$ such that $\mathbb{B} \models S_i(\overline{w}, b)$ for some $i \in [n]$, and since $\mathbb{B}$ satisfies ϕ_3, it follows that every element of $\mathbb{B}$ that belongs to some tuple of E belongs to X. Also, since $\mathbb{B}$ satisfies ϕ_5, the $\{E\}$-reduct of the substructure $\mathbb{X}$ of $\mathbb{B}$ with domain X is a disjoint union of directed paths and directed cycles. Similarly as above, one can note that if there is no directed cycle, then all substructures of $\mathbb{B}$ satisfy ϕ_2, and hence $\mathbb{B}$ hereditarily satisfies ϕ. In this case, the reduction returns the fixed τ-structure $\mathbb{F}$ that does not model Ψ. Otherwise, let $X_1, \dots, X_k$ be the subsets of X inducing a directed cycle in the $\{E\}$-reduct of $\mathbb{X}$. We claim that $k = 1$. First note that if $x \in X_i$ for some $i \in [k]$, then $\mathbb{B} \models D(x)$, because $\mathbb{B}$ satisfies ϕ_4. Hence, the union $\bigcup_{i \in [k]} X_i \subseteq D$ is linearly ordered by $<$, because $\mathbb{B}$ satisfies ϕ_6. Moreover, ϕ_6 also guarantees that there is at most one edge going backwards, i.e., an edge $(y, x) \in E$ such that $x < y$. Clearly, for every linearly ordered directed cycle X_i there is at least one edge going backwards, and so $k = 1$. We construct a τ-structure $\mathbb{A}$ with domain X_1 as follows.

- We start from the substructure of $\mathbb{B}$ with domain X_1.
- For every $i \in \{1, \dots, n\}$, add $(x_1, \dots, x_{r_i})$ to the relation R_i' whenever there is no $y \in B$ such that $S_i(x_1, \dots, x_{r_i}, y)$.
- Finally, $\mathbb{A}$ is the τ-reduct of the resulting structure.

Clearly, $\mathbb{A}$ can be computed from $\mathbb{B}$ in logarithmic space. Now, we argue that $\mathbb{A} \models \neg\Psi$ if and only if every substructure of $\mathbb{B}$ satisfies ϕ. Since $\mathbb{B} \models \phi_3 \wedge \phi_4 \wedge \phi_5 \wedge \phi_6$ and $\phi_3 \wedge \phi_4 \wedge \phi_5 \wedge \phi_6$ is a universal sentence, we have to argue that $\mathbb{A} \models \neg\Psi$ if and only if every substructure of $\mathbb{B}$ satisfies $\phi_1 \vee \phi_2$. Similarly as above, notice that a substructure $\mathbb{B}'$ of $\mathbb{B}$ satisfies ϕ_2 if and only if $B' \cap X$ is not X_1.

Let $\mathbb{B}_1$ be the substructure of $\mathbb{B}$ with vertex set $(B \setminus X) \cup X_1$. As explained above, $\mathbb{B} \in \mathrm{Her}(\phi)$ if and only if all substructures $\mathbb{B}_1'$ of $\mathbb{B}_1$ with $X_1 \subseteq B_1'$ satisfy ϕ_1. Similarly as above, we define a mapping e from substructures of $\mathbb{B}_1$ containing X_1 to expansions of $\mathbb{A}$ satisfying $R'(\bar{x}) \Rightarrow R(\bar{x})$. In this case, $e(\mathbb{B}_1')$ is the expansion of $\mathbb{A}$ where $R_i(x_1, \dots, x_{r_i})$ holds whenever there is no $y \in B_1'$ such that $S_i(x_1, \dots, x_{r_i}, y)$. Notice that in this case, the mapping e is surjective but need not be injective. However, it follows with similar arguments as before that $\mathbb{B}_1' \models \phi_1$ if and only if $e(\mathbb{B}_1') \models \neg(Q_1 x_1 \dots Q_m x_m.\psi)$. We thus conclude (via the mapping e) that $\mathbb{B}_1$ contains a structure that does not model ϕ_1 if and only if $\mathbb{A}$ has an expansion that satisfies $Q_1 x_1 \dots Q_m x_m.\psi$. It now follows that $\mathbb{A} \models \neg\Psi$ if and only if $\mathbb{B} \in \mathrm{Her}(\psi)$. Thus, $\mathrm{Her}(\phi)$ reduces in logarithmic space to $\neg\Psi$. $\square$

Theorem 4.6. *The following classes have the same computational power up to log-space equivalence, i.e., for every problem P in any of the following classes, there exists a log-space equivalent problem in each of the other ones.*

- *Monadic extensional* ESO*,*
- *extensional* ESO*, and*
- *complements of problems in* HerFO*.*

Proof. This statement is an immediate consequence of the fact that monadic extensional ESO is a fragment of extensional ESO, of Lemma 4.2, and of Theorem 4.5. $\square$

5 Monotone extensional SNP

Naturally, *extensional SNP* is the fragment of extensional ESO sentences whose first-order part is universal. Equivalently, an SNP sentence Ψ with first-order part ψ in conjunctive normal form is extensional if and only if for every existentially quantified symbol R there is a distinct symbol $R' \in R$ such that ψ contains the clause $R'(\overline{x}) \Rightarrow R(\overline{x})$, and R' does not appear in any other clause of ψ.

In this section we show that CSPs of finitely bounded structures are captured by extensional SNP (Corollary 5.5) and conclude with the proof of Theorem 5.7. For this purpose, it will be convenient to work with a slight generalization of CSPs: instead of $\mathrm{CSP}(\mathbb{B})$ for some relational structure $\mathbb{B}$, we consider $\mathrm{CSP}(\phi)$ for some universal formula ϕ, or, more generally, for some first-order theory T (see, e.g., [12, Chapter 1]).

An instance of $\mathrm{CSP}(\phi)$ (and of $\mathrm{CSP}(T)$) is a primitive positive formula ψ, and the task is to decide whether $\phi \wedge \psi$ (resp. $T \cup \{\psi\}$) is satisfiable. Equivalently, the instance is a structure $\mathbb{A}$ and the task is to decide if there is a (possibly infinite) model $\mathbb{B}$ of ϕ such that $\mathbb{A} \to \mathbb{B}$. Note that this generalizes $\mathrm{CSP}(\mathbb{B})$ for finite structure $\mathbb{B}$, since for every finite structure $\mathbb{B}$ there exists a universal sentence ϕ such that $\mathrm{CSP}(\mathbb{B})$ and $\mathrm{CSP}(\phi)$ is the same computational problem; the converse is not true in general. We first characterize monotone extensional SNP in terms of constraint satisfaction problems of universal formulas (Theorem 5.4).

5.1 Finitely bounded structures and full homomorphisms

This section contains a technical lemma that we use to characterize the expressive power of monotone extensional SNP in terms of CSPs of universal sentences (Theorem 5.4). On the way, we obtain a new (and simpler) proof of a theorem proved by Ball, Nešetřil, and Pultr [5, Theorem 3.3].

Consider a τ-structure $\mathbb{A}$. We say that two distinct vertices $a, b \in A$ are *twins* if for every $R \in \tau$ of arity r, for $(a_1, \ldots, a_r) \in A^r$ and every $i \in [r]$, it is the case that $(a_1, \ldots, a_{i-1}, a, a_{i+1}, \ldots, a_r) \in R^{\mathbb{A}}$ if and only if $(a_1, \ldots, a_{i-1}, b, a_{i+1}, \ldots, a_r) \in R^{\mathbb{A}}$. Following graph theoretic nomenclature, we say that $\mathbb{A}$ is *point-determining* if $\mathbb{A}$ contains no twins.

Notice that for every finite relational signature τ, there is an existential first-order τ-formula $\mathrm{NT}(x, y)$ such that $\mathbb{A} \models \mathrm{NT}(a, b)$ if and only if a and b are not twins in $\mathbb{A}$. Indeed, one can construct such a formula by considering the disjunction $\bigvee_{R\in\tau} \psi_R(x, y)$ where for each $R \in \tau$ of arity r, the formula $\psi_R(x, y)$ equals

$$\exists z_1, \ldots z_r \big((R(x, z_2, \ldots, z_r) \Leftrightarrow \neg R(y, z_2, \ldots, z_r)) \vee \cdots \vee (R(z_1, \ldots, z_{r-1}, x) \Leftrightarrow \neg R(z_1, \ldots, z_{r-1}, y))\big).$$

In particular, $\mathrm{NT}(x, y)$ implies the inequality $x \neq y$, and if $\mathbb{A}$ is point-determining, then $\mathbb{A} \models \mathrm{NT}(a, b)$ if and only if $a \neq b$.

We say that a structure $\mathbb{A}^*$ is a *blow-up* of a structure $\mathbb{A}$ if $\mathbb{A}$ is a substructure of $\mathbb{B}$, and every $b \in B \setminus A$ has a twin a in A. We extend this notion to a class of structures $\mathcal{C}$, and call the class $\mathcal{C}^*$ of structures $\mathbb{B}$ that are a blow-up of some $\mathbb{A} \in \mathcal{C}$ the *blow-up of* $\mathcal{C}$.

Recall that a homomorphism $f \colon \mathbb{A} \to \mathbb{B}$ is called *full* (sometimes also *strong*) if it also preserves the relation $\neg R$ for each $R \in \tau$. Notice that blow-ups and full homomorphisms are naturally related: there is a surjective full homomorphism $f \colon \mathbb{A} \to \mathbb{B}$ if and only if $\mathbb{A}$ is isomorphic to a blow-up of $\mathbb{B}$. It is straightforward to observe that for every finite structure $\mathbb{B}$, there is (up to isomorphism) a unique point-determining structure $\mathbb{B}_0$ such that $\mathbb{B}_0$ embeds into $\mathbb{B}$ and $\mathbb{B}$ is isomorphic to a blow-up of $\mathbb{B}_0$ — such a structure can be obtained by repeatedly removing twins from $\mathbb{B}$. It follows that if $\mathcal{C}$ is a hereditary class, then $\mathbb{B} \in \mathcal{C}^*$ if and only if $\mathbb{B}_0 \in \mathcal{C}$.

