Promise and Infinite-Domain Constraint Satisfaction

Antoine Mottet
Hamburg University of Technology, Research Group for Theoretical Computer Science, Germany

Abstract

Two particularly active branches of research in constraint satisfaction are the study of *promise* constraint satisfaction problems (PCSPs) with finite templates and the study of *infinite-domain* constraint satisfaction problems with ω-categorical templates. In this paper, we explore some connections between these two hitherto unrelated fields and describe a general approach to studying the complexity of PCSPs by constructing suitable infinite CSP templates. As a result, we obtain new characterizations of the power of various classes of algorithms for PCSPs, such as first-order logic, arc consistency reductions, and obtain new proofs of the characterizations of the power of the classical linear and affine relaxations for PCSPs.

2012 ACM Subject Classification Theory of computation → Complexity theory and logic

Keywords and phrases promise constraint satisfaction problems, polymorphisms, homogeneous structures, first-order logic

Digital Object Identifier 10.4230/LIPIcs.CSL.2024.41

1 Introduction

Promise constraint satisfaction problems (PCSPs) are problems of the following form: given a set of constraints on some variables, each of which coming as a pair of strong/weak constraints, determine whether the set of strong constraints is satisfiable or not even the set of weak constraints is satisfiable. Formally, such problems are parametrized by a pair $(\mathbb{A}, \mathbb{B})$ of relational structures such that $\mathbb{A}$ admits a homomorphism to $\mathbb{B}$. Typically, both structures are assumed to be finite. An instance of $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is a structure $\mathbb{X}$, whose domain are understood as variables, and whose tuples in relations correspond to constraints in the same signature as $\mathbb{A}$ and $\mathbb{B}$. The problem is to decide whether there exists a homomorphism $\mathbb{X} \to \mathbb{A}$, representing a solution to a system of strong constraints, or no homomorphism $\mathbb{X} \to \mathbb{B}$, representing the absence of a solution to a weakening of the constraints. When $\mathbb{A} = \mathbb{B}$, one recovers the classical framework of constraint satisfaction problems, for which it is known that if $\mathbb{A}$ is finite, then the associated CSP is either solvable in polynomial time or is NP-hard [44, 43, 21]. Other refined classifications are known, for instance it is known which CSPs are definable by a first-order sentence [2, 41], or by a sentence in fixpoint logic with counting [3], or solvable by Datalog programs [9], or by the basic linear relaxation [35].

The study of PCSPs started recently in the works of [6] and [18], whose motivation was to define a structurally rich framework dedicated to the study of the complexity of approximation problems such as approximate graph coloring. These problems form a framework suitable to the study of a combinatorial, or qualitative, form of approximation, compared to the usual quantitative form. This combinatorial viewpoint allows for conceptually simpler proofs of inapproximability results, such as a combinatorial version of the PCP theorem [10].

A powerful algebraic approach to the study of the complexity of PCSPs was given by [8], building on the existing algebraic tools developed in the context of constraint satisfaction. A plethora of results ensued, providing new polynomial-time algorithms solving such problems [23, 19, 27] and new tools for proving hardness [42, 34].

32nd EACSL Annual Conference on Computer Science Logic (CSL 2024).
Editors: Aniello Murano and Alexandra Silva; Article No. 41; pp. 41:1–41:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Another active branch of research in constraint satisfaction is the study of CSPs where the template is an infinite structure. Here, no dichotomy similar to the Bulatov-Zhuk theorem is possible, as every computational problem is equivalent to the CSP of an infinite structure [11, 30]. Nonetheless, a similar algebraic approach has been developed for a class of infinite structures (so called *ω-categorical structures*) that is suitable for the study the complexity of the associated CSPs. We refer to [39] and the introduction of [37] for a description of the state of the art in the area.

There is until now little to no interplay between these two directions of research. While the tractability of some PCSPs with a finite template has been shown by a reduction to infinite-domain CSPs (typically through the use of a linear relaxation or systems of linear equations over the integers), the infinite templates that arise here are not ω-categorical and are not subject to the aforementioned algebraic approach. We note that Barto [7] and Barto and Asimi [1] proved that there are finite PCSP templates $(\mathbb{A}, \mathbb{B})$ such that for every finite structure $\mathbb{C}$ such that $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ reduces to $\mathrm{CSP}(\mathbb{C})$ by a trivial reduction, then $\mathrm{CSP}(\mathbb{C})$ is NP-hard. Thus, in a sense, the use of infinite-domain CSPs is sometimes necessary.

1.1 Contributions

In this work, we show how the by-now classical tools used to study CSPs with ω-categorical templates can also be used essentially as black boxes to tackle some questions arising in the study of PCSPs with finite templates.

Descriptive Complexity of PCSPs

While algorithmic and algebraic questions concerning PCSPs have received much attention in the past years, the logical aspects pertaining to these problems, and in particular questions about their descriptive complexity, remain mostly unexplored. Here, one of the questions of interest is to determine criteria for the existence of a sentence Φ in a given logic $\mathcal{L}$ whose class of finite models separates the yes-instances from the no-instances of a PCSP. More specifically, we say that $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is *solvable* by a sentence Φ if the following holds:

- For every finite structure $\mathbb{X}$ such that $\mathbb{X}$ admits a homomorphism to $\mathbb{A}$, then $\mathbb{X} \models \Phi$,
- For every finite structure $\mathbb{X}$ such that $\mathbb{X}$ does not admit a homomorphism to $\mathbb{B}$, then $\mathbb{X} \not\models \Phi$.

▶ **Problem 1** (Separability problem for $\mathcal{L}$). *For which promise constraint satisfaction problems* $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ *does there exist a sentence* $\Phi \in \mathcal{L}$ *such that* $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ *is solvable by* Φ*?*

In the context of finite-domain CSPs, i.e., when $\mathbb{A} = \mathbb{B}$, questions of this type have been answered for several logics including first-order logic [2, 41] and some fixpoint logics [3, 9], while some important cases remain open, e.g., for the case of CSPs definable in linear Datalog or in fixpoint logics with a rank operator. Atserias and Dalmau [4] have given necessary algebraic conditions for a given PCSP to be solvable by a Datalog program, but no characterization is known at the moment.

We give an answer to Problem 1 in the case that $\mathcal{L}$ is first-order logic.

▶ **Theorem 2.** *Let* $(\mathbb{A}, \mathbb{B})$ *be a finite PCSP template. The following are equivalent:*

1. $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ *is solvable in first-order logic;*
2. $(\mathbb{A}, \mathbb{B})$ *has finite duality;*
3. *There exists a finite structure* $\mathbb{C}$ *with finite duality and such that* $\mathbb{A} \to \mathbb{C} \to \mathbb{B}$.

Interestingly, although Theorem 2 is a statement purely about finite structures, our proof uses a combination of techniques coming from the study of CSPs with infinite structures. We first construct an infinite structure $\mathbb{C}$ with the same properties as in Item 3 of Theorem 2, and using a Ramsey argument we prove that a finite factor of $\mathbb{C}$ also satisfies the required properties. This method appears to be quite flexible and underlies the proofs of the main results of this paper.

Application to the Search PCSP

Arguably, the variant of the PCSP that is most interesting from an application point of view is the following: given a structure $\mathbb{X}$ with the promise that there exists a homomorphism $\mathbb{X} \to \mathbb{A}$, compute a homomorphism $\mathbb{X} \to \mathbb{B}$. For CSPs, the search version can be solved in polynomial time, given an oracle deciding the decision version; this is not known to hold for promise CSPs.

Very little is known in general concerning the search version of promise CSPs: the tractability of the search version of $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is currently only known when there exists a structure $\mathbb{C}$ such that $\mathbb{A} \to \mathbb{C} \to \mathbb{B}$, such that $\mathrm{CSP}(\mathbb{C})$ can be solved in polynomial-time, and such that a homomorphism $\mathbb{C} \to \mathbb{B}$ can be computed in polynomial time. So far, this is only used when $\mathbb{C}$ is a finite structure.

A straightforward consequence of Theorem 2 is that if $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is solvable in first-order logic, then its search version is tractable. We are in the advantageous situation where $\mathbb{C}$ can be taken finite, but in fact the result would already follow from the existence of a suitable infinite $\mathbb{C}$. Thus, we anticipate that our methods can prove the tractability of the search version of PCSPs in much more general settings than first-order solvability.

Revisiting Arc Consistency and the Arc Consistency Reduction

Arc consistency is a common heuristic employed in constraint solvers to reduce the search space and potentially speed up the process of finding a solution to a CSP instance (or to prove that no solution exists). In this heuristic, one stores for every variable of the instance a set of possible values that this variable can take in a homomorphism $\mathbb{X} \to \mathbb{A}$, and one makes gradual refinements until a fixed point is reached.

Some PCSPs, known as width-1 PCSPs, are in fact completely solved by this heuristic: whenever the heuristic does not yield conclude the absence of a homomorphism $\mathbb{X} \to \mathbb{A}$, then there does actually exist a homomorphism $\mathbb{X} \to \mathbb{B}$. A characterization of finite width-1 CSPs was given by Feder and Vardi [29] in their seminal paper, and generalized to the case of PCSPs [8]. We reprove in Section 4 this characterization by a straightforward adaptation of the method used to characterize the power of first-order logic for PCSPs. Using the same methods, we also give in Section 4 an alternative description of the recent arc consistency reduction proposed in [28].

A current open problem in the theory of PCSPs is to characterize those templates whose PCSP is solved by a strengthening of the arc consistency heuristic, where information about tuples of variables of bounded length can be stored. This corresponds to solving Problem 1 in the case of Datalog, or the existential positive fragment of least fixpoint logic. While we do not solve this problem here, we argue in Section 4 that our method could shed some light on this problem.

Revisiting Linear and Affine Relaxations

Another heuristic to solve PCSPs is to consider linear relaxations of CSP instances, where a CSP instance is mapped to a linear program (known as the basic linear programming relaxation) or to a system of linear equations over the integers (known as the affine integer programming relaxation). The class of PCSPs that are correctly solved by these heuristic have been characterized in [8]. With the methods used in the previous sections, we give another viewpoint on these classes and reprove those characterizations.

