Complexity of Combinations of Qualitative Constraint Satisfaction Problems

Manuel Bodirsky and Johannes Greiner*

Institut für Algebra, TU Dresden
(manuel.bodirsky|johannes.greiner)@tu-dresden.de

Abstract. The CSP of a first-order theory T is the problem of deciding for a given finite set S of atomic formulas whether $T \cup S$ is satisfiable. Let T_1 and T_2 be two theories with countably infinite models and disjoint signatures. Nelson and Oppen presented conditions that imply decidability (or polynomial-time decidability) of $\mathrm{CSP}(T_1 \cup T_2)$ under the assumption that $\mathrm{CSP}(T_1)$ and $\mathrm{CSP}(T_2)$ are decidable (or polynomial-time decidable). We show that for a large class of ω-categorical theories T_1, T_2 the Nelson-Oppen conditions are not only sufficient, but also necessary for polynomial-time tractability of $\mathrm{CSP}(T_1 \cup T_2)$ (unless P=NP).

1 Introduction

Two independent proofs of the finite-domain constraint satisfaction tractability conjecture have recently been published by Bulatov and Zhuk [20, 31], settling the Feder-Vardi dichotomy conjecture. In contrast, the computational complexity of constraint satisfaction problems over infinite domains cannot be classified in general [8]. However, for a restricted class of constraint satisfaction problems that strictly all finite-domain CSPs and captures the vast majority of the problems studied in *qualitative reasoning* (see the survey article [9]) there also is a tractability conjecture (see [3–5, 17]). The situation is similar to the situation for finite-domain CSPs before Bulatov and Zhuk: there is a formal condition which provably implies NP-hardness, and the conjecture is that every other CSP in the class is in P.

For finite domain CSPs, it turned out that only few fundamentally different algorithms were needed to complete the classification; the key in both the solution of Bulatov and the solution of Zhuk was a clever combination of the existing algorithmic ideas. An intensively studied method for obtaining (polynomial-time) decision procedures for infinite-domain CSPs is the Nelson-Oppen combination method; see, e.g., [2, 30]. The method did not play any role for the classification of finite-domain CSPs, but is extremely powerful for combining algorithms for infinite-domain CSPs.

* Both authors have received funding from the European Research Council (ERC Grant Agreement no. 681988), the German Research Foundation (DFG, project number 622397), and the DFG Graduiertenkolleg 1763 (QuantLA).

In order to conveniently state what type of combinations of CSPs can be studied with the Nelson-Oppen method, we slightly generalise the notion of a CSP. The classical definition is to fix an infinite structure $\mathfrak{B}$ with finite relational signature τ; then CSP($\mathfrak{B}$) is the computational problem of deciding whether a given finite set of *atomic τ-formulas* (i.e., formulas of the form $x_1 = x_2$ or of the form $R(x_1, \ldots, x_n)$ for $R \in \tau$ and variables $x_1, \ldots, x_n$) is satisfiable in $\mathfrak{B}$. Instead of fixing a τ-structure $\mathfrak{B}$, we fix a *τ-theory* T (i.e., a set of first-order τ-sentences). Then CSP(T) is the computational problem of deciding for a given finite set S of atomic τ-formulas whether $T \cup S$ has a model. Clearly, this is a generalisation of the classical definition since CSP($\mathfrak{B}$) is the same as CSP(Th($\mathfrak{B}$)) where Th($\mathfrak{B}$) is the *first-order theory of* $\mathfrak{B}$, i.e., the set of all first-order sentences that hold in $\mathfrak{B}$. The definition for theories is *strictly* more expressive (we give an example in Section 2 that shows this).

Let T_1 and T_2 be two theories with disjoint finite relational signatures τ_1 and τ_2. We are interested in the question when CSP($T_1 \cup T_2$) can be solved in polynomial time; we refer to this problem as the *combined CSP* for T_1 and T_2. Clearly, if CSP(T_1) or CSP(T_2) is NP-hard, the CSP($T_1 \cup T_2$) is NP-hard, too. Suppose now that CSP(T_1) and CSP(T_2) can be solved in polynomial-time. In this case, there are examples where CSP($T_1 \cup T_2$) is in P, and examples where CSP($T_1 \cup T_2$) is NP-hard. Even if we know the complexity of CSP(T_1) and of CSP(T_2), a classification of the complexity of CSP($T_1 \cup T_2$) for arbitrary theories T_1 and T_2 is too ambitious (see Section 4 for a formal justification). But such a classification should be feasible at least for the mentioned class of infinite-domain CSPs for which the tractability conjecture applies.

1.1 Qualitative CSPs

The idea of *qualitative formalisms* is that reasoning tasks (e.g. about space and time) is not performed with absolute numerical values, but rather with *qualitative* predicates (such as *within*, *before*, etc.). There is no universally accepted definition in the literature that defines what a *qualitative CSP* is, but a proposal has been made in [9]; the central mathematical property for this proposal is *ω-categoricity.* A theory is called *ω-categorical* if it has up to isomorphism only one countable model. A structure is called *ω-categorical* if and only if its first-order theory is ω-categorical. Examples are $(\mathbb{Q}; <)$, Allen's Interval Algebra, and more generally all homogeneous structures with a finite relational signature (a structure $\mathfrak{B}$ is called *homogeneous* if all isomorphisms between finite substructures can be extended to an automorphism; see [6, 25]). The class of CSPs for ω-categorical theories arguably coincides with the class of CSPs for *qualitative formalisms* studied e.g. in temporal and spatial reasoning; see [9].

For an ω-categorical theory T, the complexity of CSP(T) can be studied using the universal-algebraic approach that led to the proof of the Feder-Vardi dichotomy conjecture. One of the central concepts for this approach is the concept of a *polymorphism* of a structure $\mathfrak{B}$, i.e., a homomorphism from $\mathfrak{B}^k$ to $\mathfrak{B}$ for $k \in \mathbb{N}$. It is known that the polymorphisms of a finite structure $\mathfrak{B}$ fully capture the complexity of CSP($\mathfrak{B}$) up to P-time reductions (in fact, up to Log-space reductions; see [24] for a collection of survey articles about the complexity of CSPs), and the same is true for structures $\mathfrak{B}$ with an ω-categorical theory. For an ω-categorical relational structure Γ, the relations that are primitive positive definable in Γ are uniquely determined by the polymorphisms of Γ and vice versa [15]. The possibility to use relations and polymorphisms exchangeably, to study their interplay and to combine known solutions with polymorphisms make the universal algebraic approach a versatile tool. In order to understand when we can apply the universal-algebraic approach to study the complexity of $\mathrm{CSP}(T_1 \cup T_2)$, we need to understand the following fundamental question.

Question 1: Suppose that T_1 and T_2 are theories with disjoint finite relational signatures τ_1 and τ_2. When is there an ω-categorical $(\tau_1 \cup \tau_2)$-theory T such that $\mathrm{CSP}(T)$ equals[1] $\mathrm{CSP}(T_1 \cup T_2)$?

Note that ω-categorical theories are *complete*, i.e., for every first-order sentence ϕ either T implies ϕ or T implies $\neg\phi$. In general, it is not true that $\mathrm{CSP}(T_1 \cup T_2)$ equals $\mathrm{CSP}(T)$ for a complete theory T (we present an example in Section 2).