Lemma 5.1. *Let $\mathcal{C}$ be a hereditary class of relational structures with a finite domain and signature. If $\mathcal{C}$ is finitely bounded, then the blow-up of $\mathcal{C}$ is finitely bounded. Moreover, there is a universal sentence ϕ without equalities such that a finite structure $\mathbb{A}$ satisfies ϕ if and only if $\mathbb{A} \in \mathcal{C}^*$.*

Proof. Given a finite τ-structure $\mathbb{A}$, the *diagram* of $\mathbb{A}$ is the τ-formula $\mathrm{diam}_{\mathbb{A}}$ defined as follows.

- $\mathrm{diam}_{\mathbb{A}}$ has a free variable x_a for every $a \in A$,
- for every $R \in \tau \cup \{\neq\}$ of arity r and each tuple $(a_1, \ldots, a_r)$ of A, the formula $\mathrm{diam}_{\mathbb{A}}$ contains the conjunct $R(x_{a_1}, \ldots, x_{a_r})$ if $\mathbb{A} \models R(a_1, \ldots, a_r)$,
- for every $R \in \tau$ of arity r and each tuple $(a_1, \ldots, a_r)$ of A, the formula $\mathrm{diam}_{\mathbb{A}}$ contains the conjunct $\neg R(x_{a_1}, \ldots, x_{a_r})$ if $\mathbb{A} \models \neg R(a_1, \ldots, a_r)$.

It is straightforward to observe that a τ-structure $\mathbb{B}$ satisfies $\exists \overline{x}.\, \mathrm{diam}_{\mathbb{A}}(\overline{x})$ if and only if $\mathbb{A}$ embeds into $\mathbb{B}$.

Building on the concept of the diagram of $\mathbb{A}$, we construct the *point-determining diagram* of $\mathbb{A}$, which is the formula $\mathrm{diam}^{\mathrm{pd}}_{\mathbb{A}}$ obtained from $\mathrm{diam}_{\mathbb{A}}$ by replacing each conjunct $x_a \neq x_b$ by the formula $\mathrm{NT}(x_a, x_b)$. We claim that the following statements are equivalent.

- $\mathbb{B} \models \exists \overline{x}.\, \mathrm{diam}^{\mathrm{pd}}_{\mathbb{A}}(\overline{x})$,
- $\mathbb{B}_0 \models \exists \overline{x}.\, \mathrm{diam}^{\mathrm{pd}}_{\mathbb{A}}(\overline{x})$, and
- $\mathbb{A}$ embeds into $\mathbb{B}_0$.

Since $\mathbb{B}_0$ is a point-determining structure, $\mathbb{B}_0$ satisfies $\mathrm{NT}(a, b) \Leftrightarrow a \neq b$ for every pair $a, b \in B_0$. Hence, for every tuple $\overline{b}$ of $\mathbb{B}_0$ it is the case that $\mathbb{B}_0 \models \mathrm{diam}^{\mathrm{pd}}_{\mathbb{A}}(\overline{b})$ if and only if $\mathbb{B} \models \mathrm{diam}_{\mathbb{A}}(\overline{b})$. Therefore, the last two itemized statements are equivalent, and since $\mathbb{B}_0$ embeds into $\mathbb{B}$, it follows that the second itemized statement implies the first one. Now, suppose that there is a tuple $\overline{b}$ such that $\mathbb{B} \models \mathrm{diam}^{\mathrm{pd}}_{\mathbb{A}}(\overline{b})$. Let $f \colon \mathbb{B} \to \mathbb{B}_0$ be the full homomorphism such that $f(a) = f(b)$ if and only $a = b$ or if a and b are twins. It follows from the definition of f that $\mathbb{B} \models \mathrm{NT}(a, b)$ if and only if $\mathbb{B}_0 \models \mathrm{NT}(f(a), f(b))$, and since f is a full homomorphism for every $R \in \tau$ of arity r and each r-tuple $\overline{x}$ of B it is the case that $\mathbb{B} \models R(\overline{x})$ if and only if $\mathbb{B}_0 \models R(f(\overline{x}))$. Therefore, since $\mathbb{B} \models \mathrm{diam}^{\mathrm{pd}}_{\mathbb{A}}(\overline{b})$, we conclude that $\mathbb{B}_0 \models \mathrm{diam}^{\mathrm{pd}}_{\mathbb{A}}(f(\overline{b}))$.

To show the statement of the lemma, let $\mathcal{F}$ be a finite set of bounds of $\mathcal{C}$, and let ϕ be the disjunction of the point-determining diagrams $\mathrm{diam}^{\mathrm{pd}}_{\mathbb{F}}$ ranging over all $\mathbb{F} \in \mathcal{F}$. It follows from the claim above that a finite structure $\mathbb{B}$ models $\exists \overline{x}.\phi(\overline{x})$ if and only if there is a structure $\mathbb{F} \in \mathcal{F}$ such that $\mathbb{F}$ embeds into $\mathbb{B}_0$. Equivalently, $\mathbb{B} \models \exists \overline{x}.\phi(\overline{x})$ if and only if $\mathbb{B}_0$ does not belong to $\mathcal{C}$. Since

$\mathbb{B} \in \mathcal{C}^*$ if and only if $\mathbb{B}_0 \in \mathcal{C}$, we conclude that $\mathbb{B} \models \forall x. \neg\phi(\overline{x})$ if and only if $\mathbb{B} \in \mathcal{C}^*$. This proves the moreover statement of this lemma. In particular, $\mathbb{B} \in \mathcal{C}^*$ if and only if every substructure of $\mathbb{B}$ with at most $|x|$ vertices belongs to $\mathcal{C}^*$. It thus follows that $\mathcal{C}^*$ is finitely bounded. $\square$

Before stating our intended application of this lemma, we state two implications of Lemma 5.1 of independent interest. We denote by $\mathrm{CSP}_{\mathrm{full}}(\mathbb{B})$ the class of finite structures that admit a full homomorphism to $\mathbb{B}$. *Closure under inverse full homomorphisms* is defined analogously to closure under homomorphisms.

Corollary 5.2. *Let $\mathbb{B}$ be a relational structure with finite signature. If $\mathbb{B}$ is finitely bounded, then $\mathrm{CSP}_{\mathrm{full}}(\mathbb{B})$ is finitely bounded.*

Proof. This follows from Lemma 5.1, because $\mathrm{CSP}_{\mathrm{full}}(\mathbb{B})$ is the blow-up of the age of $\mathbb{B}$. $\square$

The second application of Lemma 5.1 is a new proof of the following theorem proved by Ball, Nešetřil, and Pultr [5, Theorem 3.3]. Using their notation, given a set of structures $\mathcal{B}$, we denote the union $\bigcup_{\mathbb{B}\in\mathcal{B}} \mathrm{CSP}_{\mathrm{full}}(\mathbb{B})$ by $\mathrm{CSP}_{\mathrm{full}}(\mathcal{B})$.

Corollary 5.3. *For every finite set $\mathcal{B}$ of finite relational structures with finite signature, there is a finite set of structures $\mathcal{A}$ such that* $\mathrm{Forb}_{\mathrm{full}}(\mathcal{A}) = \mathrm{CSP}_{\mathrm{full}}(\mathcal{B})$.

Proof. Let $\mathcal{C}$ be the union of the ages of $\mathbb{B}$ over all $\mathbb{B} \in \mathcal{B}$. Notice that if m is the size of the largest structure in $\mathcal{B}$, then no structure on at least $m+1$ vertices belongs to $\mathcal{C}$, and hence $\mathcal{C}$ is finitely bounded. Since $\mathrm{CSP}_{\mathrm{full}}(\mathcal{B})$ is the blow-up of $\mathcal{C}$, there is a finite set of bounds $\mathcal{A}$ to $\mathrm{CSP}_{\mathrm{full}}(\mathcal{B})$ (Lemma 5.1).

Moreover, since $\mathrm{CSP}_{\mathrm{full}}(\mathcal{B})$ is preserved under inverse full homomorphisms, we may assume that every structure in $\mathcal{A}$ is point-determining, and thus, a structure $\mathbb{A} \in \mathcal{A}$ embeds into a structure $\mathbb{C}$ if and only if there is a full homomorphism $\mathbb{A} \to \mathbb{C}$. The claim now follows. $\square$

5.2 Monotone extensional SNP and CSPs

Given a finite relational signature τ and a first-order τ-sentence ϕ, we denote by $\mathrm{CSP}(\phi)$ the class of finite τ-structures $\mathbb{A}$ for which there is a τ-structure $\mathbb{B}$ such that $\mathbb{B} \models \phi$ and $\mathbb{A} \to \mathbb{B}$. There are first-order sentences such that $\mathrm{CSP}(\phi)$ is undecidable.[3] Nonetheless, if one restricts to a universal sentence ϕ, then $\mathrm{CSP}(\phi)$ is in NP. Indeed, suppose that for input $\mathbb{A}$ there is a structure $\mathbb{B}$ such that $\mathbb{A} \to \mathbb{B}$ and $\mathbb{B} \models \phi$. Since ϕ is universal, it follows that any substructure of $\mathbb{B}$ also models ϕ, and thus, by guessing a structure on at most $|A|$ vertices, one can verify in polynomial time whether $\mathbb{A} \in \mathrm{CSP}(\phi)$. We propose the following characterization of monotone extensional SNP.

Theorem 5.4. *For a class $\mathcal{C}$ of relational structures with finite signature, the following statements are equivalent.*

- *$\mathcal{C}$ is in monotone extensional* SNP*.*
- *There is a universal sentence ϕ such that $\mathcal{C} = \mathrm{CSP}(\phi)$.*

[3] To see this, suppose for contradiction that $\mathrm{CSP}(\phi)$ is decidable for every first-order formula ϕ. Notice that the structure $\mathbb{A}$ with one-element domain and empty interpretation of every $R \in \tau$ belongs to $\mathrm{CSP}(\phi)$ if and only if there is a non-empty model of ϕ. Hence, given a first-order formula ϕ we can check whether $\mathbb{A} \in \mathrm{CSP}(\phi)$ to decide whether ϕ is satisfied in a non-empty structure, contradicting the undecidablility of first-order logic [18].

Proof. Let τ be a finite relational signature and let $\mathcal{C}$ be a class of finite τ-structures. First, suppose that there is a monotone extensional SNP τ-sentence Φ describing the class $\mathcal{C}$. Let σ be the signature containing exactly the existentially quantified symbols from Φ and let $\tau' \subseteq \tau$ be the signature that contains the symbols R' that are extended by some $R \in \sigma$. Let $\nu \wedge \psi$ be the first-order part of Φ where ν is the conjunction of the clauses $R' \Rightarrow R$ for every $R \in \sigma$ and ψ is a universal $(\sigma \cup \tau \setminus \tau')$-formula. Let ψ' be the τ-formula obtained from ψ by replacing each existentially quantified symbol R by R'. We claim that a finite τ-structure $\mathbb{A}$ satisfies Φ if and only if $\mathbb{A} \in \mathrm{CSP}(\psi')$.