2 Definitions

A signature σ is a set of relation symbols, each of which having an arity. A σ-structure $\mathbb{A}$ consists of a set A, called its *domain*, together with an interpretation $R^{\mathbb{A}} \subseteq A^r$ for each relation symbol $R \in \sigma$ of arity $r \geq 1$. A σ-structure $\mathbb{B}$ with $B \subseteq A$ is a *substructure* of $\mathbb{A}$ if for every R in σ of arity r, we have $R^{\mathbb{B}} = R^{\mathbb{A}} \cap B^r$.[1] If $\tau \subseteq \sigma$ and $\mathbb{A}$ is a σ-structure, the τ-reduct of $\mathbb{A}$ is the τ-structure obtained by forgetting the interpretation of the symbols in $\sigma \setminus \tau$. An *expansion* of a structure $\mathbb{A}$ is a structure in a larger signature obtained by adding new relations to $\mathbb{A}$. All relational structures in this paper are at most countable and have a finite signature unless specified otherwise.

A homomorphism $h \colon \mathbb{A} \to \mathbb{B}$ between two σ-structures is a map $A \to B$ such that for all $R \in \sigma$ of arity r and $(a_1, \dots, a_r) \in R^{\mathbb{A}}$, we have $(h(a_1), \dots, h(a_r)) \in R^{\mathbb{B}}$. We use the notation $\mathbb{A} \to \mathbb{B}$ to denote the existence of a homomorphism from $\mathbb{A}$ to $\mathbb{B}$.

A cycle in a structure $\mathbb{A}$ is a set of tuples $\overline{a}^1, \dots, \overline{a}^k$ of length $r_1, \dots, r_k$, each of which appearing in a relation of $\mathbb{A}$, such that the set consisting of all entries of the tuples has size at most $\sum_i (r_i - 1)$. The smallest k for which there exists such a cycle in $\mathbb{A}$ is called the *girth* of $\mathbb{A}$. A *tree* is a structure with no cycles.

2.1 Promise Constraint Satisfaction Problems

For every two structures $\mathbb{A}, \mathbb{B}$ such that there exists a homomorphism $\mathbb{A} \to \mathbb{B}$, we define $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ as the problem of deciding, given a finite structure $\mathbb{X}$, if there exists a homomorphism $\mathbb{X} \to \mathbb{A}$ or no homomorphism $\mathbb{X} \to \mathbb{B}$; the promise is that at least one of these cases holds, and the existence of a homomorphism $\mathbb{A} \to \mathbb{B}$ ensures that at most one case holds. The pair of structures $(\mathbb{A}, \mathbb{B})$ is called the *template* of the PCSP. The search version of $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ asks to compute a homomorphism $\mathbb{X} \to \mathbb{B}$, given a finite structure $\mathbb{X}$ that is promised to admit a homomorphism to $\mathbb{A}$ (although a homomorphism $\mathbb{X} \to \mathbb{A}$ is of course not given). We define $\mathrm{CSP}(\mathbb{A})$ as $\mathrm{PCSP}(\mathbb{A}, \mathbb{A})$. Every tuple $(x_1, \dots, x_k) \in R^{\mathbb{X}}$ in an instance $\mathbb{X}$ of $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is called a *constraint*.

We say that $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is first-order solvable if there exists a first-order sentence Φ such that the following items hold for every finite structure $\mathbb{X}$:

- if $\mathbb{X}$ has a homomorphism to $\mathbb{A}$, then $\mathbb{X} \models \Phi$,
- if $\mathbb{X} \models \Phi$, then $\mathbb{X}$ admits a homomorphism to $\mathbb{B}$.

Thus, the set of finite models of Φ separates the yes-instances of $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ from the no-instances. If $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is first-order solvable, one trivially gets a logspace algorithm solving $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$.[2]

---

[1] As usual in model theory, all substructures are necessarily *induced substructures*.

[2] The truth of a first-order formula with k quantifiers can be checked by iterating over all elements of the input structure and storing k logarithmically-sized pointers.

Let $\mathcal{F}$ be a family of finite structures. A structure $\mathbb{X}$ is *$\mathcal{F}$-free* if there does not exist any $\mathbb{F} \in \mathcal{F}$ such that $\mathbb{F}$ admits a homomorphism to $\mathbb{X}$. We say that a PCSP template $(\mathbb{A}, \mathbb{B})$ has duality $\mathcal{F}$ if the following hold for every finite $\mathbb{X}$:

- if $\mathbb{X}$ admits a homomorphism to $\mathbb{A}$, then $\mathbb{X}$ is $\mathcal{F}$-free,
- if $\mathbb{X}$ is $\mathcal{F}$-free, then $\mathbb{X}$ admits a homomorphism to $\mathbb{B}$.

We say that $(\mathbb{A}, \mathbb{B})$ has *finite duality* if it has duality $\mathcal{F}$ for a finite set $\mathcal{F}$. A structure has *finite tree duality* if it has a finite duality consisting only of trees.

An operation $f\colon A^n \to A$ is a *polymorphism* of $\mathbb{A}$ if it is a homomorphism from $\mathbb{A}^n$ to $\mathbb{A}$, where $\mathbb{A}^n$ is the structure with domain A^n, the same signature as $\mathbb{A}$, and with relations

$$R^{\mathbb{A}^n} := \left\{(\overline{a}^1, \ldots, \overline{a}^r) \mid \forall j, (a^1_j, \ldots, a^r_j) \in R^{\mathbb{A}}\right\}.$$

In other words, f is a polymorphism if, whenever $\overline{a}^1, \ldots, \overline{a}^n$ are in a relation R of $\mathbb{A}$, then $f(\overline{a}^1, \ldots, \overline{a}^n)$ is in R, where f is applied componentwise to every tuple $\overline{a}^i$. We define a polymorphism of a PCSP template $(\mathbb{A}, \mathbb{B})$ similarly, as a homomorphism $\mathbb{A}^n \to \mathbb{B}$.

A polymorphism f of a structure $\mathbb{A}$ is *1-tolerant* if it satisfies that $f(\overline{a}^1, \ldots, \overline{a}^n)$ is in $R^{\mathbb{A}}$, whenever all but at most one of $\overline{a}^1, \ldots, \overline{a}^n$ is in $R^{\mathbb{A}}$. For finite-domain or ω-categorical CSPs, the following characterization of first-order solvability is known.

▶ **Theorem 3** ([36, 2, 41, 12]). *Let $\mathbb{A}$ be a finite or ω-categorical structure. The following are equivalent:*

- *$\mathbb{A}$ has finite duality,*
- *$\mathbb{A}$ has finite tree duality,*
- $\mathrm{CSP}(\mathbb{A})$ *is solvable in first-order logic,*
- *$\mathbb{A}$ has a 1-tolerant polymorphism.*

It is proven in [36] that it is possible to decide, for a finite structure $\mathbb{A}$, whether $\mathbb{A}$ has finite duality.

2.2 ω-categorical structures

An *embedding* $e\colon \mathbb{A} \to \mathbb{B}$ is an injective homomorphism such that for every relation R of arity r and every $a_1, \ldots, a_r \in A$, one has $(e(a_1), \ldots, e(a_r)) \in R^{\mathbb{B}}$ if, and only if, $(a_1, \ldots, a_r) \in R^{\mathbb{A}}$. An *automorphism* of a structure $\mathbb{A}$ is an embedding $\alpha\colon \mathbb{A} \to \mathbb{A}$ that is surjective. In other words, both α and its inverse are homomorphisms $\mathbb{A} \to \mathbb{A}$.

A structure $\mathbb{A}$ if *ω-categorical* is for all $n \geq 1$, the equivalence relation $\sim^{\mathbb{A}}_n$ on A^n defined by $x \sim^{\mathbb{A}}_n y$ iff there exists an automorphism $\alpha \in \mathrm{Aut}(\mathbb{A})$ with $\alpha(x) = y$ has finitely many equivalence classes, . These equivalence classes are called *orbits under* $\mathrm{Aut}(\mathbb{A})$. Typical examples of ω-categorical structures are the "structure with no structure" $(\mathbb{N}; =)$, for which the classes of the equivalence relation $\sim_n$ are in 1-to-1 correspondence with partitions of $\{1, \ldots, n\}$, and $(\mathbb{Q}; <)$, for which the classes of the equivalence relation $\sim_n$ correspond to weak linear orders on $\{1, \ldots, n\}$.

A structure $\mathbb{A}$ is *homogeneous* if every isomorphism $f\colon \mathbb{B} \to \mathbb{C}$ between finite substructures of $\mathbb{A}$ extends to an automorphism of $\mathbb{A}$. Thus, in a homogeneous structure, the orbits under $\mathrm{Aut}(\mathbb{A})$ are completely determined by the isomorphism types of n-element substructures of $\mathbb{A}$, or equivalently by the quantifier-free formulas with n variables up to equivalence over $\mathbb{A}$.

A countable set $\mathcal{C}$ of finite structures is said to have the *amalgamation property* if for all structures $\mathbb{X}, \mathbb{Y}_1, \mathbb{Y}_2 \in \mathcal{C}$ and embeddings $f_i\colon \mathbb{X} \to \mathbb{Y}_i$, there exist a structure $\mathbb{Z} \in \mathcal{C}$ and embeddings $e_i\colon \mathbb{Y}_i \to \mathbb{Z}$ such that $e_1 \circ f_1 = e_2 \circ f_2$. We say that $\mathbb{Z}$ is an *amalgam* over $\mathbb{Y}_1, \mathbb{Y}_2$ over $\mathbb{X}$. We say that $\mathbb{Z}$ is a *strong amalgam* if $e_1(Y_1) \cap e_2(Y_2) = e_1(f_1(X))$, and we say $\mathcal{C}$ has the *strong amalgamation property* when $\mathbb{Z}$ can always be chosen to be a strong amalgam, regardless of $\mathbb{X}, \mathbb{Y}_1, \mathbb{Y}_2$. We say that $\mathbb{Z}$ is a *free amalgam* if it is strong and no tuple containing entries from both $e_1(Y_1 \setminus f_1(X))$ and $e_2(Y_2 \setminus f_2(X))$ belongs to a relation of $\mathbb{Z}$. By a classical result of Fraïssé, for every countable class of finite structures $\mathcal{C}$ that is closed under substructures, there exists a countable homogeneous structure $\mathbb{C}$ whose finite substructures are exactly those structures that are isomorphic to a member of $\mathcal{C}$. The structure $\mathbb{C}$ is called the *Fraïssé limit* of $\mathcal{C}$.