Question 1 appears to be very difficult. However, we present a broadly applicable condition for ω-categorical theories T_1 and T_2 with infinite models that implies the existence of an ω-categorical theory T such that $\mathrm{CSP}(T_1 \cup T_2)$ equals $\mathrm{CSP}(T)$ (Proposition 1 below). The theory T that we construct has many utile properties, in particular:

1. $T_1 \cup T_2 \subseteq T$;
2. if $\phi_1(\bar{x})$ is a τ_1-formula and $\phi_2(\bar{x})$ is a τ_2-formula, both with free variables $\bar{x} = (x_1, \ldots, x_n)$, then $T \models \exists \bar{x}(\phi_1(\bar{x}) \wedge \phi_2(\bar{x}) \wedge \bigwedge_{i<j} x_i \neq x_j)$ if and only if $T_1 \models \exists \bar{x}(\phi_1(\bar{x}) \wedge \bigwedge_{i<j} x_i \neq x_j)$ and $T_2 \models \exists \bar{x}(\phi_2(\bar{x}) \wedge \bigwedge_{i<j} x_i \neq x_j)$;
3. For every $\tau_1 \cup \tau_2$ formula ϕ there exists a Boolean combination of τ_1 and τ_2 formulas that is equivalent to ϕ modulo T.

In fact, T is uniquely given by these three properties (up to equivalence of theories; see Lemma 2) and again ω-categorical, and we call it the *generic combination of* T_1 *and* T_2. Let $\mathfrak{B}_1$ and $\mathfrak{B}_2$ be two ω-categorical structures whose first-order theories have a generic combination T; then we call the (up to isomorphism unique) countably infinite model of T the *generic combination of* $\mathfrak{B}_1$ *and* $\mathfrak{B}_2$.

1.2 The Nelson-Oppen Criterion

Let T_1, T_2 be theories with disjoint finite relational signatures τ_1, τ_2 and suppose that $\mathrm{CSP}(T_1)$ is in P and $\mathrm{CSP}(T_2)$ is in P. Nelson and Oppen gave sufficient conditions for $\mathrm{CSP}(T_1 \cup T_2)$ to be solvable in polynomial time, too. Their conditions are:

---

[1] In other words: for all sets S of atomic $(\tau_1 \cup \tau_2)$-formulas, we have that $S \cup T$ is satisfiable if and only if $S \cup (T_1 \cup T_2)$ is satisfiable.

1. Both T_1 and T_2 are *stably infinite*: a τ-theory T is called *stably infinite* if for every quantifier-free τ-formula $\phi(x_1, \ldots, x_n)$, if ϕ is satisfiable over T, then there also exists an *infinite* model $\mathfrak{A}$ and elements $a_1, \ldots, a_n$ such that $\mathfrak{A} \models \phi(a_1, \ldots, a_n)$.
2. for $i = 1$ and $i = 2$, the signature τ_i contains a binary relation symbol $\neq_i$ that denotes the inequality relation, i.e., T_i implies the sentence $\forall x, y \, (x \neq_i y \Leftrightarrow \neg(x = y))$;
3. Both T_1 and T_2 are *convex* (here we follow established terminology). A τ-theory T is called *convex* if for every finite set S of atomic τ-formulas the set $T \cup S \cup \{x_1 \neq y_1, \ldots, x_m \neq y_m\}$ is satisfiable whenever $T \cup S \cup \{x_j \neq y_j\}$ is satisfiable for each $j \leq m$.

The assumption that a relation symbol denoting the inequality relation is part of the signatures τ_1 and τ_2 is often implicit in the literature treating the Nelson-Oppen method. It would be interesting to explore when it can be dropped, but we will not pursue this here. The central question of this article is the following.

Question 2. In which settings are the Nelson-Oppen conditions (and in particular, the convexity condition) not only sufficient, but also necessary for polynomial-time tractability of the combined CSP?

Again, for general theories T_1 and T_2, this is a too ambitious research goal; but we will study it for generic combinations of ω-categorical theories T_1, T_2 with infinite models. In this setting, the first condition that both T_1 and T_2 are stably infinite is trivially satisfied. The third condition on T_i, convexity, is equivalent to the existence of a binary injective polymorphism of the (up to isomorphism unique) countably infinite model of T_i (see Section 5). We mention that binary injective polymorphisms played an important role in several recent infinite-domain complexity classifications [10, 11, 27].

1.3 Results

To state our results concerning Question 1 and Question 2 we need basic terminology for permutation groups. A permutation group G on a set A is called

- *n-transitive* if for all tuples $\bar{b}, \bar{c} \in A^n$ having pairwise distinct entries there exists a permutation $g \in G$ such that $g(\bar{b}) = \bar{c}$ (where permutations are applied to tuples componentwise). G is called *transitive* if it is 1-transitive.
- *n-set-transitive* if for all subsets B, C of A with $|B| = |C| = n$ there exists a permutation $g \in G$ such that $g(B) := \{g(b) \mid b \in B\} = C$.

A structure is called n-transitive (or n-set-transitive) if its automorphism group is. The existence of generic combinations can be characterised as follows (see Section 3 for the proof).

Proposition 1. *Let $\mathfrak{B}_1$ and $\mathfrak{B}_2$ be countably infinite ω-categorical structures with disjoint relational signatures. Then $\mathfrak{B}_1$ and $\mathfrak{B}_2$ have a generic combination if and only if both $\mathfrak{B}_1$ and $\mathfrak{B}_2$ do not have algebraicity (in the model-theoretic sense; see Section 3) or at least one of $\mathfrak{B}_1$ and $\mathfrak{B}_2$ has an automorphism group which is n-transitive for all $n \in \mathbb{N}$.*

Our main result concerns Question 2 for generic combinations $\mathfrak{B}$ of countably infinite ω-categorical structures $\mathfrak{B}_1$ and $\mathfrak{B}_2$; as we mentioned before, if the generic combination exists, it is up to isomorphism unique, and again ω-categorical. Note that a structure $\mathfrak{A}$ that is 2-set-transitive gives rise to a directed graph $(A; E)$: fix two distinct elements b_1, b_2 of $\mathfrak{A}$; then two vertices c_1, c_2 are joined by a directed edge iff there exists an automorphism α with $\alpha(b_1) = c_1$ and $\alpha(b_2) = c_2$. Note that by 2-set-transitivity, it does not matter which elements b_1 and b_2 we choose, there are at most two resulting graphs and they are always isomorphic. Also note that if the structure is 2-set-transitive and not 2-transitive, then the resulting directed graph is a *tournament*, i.e., it is without loops and for any two distinct vertices a, b either $(a, b) \in E$ or $(b, a) \in E$, but not both. Examples of 2-set-transitive tournaments are the order of the rationals $(\mathbb{Q}; <)$, the countable random tournament (see, e.g., Lachlan [28]), and the countable homogeneous local order $S(2)$ (also see [23]). If $\bar{a} = (a_1, \ldots, a_n) \in B^n$ and G is a permutation group on B then $G\bar{a} := \{(\alpha(a_1), \ldots, \alpha(a_n)) \mid \alpha \in G\}$ is called the *orbit* of $\bar{a}$ (with respect to G); orbits of pairs (i.e., $n = 2$) are also called *orbitals*. Orbitals of pairs of equal elements are called *trivial*. To simplify the presentation, we introduce the following shortcut.

Definition 1. *A structure has property J if it is a countably infinite ω-categorical structure which is 2-set-transitive, but not 2-transitive, and contains binary symbols for the inequality relation and for one of the two non-trivial orbitals.*

We give some examples of structures with property J.

Example 1 *The structure* $(\mathbb{Q}; \neq, <, R_{\mathrm{mi}})$ *where*

$$R_{\mathrm{mi}} := \{(x, y, z) \in \mathbb{Q}^3 \mid x \geq y \vee x > z\}.$$

Polynomial-time tractability of the CSP of this structure has been shown in [12].

Example 2 *The structure* $(\mathbb{Q}; \neq, <, R_{ll})$ *where*

$$R_{ll} := \{(x, y, z) \in \mathbb{Q}^3 \mid x < y \vee x < z \vee x = y = z\}.$$

Polynomial-time tractability of the CSP of this structure has been shown in [13].