Let $\mathbb{B}$ be an expansion of $\mathbb{A}$ witnessing that $\mathbb{A}$ satisfies Φ. Consider the structure $\mathbb{A}'$ with domain A where $R^{\mathbb{A}'} = R^{\mathbb{A}}$ for every $R \in \tau \setminus \tau'$, and where $(R')^{\mathbb{A}'} = R^{\mathbb{B}}$ for every $R' \in \tau'$. Clearly, $\mathbb{A}' \models \psi'$ because $\mathbb{B} \models \psi$, and since $\mathbb{B} \models \nu$, it follows that the identity on A defines a homomorphism $\mathbb{A} \to \mathbb{A}'$, and thus $\mathbb{A} \in \mathrm{CSP}(\psi')$.

Conversely, suppose that $\mathbb{A} \in \mathrm{CSP}(\psi')$, and let $\mathbb{B}$ be a model of ψ' such that $\mathbb{A} \to \mathbb{B}$. Since ψ' is a universal sentence, we may assume without loss of generality that $\mathbb{B}$ is a finite structure. Moreover, since Φ is in monotone extensional SNP, it is preserved under inverse homomorphisms (see, e.g., Theorem 2.1). Thus, in order to show that $\mathbb{A}$ satisfies Φ, it suffices to show that every finite structure $\mathbb{B}$ that models ψ' also models Φ. This is immediate: consider the $\sigma \cup \tau$-expansion $\mathbb{B}'$ of $\mathbb{B}$ where there interpretation of each $R \in \sigma$ coincides with the interpretation of R' in $\mathbb{B}$; clearly, $\mathbb{B}' \models \nu$, and since $\mathbb{B} \models \psi'$, it follows that $\mathbb{B}' \models \psi$. This proves that the first item implies the second one.

Now, suppose that the second itemized statement holds, and let $\mathcal{B}$ be the class of finite structures that model ϕ. Again, since ϕ is a universal τ-sentence, a τ-structure $\mathbb{A}$ belongs to $\mathrm{CSP}(\phi)$ if and only if there is a structure $\mathbb{B} \in \mathcal{B}$ such that $\mathbb{A} \to \mathbb{B}$. In particular, every blow-up $\mathbb{B}^*$ of a structure $\mathbb{B} \in \mathcal{B}$ belongs to $\mathrm{CSP}(\phi)$ (see Section 5.1). In turn, this implies that a structure $\mathbb{A}$ belongs to $\mathrm{CSP}(\phi)$ if and only if there is an injective homomorphism $f \colon \mathbb{A} \to \mathbb{B}$ for some structure $\mathbb{B} \in \mathcal{B}^*$. Let ψ' be a universal τ-sentence without equalities that defines $\mathcal{B}^*$ (see Lemma 5.1). Consider a signature σ containing a relation symbol R^* for each $R \in \tau$, and the formula ψ obtained from ψ' by replacing each occurrence of a symbol $R \in \tau$ by R^*. Finally, let Ψ be the SNP τ-sentence with existentially quantified symbols from σ and first-order part $\nu \wedge \psi$, where ν is the conjunction of clauses $R \Rightarrow R^*$ ranging over $R \in \tau$. Similarly as above one can notice that a finite structure $\mathbb{A}$ satisfies Φ if and only if $\mathbb{A} \in \mathrm{CSP}(\phi)$: every $\mathbb{B} \in \mathcal{B}^*$ satisfies Ψ (indeed, define an expansion of $\mathbb{B}$ satisfying $\nu \wedge \psi$ by letting $R^* = R$ for each $R^* \in \sigma$), and since Ψ is in monotone SNP, we conclude that every structure $\mathbb{A}$ in $\mathrm{CSP}(\phi)$ satisfies Ψ. Conversely, if $\mathbb{B}$ satisfies Φ, then by considering the structure $\mathbb{B}'$ where $R^{\mathbb{B}'} = (R^*)^{\mathbb{B}}$ for each $R \in \tau$, we obtain a structure $\mathbb{B}'$ that satisfies ϕ, and such that $\mathbb{B} \to \mathbb{B}'$. Clearly, Φ is in monotone extensional SNP, and therefore, the second itemized statement implies the first one. □

Recall from the introduction that a structure $\mathbb{B}$ is finitely bounded if and only if there exists a finite set $\mathcal{F}$ of finite structures such that a finite structure $\mathbb{A}$ embeds into $\mathbb{B}$ if and only if no structure from $\mathcal{F}$ embeds into $\mathbb{A}$. Note that CSPs of finitely bounded structures are in NP.

Corollary 5.5. *Let $\mathbb{S}$ be a relational structure in a finite signature. If $\mathbb{S}$ is finitely bounded, then* $\mathrm{CSP}(\mathbb{S})$ *is in monotone extensional* SNP*. Moreover, if $\mathbb{S}$ has a finite set of connected bounds, then* $\mathrm{CSP}(\mathbb{S})$ *is in monotone connected extensional* SNP*.*

Proof. Since $\mathbb{S}$ is finitely bounded, there is a universal sentence ϕ such that a finite structure models ϕ if and only if $\mathbb{A}$ embeds into $\mathbb{S}$. Clearly, $\mathrm{CSP}(\mathbb{S}) \subseteq \mathrm{CSP}(\phi)$. We briefly argue that this inclusion is actually an equality. Let $\mathbb{B}$ be a model of ϕ such that there is homomorphism $f \colon \mathbb{A} \to \mathbb{B}$. Since ϕ is universal, it follows that the substructure $\mathbb{B}'$ of $\mathbb{B}$ with vertex set $f[A]$ is a finite model of ϕ such that $\mathbb{A} \to \mathbb{B}'$. We can conclude that $\mathbb{A} \in \mathrm{CSP}(\mathbb{S})$, because $\mathbb{B}'$ is a finite structure satisfying ϕ, so $\mathbb{B}'$ embeds into $\mathbb{S}$, and thus $\mathbb{A} \to \mathbb{S}$. The claim of the corollary now follows from Theorem 5.4. Finally, the moreover statement follows analogously by noticing that the construction of the formula ϕ from Lemma 5.1 satisfies that if the finite bounds of $\mathcal{C}$ are connected, then ϕ is a connected universal formula defining $\mathcal{C}^*$. □

5.3 CSPs of finitely bounded structures

Corollary 5.5 asserts that CSPs of finitely bounded structures are in monotone extensional SNP. However, there are monotone extensional SNP sentences that do not describe the CSP of a finitely bounded structure, because the class of finite models of such a sentence need not be preserved under disjoint unions. In this section we show that nonetheless, CSPs of finitely bounded structures capture the computational power of monotone extensional SNP, up to polynomial-time conjunctive-equivalence.

Theorem 2.1 guarantees that every monotone connected extensional SNP sentence Φ describes the CSP of some structure. Moreover, it follows from the proof of Theorem 5.4 that Φ describes the CSP of a finitely bounded structure.

Corollary 5.6. *Every monotone connected extensional SNP sentence Φ describes a problem of the form $\mathrm{CSP}(\mathbb{B})$ for some finitely bounded structure $\mathbb{B}$.*

In Appendix B, we include some results regarding the connected fragment of monotone extensional SNP. In particular, we show that every finite domain CSP is described by a monotone connected extensional SNP sentence (Corollary B.4), and we leave as an open question whether for every finitely bounded structure $\mathbb{B}$, there is a monotone connected extensional SNP sentence describing $\mathrm{CSP}(\mathbb{B})$. We also show that for every finitely bounded CSP there is a polynomial-time equivalent problem described by a monotone connected extensional SNP sentence (Lemma B.1); we use this fact in our next proof.

Theorem 5.7. *The following classes have the same computational power up to polynomial-time conjunctive-equivalence, i.e., for every problem P in any of the following classes, there are polynomial-time conjunctive-equivalent problems in each of the other ones.*

- *Monotone extensional* SNP*.*
- *Monotone connected extensional* SNP*.*
- CSP*s of finitely bounded structures.*

Moreover, the last two classes have the same computational power up to log-space equivalence.

Proof. It follows form Corollary 5.6 that every problem in the second itemized class belongs to the third one. It follows from Corollary 5.5 that if $\mathbb{B}$ is a finitely bounded structure, then $\mathrm{CSP}(\mathbb{B})$ is in monotone extensional SNP. By Lemma B.1 we see that if $\mathrm{CSP}(\mathbb{B})$ is in monotone extensional SNP, then there is a connected monotone extensional SNP sentence Φ which is log-space equivalent to $\mathrm{CSP}(\mathbb{B})$. This proves that the "moreover" statement holds. Clearly, every problem in the second itemized class belongs to the first itemized class of problems. Finally, by Lemma B.1 we know that for every monotone extensional SNP sentence Φ there is a polynomial-time conjunctive-equivalent monotone extensional SNP. □

6 Candidate NP-intermediate problems in extensional ESO

Here we consider the well-known computational problems of Graph Isomorphism [33] and Monotone Dualization [26]. The complexity of each of these problems has been open for more than 50 years, and some authors regard them as candidate NP-intermediate problems (see, for instance, the Wikipedia entry for candidate NP-intermediate problems). We show that for each of these problems there is an extensional ESO sentence Φ such that Φ is polynomial-time equivalent to the corresponding problem.

Graph Isomorphism. The graph isomorphism problem (GI) is one of the most studied problems in algorithmic graph theory (see, e.g., [33] for an overview on this problem). One of the major breakthroughs in this area is the quasipolynomial-time algorithm presented by Babai [4]. Despite intensive research over the past decades, the complexity of the graph isomorphism problem remains one of the problems in NP that is not known to be in P, and thus, it is considered a candidate NP-intermediate problem.

It is well known that the graph isomorphism problem is polynomial-time equivalent to the graph isomorphism completion problem. The latter takes as an input a pair of graphs $(\mathbb{G}, \mathbb{H})$ and a partial function $f'\colon U \subseteq G \to H$, and the task is to decide if there is an isomorphism $f\colon \mathbb{G} \to \mathbb{H}$ such that $f_{|U} = f'$. We now show that the graph isomorphism completion problem is in extensional ESO.