2.3 Ramsey expansions and canonical polymorphisms

An operation $f\colon A^n \to B$ is *canonical from* $\mathbb{A}$ *to* $\mathbb{B}$ if it is a homomorphism from the nth power of $(A; \sim_1^{\mathbb{A}}, \sim_2^{\mathbb{A}}, \dots)$ to $(B; \sim_1^{\mathbb{B}}, \sim_2^{\mathbb{B}}, \dots)$. In other words, f is canonical from $\mathbb{A}$ to $\mathbb{B}$ if, and only if, for all m-tuples $\overline{a}^1, \dots, \overline{a}^n$, and all $\alpha_1, \dots, \alpha_n$ automorphisms of $\mathbb{A}$, the tuples $f(\overline{a}^1, \dots, \overline{a}^m)$ and $f(\alpha_1(\overline{a}^1), \dots, \alpha_n(\overline{a}^n))$ are in the same orbit under $\operatorname{Aut}(\mathbb{B})$. Canonical functions typically arise by an application of the following result. This result uses the *Ramsey* property of some homogeneous structures; we will only use the Ramsey property as a blackbox in this paper and therefore omit the definition.

▶ **Theorem 4** ([16, 40]). *Let $\mathbb{A}$ be a homogeneous structure with the Ramsey property, let $\mathbb{B}$ be an ω-categorical structure, and let $f\colon A^n \to B$ be an arbitrary function. Then there exists $g\colon A^n \to B$ that is canonical from $\mathbb{A}$ to $\mathbb{B}$ and such that for every finite subset S of A^m, there exist $\alpha_1, \dots, \alpha_n \in \operatorname{Aut}(\mathbb{A}), \beta \in \operatorname{Aut}(\mathbb{B})$ such that $g(\overline{a}^1, \dots, \overline{a}^n) = \beta f(\alpha_1(\overline{a}^1), \dots, \alpha_n(\overline{a}^n))$ holds for all $\overline{a}^1, \dots, \overline{a}^n$ in S.*

In case $\mathbb{A} = \mathbb{B}$ in Theorem 4, we say that f *locally interpolates* g *modulo* $\operatorname{Aut}(\mathbb{A})$.

3 PCSPs solvable in First-Order Logic

It was first proven by [2] that a finite structure has finite duality if, and only if, its CSP can be defined in first-order logic. Another proof of this result was obtained by [41], who proved the following stronger statement.[3] For a first-order sentence Φ, let $\operatorname{Mod}(\Phi)$ be the class of all finite structures $\mathbb{X}$ such that $\mathbb{X} \models \Phi$.

▶ **Theorem 5** (Theorem 4.11 in [41]). *Let $\mathcal{P} \subseteq \mathcal{Q}$ be classes of structures, and Φ be a first-order sentence such that:*
- *for all finite $\mathbb{X}, \mathbb{Y}$ such that $\mathbb{X} \in \mathcal{P}$ and $\mathbb{X} \to \mathbb{Y}$, we have $\mathbb{Y} \models \Phi$,*
- *for all finite $\mathbb{X}, \mathbb{Y}$ such that $\mathbb{X} \models \Phi$ and $\mathbb{X} \to \mathbb{Y}$, we have $\mathbb{Y} \in \mathcal{Q}$.*

Then there exists an existential positive sentence Ψ such that $\mathcal{P} \subseteq \operatorname{Mod}(\Psi) \subseteq \mathcal{Q}$.

The following result, initially by Cherlin, Shelah, and Shi [22], and later improved by Hubička and Nešetřil [32], has found several applications in the study of infinite-domain CSPs in the recent years [13, 14, 15]. A structure is *connected* if it is not isomorphic to the disjoint union of two non-empty structures.

▶ **Theorem 6.** *Let $\mathcal{F}$ be a finite set of finite connected structures. There exists an ω-categorical structure $\mathbb{C}$ such that $\mathbb{C}$ has duality $\mathcal{F}$. Moreover, $\mathbb{C}$ can be chosen to have an expansion $\mathbb{C}^+$ by finitely many relations such that $\mathbb{C}^+$ is homogeneous with the Ramsey property.*

[3] To see that Theorem 5 implies Atserias's result, apply the theorem to $\mathcal{P} = \mathcal{Q} = \{\mathbb{X} \mid \mathbb{X} \not\to \mathbb{A}\}$.

In particular, $\mathbb{C}^+$ meets the hypothesis of Theorem 4, and therefore every polymorphism of $\mathbb{C}$ locally interpolates a polymorphism that is canonical with respect to $\mathrm{Aut}(\mathbb{C}^+)$. We now show that the property of being 1-tolerant is preserved under this local interpolation.

▶ **Lemma 7.** *Let $\mathbb{A}$ be an arbitrary structure. Let f be a 1-tolerant polymorphism of $\mathbb{A}$ and let Γ be a subset of $\mathrm{Aut}(\mathbb{A})$. Then every operation g that is locally interpolated by f modulo Γ is a 1-tolerant polymorphism of $\mathbb{A}$.*

Proof. Let g be locally interpolated by f modulo Γ, and let R be a relation of $\mathbb{A}$. Let $\overline{a}^1,\ldots,\overline{a}^n \in A^r$ be such that all but at most one of them are in R. By assumption, there exist $\alpha_1,\ldots,\alpha_n,\beta\in\Gamma$ such that $g(\overline{a}^1,\ldots,\overline{a}^n) = \beta f(\alpha_1(\overline{a}^1),\ldots,\alpha_n(\overline{a}^n))$. Since $\alpha_1,\ldots,\alpha_n$ are automorphisms of $\mathbb{A}$, we obtain that for all but at most one j, one has $\alpha_j(\overline{a}^j)\in R$. Since f is 1-tolerant, $f(\alpha_1(\overline{a}^1),\ldots,\alpha_n(\overline{a}^n))$ is in R, and thus $g(\overline{a}^1,\ldots,\overline{a}^n)$ is in R since β is an automorphism of $\mathbb{A}$. ◀

In the following, recall that a first-order formula is *primitive positive* if it only consists of existential quantifications, conjunctions, and atomic formulas only. Every primitive positive formula $\varphi(x_1,\ldots,x_n)$ without equalities corresponds in a canonical way to a relational structure $\mathbb{A}$, its *canonical database*, whose domain is the set of variables of the formula, and whose relations are determined by the conjuncts of the formula.

▶ **Theorem 2.** *Let $(\mathbb{A},\mathbb{B})$ be a finite PCSP template. The following are equivalent:*

1. $\mathrm{PCSP}(\mathbb{A},\mathbb{B})$ *is solvable in first-order logic;*
2. $(\mathbb{A},\mathbb{B})$ *has finite duality;*
3. *There exists a finite structure $\mathbb{C}$ with finite duality and such that $\mathbb{A}\to\mathbb{C}\to\mathbb{B}$.*

Proof. (1) implies (2). This is an immediate consequence of Theorem 5. Let $\mathcal{P}$ be the class of finite structures that do not admit a homomorphism to $\mathbb{B}$, and $\mathcal{Q}$ be the class of finite structures that do not admit a homomorphism to $\mathbb{A}$. Let Φ be a first-order sentence proving that $\mathrm{PCSP}(\mathbb{A},\mathbb{B})$ is in FO. Then $\mathcal{P}\subseteq \mathrm{Mod}(\neg\Phi)\subseteq\mathcal{Q}$ holds by definition. Moreover, if $\mathbb{X}\in\mathcal{P}$ and $\mathbb{X}\to\mathbb{Y}$, then $\mathbb{Y}\not\to\mathbb{B}$, hence $\mathbb{Y}\models\neg\Phi$. Similarly, if $\mathbb{X}\to\mathbb{Y}$ and $\mathbb{X}\models\neg\Phi$, then $\mathbb{X}$ does not admit a homomorphism to $\mathbb{A}$, so $\mathbb{Y}$ does not admit a homomorphism to $\mathbb{A}$, i.e., $\mathbb{Y}\in\mathcal{Q}$. Thus, Theorem 5 applies, and there exists an existential positive formula Ψ such that $\mathcal{P}\subseteq\mathrm{Mod}(\Psi)\subseteq\mathcal{Q}$, and Ψ is equivalent to a disjunction $\bigvee\Psi_i$ where each Ψ_i is a primitive positive sentence. Moreover, each Ψ_i can be assumed without loss of generality to not contain any equalities.[4] Let $\mathcal{F}$ be the set of canonical databases for each Ψ_i. For any finite $\mathbb{X}$, if there exists $\mathbb{F}\in\mathcal{F}$ such that $\mathbb{F}\to\mathbb{X}$, then $\mathbb{X}\models\Psi$, so that $\mathbb{X}\in\mathcal{Q}$ and $\mathbb{X}$ does not admit a homomorphism to $\mathbb{A}$. If $\mathbb{X}$ does not admit a homomorphism to $\mathbb{B}$, then $\mathbb{X}\in\mathcal{P}$ so that $\mathbb{X}\models\Psi$, and therefore there is $\mathbb{F}\in\mathcal{F}$ such that $\mathbb{F}\to\mathbb{X}$. Thus, $\mathcal{F}$ forms a duality for $(\mathbb{A},\mathbb{B})$.