Further examples of structure with property J come from expansions of the countable random tournament and the countable homogeneous local order mentioned above. The proof of the following theorem can be found in Section 5.

Theorem 3. *Let $\mathfrak{B}$ be the generic combination of two structures $\mathfrak{B}_1$ and $\mathfrak{B}_2$ with property J such that* $\mathrm{CSP}(\mathfrak{B}_1)$ *and* $\mathrm{CSP}(\mathfrak{B}_2)$ *are in P. Then one of the following applies:*

- $\mathrm{Th}(\mathfrak{B}_1)$ *or* $\mathrm{Th}(\mathfrak{B}_2)$ *is not convex; in this case,* $\mathrm{CSP}(\mathfrak{B})$ *is NP-hard.*
- *Each of* $\mathrm{Th}(\mathfrak{B}_1)$ *and* $\mathrm{Th}(\mathfrak{B}_2)$ *is convex, and* $\mathrm{CSP}(\mathfrak{B})$ *is in P.*

In other words, either the Nelson-Oppen conditions apply, and $\mathrm{CSP}(\mathfrak{B})$ is in P, or otherweise $\mathrm{CSP}(\mathfrak{B})$ is NP-complete.

Example 4 *Let $\mathfrak{B}_1$ be the relational structure $(\mathbb{Q}; <, \neq, R_{\text{mi}})$ where R_{mi} is defined as above. Let $\mathfrak{B}_2 := (\mathbb{Q}; \prec, \not\approx)$ where $\prec$ also denotes the strict order of the rationals, and $\not\approx$ also denotes the inequality relation (we chose different symbols than $<$ and $\neq$ to make the signatures disjoint). It is easy to see that $\mathfrak{B}_1$ and $\mathfrak{B}_2$ satisfy the assumptions of Proposition 1, so they have a generic combination $\mathfrak{B}$. It is also easy to see that $\mathfrak{B}_1$ and $\mathfrak{B}_2$ are 2-set-transitive, but not 2-transitive. We have already mentioned that they also have polynomial-time tractable CSPs. However, $\mathfrak{B}_1$ does not have a convex theory, and hence our result implies that the CSP of the combined structure is NP-complete (we invite the reader to find an NP-hardness proof without using our theorem!).*

A structure $\mathfrak{B}_1$ is called a *reduct* of a structure $\mathfrak{B}_2$, and $\mathfrak{B}_2$ is called an *expansion* of $\mathfrak{B}_1$, if $\mathfrak{B}_1$ is obtained from $\mathfrak{B}_2$ by dropping some of the relations of $\mathfrak{B}_1$. If $\mathfrak{B}_1$ is a reduct of $\mathfrak{B}_2$ with the signature τ then we write $\mathfrak{B}_2^\tau$ for $\mathfrak{B}_1$. An expansion $\mathfrak{B}_2$ of $\mathfrak{B}_1$ is called a *first-order expansion* if all additional relations in $\mathfrak{B}_2$ have a first-order definition in $\mathfrak{B}_1$. A structure $\mathfrak{B}_1$ is called a *first-order reduct* if $\mathfrak{B}_1$ is a reduct of a first-order expansion of $\mathfrak{B}_2$. Note that if a structure $\mathfrak{B}$ is 2-set-transitive then so is every first-order reduct of $\mathfrak{B}$ (since its automorphism group contains the automorphisms of $\mathfrak{B}$).

The CSPs for first-order reducts of $\mathbb{Q}$ have been called *temporal CSPs*; their computational complexity has been classified completely [12]. There are many interesting polynomial-time tractable temporal CSPs that have non-convex theories, which makes temporal CSPs a particularly interesting class for understanding the situation where the Nelson-Oppen conditions do not apply. Generic combinations of temporal CSPs are isomorphic to first-order reducts of the *countable random permutation* introduced in [22] and studied in [29]; a complexity classification of the CSPs of all reducts of the random permutation (as e.g. in [10,12,27] for simpler structures than the random permutation) is out of reach for the current methods (in particular, the classification method via a reduction to the finite-domain CSP dichotomy from [14] cannot be applied).

Examples of ω-categorical structures with 2-transitive automorphism groups can be found in phylogenetic analysis; see [10]. A generic combination of a structure with a 2-transitive automorphism with $(\mathbb{Q}; <)$ is no longer 2-transitive, but still 2-set-transitive (this will become obvious from the results in Section 3). So any 2-transitive structure without algebraicity can be used to produce further interesting examples that satisfy the conditions of Theorem 3.

2 Combinations of CSPs

We already mentioned that our definition of CSPs for theories is a strict generalisation of the notion of CSPs for structures, and this will be clarified by the following proposition which is an immediate consequence of Proposition 2.4.6 in [6].

Proposition 2. *Let T be a first-order theory with finite relational signature. Then there exists a structure $\mathfrak{B}$ such that* $\text{CSP}(\mathfrak{B}) = \text{CSP}(T)$ *if and only if T*

has the Joint Homomorphism Property (JHP)*, that is, for any two models* $\mathfrak{A}$*,* $\mathfrak{B}$ *of* T *there exists a model* $\mathfrak{C}$ *of* T *such that both* $\mathfrak{A}$ *and* $\mathfrak{B}$ *homomorphically map to* $\mathfrak{C}$*.*

Example 5 *A simple example of two theories* T_1, T_2 *with the JHP such that* $T_1 \cup T_2$ *does not have the JHP is given by*

$$T_1 := \{\forall x, y\ ((O(x) \wedge O(y)) \Rightarrow x = y)\}$$
$$T_2 := \{\forall x.\ \neg(P(x) \wedge Q(x)\}$$

Suppose for contradiction that $T_1 \cup T_2$ *has the JHP. Note that*

$$T_1 \cup T_2 \cup \{\exists x(O(x) \wedge P(x))\} \quad \textit{and} \quad T_1 \cup T_2 \cup \{\exists y(O(y) \wedge Q(y))\}$$

are satisfiable. The JHP implies that

$$T_1 \cup T_2 \cup \{\exists x(O(x) \wedge P(x)), \exists y(O(y) \wedge Q(y))\}$$

has a model $\mathfrak{A}$*, so* $\mathfrak{A}$ *has elements* u, v *satisfying* $O(u) \wedge O(v) \wedge P(u) \wedge Q(v)$*. Since* $\mathfrak{A} \models T_1$ *we must have* $u = v$*, and so* $\mathfrak{A}$ *does not satisfy the sentence* $\forall x.\ \neg(P(x) \wedge Q(x))$ *from* T_2*, a contradiction.*

3 Generic Combinations

For general theories T_1, T_2 even the question whether $T_1 \cup T_2$ has the JHP might be a difficult question. But if both T_1 and T_2 are ω-categorical with a countably infinite model that *does not have algebraicity*, then $T_1 \cup T_2$ always has the JHP (a consequence of Lemma 1 below). A structure $\mathfrak{B}$ (and its first-order theory) *does not have algebraicity* if for all first-order formulas $\phi(x_0, x_1, \ldots, x_n)$ and all elements $a_1, \ldots, a_n \in B$ the set $\{a_0 \in B \mid \mathfrak{B} \models \phi(a_0, a_1, \ldots, a_n)\}$ is either infinite or contained in $\{a_1, \ldots, a_n\}$; otherwise, we say that the structure *has algebraicity*.