Let U_1 and U_2 be unary relation symbols, E a binary relation symbol, and I and I' binary relation symbols that will encode isomorphisms and partial isomorphisms, respectively. Here, the input signature is $\tau := \{U_1, U_2, E, I'\}$. Let $\mathrm{ISO}(I)$ be a first-order $\{U_1, U_2, E\}$-formula stating that I defines an isomorphism from the graph with vertices in U_1 to the graph with vertex set U_2, e.g.,

$$\begin{aligned}\mathrm{ISO}(I) := {} & \forall x, y\ (I(x,y) \Rightarrow U_1(x) \wedge U_2(x)) \\ & \wedge \forall x, y, z\ \big((I(x,y) \wedge I(x,z) \Rightarrow y = z) \wedge (I(y,z) \wedge I(y,z) \Rightarrow x = y)\big) \\ & \wedge \forall x \exists y\ \big((U_1(x) \Rightarrow I(x,y)) \wedge (U_2(x) \Rightarrow I(y,x))\big) \\ & \wedge \forall x_1, x_2, y_1, y_2\ \big((E(x_1,y_1) \wedge I(x_1,x_2) \wedge I(y_1,y_2)) \Rightarrow E(y_1,y_2)\big).\end{aligned}$$

Hence, if ϕ is a first-order $\{U_1, U_2, E\}$-sentence stating that E is a symmetric irreflexive relation that connects only vertices within U_1 or within U_2, we get that Graph Isomorphism Completion can be expressed by the extensional ESO sentence

$$\Psi := \exists I.\ \forall x, y\big((I'(x,y) \Rightarrow I(x,y)) \wedge \mathrm{Iso}(I) \wedge \phi\big).$$

Clearly, every GI completion instance X can be encoded as a τ-structure $\mathbb{A}$ such that $\mathbb{A} \models \Psi$ if and only if X is a yes-instance. Also notice that if an input τ-structure $\mathbb{A}$ is such that I' does not define a partial isomorphism from the graph with vertex set $U_1(\mathbb{A})$ to the graph with vertex set $U_2(\mathbb{A})$, then $\mathbb{A} \not\models \Psi$. In this way, we obtain a polynomial-time equivalence between the graph isomorphism completion problem and deciding Ψ.

Monotone Dualization. We assume basic familiarity with propositional logic (see, e.g., [25, Section 2]). Our presentation of the monotone dualization problem is similar to the presentation given in [26], and we refer to that survey for a detailed exposition. The *dual* of a boolean function $f\colon \{0,1\}^n \to \{0,1\}$ is the function $f^d(x_1, \ldots, x_n) := \bar{f}(\bar{x}_1, \ldots, \bar{x}_n)$. In particular, if ϕ is a CNF formula representing f (where the constant 0 is represented by an empty clause, and the constant 1 is represented by an empty set of clauses), then the dual f^d of f can be easily obtained by interchanging the symbols $\wedge$ and $\vee$. Computing the CNF of the dual of a boolean function f

presented in CNF traces back to the middle of the last century (see, e.g, [43]), and it is one of few problems whose complexity remains unclassified [26].

A boolean function $f\colon \{0,1\}^n \to \{0,1\}$ is *monotone* if $f(\bar{x}) \leq f(\bar{y})$ whenever $x_i \leq y_i$ for each argument $i \in [n]$. Equivalently, a boolean function is monotone if it can be represented by a boolean formula ϕ in CNF with no negative literals.

Monotone Dual

INPUT: Functions f and g represented by formulas ϕ and ψ in CNF without negative literals.

QUESTION: If g the dual of f?

Clearly, this problem is in coNP: if there is an evaluation $v_1, \dots, v_x$ of the variables $x_1, \dots, x_n$ such that $\phi^d(v_1, \dots, v_n)$ is true, and $\psi(v_1, \dots, v_n)$ is false, then g is not the dual of f (here, ϕ^d is the boolean formula in DNF obtained by swapping the symbols $\land$ and $\lor$ in ϕ). This problem can be solved in quasi-polynomial time [31], and according to [26] is has been open for over 40 years now, whether it can be solved in polynomial time. We construct an extensional ESO sentence Φ such that deciding Φ is polynomial-time equivalent to the complement of Monotone Dual, i.e., deciding whether g and f are not duals.

To begin with, we will consider a boolean formula ϕ in CNF or DNF, and without negative literals, empty clauses, and empty monomials. We will encode such formulas as structures $F(\phi)$ with the binary signature $\{C, EQ\}$ as follows. The vertices are occurrences of variables (i.e., for each clause c and each variable x in c, there is a vertex $v_{x,c}$), and the interpretations of C and EQ are equivalence relations where

- each class of C consists of all occurrences of a variable appearing in the same clause, and
- each class of EQ consists of the occurrences of the same variable.

In symbols, $(v_{x,c}, v_{y,c'}) \in C$ iff $c = c'$, and $(v_{x,c}, v_{y,c'}) \in EQ$ iff $x = y$. In particular, if ϕ is in CNF and ϕ^d is the dual DNF formula, then $F(\phi) = F(\phi^d)$. Also, for every $\{C, EQ\}$-structure $\mathbb{A}$ where C and EQ are equivalence relations, there is a monotone CNF formula ϕ such that $F(\phi) = \mathbb{A} = F(\phi^d)$.

We now consider a unary predicate U_1 that will encode an evaluation of the variables in the formula ϕ. In order for U_1 to encode a valid evaluation we must impose the first-order restriction stating that if $(x,y) \in EQ$, then $x \in U_1$ if and only if $y \in U_1$ (i.e., different occurrence of the same variable must receive the same value). So let δ be the (universal) first-order $\{U_1, C, EQ\}$-sentence stating that C and EQ are equivalence relations and that U_1 is preserved by the relation EQ in the sense explained above. Clearly, structures that satisfy δ are in one-to-one correspondence with evaluations of formulas in CNF (and in DNF) without negative literals, empty clauses, and empty monomials.

Consider the following first-order $\{U_1, C, EQ\}$-sentences

$$\texttt{true}_\texttt{C} := \forall x \exists y \; (C(x,y) \land U_1(y)) \quad \text{and} \quad \texttt{true}_\texttt{D} := \exists x \forall y \; (C(x,y) \Rightarrow U_1(y)).$$

Observe that if ϕ is a monotone CNF formula with variables X, then an evaluation $e\colon X \to \{0,1\}$ makes ϕ true if and only if the structure $(F(\phi), U_1 := e^{-1}(1))$ satisfies $\texttt{true}_\texttt{C}$. Similarly, e makes the dual formula ϕ^d in DNF true if and only if $(F(\phi), U_1 = e^{-1}(1)) \models \texttt{true}_\texttt{D}$. In fact, the $\{U_1, C, EQ\}$-expansions of $F(\phi)$ that satisfy $\texttt{true}_\texttt{C}$ (resp. $\texttt{true}_\texttt{D}$) are in one-to-one correspondence with the satisfying evaluations of ϕ (resp. of ϕ^d).

To obtain an extensional ESO sentence that is polynomial-time equivalent to the problem Monotone Dual, we need three more unary symbols U_f, U_g, and U_1'. Let ρ be a first-order sentence stating the following:

- the interpretations of U_f and U_g are total and disjoint, i.e., every vertex belongs to exactly one of U_f and U_g,
- if a pair of vertices x and y are connected by an edge in C, then x and y both belong to U_f or both to U_g, and
- for each vertex $x \in U_f$ there is a vertex $y \in U_g$ such that $xy \in EQ$, and vice versa.

Given a pair of CNF formulas ϕ and ψ on the same set of variables without negative literals and empty clauses, we construct a structure $F(\phi, \psi)$ with signature $\{U_f, U_g, U_1', C, EQ\}$ as follows. This structure is obtained from the disjoint union $F(\phi) + F(\psi)$ where we colour the vertices in $F(\phi)$ with U_f and the vertices in $F(\psi)$ with U_g, the interpretation of U_1' is empty, and we modify the interpretation of EQ by connecting all pairs of vertices x and y that represent the same variable (in particular, $EQ(F(\phi)) \cup EQ(F(\psi)) \subseteq EQ(F(\phi, \psi))$). Clearly, $F(\phi, \psi) \models \rho \wedge \delta$.

Finally, for $\mathtt{X} \in \{\mathtt{C}, \mathtt{D}\}$ and $h \in \{f, g\}$, we denote by $\mathtt{true}_{\mathtt{X}}(U_h)$ the restriction of $\mathtt{true}_{\mathtt{X}}$ to the substructure with domain U_h (which can be written as a first-order sentence with symbols in $\{C, EQ, U_1, U_h\}$). With these sentences we define the extensional ESO sentence

$$\Gamma := \exists U_1 \big((U_1' \Rightarrow U_1) \wedge \rho \wedge \delta \\ \wedge \left((\mathtt{true}_{\mathtt{D}}(U_f) \wedge \neg \mathtt{true}_{\mathtt{C}}(U_g)) \vee (\neg \mathtt{true}_{\mathtt{D}}(U_f) \wedge \mathtt{true}_{\mathtt{C}}(U_g)) \right) \big).$$

A structure of the form $F(\phi, \psi)$ satisfies Γ if and only if there is an evaluation of the variables that either make exactly one of ϕ^d or ψ true, and the other one false. Therefore, on an input (ϕ, ψ) to Monotone Dual, where ϕ and ψ are formulas CNF formulas without negative literals and empty clauses that represent the Boolean functions f and g, respectively, we have that g is not the dual of f if and only if $F(\phi, \psi)$ satisfies Γ. This gives us a polynomial-time reduction from Monotone Dual to deciding $\neg\Gamma$.

Now we show that conversely, deciding $\neg\Gamma$ reduces in polynomial time to Monotone Dual. First, fix a pair of formulas ϕ_X and ψ_X in monotone CNF over the same set of variable such that ψ_X is not the dual of ϕ_X. The reduction works as follows. Given a $\{U_1', U_f, U_g, C, EQ\}$-structure $\mathbb{A}$ we first first verify whether $\mathbb{A} \models \rho \wedge \delta$, and if not, the reduction returns the pair (ϕ_X, ψ_X). Otherwise, we construct ϕ and ψ as follows. The variables of ϕ and of ψ are the EQ-equivalence classes, and for each C-equivalence class Y in $U_f(\mathbb{A})$ (resp. in $U_g(\mathbb{A})$) we add a clause in ϕ (resp. in ψ) containing the EQ-equivalence classes that intersect Y. Finally, if some vertex v is colored with U_1', then we remove all clauses from ϕ and ψ that contain v. It readily follows from the arguments above that $\mathbb{A} \models \Gamma$ if and only if ϕ is not the dual of ψ. Hence, the reduction that maps $\mathbb{A}$ to (ϕ_X, ψ_X) if $\mathbb{A} \not\models \delta \wedge \rho$, and to (ϕ, ψ) otherwise, is a polynomial-time reduction from deciding $\neg\Gamma$ to the Monotone Dual problem. Both reductions together show that Monotone Dual is polynomial-time equivalent to deciding $\neg\Gamma$.