(2) implies (3). Let $\mathcal{F}$ be a duality for $(\mathbb{A},\mathbb{B})$. Without loss of generality, we can assume that $\mathcal{F}$ consists of connected structures. Indeed, suppose that $\mathbb{F}\in\mathcal{F}$ is isomorphic to a disjoint union of non-empty structures $\mathbb{F}_1,\mathbb{F}_2$. Since $\mathbb{A}$ is $\mathcal{F}$-free, there is no homomorphism $\mathbb{F}\to\mathbb{A}$ and therefore one of $\mathbb{F}_1$ or $\mathbb{F}_2$ does not admit a homomorphism to $\mathbb{A}$, say without loss of generality that $\mathbb{F}_1$ does not. Consider $\mathcal{F}' := (\mathcal{F}\cup\{\mathbb{F}_1\})\setminus\{\mathbb{F}\}$, which we prove is a duality for $(\mathbb{A},\mathbb{B})$. Suppose that $\mathbb{X}$ admits a homomorphism to $\mathbb{A}$. Then $\mathbb{X}$ is $\mathcal{F}$-free and also $\mathbb{F}_1$-free since $\mathbb{F}_1\not\to\mathbb{A}$, so that $\mathbb{X}$ is $\mathcal{F}'$-free. If $\mathbb{X}$ is $\mathcal{F}'$-free, then $\mathbb{X}$ is $\mathcal{F}$-free, so that $\mathbb{X}$ admits a homomorphism to $\mathbb{B}$.

[4] An equality in any Ψ_i can be removed by merging the corresponding variables.

By Theorem 6, there exists an ω-categorical structure $\mathbb{C}$ that has duality $\mathcal{F}$. Since $\mathbb{A}$ is $\mathcal{F}$-free, we have $\mathbb{A} \to \mathbb{C}$. Since every finite substructure of $\mathbb{C}$ is $\mathcal{F}$-free, there is a homomorphism from every finite substructure of $\mathbb{C}$ to $\mathbb{B}$. By compactness, there exists a homomorphism $\mathbb{C} \to \mathbb{B}$.

Since $\mathbb{C}$ is ω-categorical and has finite duality, there is an $f\colon \mathbb{C}^n \to \mathbb{C}$ that is a 1-tolerant polymorphism of $\mathbb{C}$ by Theorem 3. Moreover, $\mathbb{C}$ admits a Ramsey expansion $\mathbb{C}^+$. By applying Theorem 4 to $\mathbb{C}^+$, f locally interpolates an operation g modulo $\mathrm{Aut}(\mathbb{C}^+)$ that is canonical with respect to $\mathbb{C}^+$. By Lemma 7, g is a 1-tolerant polymorphism of $\mathbb{C}$.

Consider the structure $\mathbb{C}' := \mathbb{C}/\mathrm{Aut}(\mathbb{C}^+)$ whose domain consists of the classes of the equivalence relation $\sim_1$ induced by $\mathrm{Aut}(\mathbb{C}^+)$, and such that for every relation symbol R of arity r in the signature of $\mathbb{C}$, one has $(O_1,\dots,O_r) \in R^{\mathbb{C}'}$ if, and only if, there exist $a_1 \in O_1,\dots,a_r \in O_r$ such that $(a_1,\dots,a_r) \in R^{\mathbb{C}}$. Then g induces a 1-tolerant polymorphism of $\mathbb{C}'$: define $\tilde{g}(O_1,\dots,O_n)$ to be the $\sim_1$-class of $g(a_1,\dots,a_n)$, for arbitrary $a_1 \in O_1,\dots,a_n \in O_n$. Since g is canonical with respect to $\mathbb{C}^+$, the definition of $\tilde{g}$ does not depend on the chosen elements $a_1,\dots,a_n$. One readily checks that $\tilde{g}$ thus defined is a 1-tolerant polymorphism of $\mathbb{C}'$.

Moreover, let h be a homomorphism $\mathbb{C} \to \mathbb{B}$. By Theorem 4 applied with $\mathbb{C}^+$ and $\mathbb{B}$, there exists a homomorphism $h'\colon \mathbb{C} \to \mathbb{B}$ that is canonical from $\mathbb{C}^+$ to $\mathbb{B}$. Similarly as above, h' induces a homomorphism $\tilde{h}'$ from $\mathbb{C}'$ to $\mathbb{B}$. Thus we get that $\mathbb{A} \to \mathbb{C}' \to \mathbb{B}$. Moreover $\mathbb{C}'$ has a 1-tolerant polymorphism, so by Theorem 3, $\mathbb{C}'$ has finite duality.

(3) implies (1). By Theorem 3, $\mathrm{CSP}(\mathbb{C})$ can be defined by a first-order sentence Φ. This sentence shows that $\mathrm{PCSP}(\mathbb{A},\mathbb{B})$ is solvable in first-order logic. ◀

An anonymous reviewer of this paper provided another proof of the implication from (2) to (3) in Theorem 2 using the sparse incomparability lemma (see, e.g., [29]) to show directly that the duality $\mathcal{F}$ can be taken to consist of trees (which in our case follows from an application of Theorem 3), and then using the fact that finite families of trees admit a finite dual structure [38]. Namely, given an arbitrary finite duality $\mathcal{F}$ for $(\mathbb{A},\mathbb{B})$, consider the family $\mathcal{G}$ consisting of homomorphic images of structures from $\mathcal{F}$ and that are trees. If $\mathbb{X}$ admits a homomorphism to $\mathbb{A}$, then it is $\mathcal{F}$-free and therefore $\mathcal{G}$-free. Suppose now that $\mathbb{X} \not\to \mathbb{B}$. By the sparse incomparability lemma, one can find a structure $\mathbb{X}'$ such that $\mathbb{X}' \to \mathbb{X}$ and $\mathbb{X}' \not\to \mathbb{B}$, and $\mathbb{X}'$ has girth larger than the size of any structure in $\mathcal{F}$. Since $\mathbb{X}' \not\to \mathbb{B}$, there exists $\mathbb{F} \in \mathcal{F}$ and a homomorphism $h\colon \mathbb{F} \to \mathbb{X}'$, and the image of $\mathbb{F}$ under h must be a tree, which implies that $\mathbb{X}'$ is not $\mathcal{G}$-free. Since $\mathbb{X}' \to \mathbb{X}$, $\mathbb{X}$ is not $\mathcal{G}$-free either.

Thus, the construction of an infinite structure $\mathbb{C}$ as in our proof of Theorem 2 is not necessary; however our method here applies to a wider setting as the next sections show.

We obtain as a corollary to Theorem 2 a characterization of the pair $(\mathbb{A},\mathbb{B})$ that have finite duality, in the case that $\mathbb{A}$ is a digraph containing a directed cycle.

▶ **Corollary 8.** *Let $(\mathbb{A},\mathbb{B})$ be a PCSP template where $\mathbb{A}$ is a digraph containing a directed cycle. Then $(\mathbb{A},\mathbb{B})$ has finite duality if, and only if, $\mathbb{B}$ contains a loop.*

Proof. If $\mathbb{B}$ has a loop, then the empty set is a duality for $\mathrm{PCSP}(\mathbb{A},\mathbb{B})$. Suppose now that $(\mathbb{A},\mathbb{B})$ has finite duality. By Theorem 2, there exists a finite $\mathbb{C}$ with finite duality and such that $\mathbb{A}$ admits a homomorphism to $\mathbb{C}$ and $\mathbb{C} \to \mathbb{B}$. Since $\mathbb{A}$ has a directed cycle, so does $\mathbb{C}$.

Since $\mathbb{C}$ has finite duality, its duality must consist of trees. Note that every orientation of a tree admits a homomorphism to $\mathbb{C}$ since $\mathbb{C}$ contains a directed cycle, and therefore its duality must be empty. This implies in particular that $\mathbb{C}$ has a loop, and so does $\mathbb{B}$ since $\mathbb{C} \to \mathbb{B}$. ◀

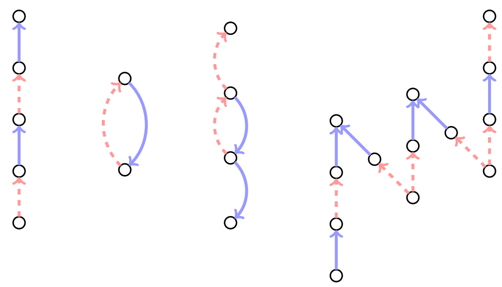

Figure 1 Illustration of the structures $\mathbb{A}$, $\mathbb{C}$, $\mathbb{B}$, and $\mathbb{F}_2$ (from left to right) in Proposition 9. The red dashed arcs correspond to pairs in the relation R, while blue solid arcs correspond to pairs in the relation B.

We conclude this section by showing that there are proper examples of PCSP templates with finite duality.

▶ **Proposition 9.** *There exists a PCSP template* $(\mathbb{A}, \mathbb{B})$ *with finite duality such that neither* $\mathbb{A}$ *nor* $\mathbb{B}$ *has finite duality.*

Proof. Consider the structures in a binary relational signature with two binary symbols R and B displayed in Figure 1. Then $\mathbb{A} \to \mathbb{C} \to \mathbb{B}$ and $\mathbb{C}$ has the duality that consists of the following structures: an R-path of length 2, a B-path of length 2, a vertex with incoming B- and R-edges, and a vertex with outgoing B- and R-edges.

We show that the structure $\mathbb{A}$ does not have finite duality. Let $\mathbb{P}$ be an R-edge followed by a B-edge. We use the notation $+$ and $-$ to denote the obvious amalgamation of copies of $\mathbb{P}$. For every $n \geq 1$, consider the structure $\mathbb{F}_n$ defined by taking $\mathbb{P}+(\mathbb{P}-\mathbb{P})+\cdots+(\mathbb{P}-\mathbb{P})+\mathbb{P}+\mathbb{P}$, with a total of $2n+3$ copies of $\mathbb{P}$, and removing the first vertex and its adjacent R-edge and the last vertex and its adjacent B-edge ($\mathbb{F}_2$ is showed in Figure 1). No $\mathbb{F}_n$ admits a homomorphism to $\mathbb{A}$, although every structure obtained by removing an edge does. Thus, every $\mathbb{F}_n$ must be in a duality for $\mathbb{A}$. Since all the structures $\mathbb{F}_n$ are homomorphically incomparable, we obtain that $\mathbb{A}$ does not have finite duality.