It is a well-known fact from model theory that the concept of having no algebraicity is closely related to the concept of *strong amalgamation* (see [25], page 138f). The *age* of a relational τ-structure $\mathfrak{B}$ is the class of all finite τ-structures that embed into $\mathfrak{B}$. A class $\mathcal{K}$ of structures has the *amalgamation property* if for all $\mathfrak{A}, \mathfrak{B}_1, \mathfrak{B}_2 \in \mathcal{K}$ and embeddings $f_i \colon \mathfrak{A} \to \mathfrak{B}_i$, for $i = 1$ and $i = 2$, there exist $\mathfrak{C} \in \mathcal{K}$ and embeddings $g_i \colon \mathfrak{B}_i \to \mathfrak{C}$ such that $g_1 \circ f_1 = g_2 \circ f_2$. It has the *strong amalgamation property* if additionally $g_1(B_1) \cap g_2(B_2) = g_1(f_1(A)) = g_2(f_2(A))$. If $\mathcal{K}$ is a class of structures with finite relational signature which is closed under isomorphism, substructures, and has the amalgamation property, then there exists an (up to isomorphism unique) countable homogeneous structure $\mathfrak{B}$ whose age is $\mathcal{K}$ (see [26]). Moreover, in this case $\mathfrak{B}$ has no algebraicity if and only if $\mathcal{K}$ has the strong amalgamation property (see, e.g., [21]). The significance of strong amalgamation in the theory of combining decision procedures has already been pointed out by Bruttomesso, Ghilardi, and Ranise [19]. By the theorem of Ryll-Nardzewski, Engeler, and Svenonius (see [25]) a homogeneous structure with finite relational signature is ω-categorical, and the expansion of an ω-categorial structure by all first-order definable relations is homogeneous.

When $\mathcal{K}$ is a class of structures, we write $I(\mathcal{K})$ for the class of all structures isomorphic to a structure in $\mathcal{K}$. Let τ_1 and τ_2 be disjoint relational signatures, and let $\mathcal{K}_i$ be a class of finite τ_i-structures, for $i \in \{1,2\}$. Then $\mathcal{K}_1 * \mathcal{K}_2$ denotes the class of $(\tau_1 \cup \tau_2)$-structures given by $\{\mathfrak{A} \mid \mathfrak{A}^{\tau_1} \in I(\mathcal{K}_1) \text{ and } \mathfrak{A}^{\tau_2} \in I(\mathcal{K}_2)\}$. If B is a set and $n \in \mathbb{N}$, we write $B^{(n)}$ for the set of tuples from B^n with pairwise distinct entries.

Lemma 1. *Let T_1 and T_2 be ω-categorical theories with disjoint relational signatures τ_1 and τ_2, with infinite models without algebraicity. Then there exists an ω-categorical model $\mathfrak{B}$ of $T_1 \cup T_2$ without algebraicity such that*

$$\text{for all } k \in \mathbb{N}, \bar{a}, \bar{b} \in B^{(k)}: \quad \operatorname{Aut}(\mathfrak{B}^{\tau_1})\bar{a} \cap \operatorname{Aut}(\mathfrak{B}^{\tau_2})\bar{b} \neq \emptyset \tag{1}$$

$$\text{and for all } k \in \mathbb{N}, \bar{a} \in B^{(k)}: \quad \operatorname{Aut}(\mathfrak{B}^{\tau_1})\bar{a} \cap \operatorname{Aut}(\mathfrak{B}^{\tau_2})\bar{a} = \operatorname{Aut}(\mathfrak{B})\bar{a}\,. \tag{2}$$

The proof works via expansion with all first-order definable relations and a Fraïssé-limit. It can be found in the extended version [7].

Note that by the facts on ω-categorical structures mentioned above, the properties (1) and (2) for $\mathfrak{B}$, $\mathfrak{B}^{\tau_1}$, $\mathfrak{B}^{\tau_2}$ are equivalent to items (2) and (3) in Section 1.1 for $T = \operatorname{Th}(\mathfrak{B})$, $T_1 = \operatorname{Th}(\mathfrak{B}^{\tau_1})$, $T_2 = \operatorname{Th}(\mathfrak{B}^{\tau_2})$, respectively. Lemma 1 motivates the following definition.

Definition 2 (Generic Combination). *Let $\mathfrak{B}_1$ and $\mathfrak{B}_2$ be countably infinite ω-categorical structures with disjoint relational signatures τ_1 and τ_2, and let $\mathfrak{B}$ be a model of* $\operatorname{Th}(\mathfrak{B}_1) \cup \operatorname{Th}(\mathfrak{B}_2)$. *If $\mathfrak{B}$ satisfies item* (1) *then we say that $\mathfrak{B}$ is a* free combination *of $\mathfrak{B}_1$ and $\mathfrak{B}_2$. If $\mathfrak{B}$ satisfies both item* (1) *and item* (2) *then we say that $\mathfrak{B}$ is a* generic combination *(or* random combination*; see [1]) of $\mathfrak{B}_1$ and $\mathfrak{B}_2$.*

The following can be shown via a back-and-forth argument.

Lemma 2. *Let $\mathfrak{B}_1$ and $\mathfrak{B}_2$ be countable ω-categorical structures. Then up to isomorphism, there is at most one generic combination of $\mathfrak{B}_1$ and $\mathfrak{B}_2$.*

In later proofs we need the following lemma.

Lemma 3 (Extension Lemma). *For $i = 1$ and $i = 2$, let $\mathfrak{B}_i$ be an ω-categorical structure with signature τ_i such that $\mathfrak{B}_1$ and $\mathfrak{B}_2$ have a generic combination. Let $\bar{a}$, $\bar{b}_1$, $\bar{b}_2$ be tuples such that the tuples $(\bar{a}, \bar{b}_1)$ and $(\bar{a}, \bar{b}_2)$ have pairwise distinct entries and equal length. Then there exist $\alpha_i \in \operatorname{Aut}(\mathfrak{B}_i, \bar{a})$ such that $\alpha_2(\alpha_1(\bar{b}_1)) = \bar{b}_2$.*

Proof. By the definition of free combinations (Property (1)) there exist $\alpha \in \operatorname{Aut}(\mathfrak{B}_1)$ and $\beta \in \operatorname{Aut}(\mathfrak{B}_2)$ such that $\beta(\alpha(\bar{a}, \bar{b}_1)) = (\bar{a}, \bar{b}_2)$. Note that $\alpha(\bar{a})$ lies in the same orbit as $\bar{a}$ both with respect to $\mathfrak{B}_1$ and with respect to $\mathfrak{B}_2$, so by Property (2) of generic combinations there exists an automorphism $\delta \in \operatorname{Aut}(\mathfrak{B})$ that maps $\alpha(\bar{a})$ to $\bar{a}$. Then $\alpha_1 := \delta \circ \alpha$ and $\alpha_2 := \beta \circ \delta^{-1}$ have the desired properties.

We now prove Proposition 1 that we already stated in the introduction, and which states that two countably infinite ω-categorical structures with disjoint relational signatures have a generic combination if and only if both have no algebraicity, or at least one of the structures has an automorphism group which is n-transitive for all $n \in \mathbb{N}$. Note that the countably infinite structures whose automorphism group is n-transitive for all $n \in \mathbb{N}$ are precisely the structures that are isomorphic to a first-order reduct of $(\mathbb{N}; =)$.

Proof. If both $\mathfrak{B}_1$ and $\mathfrak{B}_2$ do not have algebraicity then the existence of an ω-categorical generic combination follows from Lemma 1. If on the other hand $\mathfrak{B}_1$ is n-transitive for all n then an ω-categorical generic combination trivially exists (it will be a first-order expansion of $\mathfrak{B}_2$). The case that $\mathfrak{B}_2$ is n-transitive for all n is analogous.