Conjecture. We have seen that surjective finite-domain CSPs (Section 3.1) and CSPs of finitely bounded structures (Corollary 5.5) are expressible in ESO. Moreover, we proved that extensional ESO has the same computational power as HerFO (Theorem 4.6), and so, the tractability meta-problem for extensional ESO is undecidable (Corollary 4.4). Finally, we argued in this section that extensional ESO can express problems that are polynomial-time equivalent to Graph Isomorphism, and to the complement of Monotone Dualization. These observations make us believe that the following is true.

Conjecture 6.1. *There is an extensional ESO sentence Φ such that deciding Φ is NP-intermediate.*

7 Limitations of the computational power

In this section we show that extensional ESO is not NP-rich (unless E = NE), i.e., that there are sets L in NP which are not polynomial-time equivalent to a problem in extensional ESO. To do so, we study disjunctive self-reducibility; a notion in computational complexity that traces back to [44, 45].

7.1 Disjunctive self-reducible sets

Intuitively speaking, a set L is *self-reducible* if deciding whether x belongs to L reduces in polynomial-time to deciding L for "smaller" instances than x. Moreover, it is *disjunctive self-reducible* if there is a deterministic polynomial-smaller instances $x_1, \ldots, x_n$ such that $x \in L$ if and only if $x_i \in L$ for some $i \in [n]$. A typical example of a disjunctive self-reducible set is SAT: a boolean formula $\phi(x_1, \ldots, x_n)$ is satisfiable if and only if at $\phi(0, x_2, \ldots, x_n)$ and $\phi(1, x_2, \ldots, x_n)$ is satisfiable. This concept has played a major role in the complexity theory literature on the class NP, see, e.g., [3, 35].

Let us formally define disjunctive self-reducible sets (also see also [44, 45]). A polynomial-time computable partial order $<$ on $\{0,1\}^*$ is *OK* if and only if there is a polynomial p such that

- every strictly decreasing chain $x_1 < \cdots < x_n$ of words from $\{0,1\}^*$ is shorter than $p(x_n)$, i.e., $n \leq p(x_n)$.
- for all $x, y \in \{0,1\}^*$, $x < y$ implies $|x| \leq p(|y|)$.

A set $L \subseteq \{0,1\}^*$ is *self-reducible* if there are an OK partial order $<$ on $\{0,1\}^*$ and a deterministic polynomial-time oracle machine M that accepts L and asks its oracle queries about words smaller than the input x, i.e., on input x, if M asks a query to its oracle for L about y, then $y < x$. We say that L is *disjunctive self-reducible* if on every input x the machine M either

- decides whether $x \in L$ without calling the oracle, or
- computes queries $x_1, \ldots, x_n \in \{0,1\}^*$ strictly smaller than x (with respect to $<$) such that

$$x \in L \text{ if and only if } \{x_1, \ldots, x_n\} \cap L \neq \varnothing.$$

Clearly, every set in P is disjunctive self-reducible, and it follows from [38] that every disjunctive self-reducible set is in NP.

Lemma 7.1. *If a class of structures $\mathcal{C}$ is expressible in extensional* ESO*, then $\mathcal{C}$ is disjunctive self-reducible. In particular,* CSPs *of finitely bounded structures are disjunctive-self-reducible.*

Proof. We prove the claim for the case that Φ has one existentially quantified relation symbol R or arity r — the general case follows with similar arguments. Let R' be the input symbol such that Φ contains the clause $R'(x) \Rightarrow R(x)$ and R' does not appear in any other clause of Φ. Given an input τ-structure $\mathbb{A}$ we first test whether the expansion $\mathbb{A}'$ with $R^{\mathbb{A}'} = R'$ satisfies the first-order part of Φ. If yes, we accept, and otherwise we return the structures obtained from $\mathbb{A}$ by adding one new tuple to the interpretation of R'. It is clear that the order $\mathbb{B} < \mathbb{A}$ if $A = B$, $S^{\mathbb{A}} = S^{\mathbb{B}}$ for every $S \in \tau \setminus \{R'\}$, and $(R')^{\mathbb{A}'} \subseteq (R')^{\mathbb{B}'}$ is an OK poset witnessing that the previous reduction is a disjunctive self-reduction. $\square$

In the rest of this section we study disjunctive self-reducible sets. In particular, we show that the class of disjunctive self-reducible sets is not NP-rich, unless E = NE (Theorem 7.7); and also prove that there are NP-intermediate disjunctive self-reducible sets unless P = NP (Theorem 7.3). The former, together with Lemma 7.1 implies that extensional ESO does not have the full computational power of NP (Corollary 7.8), unless E = NE.

7.2 NP-intermediate disjunctive self-reducible sets

We consider the following family of disjunctive self-reducible sets.

Lemma 7.2. *Let τ be a finite relational signature and $\mathcal{F}$ a (not necessarily finite) class of finite τ-structures. If $\mathcal{F}$ is in* P*, then the class of structures $\mathbb{A}$ that contain some substructure $\mathbb{F}$ from $\mathcal{F}$ is disjunctive self-reducible.*

Proof. Given an input structure $\mathbb{A}$, we first test whether $\mathbb{A} \in \mathcal{F}$. If yes, we accept, otherwise, we return the substructures obtained from $\mathbb{A}$ by removing exactly one element from the domain. Here, the order $\mathbb{A} < \mathbb{B}$ defined by $\mathbb{A}$ being a substructure of $\mathbb{B}$ is clearly an OK poset for which the previous reduction is a disjunctive self-reduction. □

For our next result we will use a construction of coNP-intermediate CSPs from Bodirsky and Grohe [13] which uses the following set of tournaments. The *Henson set* is the set $\mathcal{T}$ of tournaments $\mathbb{T}_n$ defined for positive integer $n \geq 5$ as follows. The vertex set of $\mathbb{T}_n$ is $[n]$ and it contains the edges

- $(1, n)$,
- $(i, i+1)$ for $i \in [n-1]$, and
- (j, i) for $j > i+1$ and $(j, i) \neq (n, 1)$.

It is straightforward to observe that, for any fixed (possibly infinite) set of tournaments $\mathcal{F}$ (such as $\mathcal{T}$), the class of oriented graphs that do not embed any tournament from $\mathcal{F}$ is of the form $\mathrm{CSP}(\mathbb{B}_{\mathcal{F}})$ for some countably infinite structure $\mathbb{B}_{\mathcal{F}}$ (we may even choose $\mathbb{B}_{\mathcal{F}}$ to be homogeneous). In the concrete situation of the class $\mathcal{T}$ defined above, $\mathrm{CSP}(\mathbb{B}_{\mathcal{T}})$ is an example of coNP-complete CSP (see, e.g., [12, Proposition 13.3.1]).

Theorem 7.3. *There is no* P*- versus* NP*-complete dichotomy for disjunctive self-reducible sets.*

Proof. We use that there exists a subset $\mathcal{T}_0$ of the Henson set $\mathcal{T}$ such that there exists a linear-time Turing machine G computing a function $g \colon \mathbb{N} \to \mathbb{N}$ (where the input to G is represented in unary) with

- $\mathcal{T}_0 = \{\mathbb{T}_n \mid g(n) \text{ is even}\}$,
- $\mathrm{CSP}(\mathbb{B}_{\mathcal{T}_0})$ is not in P and not coNP-hard, unless P = coNP [13].

Clearly, $\mathcal{T}_0$ is in P: given an input digraph D verify whether $g(|D|)$ is even (if not, reject), and if this is the case, verify whether D belongs to the Henson set $\mathcal{T}$ — which is even in first-order logic [14]. Since $\mathrm{CSP}(\mathbb{B}_{\mathcal{T}_0})$ is the class of $\mathcal{T}_0$-free loopless oriented graphs, it follows from Lemma 7.2 that $\mathrm{CSP}(\mathbb{B}_{\mathcal{T}_0})$ is disjunctive self-reducible. Note that this implies the unconditional non-dichotomy result stated in the theorem: if P = NP, then there is no dichotomy, and if P is different from NP, then there are NP-intermediate disjunctive self-reducible sets. □

7.3 Search reduces to decision

For a polynomial-time computable $R \subseteq (\{0,1\}^*)^2$, the *search problem for R* is the computational problem to compute for a given $u \in \{0,1\}^*$ some $v \in \{0,1\}^*$ such that $(u, v) \in R$; if there is no such v, the output can be arbitrary (even undefined). Clearly, the search problem for R is at least as hard deciding the language

$$L_R := \{u \in \{0,1\}^* \mid \text{there is } v \in \{0,1\}^* \text{ such that } (u, v) \in R\}.$$

For a given R, we say that *search for R reduces to decision* if the search problem for R can be solved in polynomial time given an oracle for the decision of L. For a given language L in NP, we say that *search for L reduces to decision* if there exists a polynomial-time computable $R \subseteq (\{0,1\}^*)^2$ such that $L = L_R$ and search for R reduces to decision.

In particular, for every disjunctive self-reducible set L search for L reduces to decision. Indeed, suppose that M is a Turing machine witnessing that L is disjunctive self-reducible and $<$ is the corresponding OK poset. For each $u \in L$ consider witnesses of the form $v_1 \dots v_k$ where $u = v_1$, $v_i < v_{i+1}$ for each $i \in [k-1]$, v_k is accepted by M without calling the oracle, and v_{i+1} is one of the queries that M computes on input v_i before calling the oracle. Clearly, the set R of such tuples $(u, v_1 \dots v_k)$ is decidable in polynomial-time, and search reduces to decision for R.