The proof that $\mathbb{B}$ does not have finite duality is similar, where this time one defines $\mathbb{P}$ to be an R-path or a B-path of length 2. ◀

Non-sufficient conditions

For finite-domain CSPs, a number of other conditions are known to be equivalent to the fact that $\mathbb{A}$ has finite duality:

- as mentioned, the existence of a 1-tolerant polymorphism of $\mathbb{A}$,
- the connectivity of a specific graph $L(\mathbb{G}, \mathbb{A}')$ for some retract $\mathbb{A}'$ of $\mathbb{A}$ and all finite structures $\mathbb{G}$ [20],
- the fact that there exists a retract $\mathbb{A}'$ of $\mathbb{A}$ is such that $(\mathbb{A}')^2$ dismantles to its diagonal [36].

It is not clear what could be generalizations of the last two items in the case of promise templates $(\mathbb{A}, \mathbb{B})$. However, the first item has a clear candidate for a generalization, namely the existence of a 1-tolerant polymorphism of $(\mathbb{A}, \mathbb{B})$, i.e., a map $f \colon A^n \to B$ such that for every relation symbol R, and every $\overline{a}^1, \ldots, \overline{a}^n$ such that all but at most one are in $R^{\mathbb{A}}$, then $f(\overline{a}^1, \ldots, \overline{a}^n)$ is in $R^{\mathbb{B}}$.

We remark that a 1-tolerant polymorphism of arity n of a structure $\mathbb{C}$ can be composed with itself to obtain a 1-tolerant polymorphism of any arity $m \geq n$. Thus, in the case of CSPs, first-order solvability can also be characterized by the existence of 1-tolerant polymorphisms of all but finitely many arities.

▶ **Proposition 10.** *Let* $(\mathbb{A},\mathbb{B})$ *be a PCSP template. If* $(\mathbb{A},\mathbb{B})$ *has finite duality, then it has a* 1*-tolerant polymorphism. However, there exists a PCSP template with* 1*-tolerant polymorphisms of all arities but not finite duality.*

Proof. Suppose that $(\mathbb{A},\mathbb{B})$ has finite duality. By Theorem 2, there exists a finite structure $\mathbb{C}$ with finite duality and homomorphisms $h\colon \mathbb{A}\to\mathbb{C}$ and $g\colon\mathbb{C}\to\mathbb{B}$. By Theorem 3, $\mathbb{C}$ has a 1-tolerant polymorphism $f\colon\mathbb{C}^n\to\mathbb{C}$. The composition $g\circ f\circ h$ is then a 1-tolerant polymorphism of $(\mathbb{A},\mathbb{B})$.

Consider now $\mathbb{A} := K_c$ and $\mathbb{B} := K_{c^2}$ to be complete graphs on c and c^2 vertices, for any $c\geq 2$. Let $n\geq 1$. We simply name the vertices of $\mathbb{B}$ by pairs (a,b) of elements of $\mathbb{A}$. Then the map $f\colon A^n\to B$ defined by $f(a_1,\dots,a_n) := (a_1,a_2)$ is a 1-tolerant n-ary polymorphism of $(\mathbb{A},\mathbb{B})$: if $(a_1,b_1),\dots,(a_n,b_n)$ are pairs such that $a_j\neq b_j$ for all but at most one j, then $(a_1,a_2)\neq(b_1,b_2)$. However, it follows from Corollary 8 that $(\mathbb{A},\mathbb{B})$ does not have finite duality, since $\mathbb{A}$ has a directed cycle of length 2 and $\mathbb{B}$ does not have a loop. ◀

4 Local Consistency for PCSPs

We now apply the same reasoning as in the previous section to characterize the power of the arc consistency reduction, recently introduced in [28, 33]. For this reduction, we need the following concepts.

Arc consistency is a polynomial-time algorithm that takes as input a structure $\mathbb{X}$, as an instance of CSP($\mathbb{A}$), and that computes for every $x\in X$ a set $P_x\subseteq A$, and for every constraint C of the form $(x_1,\dots,x_n)\in R^{\mathbb{X}}$ a set $Q_C\subseteq R^{\mathbb{A}}$ such that:

1. for every constraint C whose ith variable is $x_i\in X$, the ith projection of Q_C is equal to P_{x_i},
2. for every homomorphism $h\colon\mathbb{X}\to\mathbb{A}$ and every $x\in X$, we have $h(x)\in P_x$. In particular, if P_x is empty then $\mathbb{X}$ does not admit a homomorphism to $\mathbb{A}$.

A minion is a functor $\mathscr{M}$ from finite sets to sets: for every finite set X, one has a set $\mathscr{M}X$ and for every function $\sigma\colon X\to Y$, one has a function $\mathscr{M}\sigma\colon\mathscr{M}X\to\mathscr{M}Y$, such that $\mathscr{M}\mathrm{id}_X=\mathrm{id}_{\mathscr{M}X}$ for all X and $\mathscr{M}(\sigma\circ\tau)=\mathscr{M}\sigma\circ\mathscr{M}\tau$ whenever the composition of σ and τ is well defined.

A *minor identity* is a formal statement of the form $f^\sigma\approx g^\tau$ where f is a symbol of type X, g is a symbol of type Y, $\sigma\colon X\to Z$ and $\tau\colon Y\to Z$ are functions, and X,Y,Z are arbitrary finite sets. A *minor condition* is a set Σ of minor identities. A minor condition Σ is *satisfied* in $\mathscr{M}$ if the symbols in Σ of any type X can be mapped to elements of $\mathscr{M}X$ such that if $f^\sigma\approx g^\tau$ is in Σ, then $(\mathscr{M}\sigma)(f)=(\mathscr{M}\tau)(g)$. The set $\mathrm{Pol}(\mathbb{A},\mathbb{B})$ can be seen to be such a minion $\mathscr{M}$, where $\mathscr{M}X$ consists of the homomorphisms $\mathbb{A}^X\to\mathbb{B}$, and the functions $\mathscr{M}\sigma$ are obtained by identifying arguments of such homomorphisms according to σ.

The problem $\mathrm{PMC}_N(\mathscr{M})$ is the problem taking as input a minor condition Σ whose symbols have sorts of size at most N, whose yes-instances are those Σ that can be satisfied in *every* minion, and whose no-instances are those that are not satisfiable in $\mathscr{M}$. It is known that for N large enough only depending on the size of $\mathbb{A}$ and the size of its relations, $\mathrm{PMC}_N(\mathrm{Pol}(\mathbb{A},\mathbb{B}))$ and PCSP($\mathbb{A},\mathbb{B}$) are equivalent under logspace reductions [8]. In the following, we drop the subscript and always take N large enough for this equivalence to hold. The arc consistency reduction described below is a complete reduction from PCSP($\mathbb{A},\mathbb{B}$) to PMC($\mathscr{M}$) for a minion $\mathscr{M}$, although it is not necessarily sound; by the previous sentence, this is essentially the same as trying to reduce from one PCSP to another.

This reduction applied to an input $\mathbb{X}$ of PCSP($\mathbb{A},\mathbb{B}$), and whose output is an instance of PMC($\mathscr{M}$), works as follows:

- First, apply the arc consistency algorithm to $\mathbb{X}$, seen as an instance of $\mathrm{CSP}(\mathbb{A})$; one obtains a family of subsets $(P_x)_{x\in X}$ and (Q_C) satisfying the arc consistency condition above,
- Associate a function symbol f_x with every $x \in X$, whose arguments are labelled by the elements in P_x; associate a function symbol f_C with every constraint $C = R(x_1, \ldots, x_k)$, whose arguments are labelled by the tuples in Q_C;
- Output the minor condition containing the identities $f_{x_i} \approx f_C^\sigma$ where $\sigma\colon Q_C \to P_{x_i}$ is the projection on the ith component, as an input to $\mathrm{PMC}(\mathscr{M})$.

We turn the set of instances $\mathbb{X}$ that are not rejected by this reduction into a class with the amalgamation property. The additional symbols $H_{(P,f)}$ are indexed by pairs (P, f) where P is a non-empty subset of A or of $R^{\mathbb{A}}$ for some R, and f is an element of $\mathscr{M}P$. Let $\mathbb{X}$ be an arbitrary finite structure such that there exist a family of subsets $(P_x)_{x\in X}$ and (Q_C) satisfying the arc consistency condition (Item 1) together with a map ξ witnessing the satisfiability of the corresponding minor condition Σ in $\mathscr{M}$. Let $\mathbb{X}^*$ be the expansion of $\mathbb{X}$ where:

- $x \in H_{(P,f)}$ iff $P = P_x$ and $\xi(f_x) = f$,
- for every constraint $(x_1, \ldots, x_n) \in R^{\mathbb{X}}$, we let $(x_1, \ldots, x_n) \in H_{(Q,g)}$ iff $Q_C = Q$ and $\xi(f_C) = g$.

We say that $\mathbb{X}^*$ is a *valid encoding*, witnessed by the sets $(P_x), (Q_C)$ and the map ξ. Let $\mathcal{C}_{red}$ be the class of valid encodings.

▶ **Proposition 11.** *$\mathcal{C}_{red}$ is closed under substructures and has the strong amalgamation property.*

Proof. The fact that $\mathcal{C}_{red}$ is closed under substructures is readily checked.

Let now $\mathbb{X}^*, \mathbb{Y}_1^*, \mathbb{Y}_2^* \in \mathcal{C}_{red}$ and embeddings $f_i\colon \mathbb{X}^* \to \mathbb{Y}_i^*$. Without loss of generality, we can suppose that $X \subseteq Y_1, Y_2$ and that f_i is the inclusion map. Take $\mathbb{Z}^*$ to be the free amalgam of $\mathbb{Y}_1^*$ and $\mathbb{Y}_2^*$ over $\mathbb{X}^*$. We write $\mathbb{Z}$ for the reduct of $\mathbb{Z}^*$ to the signature of $\mathbb{A}$. Let $(P_y^i), (Q_C^i), \xi^i$ be the witnesses for the fact that $\mathbb{Y}_i^*$ is a valid encoding, for $i \in \{1, 2\}$.

Let $P_z := P_z^i$ if $z \in Y_i$; this is well defined since if $z \in Y_1 \cap Y_2 = X$, then we have $P_z^1 = P_z^2$. If C is a constraint $(z_1, \ldots, z_n) \in R^{\mathbb{Z}}$, then by definition $\{z_1, \ldots, z_n\} \subseteq Y_i$ for some $i \in \{1, 2\}$. Let $Q_C := Q_C^i$, and note again that if $\{z_1, \ldots, z_n\} \subseteq Y_1 \cap Y_2$ then $Q_C^1 = Q_C^2$. Define similarly ξ, where Σ is the minor condition arising from the sets (P_z) and (Q_C), by defining $\xi(f_z) := \xi^i(f_z)$ and $\xi(f_C) := \xi^i(f_C)$ for a suitable i.