For the converse direction, let $\mathfrak{B}$ be the generic combination of the τ_1-structure $\mathfrak{B}_1$ and the τ_2-structure $\mathfrak{B}_2$. Recall that $\mathfrak{B}^{\tau_i}$ is isomorphic to $\mathfrak{B}_i$, for $i \in \{1, 2\}$. By symmetry between $\mathfrak{B}_1$ and $\mathfrak{B}_2$, we will assume towards a contradiction that $\mathfrak{B}^{\tau_1}$ has algebraicity and $\text{Aut}(\mathfrak{B}^{\tau_2})$ is not n-transitive for some $n \in \mathbb{N}$. Choose n to be smallest possible, so that $\text{Aut}(\mathfrak{B}^{\tau_2})$ is not n-transitive. Therefore there exist tuples $(b_0, \dots, b_{n-1})$ and $(c_0, \dots, c_{n-1})$, each with pairwise distinct entries, that are in different orbits with respect to $\text{Aut}(\mathfrak{B}^{\tau_2})$. By the minimality of n, there exists $\alpha \in \text{Aut}(\mathfrak{B}^{\tau_2})$ such that $\alpha(b_1, \dots, b_{n-1}) = (c_1, \dots, c_{n-1})$. Algebraicity of $\mathfrak{B}^{\tau_1}$ implies that there exists a first-order τ_1-formula $\phi(x_0, x_1, \dots, x_m)$ and pairwise distinct elements $a_1, \dots, a_m$ of $\mathfrak{B}$ such that $\phi(x, a_1, \dots, a_m)$ holds for precisely one element $x = a_0$ other than $a_1, \dots, a_m$ in $\mathfrak{B}$. By adding unused extra variables to ϕ we can assume that $m \geq n - 1$. Choose elements $b_n, \dots, b_m$ such that the entries of $(b_0, \dots, b_{n-1}, b_n, \dots, b_m)$ are pairwise distinct and define $c_i := \alpha(b_i)$ for $i \in \{n, \dots, m\}$. Since $\mathfrak{B}$ is a free combination, there exist tuples $(b'_0, \dots, b'_m)$, $(c'_0, \dots, c'_m)$ and $\beta_1, \gamma_1 \in \text{Aut}(\mathfrak{B}^{\tau_1})$ and $\beta_2, \gamma_2 \in \text{Aut}(\mathfrak{B}^{\tau_2})$ such that

$$\beta_2(b_0, \dots, b_m) = (b'_0, \dots, b'_m), \qquad \beta_1(b'_0, \dots, b'_m) = (a_0, \dots, a_m),$$
$$\gamma_2(c_0, \dots, c_m) = (c'_0, \dots, c'_m), \qquad \gamma_1(c'_0, \dots, c'_m) = (a_0, \dots, a_m).$$

Because $\gamma_1^{-1} \circ \beta_1 \in \text{Aut}(\mathfrak{B}^{\tau_1})$ and $\gamma_2 \circ \alpha \circ \beta_2^{-1} \in \text{Aut}(\mathfrak{B}^{\tau_2})$ both map $(b'_1, \dots, b'_m)$ to $(c'_1, \dots, c'_m)$, and due to the second condition for generic combinations, there exists $\mu \in \text{Aut}(\mathfrak{B})$ such that $\mu(b'_1, \dots, b'_m) = (c'_1, \dots, c'_m)$. Since any operation in $\text{Aut}(\mathfrak{B}^{\tau_1})$ preserves ϕ, we have $\gamma_1 \circ \mu \circ \beta_1^{-1}(a_0, \dots, a_m) = (a_0, \dots, a_m)$. Therefore μ must map b'_0 to c'_0. Hence, $\gamma_2^{-1} \circ \mu \circ \beta_2 \in \text{Aut}(\mathfrak{B}^{\tau_2})$ maps $(b_0, \dots, b_{n-1})$ to $(c_0, \dots, c_{n-1})$, contradicting our assumption that they lie in different orbits with respect to $\text{Aut}(\mathfrak{B}^{\tau_2})$.

4 Difficulties for a General Complexity Classification

Let T_1 and T_2 be ω-categorical theories with disjoint finite relational signatures such that $\text{CSP}(T_1)$ is in P and $\text{CSP}(T_2)$ is in P. The results in this section suggest that in general we cannot hope to get a classification of the complexity of $\mathrm{CSP}(T_1 \cup T_2)$. We use the result from [8] that there are homogeneous directed graphs $\mathfrak{B}$ such that $\mathrm{CSP}(\mathfrak{B})$ is undecidable. There are even homogeneous directed graphs $\mathfrak{B}$ such that $\mathrm{CSP}(\mathfrak{B})$ is *coNP-intermediate*, i.e., in coNP, but neither coNP-hard nor in P [8] (unless P = coNP). All of the homogeneous graphs $\mathfrak{B}$ used in [8] can be described by specifying a set of finite tournaments $\mathcal{T}$. Let $\mathcal{C}$ be the class of all finite directed loopless graphs $\mathfrak{A}$ such that no tournament from $\mathcal{T}$ embeds into $\mathfrak{A}$. It can be checked that $\mathcal{C}$ is a strong amalgamation class; the Fraïssé-limits of those classes are called the *Henson digraphs*.

Proposition 3. *For every Henson digraph $\mathfrak{B}$ there exist ω-categorical convex theories T_1 and T_2 with disjoint finite relational signatures such that $\mathrm{CSP}(T_1)$ is in P, $\mathrm{CSP}(T_2)$ is in P, and $\mathrm{CSP}(T_1 \cup T_2)$ is polynomial-time Turing equivalent to $\mathrm{CSP}(\mathfrak{B})$.*

The proof is omitted for reasons of space, but can be found in [7]. Note that the Nelson-Oppen conditions do not apply here because it is crucial for our construction that T_1 does not contain a symbol for inequality. We mention that another example of two theories such that $\mathrm{CSP}(T_1)$ and $\mathrm{CSP}(T_2)$ are decidable but $\mathrm{CSP}(T_1 \cup T_2)$ is not can be found in [18].

5 On the Necessity of the Nelson-Oppen Conditions

In this section we introduce a large class of ω-categorical theories where the condition of Nelson and Oppen (the existence of binary injective polymorphisms) is not only a sufficient, but also a necessary condition for the polynomial-time tractability of generic combinations (unless P = NP); in particular, we prove Theorem 3 from the introduction. We need the following characterisation of convexity of ω-categorical theories.

Theorem 6 (Lemma 6.1.3 in [6]).
Let $\mathfrak{B}$ be an ω-categorical structure and let T be its first-order theory. Then the following are equivalent.

- *T is convex;*
- *$\mathfrak{B}$ has a binary injective polymorphism.*

Moreover, if $\mathfrak{B}$ contains the relation $\neq$, these conditions are also equivalent to the following.

- *for every finite set S of atomic τ-formulas such that $S \cup T \cup \{x_1 \neq y_1\}$ is satisfiable and $S \cup T \cup \{x_2 \neq y_2\}$ is satisfiable, then $T \cup S \cup \{x_1 \neq y_1, x_2 \neq y_2\}$ is satisfiable, too.*

The following well-known fact easily follows from many published results, e.g., from the results in [16]. An operation $f \colon B^k \to B$ is called *essentially unary* if there exists an $i \leq k$ and a function $g \colon B \to B$ such that $f(x_1, \dots, x_k) = g(x_i)$ for all $x_1, \dots, x_k \in B$. The operation f is called *essential* if it is not essentially unary.

Proposition 4 (see [16]). *Let $\mathfrak{B}$ be an infinite ω-categorical structure with finite relational signature containing the relation $\neq$ and such that all polymorphisms of $\mathfrak{B}$ are essentially unary. Then* CSP($\mathfrak{B}$) *is NP-hard.*

Hence, we want to show that the existence of an essential polymorphism of the generic combination of two countably infinite ω-categorical structures $\mathfrak{B}_1$ and $\mathfrak{B}_2$ implies the existence of a binary injective polymorphism. The key technical result, which we prove at the end of this section, is the following proposition.