Lemma 7.4. *Let L and L' be sets such that search reduces to decision for L. If there is a polynomial-time disjunctive reduction from L' to L, and a polynomial-time Turing reduction from L to L', then search reduces to decision for L'.*

Proof. Let R be a binary relation such that $L = L_R$ and search reduces to decision for R, and let f be a polynomial-time computable function proving that there is a disjunctive polynomial-time reduction from L' to L, i.e., $x \in L'$ if and only if $f(x) \cap L \neq \varnothing$. Note that a witness for $x \in L'$ is a pair (u, v) such that $u \in f(x)$ and $(u, v) \in R$. Since there is a polynomial-time Turing reduction from $L = L_R$ to L', and search reduces to decision for R, there is a polynomial-time deterministic oracle machine M that solves the search problem for R using L' as an oracle. Let M' be the deterministic polynomial-time oracle machine that on input x does the following:

- it first computes $f(x)$;
- for each $u \in f(x)$ it simulates M to find a witness v such that $(u, v) \in R$;
- if it finds such a v for some $u \in f(x)$, then M' outputs (u, v);
- otherwise it rejects.

It clearly follows that M' is an oracle machine proving that search reduces to decision for L'. □

It is now straightforward to observe that search reduces to decision for every problem in any of the classes considered so far.

Theorem 7.5. *Search reduces to decision for every problem P in any of the following classes.*

- CSP*s of finitely bounded structures.*
- *Monotone extensional* SNP*.*
- *Extensional* ESO*.*

In particular, neither of these classes in disjunctive NP*-rich unless* EE = NEE*.*

Proof. We already argued that search reduces to decision for every disjunctive self-reducible sets, and it follows from Lemmas 7.1 and 7.2 that any problem in these classes is disjunctive self-reducible. To see that the last claim of this theorem holds, we use Lemma 7.4 to note that if L is polynomial-time disjunctive equivalent to L' where L' belongs to one of the itemized classes, then search reduces to decision for L. The claim now follows, because Ballare and Goldwasser [10] proved that there exists a language $L \subseteq \{0,1\}^*$ in NP such that search does not reduce to decision for L unless EE = NEE. □

7.4 Promise problems

Several NP-complete problems become tractable given a strong enough promise. For instance, k-colourability is decidable in polynomial time, given the promise that the input graph G is *perfect*, i.e., the chromatic number of H equals the largest complete subgraph of H for every induced subgraph H of G. Hence, to decide whether G is k-colourable it suffices to verify whether G is K_{k+1}-free, which is clearly in P. In this example, the promise is even verifiable in polynomial-time [21], but this need not be the case in general. A prominent example being the problem 1-in-3 SAT, which is tractable given the promise that the input formula ϕ either has a 1-in-3 solution, or does not even have a not-all-equal solution [19]. In contrast, some problems remain hard even given some non-trivial promise. A typical example is 3-colourability which is NP-hard given the promise that the input graph G is either 3-colourable or not even 5-colourable [9].

A general framework for studying hard promises was introduced in [46]. Formally, a (decidable) *promise problem* is a pair of decidable sets (Q, R) where Q is the *promise*, and R is the *property* that needs to be decided given the promise Q. A deterministic Turing machine M *solves* the promise problem (Q, R) if for every input x the following holds

$$\text{if } x \in Q, \text{ then } M \text{ accepts } x \text{ if and only if } x \in R.$$

In this case, we say that the language L accepted by M is a *solution* to (Q, R). In particular, notice that R and $Q \cap R$ are solutions to (Q, R). Going back to our previous examples, if Q is the class of perfect graphs, and R the class of k-colourable graphs, then the class of K_{k+1}-free graphs is a polynomial-time solution to (Q, R).

A promise problem (Q, R) is in P if (Q, R) has a solution L in P, and (Q, R) is *NP-hard* if every solution to (Q, R) is NP-hard. Note that for every promise problem (Q, R), we have that R is a solution to (Q, R), and hence if R is in P, then (Q, R) is in P. A promise problem (Q, R) is *polynomial-time Turing-reducible*[4] to a promise problem (S, T) if for every solution A of (S, T) there is a solution B of (Q, T) such that B is polynomial-time Turing-reducible to A. Notice that if L is a decidable set, then the only solution to $(\{0, 1\}^*, L)$ is L, and hence, we say that a set L is polynomial-time Turing-reducible to a promise problem (Q, R) if $(\{0, 1\}^*, L)$ is polynomial-time Turing-reducible to (Q, R).

Natural promise problems trace back to at least 1984, and they arise in the context of public-key cryptography [27]. Given a decidable set $L \subseteq \Sigma^*$, for some finite alphabet Σ, the *natural promise problem* for L is the promise problem

$$(L \times (\Sigma^* \setminus L) \ \uplus \ (\Sigma^* \setminus L) \times L, \ L \times \Sigma^*) \, .$$

That is, an instance to the problem is a pair (x, y), the promise is that exactly one of x or y belongs to L, and the property to decide is whether $x \in L$. Clearly, the natural promise problem for L is at most as hard as L, and thus, if L is in P then L and the natural promise problem for L are trivially polynomial-time equivalent. In [27], the authors prove that the natural promise problem for SAT is NP-hard, and it follows from [46] that L and its natural promise problem are polynomial-time Turing-equivalent for every NP-complete set L. In particular, it follows from the finite domain dichotomy conjecture [20, 48] that every finite domain CSP is polynomial-time Turing-equivalent to its natural promise problem – this will also follow from our next theorem without using the finite domain dichotomy. The graph isomorphism problem (which is not known to belong to P and most likely not NP-complete) is polynomial-time equivalent to its natural promise problem [46].

[4] Many-one, disjunctive, and conjunctive reductions are defined analogously, but are not needed for this paper.

Theorem 7.6. *For every L in any of the following classes, the natural promise problem for L is polynomial-time Turing equivalent to L.*

- CSP*s of finitely bounded structures.*
- *Monotone extensional* SNP*.*
- *Extensional* ESO*.*

Proof. Clearly, every set L is at least as hard as its natural promise problem. On the other hand, Theorem 3 in [46] asserts that if L is disjunctive self-reducible, then there is a polynomial-time Turing-reduction from L to the natural promise problem L. By Lemmas 7.1 and 7.2 we know that any problem in L in any of the itemized classes is disjunctive self-reducible, and so, the natural promise problem for L is polynomial-time Turing equivalent to L. □

7.5 Failure of NP-richness

It follows from results in [46] that there are problems in NP which are not polynomial-time disjunctive-equivalent to any disjunctive self-reducible set, unless $\mathrm{E} = \mathrm{NE}$.

Theorem 7.7 (essentially in [46])**.** *The class of disjunctive self-reducible sets is not disjunctive* NP*-rich, unless* $\mathrm{E} = \mathrm{NE}$*.*

Proof. Theorem 3 in [46] asserts that if L is disjunctive self-reducible, then the natural promise problem for L is polynomial-time Turing-equivalent to L, and Theorem 4 in [46] asserts that the property "L and the natural promise problem for L are polynomial-time Turing equivalent" is preserved under polynomial time disjunctive-equivalence. Finally, Theorem 7 in [46] guarantees that if $\mathrm{E} \neq \mathrm{NE}$, then there is a set $L \in \mathrm{NP} \setminus \mathrm{P}$ such that the natural promise problem for L is in P, i.e., L and its natural promise problem are not polynomial-time Turing-equivalent. □

Corollary 7.8. *None of the following classes is* NP*-rich, unless* $\mathrm{E} = \mathrm{NE}$*.*

- CSP*s of finitely bounded structures.*
- *Monotone extensional* SNP*.*
- *Extensional* ESO*.*

8 Conclusion and Open Problems

We proved that the complement of every problem in extensional ESO is polynomial-time equivalent to a problem in HerFO, and vice versa (Theorem 4.6). We also showed that there are problems in coNP that are not polynomial-time equivalent to problems in Extensional ESO, unless NE=E (Corollary 7.8). Our results also show that there are subclasses of NP that are believed to have no P versus NP-complete dichotomy, while the previous technique of showing this by establishing that the class is NP-rich does not work. For instance for proving that CSPs of finitely bounded structures do not have a complexity dichotomy, we therefore need a different approach.

Besides Conjecture 6.1, a number of open problems are left for future research.

1. Is every disjunctive self-reducible set polynomial-time (disjunctive-) equivalent to an extensional ESO sentence (see Figure 1)?

2. Is it true that for every extensional ESO sentence Φ there is a finitely bounded structure $\mathbb{A}$ such that Φ and CSP($\mathbb{A}$) are polynomial-time equivalent (see again Figure 1)?

3. Are there NP-intermediate CSPs of finitely bounded structures (assuming P $\neq$ NP)?

4. Is there a finitely bounded structure $\mathbb{A}$ such that CSP($\mathbb{A}$) is polynomial-time equivalent to the graph isomorphism problem?

5. Is every CSP of a finitely bounded structure in connected monotone extensional SNP? (Compare with Corollary 5.6 and the moreover statement in Theorem 5.7).

Acknowledgments

The authors are thankful to Antoine Amarilli for a guided tour through Constraint Topological Sorting.

9 References

[1] J. Alvarado, S. Dantas, and D. Rautenbach. Sandwiches missing two ingredients of order four. *Annals of Operations Research*, 280:47–63, 2019.

[2] A. Amarilli and C. Paperman. Topological Sorting with Regular Constraints. In I. Chatzigiannakis, C. Kaklamanis, D. Marx, and D. Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 115:1–115:14, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[3] K. Ambos-Spies and J. Kämper. On disjunctive self-reducibility. In E. Börger, H. K. Büning, and M. M. Richter, editors, *CSL '88, 2nd Workshop on Computer Science Logic, Duisburg, Germany, October 3-7, 1988, Proceedings*, volume 385 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 1988.

[4] L. Babai. Graph isomorphism in quasipolynomial time. *arXiv:1512.03547*, 2015.

[5] R. N. Ball, J. Nešetřil, and A. Pultr. Dualities in full homomorphisms. *European Journal of Combinatorics*, 31:106–199, 2010.

[6] J. Bang-Jensen, J. Huang, and X. Zhu. Completing orientations of partially oriented graphs. *Journal of Graph Theory*, 87(3):285–304, 2017.

[7] M. Bannach, F. Chudigiewitsch, and T. Tantau. On the Descriptive Complexity of Vertex Deletion Problems. In R. Královič and A. Kučera, editors, *49th International Symposium on Mathematical Foundations of Computer Science (MFCS 2024)*, volume 306 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:14, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[8] A. Barsukov and F. R. Madelaine. On guarded extensions of MMSNP. In G. D. Vedova, B. Dundua, S. Lempp, and F. Manea, editors, *Unity of Logic and Computation - 19th Conference on Computability in Europe, CiE 2023, Batumi, Georgia, July 24-28, 2023, Proceedings*, volume 13967 of *Lecture Notes in Computer Science*, pages 202–213. Springer, 2023.