We prove that the family of sets $(P_z), (Q_C)$ satisfies Item 1. If C is a constraint $(z_1, \ldots, z_n) \in R^{\mathbb{Z}}$, then by definition $\{z_1, \ldots, z_n\} \subseteq Y_i$ for some i. Thus, since $(P_y^i), (Q_C^i)$ satisfies Item 1, the projection of Q_C^i to each component coincides with the corresponding P_z^i, and we are done.

Similarly, it is easy to see that ξ witnesses that Σ is satisfiable in $\mathscr{M}$. Thus, $\mathbb{Z}^*$ is a valid encoding, witnessed by the sets $(P_z), (Q_C)$, and ξ. ◀

Consider the Fraïssé limit $\mathrm{AC}(\mathbb{A}, \mathscr{M})^*$ of $\mathcal{C}_{red}$ and its reduct $\mathrm{AC}(\mathbb{A}, \mathscr{M})$ to the signature of $\mathbb{A}$. One obtains a structure that gives a classification of PCSPs for which the arc consistency reduction correctly solves the problem.

▶ **Theorem 12.** *Let $(\mathbb{A}, \mathbb{B})$ be a PCSP template. The following hold:*

- *There exists a homomorphism $\mathbb{A} \to \mathrm{AC}(\mathbb{A}, \mathscr{M})$,*
- *For every finite structure $\mathbb{B}$, arc consistency is a correct reduction from $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ to $\mathrm{PMC}(\mathscr{M})$ if, and only if, there exists a homomorphism $\mathrm{AC}(\mathbb{A}, \mathscr{M}) \to \mathbb{B}$.*

Proof. Since $\mathbb{A}$ is not rejected by the arc consistency reduction, there exists a valid encoding $\mathbb{A}^* \in \mathcal{C}_{red}$, and therefore $\mathbb{A}^*$ embeds into $\mathrm{AC}(\mathbb{A}, \mathscr{M})^*$. It follows that $\mathbb{A}$ embeds into $\mathrm{AC}(\mathbb{A}, \mathscr{M})$.

If arc consistency is a correct reduction, then all the structures that have a valid encoding in $\mathcal{C}_{red}$ admit a homomorphism to $\mathbb{B}$; thus, a compactness argument gives that $\mathrm{AC}(\mathbb{A}, \mathscr{M})$ admits a homomorphism to $\mathbb{B}$.

For the other direction, we remark that the reduction is always *complete*, i.e., every yes-instance of $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is mapped to a yes-instance of $\mathrm{PMC}(\mathscr{M})$. Conversely, suppose that the minor condition Σ that is computed by the algorithm on an instance $\mathbb{X}$ admits a solution in $\mathscr{M}$. We give a homomorphism $\mathbb{X} \to \mathbb{B}$ as follows. Let $\mathbb{X}^* \in \mathcal{C}_{red}$ be a valid encoding, which exists since $\mathbb{X}$ is not rejected by the reduction. We get an embedding $\mathbb{X}^* \to \mathrm{AC}(\mathbb{A}, \mathscr{M})^*$, giving a homomorphism $\mathbb{X} \to \mathrm{AC}(\mathbb{A}, \mathscr{M})$, which can then be composed with the homomorphism $\mathrm{AC}(\mathbb{A}, \mathscr{M}) \to \mathbb{B}$ that exists by assumption, from which we get a homomorphism $\mathbb{X} \to \mathbb{B}$. ◀

Note that if the number of elements in $\mathscr{M}P$ is bounded for every subset P of A or of a relation of A, then $\mathrm{AC}(\mathbb{A}, \mathscr{M})^*$ is homogeneous in a finite language, and is therefore ω-categorical. By [32, Theorem 2.11], the expansion of $\mathrm{AC}(\mathbb{A}, \mathscr{M})^*$ by a free linear order is a Ramsey structure. Let $\mathscr{G}$ be its automorphism group. Thus, the existence of a homomorphism $\mathrm{AC}(\mathbb{A}, \mathscr{M}) \to \mathbb{B}$ is equivalent, by Theorem 4, to the existence of a homomorphism $\mathrm{AC}(\mathbb{A}, \mathscr{M})/\mathscr{G} \to \mathbb{B}$, which can be effectively tested when $\mathrm{AC}(\mathbb{A}, \mathscr{M})/\mathscr{G}$ is finite. The domain of $\mathbb{D} := \mathrm{AC}(\mathbb{A}, \mathscr{M})/\mathscr{G}$ consists, by homogeneity, of pairs (P, f) where P is a non-empty subset of A and f is in $\mathscr{M}P$. Moreover, if R is an n-ary relation symbol, then $((P_1, f_1), \ldots, (P_n, f_n)) \in R^{\mathbb{D}}$ if, and only if, there exists a $g \in \mathscr{M}(R^{\mathbb{A}} \cap (P_1 \times \cdots \times P_n))$ such that $(\mathscr{M}\sigma_i)(g) = f_i$ holds for all $i \in \{1, \ldots, n\}$, where $\sigma_i \colon R^{\mathbb{A}} \cap (P_1 \times \cdots \times P_n) \to P_i$ is the ith projection. Provided that the elements of $\mathscr{M}$ can be presented to an algorithm in some way, then the structure $\mathbb{D}$ can be computed according to the definition above, and the existence of a homomorphism $\mathbb{D} \to \mathbb{B}$ can be tested. This gives a decision procedure to check whether the arc consistency reduction to $\mathscr{M}$ solves a given PCSP.

It is noted in [28] that arc consistency as an algorithm solving $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ can be seen as a reduction from $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ to $\mathrm{PMC}(\mathscr{M}_0)$, where $\mathscr{M}_0$ is the minion consisting of all operations of finite arity on a 1-element set (i.e., for every n, $\mathscr{M}_0$ contains a single function of arity n). In this case, the elements of the finite structure $\mathbb{D}$ correspond exactly to non-empty subsets of $\mathbb{A}$, and we obtain another proof of the characterization of the power of the arc consistency algorithm for PCSPs by means of the powerset structure defined by Feder and Vardi [29].

The arc consistency procedure can be generalized to the *k-consistency algorithm* by computing sets $P_{x_1,\ldots,x_k} \subseteq A^k$ for every k-tuple of elements from $\mathbb{X}$, and asking for similar conditions as in Item 1. We say that a PCSP template $(\mathbb{A}, \mathbb{B})$ has *width k* if every instance $\mathbb{X}$ of $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ that is not rejected by the k-consistency algorithm (when seeing $\mathbb{X}$ as an instance of $\mathrm{CSP}(\mathbb{A})$) admits a homomorphism to $\mathbb{B}$.

When a PCSP template has bounded width (i.e., has width k for some k) then the corresponding PCSP can be solved in polynomial time. It is hitherto not known whether the search version of the PCSP can then be solved in polynomial time. The existence of a structure $\mathbb{C}$ that would allow us to follow the same line of reasoning as for arc consistency is open. Atserias and Toruńczyk [5] proved that the class of locally consistent systems of linear equations over $\mathbb{Z}_2$ cannot be turned into a class with the amalgamation property using an expansion by finitely many relations (i.e., this class is not *homogenizable*). However, in order to obtain a structure $\mathbb{A}'$ playing the role of $\mathrm{AC}(\mathbb{A}, \mathscr{M}_0)$ in Theorem 12 for any $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$

that is solvable by, say, 3-consistency, it is plausible that one only needs the number of orbits of elements of $\mathbb{A}'$ to be finite, a much weaker condition than homogeneity in a finite relational language, or even ω-categoricity.

5 The Basic Linear Programming Relaxation

The *basic linear programming relaxation* of an instance $\mathbb{X}$ of CSP($\mathbb{A}$) is the following linear program with variables $\lambda_x(a)$ for $x \in \mathbb{X}$ and $a \in \mathbb{A}$:

$$\begin{array}{rcll} \sum_{a\in A} \lambda_x(a) & = & 1 & \text{for all } x \in X \\ \sum_{\overline{a}\in R^{\mathbb{A}}} \lambda_C(\overline{a}) & = & 1 & \text{for all constraints } C := \overline{y} \in R^{\mathbb{X}} \\ \sum_{\overline{a}: a_i = b} \lambda_C(\overline{a}) & = & \lambda_x(b) & \text{for all } x,\ b, \text{ and } \overline{y} \text{ s.t. } y_i = x \\ \lambda_x(a), \lambda_C(\overline{a}) & \geq & 0 & \text{for all variables} \end{array} \qquad (\mathrm{BLP}(\mathbb{X},\mathbb{A}))$$

Note that if $\mathbb{X}$ admits a homomorphism to $\mathbb{A}$, then the λ's can be taken to have values in $\{0,1\}$ and to describe completely a homomorphism $\mathbb{X} \to \mathbb{A}$. However, there can be proper solutions to the program BLP($\mathbb{X}, \mathbb{A}$) that do not correspond to any homomorphism.

We say that PCSP($\mathbb{A}, \mathbb{B}$) is *solvable by BLP* if whenever BLP($\mathbb{X}, \mathbb{A}$) has a solution for a given $\mathbb{X}$, then $\mathbb{X}$ admits a homomorphism to $\mathbb{B}$. Note that if PCSP($\mathbb{A}, \mathbb{B}$) is solvable by BLP, then it is in particular solvable in polynomial time. It is proven in [8] that PCSP($\mathbb{A}, \mathbb{B}$) is solvable by BLP if, and only if, a certain structure LP($\mathbb{A}$) whose domain is the set of probability distributions on A admits a homomorphism to $\mathbb{B}$. It is unknown whether it is always true that a PCSP($\mathbb{A}, \mathbb{B}$) that is solvable by BLP is also polynomially solvable in its search variant. Moreover, it is not known whether the BLP-solvability of a given PCSP is decidable.