Proposition 5. *Let $\mathfrak{B}_1, \mathfrak{B}_2$ be ω-categorical structures with generic combination $\mathfrak{B}$ so that*

- *each of $\mathfrak{B}_1$ and $\mathfrak{B}_2$ has a relation symbol that denotes the relation $\neq$;*
- *$\mathfrak{B}$ has a binary essential polymorphism;*
- *$\mathfrak{B}_1$ is 2-set-transitive; and*
- *$\mathfrak{B}_2$ is 1-transitive and contains a binary antisymmetric irreflexive relation.*

Then $\mathfrak{B}_1$ must have a binary injective polymorphism.

To apply Proposition 5, we therefore need to prove the existence of binary essential polymorphisms of generic combinations $\mathfrak{B}$. For this, we use an idea that first appeared in [12] and was later generalized in [6], based on the following concept. A permutation group G on a set B has the *orbital extension property (OEP)* if there is an orbital O such that for all $b_1, b_2 \in B$ there is an element $c \in B$ where $(b_1, c) \in O$ and $(b_2, c) \in O$. The relevance of this property comes from the following lemma.

Lemma 4 (Kára's Lemma; see [6], Lemma 5.3.10). *Let $\mathfrak{B}$ be a structure with an essential polymorphism and an automorphism group with the OEP. Then $\mathfrak{B}$ must have a binary essential polymorphism.*

To apply this lemma to the generic combination $\mathfrak{B}$ of $\mathfrak{B}_1$ and $\mathfrak{B}_2$, we have to verify that Aut($\mathfrak{B}$) has the OEP.

Lemma 5. *Any 2-set-transitive permutation group action on a set with at least 3 elements has the OEP.*

Lemma 6. *Let $\mathfrak{B}$ be a generic combination of two ω-categorical structures $\mathfrak{B}_1$ and $\mathfrak{B}_2$ with the OEP. Then $\mathfrak{B}$ has the OEP.*

Now, we proof Theorem 3. Property J from Definition 1 is needed in order to apply Proposition 5 twice.

Proof (Proof of Theorem 3). If all polymorphisms of $\mathfrak{B}$ are essentially unary then Proposition 4 shows that CSP($\mathfrak{B}$) is NP-hard. Otherwise, $\mathfrak{B}$ has a binary essential polymorphism by Lemma 4, because $\mathfrak{B}$ has the OEP by Lemma 5 and Lemma 6. Property J implies that $\mathfrak{B}_i$, for $i = 1$ and $i = 2$, is 2-set-transitive and contains a binary relation symbol that denotes $\neq$ and a binary relation symbol that denotes the orbital of $\mathfrak{B}_i$, which is a binary antisymmetric irreflexive relation. Thus, $\mathfrak{B}_1$ and $\mathfrak{B}_2$ satisfy the assumptions of Proposition 5. It follows that $\mathfrak{B}_1$ has a binary injective polymorphism. By Theorem 6, this shows that $\text{Th}(\mathfrak{B}_1)$ is convex. Since we have the same assumptions on $\mathfrak{B}_1$ and on $\mathfrak{B}_2$, we can use Proposition 5 again to show that also $\text{Th}(\mathfrak{B}_2)$ is convex. Now, the Nelson-Oppen combination procedure implies that $\text{CSP}(\mathfrak{B})$ is in P.

Proof (Proof of Proposition 5). Since $\mathfrak{B}_1$ and $\mathfrak{B}_2$ have a generic combination, by Proposition 1 either both $\mathfrak{B}_1$ and $\mathfrak{B}_2$ have no algebraicity or at least one of $\mathfrak{B}_1, \mathfrak{B}_2$ is n-transitive for all $n \in \mathbb{N}$. The structure $\mathfrak{B}_2$ is not 2-transitive. Suppose that $\mathfrak{B}_1$ is n-transitive for all $n \in \mathbb{N}$. Since $\mathfrak{B}$ has a binary essential polymorphism, so has $\mathfrak{B}_1$. Since $\mathfrak{B}_1$ also contains a symbol that denotes the relation $\neq$, it must also have a binary injective polymorphism (see [11]) and we are done. So we assume in the following that both $\mathfrak{B}_1$ and $\mathfrak{B}_2$ do not have algebraicity. Since $\mathfrak{B}_1$ and $\mathfrak{B}_2$ are isomorphic to reducts of $\mathfrak{B}$, we may assume that they actually are reducts of $\mathfrak{B}$. Let ϕ be a primitive positive formula over the signature of $\mathfrak{B}_1$ and suppose that $\phi \wedge x_1 \neq y_1$ has a satisfying assignment s_1 over $\mathfrak{B}_1$ and $\phi \wedge x_2 \neq y_2$ has a satisfying assignment s_2 over $\mathfrak{B}_2$. By Theorem 6 it suffices to show that in $\mathfrak{B}_1$ there exists a satisfying assignment to

$$\phi \wedge x_1 \neq y_1 \wedge x_2 \neq y_2. \tag{3}$$

If $s_1(x_2) \neq s_1(y_2)$ or if $s_2(x_1) \neq s_2(y_1)$ then there is nothing to be shown, so we assume that this is not the case. Let f be the binary essential polymorphism of $\mathfrak{B}$. Then there are $a_1, a_2, a_3, b_1, b_2, b_3 \in B$ such that $f(a_2, b_1) \neq f(a_3, b_1)$ and $f(a_1, b_2) \neq f(a_1, b_3)$. It is easy to see that then there also exist elements $u_1, u_2, v_1, v_2 \in B$ such that $f(u_1, v_1) \neq f(u_2, v_1)$ and $f(u_1, v_1) \neq f(u_1, v_2)$ (choose $u_1 = a_1$, $v_1 = b_1$ and suitable $u_2 \in \{a_2, a_3\}$, $v_2 \in \{b_2, b_3\}$). Note that in particular $u_1 \neq u_2$ and $v_1 \neq v_2$. By the 2-set-transitivity of $\mathfrak{B}_1$, there exist $\alpha_1, \alpha_2 \in \text{Aut}(\mathfrak{B}_1)$ such that

$$\alpha_1(\{s_1(x_1), s_1(y_1)\}) = \{u_1, u_2\} \quad \text{and} \quad \alpha_2(\{s_2(x_2), s_2(y_2)\}) = \{v_1, v_2\}.$$

By renaming variables if necessary we may assume that $\alpha_1(s_1(x_1), s_1(y_1)) = (u_1, u_2)$ and $\alpha_2(s_2(x_2), s_2(y_2)) = (v_1, v_2)$.
Note that $|s_1(\{x_1, y_1, x_2, y_2\})|, |s_2(\{x_1, y_1, x_2, y_2\})| \in \{2, 3\}$.

Case 1. $|s_1(\{x_1, y_1, x_2, y_2\})| = |s_2(\{x_1, y_1, x_2, y_2\})| = 3$. In other words, $s_1(x_2) = s_1(y_2) \notin \{s_1(x_1), s_1(y_1)\}$ and $s_2(x_1) = s_2(y_1) \notin \{s_2(x_2), s_2(y_2)\}$.