[9] L. Barto, J. Bulín, A. Krokhin, and J. Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4), July 2021.

[10] M. Bellare and S. Goldwasser. The complexity of decision versus search. *SIAM J. Comput.*, 23(1):97–119, 1994.

[11] M. Bienvenu, B. ten Cate, C. Lutz, and F. Wolter. Ontology-Based Data Access: A Study through Disjunctive Datalog, CSP, and MMSNP. *ACM Trans. Database Syst.*, 39(4):33:1–33:44, 2014.

[12] M. Bodirsky. *Complexity of Infinite-Domain Constraint Satisfaction*. Lecture Notes in Logic (52). Cambridge University Press, Cambridge, United Kingdom; New York, NY, 2021.

[13] M. Bodirsky and M. Grohe. Non-dichotomies in constraint satisfaction complexity. pages 184 –196. Springer Verlag, July 2008.

[14] M. Bodirsky and S. Guzmán-Pro. Hereditary first-order logic: the tractable quantifier prefixes, 2025. Preprint arXiv:2411.10860v2.

[15] M. Bodirsky and S. Guzmán-Pro. The generic circular triangle-free graph. *Journal of Graph Theory*, 109(4):426–445, 2025.

[16] M. Bodirsky, S. Knäuer, and F. Starke. ASNP: A tame fragment of existential second-order logic. In *CiE 2020, Proceedings*, volume 12098 of *Lecture Notes in Computer Science*, pages 149–162. Springer, 2020.

[17] J. A. Bondy and U. S. R. Murty. *Graph Theory*. Springer, Berlin, 2008.

[18] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.

[19] J. Brakensiek and V. Guruswami. Promise constraint satisfaction: Algebraic structure and a symmetric boolean dichotomy. *SIAM J. Comput.*, 50(6):1663–1700, 2021.

[20] A. A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17*, pages 319–330, 2017.

[21] M. Chudnovsky, N. Roberston, P. Seymour, and R. Thomas. The strong perfect graph theorem. *Annals of Mathematics*, 164:51–229, 2006.

[22] S. Dantas, C. M. de Figueiredo, M. V. da Silva, and R. B. Teixeira. On the forbidden induced subgraph sandwich problem. *Discrete Applied Mathematics*, 159:1717–1725, 2015.

[23] S. Dantas, C. M. de Figueiredo, P. Petito, and R. B. Teixeira. A general method for forbidden induced subgraph sandwich problem np-completeness. *Electronic Notes in Theoretical Computer Science*, 346:393–400, 2019.

[24] C. M. de Figueiredo and S. Spirkl. Sandwich and probe problems for excluding paths. *Discrete Applied Mathematics*, 251:146–154, 2018.

[25] T. Eiter, G. Gottlob, and K. Makino. New results on monotone dualization and generating hypergraph transversals. *SIAM Journal on Computing*, 32(2):514–537, 2003.

[26] T. Eiter, K. Makino, and G. Gottlob. Computational aspects of monotone dualization: A brief survey. *Discrete Applied Mathematics*, 156(11):2035–2049, 2008. In Memory of Leonid Khachiyan (1952 - 2005).

[27] S. Even, A. L. Selman, and Y. Yacobi. The complexity of promise problems with applications to public-key cryptography. *Information and Control*, 61(2):159–173, 1984.

[28] R. Fagin. Generalized first-order spectra and polynomial-time recognizable sets. *Complexity of Computation*, 7:43–73, 1974.

[29] T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1999.

[30] F. Fomin, P. Golovach, and D. Thilikos. On the parameterized complexity of graph modification to first-order logic properties. *Theory of Computing Systems*, 64:251–271, 2020.

[31] M. L. Fredman and L. Khachiyan. On the complexity of dualization of monotone disjunctive normal forms. *Journal of Algorithms*, 21(3):618–628, 1996.

[32] M. Golumbic, H. Kaplan, and R. Shamir. Graph Sandwich Problems. *Journal of Algorithms*, 19:449–473, 1995.

[33] M. Grohe and P. Schweitzer. The graph isomorphism problem. *Commun. ACM*, 63(11):128–134, Oct. 2020.

[34] E. Hemaspaandra, A. V. Naik, M. Ogihara, and A. L. Selman. P-selective sets and reducing search to decision vs self-reducibility. *Journal of Computer and System Sciences*, 5:194–209, 1996.

[35] L. A. Hemaspaandra. *The Power of Self-Reducibility: Selectivity, Information, and Approximation*, pages 19–47. Springer International Publishing, Cham, 2020.

[36] W. Hodges. *A shorter model theory.* Cambridge University Press, Cambridge, 1997.

[37] N. Immerman. *Descriptive Complexity.* Graduate Texts in Computer Science, Springer, New York, 1998.

[38] K.-I. Ko. On self-reducibility and weak p-selectivity. *Journal of Computer and System Sciences*, 26:209–221, 1983.

[39] G. Kun and J. Nešetřil. Dichotomy for orderings?, 2025.

[40] A. H. Lachlan and R. E. Woodrow. Countable ultrahomogeneous undirected graphs. *Transactions of the AMS*, 262(1):51–94, 1980.

[41] R. E. Ladner. On the structure of polynomial time reducibility. *J. ACM*, 22(1):155–171, Jan. 1975.

[42] D. Macpherson. A survey of homogeneous structures. *Discrete Mathematics*, 311(15):1599–1634, 2011. Infinite Graphs: Introductions, Connections, Surveys.

[43] K. Maghout. Sur la determination des nombres de stabilite et du nombre chromatique dun graphe. *COMPTES RENDUS HEBDOMADAIRES DES SEANCES DE L ACADEMIE DES SCIENCES*, 248(25):3522–3523, 1959.

[44] A. Meyer and M. Patersons. With what frequency are apparently intractable problems difficult?, 1979. Tech. Report MIT/LCS/TM-126, Massachusetts Institute of Technology, Cambridge, MA.

[45] C.-P. Schnorr. Optimal Algorithms for Self-Reducible Problems. In *Proceedings of the Third International Colloquium on Automata, Languages and Programming*, pages 322–337, 1976.

[46] A. L. Selman. Promise problems for complexity classes. *Information and Computation*, 78:87–98, 1988.

[47] H. Straubing. First-order logic and aperiodic languages: a revisionist history. *ACM SIGLOG News*, 5(3):4–20, July 2018.

[48] D. N. Zhuk. A proof of CSP dichotomy conjecture. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17*, pages 331–342, 2017. https://arxiv.org/abs/1704.01914.

A A remark concerning the definition of extensional ESO

We point out that if we remove the condition "and R' does not appear anywhere else in Ψ" in the definition of extensional ESO (in Section 3), we obtain essentially the same logic as ESO.

Observation A.1. *For every* ESO *τ-sentence Φ, there are a signature $\tau^* \supseteq \tau$ and an* ESO *τ^*-sentence Ψ such that*

- *a τ^*-structure $\mathbb{A}$ satisfies Ψ if and only if $\mathbb{A}$ is a τ-structure and $\mathbb{A} \models \Phi$,*
- *if Φ is monotone (resp. in* SNP*), then Ψ is monotone (resp. in* SNP*), and*
- *for every existentially quantified symbol R of Ψ, there is a distinct symbol $R' \in \tau^*$ such that Ψ contains the clause $\forall \overline{x}. R'(\overline{x}) \Rightarrow R(\overline{x})$.*

Proof. Let ϕ be the first-order part of Φ; we assume ϕ is in conjunctive normal form. We first consider the case where Φ holds on every finite τ-structure. In this case, let ψ be the first-order τ-sentence obtained from ϕ by removing all clauses that contain an existentially quantified symbol from Φ. Since Φ holds on all finite τ-structures, ψ does as well. Hence, Φ and ψ are equivalent on finite structures. Moreover, if Φ is monotone (resp. universal) then so is ψ, and ψ satisfies the third itemized statement because ψ has no quantified symbols.

Let σ be the set of existentially quantified symbols in Φ. Consider the relational signature τ^* that contains τ and additionally a symbol R' of arity r for every $R \in \sigma$ of arity r. Let m be the maximum arity of the quantified symbols of Φ, and define ψ as

$$\psi := \phi \wedge \forall x_1, \dots, x_m \left(\bigwedge_{R \in \sigma} \neg R'(x_1, \dots, x_r) \wedge \left(\bigwedge_{R \in \sigma} R'(x_1, \dots, x_r) \Rightarrow R(x_1, \dots, x_r) \right) \right).$$

Note that this formula can be simplified; the purpose of writing it in the given redundant form is to make sure that the formula Ψ that we define next meets the syntactic requirements from the observation. We now define Ψ as the ESO τ^*-formula obtained from Φ be replacing ϕ by ψ. Since $\tau \subseteq \tau^*$, every τ-structure $\mathbb{A}$ is a τ^*-structure. It is also straightforward to observe that $\mathbb{A}$ satisfies Φ if and only if it satisfies Ψ. Hence, the forward implication of the first itemized statement holds for Φ and Ψ. To see that the converse also holds consider a τ^*-structure $\mathbb{B}$. Notice that if the interpretation of R' in $\mathbb{B}$ is not empty for some $R \in \sigma$, then $\mathbb{B} \not\models \Psi$. Otherwise, $\mathbb{B}$ is a τ-structure, and it follows from the definition of Ψ that $\mathbb{B} \models \Phi$ if and only if $\mathbb{B} \models \Psi$. We thus conclude that the first itemized statement holds for Φ and Ψ. The second and third itemized statement hold by the definition of Ψ. $\square$

B Monotone connected extensional SNP

The following lemma is used in our proof of Theorem 5.7.