The λ's represent probability distributions on $\mathbb{A}$ and $R^{\mathbb{A}}$ that are required to be consistent; one sees that the supports of any solution to BLP($\mathbb{X}, \mathbb{A}$) forms a family of sets P_x, Q_C satisfying the arc consistency condition. Thus, the BLP relaxation is more powerful than arc consistency.

The approach used in the previous sections can be applied here, by encoding the probability distributions λ's arising from a solution to BLP($\mathbb{X}, \mathbb{A}$) as suitable relations over $\mathbb{X}$ using additional symbols $P_{\overline{a},q}$ for tuples a from A and $q \in \mathbb{Q} \cap [0,1]$ to encode the probability distributions λ's. This gives a class $\mathcal{C}_{\mathrm{BLP}}$ of structures that has the amalgamation property. Let $\mathbb{C}^*_{\mathrm{BLP}}$ be the Fraïssé limit of $\mathcal{C}_{\mathrm{BLP}}$, and let $\mathbb{C}_{\mathrm{BLP}}$ be its reduct of the signature of $\mathbb{A}$.

▶ **Proposition 13.** *Let* $(\mathbb{A}, \mathbb{B})$ *be a PCSP template. The following are equivalent:*

- PCSP($\mathbb{A}, \mathbb{B}$) *is solvable by BLP,*
- *there exists a homomorphism* $\mathbb{C}_{\mathrm{BLP}} \to \mathbb{B}$*.*

We note that infinitely many new predicates are required for this encoding, and therefore $\mathbb{C}_{\mathrm{BLP}}$ is not ω-categorical, just as is the case of LP($\mathbb{A}$) [8]. In fact, LP($\mathbb{A}$) can be seen to be homomorphically equivalent to $\mathbb{C}_{\mathrm{BLP}}$, using the natural correspondence between probability distributions on A and the orbits of $\mathbb{C}^*_{\mathrm{BLP}}$, which are described by the unary predicates $P_{a,q}$ for $a \in A$ and $q \in \mathbb{Q} \cap [0,1]$.

However, for every $N \geq 1$, one can consider the class $\mathcal{C}^{(N)}_{\mathrm{BLP}}$ of structures $\mathbb{X}$ endowed with rational probability distributions where no denominator is greater than N. Every $\mathcal{C}^{(N)}_{\mathrm{BLP}}$ has the amalgamation property, and therefore a Fraïssé limit $\mathbb{C}^{*,N}_{\mathrm{BLP}}$. Since the language is now finite, each $\mathbb{C}^{*,N}_{\mathrm{BLP}}$ is ω-categorical, and for every $N < M$ we have embeddings $\mathbb{C}^{*,N}_{\mathrm{BLP}} \hookrightarrow \mathbb{C}^{*,M}_{\mathrm{BLP}}$. We thus get the following refinement of Proposition 13.

▶ **Proposition 14.** *Let* $(\mathbb{A}, \mathbb{B})$ *be a PCSP template. The following are equivalent:*
- $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ *is solvable by BLP,*
- $\mathbb{A} \to \mathbb{C}^1_{\mathrm{BLP}} \to \mathbb{C}^2_{\mathrm{BLP}} \to \cdots \to \mathbb{C}_{\mathrm{BLP}} \to \mathbb{B}$.

Once more, [32, Theorem 2.11] applies and gives that each $\mathbb{C}^N_{\mathrm{BLP}}$ admits a homogeneous Ramsey expansion by finitely many relations. It is not immediately clear if this observation and Proposition 14 can be used to derive a decision procedure for solvability of a PCSP by BLP.

6 The Affine Integer Relaxation

Similarly as in Section 5, we can obtain a limit structure characterizing the power of the so-called affine integer relaxation (AIP) [8]. Given an input $\mathbb{X}$ to $\mathrm{CSP}(\mathbb{A})$, the system $\mathrm{AIP}(\mathbb{X}, \mathbb{A})$ is like $\mathrm{BLP}(\mathbb{X}, \mathbb{A})$ except that the variables are integer-valued. If $\mathbb{X} \to \mathbb{A}$, then the same $\{0, 1\}$ solution to $\mathrm{BLP}(\mathbb{X}, \mathbb{A})$ is a solution to $\mathrm{AIP}(\mathbb{X}, \mathbb{A})$, and we say that AIP solves $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ if whenever a solution to $\mathrm{AIP}(\mathbb{X}, \mathbb{A})$ exists, then $\mathbb{X} \to \mathbb{B}$. The power of AIP to solve PCSPs has been characterized in [8] by means of the existence of a structure $\mathrm{IP}(\mathbb{A})$, similarly as $\mathrm{LP}(\mathbb{A})$ characterizes the power of BLP.

The class of structures $\mathbb{X}$ for which $\mathrm{AIP}(\mathbb{X}, \mathbb{A})$ has a solution can be expanded by relations encoding, for each $\mathbb{X}$, a possible solution to $\mathrm{AIP}(\mathbb{X}, \mathbb{A})$. The resulting class $\mathcal{C}_{\mathrm{AIP}}$ of structures has the amalgamation property (where the amalgam is always free), and therefore it has a Fraïssé limit $\mathbb{C}^*_{\mathrm{AIP}}$, whose reduct $\mathbb{C}_{\mathrm{AIP}}$ to the signature of $\mathbb{A}$ characterizes solvability by AIP.

▶ **Proposition 15.** *Let* $(\mathbb{A}, \mathbb{B})$ *be a PCSP template. The following are equivalent:*
- $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ *is solvable by AIP,*
- *there exists a homomorphism* $\mathbb{C}_{\mathrm{AIP}} \to \mathbb{B}$.

7 Further connections to infinite-domain CSPs

We conclude this paper by hinting at further connections between algorithms solving PCSPs and infinite-domain CSPs.

Sandwiches and Monotone Algorithms

All current polynomial-time algorithms for solving or reducing PCSPs feature the use of algorithms solving CSPs either via a trivial reduction, via computationally simple many-one reductions, or as oracles.

In the first case, one solves $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ through a "trivial" reduction to a problem of the form $\mathrm{CSP}(\mathbb{C})$, where the reduction does not transform the input. By definition, this "do-nothing" reduction is valid if, and only if, there exists a homomorphism from $\mathbb{A}$ to $\mathbb{C}$ and from $\mathbb{C}$ to $\mathbb{B}$, which we denote by $\mathbb{A} \to \mathbb{C} \to \mathbb{B}$. Because of this characterization, characterizing when a "do-nothing" reduction is a valid reduction from $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ to $\mathrm{CSP}(\mathbb{C})$ relies on studying the templates that are *sandwiched* between $\mathbb{A}$ and $\mathbb{B}$ in the homomorphism preorder. Barto [7] and Barto and Asimi [1] have showed examples of problems of the form $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ where such a $\mathbb{C}$ with $\mathrm{CSP}(\mathbb{C})$ polynomial-time tractable exists but cannot be taken to be a finite structure. As explained in Sections 5 and 6, the tractability of PCSPs that are solvable by relaxations like the basic linear relaxation and integer affine relaxation can also be explained by sandwiches, taking $\mathbb{C}$ to be a reduct of a given structure whose domain consists of tuples of rational or integer numbers.

A second type of algorithms solving PCSPs consists in having a computationally simple, but non-trivial, reduction to a CSP. The power of reductions known as *gadget reductions* is completely classified (even where the target problem is itself a PCSP). As in the case of CSPs, the existence of a gadget reduction between two PCSPs is equivalent to the existence of a certain type of map between the sets of polymorphisms of the corresponding templates. A more powerful type of reductions, called *k-reductions* or *k-consistency reductions*, has emerged recently [33, 28] and the computational power of such reductions is still unclear.

The computationally more powerful algorithms leverage algorithms for CSPs as blackboxes, mainly using the solvability of linear programming over $\mathbb{Q}$ or of linear diophantine equations over $\mathbb{Z}$. This is for example the case of CLAP [23], BLP+AIP [19], and cohomological consistency [27].

For each of the three types of algorithms or reductions proposed above, it is in general hard to characterize the power of the respective methods, and in particular it is hard to prove that a given approach does *not* solve a given PCSP, even in concrete cases. For example, considerable effort has been put recently into proving that specific polynomial-time algorithms do not solve $\mathrm{PCSP}(K_s, K_c)$ for $c \geq s > 2$ [25, 26], which is conjectured to be NP-hard and therefore should not be solvable by *any* polynomial-time algorithm if $\mathsf{P} \neq \mathsf{NP}$. Beyond the main result of this paper, the thesis we put forward here is that the tools from various fields of logic can be used to study these questions in a more general setting than Theorem 2.

We note that by allowing arbitrary logspace reductions, one can prove the tractability of *every* PCSP by a reduction to a CSP (potentially with an infinite template).

▶ **Observation 16.** *For every PCSP template* $(\mathbb{A}, \mathbb{B})$ *such that* $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ *is in* P*, there exists a structure* $\mathbb{C}$ *such that* $\mathrm{CSP}(\mathbb{C})$ *is in* P *and such that* $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ *reduces to* $\mathrm{CSP}(\mathbb{C})$ *in logspace.*

Proof. Let L be the set of instances accepted by a given polynomial-time algorithm solving $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$. By [11, Theorem 1], there exists a structure $\mathbb{C}$ such that L admits a logspace reduction to $\mathrm{CSP}(\mathbb{C})$, and $\mathrm{CSP}(\mathbb{C})$ admits a polynomial-time Turing reduction to L (and is therefore in P). Since $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ reduces to L (by a trivial reduction), we have that $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ reduces to $\mathrm{CSP}(\mathbb{C})$. ◀

We mention that while every finite-domain PCSP reduces to a problem in NP, this is not the case for infinite-domain CSPs, even for decidable ones (as there exist, e.g., NEXPTIME-complete CSPs [31]). Thus, it is probable that structural restrictions can be imposed on the infinite templates appearing in Observation 16, so that their CSPs are still able to "solve" all the finite-domain PCSPs. While the assumption that $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is in P is difficult to use, it seems that assuming that the tractability of $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ comes from specific polynomial-time algorithms *does* allow us to provide better constructions for $\mathbb{C}$, as in the cases explored in the previous sections.