By the transitivity of $\text{Aut}(\mathfrak{B}_1)$ there exist $\beta_1, \beta_2 \in \text{Aut}(\mathfrak{B}_1)$ such that $\beta_1(s_1(x_2)) = u_1$ and $\beta_2(s_2(y_1)) = v_1$. We can choose $\beta_1 \in \text{Aut}(\mathfrak{B}_1)$ such that $\beta_1(s_1(x_1)), \beta_1(s_1(y_1))$ are distinct from $\alpha_1(s_1(x_2))$ and u_2: to see this, note that $\text{Aut}(\mathfrak{B}_1, u_1)$ has no finite orbits other than $\{u_1\}$ because $\mathfrak{B}_1$ has no algebraicity, and by Neumann's lemma (see e.g. [25], page 141, Corollary 4.2.2) there exists a $g \in \text{Aut}(\mathfrak{B}_1, u_1)$ such that

$$g(\{\beta_1(s_1(x_1)), \beta_1(s_1(x_1))\}) \cap \{\alpha_1(s_1(x_2)), u_2\} = \emptyset.$$

We can thus replace β by $g \circ \beta$. Hence, $u_1, u_2, \beta_1(s_1(x_1)), \alpha_1(s_1(x_2)), \beta_1(s_1(y_1))$ are pairwise distinct. Likewise, we can choose $\beta_2 \in \text{Aut}(\mathfrak{B}_2)$ such that v_1, v_2, $\beta_2(s_2(x_2))$,$\alpha_2(s_2(x_1))$, $\beta_2(s_2(y_2))$ are pairwise distinct.

Let R be the binary antisymmetric irreflexive relation of $\mathfrak{B}_2$, choose any $(a,b) \in R$, and let $\alpha \in \mathrm{Aut}(\mathfrak{B}_2)$ be such that $\alpha(a) = b$. Define $c := \alpha(b)$ and note that $c \neq a$ since otherwise $(a,b), (b,a) \in R$ contrary to our assumptions. Since $\mathfrak{B}$ is a generic combination and $\mathfrak{B}_1$, $\mathfrak{B}_2$ are transitive, $\mathfrak{B}$ is transitive as well and we can choose a, b, c disjoint from u_1, u_2, v_1, v_2 by Neumanns Lemma as above. Then the Extension Lemma (Lemma 3) asserts the existence of elements u_3, u_4, u_5 and automorphisms $\delta_{i,1} \in \mathrm{Aut}(\mathfrak{B}_i)$, for $i \in \{1,2\}$, such that

$$\begin{aligned} \delta_{1,1}(u_1, u_2, u_3, u_4, u_5) &= (u_1, u_2, \beta_1(s_1(x_1)), \alpha_1(s_1(x_2)), \beta_1(s_1(y_1))) \\ \text{and } \delta_{2,1}(u_1, u_2, u_3, u_4, u_5) &= (u_1, u_2, a, b, c)\,. \end{aligned}$$

Similarly, there are elements v_3, v_4, v_5 and $\delta_{i,2} \in \mathrm{Aut}(\mathfrak{B}_i)$, for $i \in \{1,2\}$, such that

$$\begin{aligned} \delta_{1,2}(v_1, v_2, v_3, v_4, v_5) &= (v_1, v_2, \beta_2(s_2(x_2)), \alpha_2(s_2(x_1)), \beta_2(s_2(y_2))) \\ \text{and } \delta_{2,2}(v_1, v_2, v_3, v_4, v_5) &= (v_1, v_2, a, b, c)\,. \end{aligned}$$

See Figure 1. If $f(u_4, v_3) \neq f(u_4, v_5)$, then

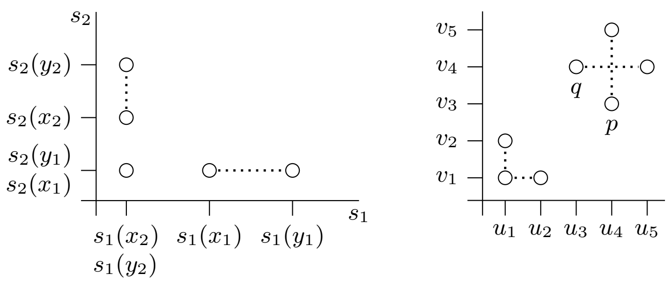

Fig. 1. An illustration of the first case in the proof of Proposition 5. Dashed edges indicate (potential) inequalities between function values.

$$s := f(\delta_{1,1}^{-1}\alpha_1 s_1, \delta_{1,2}^{-1}\beta_2 s_2)$$

is a solution to (3):

$$\begin{aligned} s(x_1) &= f(\delta_{1,1}^{-1}\alpha_1 s_1(x_1), \delta_{1,2}^{-1}\beta_2 s_2(x_1)) = f(\delta_{1,1}^{-1}(u_1), \delta_{1,2}^{-1}(v_1)) = f(u_1, v_1) \\ &\neq f(u_2, v_1) = f(\delta_{1,1}^{-1}(u_2), \delta_{1,2}^{-1}(v_1)) = f(\delta_{1,1}^{-1}\alpha_1 s_1(y_1), \delta_{1,2}^{-1}\beta_2 s_2(y_1)) = s(y_1) \\ s(x_2) &= f(\delta_{1,1}^{-1}\alpha_1 s_1(x_2), \delta_{1,2}^{-1}\beta_2 s_2(x_2)) = f(u_4, v_3) \\ &\neq f(u_4, v_5) = f(\delta_{1,1}^{-1}\alpha_1 s_1(x_2), v_5) = f(\delta_{1,1}^{-1}\alpha_1 s_1(y_2), \delta_{1,2}^{-1}\beta_2 s_2(y_2)) = s(y_2). \end{aligned}$$

So let us assume that $p := f(u_4, v_3) = f(u_4, v_5)$. If $f(u_3, v_4) \neq f(u_5, v_4)$, then

$$s := f(\delta_{1,1}^{-1}\beta_1 s_1, \delta_{1,2}^{-1}\alpha_2 s_2)$$

is a solution to (3), by similar reasoning as above. Thus, we also assume that $q := f(u_3, v_4) = f(u_5, v_4)$. As $(a, b) \in R$, $(b, c) \in R$, $\delta_{2,1}, \delta_{2,2} \in \mathrm{Aut}(\mathfrak{B}_2)$, and f preserves R,

$$(p, q) = (f(u_4, v_3), f(u_5, v_4)) = (f(\delta_{2,1}^{-1}(b), \delta_{2,2}^{-1}(a)), f(\delta_{2,1}^{-1}(c), \delta_{2,2}^{-1}(b))) \in R.$$

Similarly,

$$(q, p) = (f(u_3, v_4), f(u_4, v_5)) = (f(\delta_{2,1}^{-1}(a), \delta_{2,2}^{-1}(b)), f(\delta_{2,1}^{-1}(b), \delta_{2,2}^{-1}(c))) \in R.$$

Hence, both $(p, q) \in R$ and $(q, p) \in R$, contradicting our assumptions.

Case 2. $|s_1(\{x_1, y_1, x_2, y_2\})| = |s_2(\{x_1, y_1, x_2, y_2\})| = 2$.

Case 2a. $s_1(x_2) = s_1(y_2) = s_1(x_1)$ and $s_2(x_1) = s_2(y_1) = s_2(x_2)$. In this case, it is easy to verify that

$$s := f(\alpha_1(s_1), \alpha_2(s_2))$$

is a solution to (3).

Case 2b. $s_1(x_2) = s_1(y_2) = s_1(y_1)$ and $s_2(x_1) = s_2(y_1) = s_2(x_2)$. This case can be proven similarly to Case 1 and is written out in [7].

Case 2c. $s_1(x_2) = s_1(y_2) = s_1(x_1)$ and $s_2(x_1) = s_2(y_1) = s_2(y_2)$. This case can be shown analogously to case 2b (swap x and y).

Case 2d. $s_1(x_2) = s_1(y_2) = s_1(y_1)$ and $s_2(x_1) = s_2(y_1) = s_2(y_2)$. This case can be shown analogously to case 2a (swap x and y).

Case 3. $|s_1(\{x_1, y_1, x_2, y_2\})| = 3$ and $|s_2(\{x_1, y_1, x_2, y_2\})| = 2$.

Case 3a. $s_1(x_2) = s_1(y_2) \notin \{s_1(x_1), s_1(y_1)\}$ and $s_2(x_1) = s_2(y_1) = s_2(x_2)$. The proof is similar to the first and to the second case.