Lemma B.1. *For every extensional* SNP *τ-sentence Φ and a binary symbol E' not in τ there is a connected extensional* SNP *$(\tau \cup \{E'\})$-sentence Ψ such that the following statements hold.*

- *There is a polynomial-time reduction from Φ to Ψ, and a polynomial-time conjunctive reduction from Ψ to Φ.*
- *If Φ is preserved under disjoint unions, then Φ and Ψ are log-space equivalent.*

- *If Φ is in monotone* SNP*, then Ψ is in monotone* SNP*.*

Proof. We follow almost the same proof as the one for Proposition 1.4.11 in [12] — we thus choose to refer the reader to that proof for any missing details in this one. Let τ be a finite relational signature and let Φ be an extensional SNP τ-sentence of the form

$$\exists R_1, \dots, R_k \forall x_1, \dots, x_n (\nu \wedge \phi)$$

where ν consists of the conjuncts $R_i' \Rightarrow R_i$ for every $i \in [k]$ and ϕ is a quantifier-free $(\tau \cup \{R_1, \dots, R_k\})$-formula in conjunctive normal form. We consider two new binary relation symbols E' and E and the signature $\tau' := \tau \cup \{E'\}$. We define the uniform SNP τ'-sentence

$$\Psi := \exists R_1, \dots, R_k, E\, \forall x_1, \dots, x_n, z_1, z_2, z_3 (\nu \wedge E'(z_1, z_2) \Rightarrow E(z_1, z_2) \wedge \psi)$$

where ψ is the conjunction of the following clauses:

- $\neg E(z_1, z_2) \vee \neg E(z_2, z_3) \vee E(z_1, z_3)$ (i.e., E is a transitive relation),
- $\neg E(z_1, z_2) \vee E(z_2, z_1)$ (i.e., E is a symmetric relation),
- for every $R \in \tau$ of arity r and all $i, j \in [r]$ where $i \neq j$, the formula ψ contains the conjunct $R(x_1, \dots, x_r) \Rightarrow E(x_i, x_j)$, and
- for each clause ϕ' of ϕ with free variables $y_1, \dots, y_m \subseteq \{x_1, \dots, x_n\}$, the formula ψ contains the conjunct
$$\phi' \vee \bigvee_{i,j \in [m]} \neg E(y_i, y_j).$$

Clearly, Ψ is connected. Since Φ is in extensional SNP, Ψ is in extensional SNP as well. It is not hard to verify that a τ'-structure $(\mathbb{B}, E')$ satisfies Ψ if and only if the τ-reduct $\mathbb{C}$ of every connected component $(\mathbb{C}, E')$ of $(\mathbb{B}, E')$ satisfies Φ. Thus, there is a polynomial-time conjunctive reduction from Ψ to Φ. If Φ is preserved under disjoint unions, then $(\mathbb{B}, E') \mapsto \mathbb{B}$ is a polynomial-time (and log-space) reduction from Ψ to Φ. For the converse reduction, on input structure $\mathbb{A}$ consider the $(\tau \cup \{E'\})$-structure $\mathbb{B} := (\mathbb{A}, A^2)$ and notice that $\mathbb{A} \models \Phi$ if and only if $\mathbb{B} \models \Psi$. Hence, Φ reduces in logarithmic space to Ψ. Finally, it follows from the definition of Ψ that if Φ is monotone, then Ψ is monotone as well. $\square$

Finite-domain CSPs

This appendix contains a lemma which in particular implies that every finite domain CSP is expressible in connected monotone extensional SNP. Given a relational τ-structure $\mathbb{A}$, we denote by $\mathrm{Age}(\mathbb{A})$ the class of finite structures that embed into $\mathbb{A}$; equivalently, the class of finite substructures of $\mathbb{A}$, up to isomorphism. A set of finite structures $\mathcal{F}$ is called a set of *bounds* of $\mathbb{A}$, if for every finite structure $\mathbb{B}$ it is the case that $\mathbb{B} \in \mathrm{Age}(\mathbb{A})$ if and only if $\mathbb{B}$ is $\mathcal{F}$-free. We say that $\mathcal{F}$ is a set of *connected* bounds if every structure in $\mathcal{F}$ is connected. The *Gaifman graph* of $\mathbb{A}$ is the undirected graph $G(\mathbb{A})$ with vertex set A and there is an edge ab if a and b belong to some common tuple in $R^{\mathbb{A}}$ for some $R \in \tau$. An *induced path* in $\mathbb{A}$ is a sequence of vertices $a_1, \dots, a_n$ that induce a path in the Gaifman graph of $\mathbb{A}$. In this case, we say that the path is an *$a_1 a_n$-path*, and that the *length* of this path is n. It is straightforward to observe that for every finite relational signature τ and every positive integer N, there is a finite set of τ-structures $\mathcal{P}_N$ such that a τ-structure $\mathbb{A}$ has no induced path of length N if and only if $\mathbb{A}$ is $\mathcal{P}_N$-free.

Lemma B.2. *Let $\mathbb{A}$ be a finitely bounded structure. If there is a positive integer N such that every induced path in $\mathbb{A}$ has length strictly less than N, then there is a structure $\mathbb{B}$ such that*

- $\mathrm{Age}(\mathbb{A}) \subseteq \mathrm{Age}(\mathbb{B})$,
- *a structure $\mathbb{C}$ belongs to* $\mathrm{Age}(\mathbb{B})$ *if and only if every connected component of $\mathbb{C}$ belongs to* $\mathrm{Age}(\mathbb{A})$, *and*
- *there is a finite set of connected bounds for* $\mathrm{Age}(\mathbb{B})$.

Proof. Let $\mathcal{F}$ be a finite set of bounds of $\mathbb{A}$, and let $\mathcal{P}_N$ be as in the paragraph above. Let $\mathcal{F}_C$ be the set of connected structures in $\mathcal{F}$, and for every disconnected structure $\mathbb{F} \in \mathcal{F}$ where $\mathbb{F} := \biguplus_{i=1}^{n} \mathbb{F}_n$ and each $\mathbb{F}_n$ is a connected structure, let $\mathcal{C}(\mathbb{F})$ be the class of finite τ-structures $\mathbb{C}$ such that

- $\mathbb{F}$ embeds into $\mathbb{C}$, and
- for all distinct $i, j \in \{1, \ldots, n\}$, $a \in F_i$, and $b \in F_j$, there is an induced ab-path in $\mathbb{C}$ of length strictly less than N.

Notice that $\mathcal{C}(\mathbb{F})$ has, up to isomorphism, a finite set of minimal structures $\mathbb{M}$ with respect to the embedding order. Indeed, let $\mathbb{M}$ be such a minimal structure, and identify $\mathbb{F}$ with a substructure of $\mathbb{M}$ (isomorphic to $\mathbb{F}$). For every $a, b \in F$ that belong to different components of $\mathbb{F}$, let P_{ab} be the set of vertices witnessing that there is an induced ab-path of length less than N in $\mathbb{M}$. By the minimality of $\mathbb{M}$, it must be the case that every element of $\mathbb{M}$ belongs to F or to P_{ab} for some $a, b \in F$ (that belong to different components of $\mathbb{F}$). Hence, the cardinality of M is bounded by $|F|$, plus $|F|^2$ times the maximum number m of vertices in an induced path of length $N-1$ — notice that the last m only depends on the signature of $\mathbb{M}$ and on N, so it does not depend on $\mathbb{M}$.

We denote this finite set of minimal structures in $\mathcal{C}(\mathbb{F})$ by $\min(\mathcal{C}(\mathbb{F}))$. Finally, let

$$\mathcal{F}^* := \mathcal{F}_C \cup \mathcal{P}_N \cup \bigcup_{\mathbb{F} \in \mathcal{F} \setminus \mathcal{F}_C} \min(\mathcal{C}(\mathbb{F})).$$

Since $\mathcal{F}^*$ consists of connected structures, the class of $\mathcal{F}^*$-free structures is closed under disjoint unions, and so, there is a structure $\mathbb{B}$ whose age is the class of finite $\mathcal{F}^*$-free structures (see, e.g., [12, Proposition 2.3.1]). We claim that $\mathbb{B}$ satisfies the itemized statements of the lemma. By construction, it satisfies the third one. To show the first one, let $\mathbb{C}$ be a finite τ-structure that embeds into $\mathbb{A}$. In particular, $\mathbb{C}$ is $\mathcal{F}$-free. Moreover, $\mathbb{A}$ is $\mathcal{P}_N$-free because $\mathbb{A}$ has no induced path of length N, and thus $\mathbb{C}$ is also $\mathcal{P}_N$-free. Hence, $\mathbb{C}$ is $\mathcal{F}^*$-free, and so $\mathbb{C} \in \mathrm{Age}(\mathbb{B})$.

To show the second itemized statement, first observe that if every connected component of a finite structure $\mathbb{C}$ belongs to $\mathrm{Age}(\mathbb{A})$, then $\mathbb{C}$ is a disjoint unions of structures that belong to $\mathrm{Age}(\mathbb{A}) \subseteq \mathrm{Age}(\mathbb{B})$. Since $\mathbb{B}$ has a set of connected bounds, its age is closed under disjoint unions, and so $\mathbb{C} \in \mathrm{Age}(\mathbb{B})$. For the converse implication, it suffices to prove that every connected structure $\mathbb{C}$ in the age of $\mathbb{B}$ belongs to the age of $\mathbb{A}$. Since $\mathbb{C}$ is an $\mathcal{F}^*$-free structure, it contains no induced path of length larger than N. Now, suppose for contradiction that $\mathbb{C} \notin \mathrm{Age}(\mathbb{A})$, so there is an embedding of some $\mathbb{F} \in \mathcal{F}$ into $\mathbb{C}$. Moreover, since $\mathcal{F}_C \subseteq \mathcal{F}^*$, it must be the case that $\mathbb{F}$ is not connected. But since $\mathbb{C}$ is connected and has no path of length larger than N, we must have $\mathbb{C} \in \mathcal{C}(\mathbb{F})$. This implies that $\mathbb{C}$ contains some structure from $\min(\mathcal{C}(\mathbb{F})) \subseteq \mathcal{F}^*$, contradicting the assumption that $\mathbb{C}$ is $\mathcal{F}^*$-free. The claim now follows. $\square$

Lemma B.3. *Let $\mathbb{A}$ be a finitely bounded structure. If there is a positive integer N such that every induced path in $\mathbb{A}$ has length strictly less than N, then* $\mathrm{CSP}(\mathbb{A})$ *is in monotone connected extensional* SNP.

Proof. Let $\mathbb{B}$ be the structure from Lemma B.2 for $\mathbb{A}$. It follows from the second itemized property of $\mathbb{B}$ that $\mathrm{CSP}(\mathbb{A}) = \mathrm{CSP}(\mathbb{B})$. Since $\mathbb{B}$ has a finite set of connected bounds by the third itemized property of $\mathbb{B}$, the claim of the lemma follows from Corollary 5.5. $\square$

Corollary B.4. *If $\mathbb{A}$ is a finite structure, then $\mathrm{CSP}(\mathbb{A})$ is in monotone connected extensional SNP.*