Case 3b. $s_1(x_2) = s_1(y_2) \notin \{s_1(x_1), s_1(y_1)\}$ and $s_2(x_1) = s_2(y_1) = s_2(y_2)$. The proof is analogous to Case 3a.

Case 4. $|s_1(\{x_1, y_1, x_2, y_2\})| = 2$ and $|s_2(\{x_1, y_1, x_2, y_2\})| = 3$. This case is symmetric to Case 3 (swap s_1 and s_2).

6 Conclusion and Future Work

For many theories T_1 and T_2 we have shown that the Nelson-Oppen conditions are not only a sufficient, but also a necessary condition for the polynomial-time tractability of the combined constraint satisfaction problem $\mathrm{CSP}(T_1 \cup T_2)$. Our results imply for example the following complexity classification for combinations of temporal CSPs.

Corollary 1. *Let $\mathfrak{B}_1$ and $\mathfrak{B}_2$ be two first-order expansions of $(\mathbb{Q}; <, \neq)$; rename the relations of $\mathfrak{B}_1$ and $\mathfrak{B}_2$ so that $\mathfrak{B}_1$ and $\mathfrak{B}_2$ have disjoint signatures. Then $\mathrm{CSP}(\mathrm{Th}(\mathfrak{B}_1) \cup \mathrm{Th}(\mathfrak{B}_2))$ is in P if $\mathrm{CSP}(\mathfrak{B}_1)$ and $\mathrm{CSP}(\mathfrak{B}_2)$ are in P and if both $\mathrm{Th}(\mathfrak{B}_1)$ and $\mathrm{Th}(\mathfrak{B}_2)$ are convex. Otherwise, $\mathrm{CSP}(\mathrm{Th}(\mathfrak{B}_1) \cup \mathrm{Th}(\mathfrak{B}_2))$ is NP-hard.*

This follows from Proposition 1 which characterises the existence of a generic combination of T_1 and T_2, and from Theorem 3 which classifies the computational complexity of the generic combination.

It would be interesting to show our complexity result for even larger classes of ω-categorical theories T_1 and T_2. It would also be interesting to drop the assumption that the signatures of T_1 and T_2 contain a symbol for the inequality relation.

References

1. N. Ackerman, C. Freer, and R. Patel. Invariant measures concentrated on countable structures. *Forum of Mathematics Sigma*, 4, 2016.
2. F. Baader and K. Schulz. Combining constraint solving. *H. Comon, C. March, and R. Treinen, editors, Constraints in Computational Logics*, 2001.
3. L. Barto, M. Kompatscher, M. Olšák, M. Pinsker, and T. V. Pham. The equivalence of two dichotomy conjectures for infinite domain constraint satisfaction problems. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science – LICS'17*, 2017. Preprint arXiv:1612.07551.
4. L. Barto, J. Opršal, and M. Pinsker. The wonderland of reflections. *Israel Journal of Mathematics*, 2017. To appear. Preprint arXiv:1510.04521.
5. L. Barto and M. Pinsker. The algebraic dichotomy conjecture for infinite domain constraint satisfaction problems. In *Proceedings of the 31th Annual IEEE Symposium on Logic in Computer Science – LICS'16*, pages 615–622, 2016. Preprint arXiv:1602.04353.
6. M. Bodirsky. Complexity classification in infinite-domain constraint satisfaction. Mémoire d'habilitation à diriger des recherches, Université Diderot – Paris 7. Available at arXiv:1201.0856, 2012.
7. M. Bodirsky and J. Greiner. Complexity of combinations of qualitative constraint satisfaction problems. Preprint arXiv:1801.05965, 2018.
8. M. Bodirsky and M. Grohe. Non-dichotomies in constraint satisfaction complexity. In L. Aceto, I. Damgard, L. A. Goldberg, M. M. Halldórsson, A. Ingólfsdóttir, and I. Walukiewicz, editors, *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, Lecture Notes in Computer Science, pages 184 –196. Springer Verlag, July 2008.
9. M. Bodirsky and P. Jonsson. A model-theoretic view on qualitative constraint reasoning. *Journal of Artificial Intelligence Research*, 58:339–385, 2017.
10. M. Bodirsky, P. Jonsson, and T. V. Pham. The Complexity of Phylogeny Constraint Satisfaction Problems. *ACM Transactions on Computational Logic (TOCL)*, 18(3), 2017. An extended abstract appeared in the conference STACS 2016.
11. M. Bodirsky and J. Kára. The complexity of equality constraint languages. *Theory of Computing Systems*, 3(2):136–158, 2008. A conference version appeared in the proceedings of Computer Science Russia (CSR'06).
12. M. Bodirsky and J. Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2):1–41, 2009. An extended abstract appeared in the Proceedings of the Symposium on Theory of Computing (STOC).
13. M. Bodirsky and J. Kára. A fast algorithm and Datalog inexpressibility for temporal reasoning. *ACM Transactions on Computational Logic*, 11(3), 2010.

14. M. Bodirsky and A. Mottet. Reducts of finitely bounded homogeneous structures, and lifting tractability from finite-domain constraint satisfaction. In *Proceedings of the 31th Annual IEEE Symposium on Logic in Computer Science – LICS'16*, pages 623–632, 2016. Preprint available at ArXiv:1601.04520.
15. M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.
16. M. Bodirsky and M. Pinsker. Topological Birkhoff. *Transactions of the American Mathematical Society*, 367:2527–2549, 2015.
17. M. Bodirsky, M. Pinsker, and A. Pongrácz. Projective clone homomorphisms. Accepted for publication in the Journal of Symbolic Logic, Preprint arXiv:1409.4601, 2014.
18. M. P. Bonacina, S. Ghilardi, E. Nicolini, S. Ranise, and D. Zucchelli. *Decidability and Undecidability Results for Nelson-Oppen and Rewrite-Based Decision Procedures*, pages 513–527. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
19. R. Bruttomesso, S. Ghilardi, and S. Ranise. Quantifier-free interpolation in combinations of equality interpolating theories. *ACM Trans. Comput. Log.*, 15(1):5, 2014.
20. A. A. Bulatov. A dichotomy theorem for nonuniform CSPs. Accepted for publication at FOCS 2017, arXiv:1703.03021, 2017.
21. P. J. Cameron. *Oligomorphic permutation groups.* Cambridge University Press, Cambridge, 1990.
22. P. J. Cameron. Homogeneous permutations. *Electronic Journal of Combinatorics*, 9(2), 2002.
23. G. Cherlin. The classification of countable homogeneous directed graphs and countable homogeneous n-tournaments. *AMS Memoir*, 131(621), January 1998.
24. N. Creignou, P. G. Kolaitis, and H. Vollmer, editors. *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar]*, volume 5250 of *Lecture Notes in Computer Science*. Springer, 2008.
25. W. Hodges. *Model theory.* Cambridge University Press, 1993.
26. W. Hodges. *A shorter model theory.* Cambridge University Press, Cambridge, 1997.
27. M. Kompatscher and T. V. Pham. A Complexity Dichotomy for Poset Constraint Satisfaction. In *34th Symposium on Theoretical Aspects of Computer Science (STACS 2017)*, volume 66 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47:1–47:12, 2017.
28. A. H. Lachlan. Countable homogeneous tournaments. *TAMS*, 284:431–461, 1984.
29. J. Linman and M. Pinsker. Permutations on the random permutation. *Electronic Journal of Combinatorics*, 22(2):1–22, 2015.
30. G. Nelson and D. C. Oppen. Fast decision procedures based on congruence closure. *J. ACM*, 27(2):356–364, 1980.
31. D. Zhuk. The Proof of CSP Dichotomy Conjecture. Accepted for publication at FOCS 2017, arXiv:1704.01914, 2017.