The power of the sandwich approach

Provided that $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is solvable by a "natural" algorithm, we even obtain a structure $\mathbb{C}$ such that $\mathrm{CSP}(\mathbb{C})$ is in P and such that $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ reduces to $\mathrm{CSP}(\mathbb{C})$ by the do-nothing reduction, i.e., $\mathbb{A} \to \mathbb{C} \to \mathbb{B}$. We call an algorithm M *natural* if it satisfies the following conditions:

1. for every finite $\mathbb{X}$, M accepts the input $\mathbb{X}$ iff M accepts all the connected components of $\mathbb{X}$,
2. for every finite $\mathbb{X}, \mathbb{Y}$ such that M accepts $\mathbb{Y}$ and such that $\mathbb{X}$ admits a homomorphism to $\mathbb{Y}$, M accepts $\mathbb{X}$.

Let us call M *monotone* if it satisfies Item 2.[5] If $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is solvable by a monotone polynomial-time algorithm M, then there exists a polynomial-time algorithm M' satisfying Item 1, making polynomially many calls to M, and such that M' solves $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$. Indeed, on input $\mathbb{X}$, M' simply calls M on all the connected components of $\mathbb{X}$ and accepts $\mathbb{X}$ if all its connected components are accepted by M. Note that M' still solves $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ since for every $\mathbb{X}$, whether $\mathbb{X}$ homomorphically maps to $\mathbb{A}$ or $\mathbb{B}$ only depends on whether all its connected components do.

▶ **Observation 17.** *Let $(\mathbb{A}, \mathbb{B})$ be a PCSP template. The following are equivalent:*

1. *There exists a structure $\mathbb{C}$ such that $\mathbb{A} \to \mathbb{C} \to \mathbb{B}$ and such that $\mathrm{CSP}(\mathbb{C})$ is in* P*,*
2. $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ *is solvable by a monotone polynomial-time algorithm.*

Proof. The implication from Item 1 to Item 2 is immediate, as any algorithm solving $\mathrm{CSP}(\mathbb{C})$ must be natural, and $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ reduces to $\mathrm{CSP}(\mathbb{C})$ by a trivial reduction. This gives a monotone algorithm solving $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$.

Suppose now that $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ is solvable by a monotone polynomial-time algorithm M. By the argument above Observation 17, one can even assume that M is natural. Let $\mathbb{C}$ be the disjoint union of all the finite structures $\mathbb{X}$ such that M accepts $\mathbb{X}$. Since M needs to accept $\mathbb{A}$, we have $\mathbb{A} \to \mathbb{C}$. Moreover, since every accepted structure admits a homomorphism to $\mathbb{B}$, a compactness argument gives that $\mathbb{C} \to \mathbb{B}$.

Note that for every finite $\mathbb{X}$, M accepts $\mathbb{X}$ iff $\mathbb{X} \to \mathbb{C}$. Indeed, if $\mathbb{X}$ is accepted by M then it is even an induced substructure of $\mathbb{C}$. Conversely, if $\mathbb{X}$ has a homomorphism to $\mathbb{M}$, every connected component of $\mathbb{X}$ admits a homomorphism to a structure $\mathbb{Y}$ such that $\mathbb{Y}$ is accepted by M. By assumption, this means that $\mathbb{X}$ is accepted by M. Thus, $\mathrm{CSP}(\mathbb{C})$ is solved by M.

Finally, note that the trivial reduction $\mathbb{X} \mapsto \mathbb{X}$ is a reduction from $\mathrm{PCSP}(\mathbb{A}, \mathbb{B})$ to $\mathrm{CSP}(\mathbb{C})$. ◀

It was conjectured in [17] that every tractable PCSP must sandwich a tractable CSP. By Observation 17, this is equivalent to conjecturing that every polynomial-time tractable PCSP can be solved by a *monotone* algorithm running in polynomial time. Since all the known algorithms for PCSPs satisfy the condition of being natural, we see that the sandwich approach, although looking at first sight more limited than the other two mentioned approaches, is in fact currently also the most general.

8 Conclusion

We have provided in Sections 4–7 results relating the power of certain algorithms and reductions to solve PCSPs that has already been studied in the literature [8, 28, 19] providing characterizations of the applicability of a given algorithm by properties of the polymorphisms of the PCSP templates. We give here a logical take on the problem of characterizing the power of these algorithms. The general approach that we give has the advantage that it is fairly automatic: given the description of an algorithm, it is quite immediate to encode the inputs of the algorithm that are accepted as a class of relational structures, and study the potential generic objects for such a class. Remarkably, the powerful Ramsey theorem of [32] provides “out-of-the-box” strong combinatorial properties for these objects that can be potentially be used to prove tractability properties for the search variant of PCSPs, as well as decidable conditions for the applicability of a given algorithm.

[5] This can be seen as a special case of the *monotone reductions* described in [34].

The algorithms that arise as "higher levels" of algorithms described here (e.g., k-consistency as a higher level of arc consistency, kth level of the Sherali-Adams hierarchy as a "higher level" version of BLP) escape both the algebraic methods and the methods presented here. We note that [24] have algebraic characterizations of the instances that are accepted by a given algorithm, however this does not answer the question of which PCSP templates have the property that *all* their instances that are accepted by this algorithm have a solution.

We are thankful for the numerous comments by anonymous reviewers that helped improve the quality of this paper.

References

1. Kristina Asimi and Libor Barto. Finitely tractable promise constraint satisfaction problems. In Filippo Bonchi and Simon J. Puglisi, editors, *46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021, August 23-27, 2021, Tallinn, Estonia*, volume 202 of *LIPIcs*, pages 11:1–11:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.MFCS.2021.11.
2. Albert Atserias. On digraph coloring problems and treewidth duality. *Eur. J. Comb.*, 29(4):796–820, 2008. doi:10.1016/j.ejc.2007.11.004.
3. Albert Atserias, Andrei A. Bulatov, and Anuj Dawar. Affine systems of equations and counting infinitary logic. *Theor. Comput. Sci.*, 410(18):1666–1683, 2009. doi:10.1016/j.tcs.2008.12.049.
4. Albert Atserias and Víctor Dalmau. Promise constraint satisfaction and width. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9–12, 2022*, pages 1129–1153. SIAM, 2022. doi:10.1137/1.9781611977073.48.
5. Albert Atserias and Szymon Toruńczyk. Non-homogenizable classes of finite structures. In Jean-Marc Talbot and Laurent Regnier, editors, *25th EACSL Annual Conference on Computer Science Logic, CSL 2016, August 29 – September 1, 2016, Marseille, France*, volume 62 of *LIPIcs*, pages 16:1–16:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.CSL.2016.16.
6. Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\epsilon)$-Sat is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017. doi:10.1137/15M1006507.
7. Libor Barto. Promises make finite (constraint satisfaction) problems infinitary. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, pages 1–8. IEEE, 2019. doi:10.1109/LICS.2019.8785671.
8. Libor Barto, Jakub Bulin, Andrei A. Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):28:1–28:66, 2021. doi:10.1145/3457606.
9. Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *J. ACM*, 61(1):3:1–3:19, 2014. doi:10.1145/2556646.
10. Libor Barto and Marcin Kozik. Combinatorial gap theorem and reductions between promise CSPs. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9–12, 2022*, pages 1204–1220. SIAM, 2022. doi:10.1137/1.9781611977073.50.
11. Manuel Bodirsky and Martin Grohe. Non-dichotomies in constraint satisfaction complexity. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part II – Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations*, volume 5126 of *Lecture Notes in Computer Science*, pages 184–196. Springer, 2008. doi:10.1007/978-3-540-70583-3_16.
12. Manuel Bodirsky, Martin Hils, and Barnaby Martin. On the scope of the universal-algebraic approach to constraint satisfaction. In *Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11-14 July 2010, Edinburgh, United Kingdom*, pages 90–99. IEEE Computer Society, 2010. doi:10.1109/LICS.2010.13.

13 Manuel Bodirsky, Florent R. Madelaine, and Antoine Mottet. A proof of the algebraic tractability conjecture for monotone monadic SNP. *SIAM J. Comput.*, 50(4):1359–1409, 2021. doi:10.1137/19M128466X.

14 Manuel Bodirsky, Antoine Mottet, Miroslav Olsak, Jakub Opršal, Michael Pinsker, and Ross Willard. ω-categorical structures avoiding height 1 identities. *Trans. Amer. Math. Soc.*, 374:327–350, 2021.

15 Manuel Bodirsky, Antoine Mottet, Miroslav Olšák, Jakub Opršal, Michael Pinsker, and Ross Willard. Topology is relevant (in a dichotomy conjecture for infinite-domain constraint satisfaction problems). In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, pages 1–12. IEEE, 2019. doi:10.1109/LICS.2019.8785883.

16 Manuel Bodirsky, Michael Pinsker, and Todor Tsankov. Decidability of definability. *The Journal of Symbolic Logic*, 78(4):1036–1054, 2013. doi:10.2178/jsl.7804020.

17 Joshua Brakensiek and Venkatesan Guruswami. An algorithmic blend of LPs and ring equations for promise CSPs. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 436–455. SIAM, 2019. doi:10.1137/1.9781611975482.28.

18 Joshua Brakensiek and Venkatesan Guruswami. Promise constraint satisfaction: Algebraic structure and a symmetric boolean dichotomy. *SIAM J. Comput.*, 50(6):1663–1700, 2021. doi:10.1137/19M128212X.

19 Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Zivný. The power of the combined basic linear programming and affine relaxation for promise constraint satisfaction problems. *SIAM J. Comput.*, 49(6):1232–1248, 2020. doi:10.1137/20M1312745.

20 Raimundo Briceno, Andrei Bulatov, Víctor Dalmau, and Benoit Larose. Dismantlability, connectedness, and mixing in relational structures. *J. Comb. Theory, Ser. B*, 147:37–70, 2021. doi:10.1016/j.jctb.2020.10.001.

21 Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 319–330. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.37.

22 Gregory Cherlin, Saharon Shelah, and Niandong Shi. Universal graphs with forbidden subgraphs and algebraic closure. *Advances in Applied Mathematics*, 22:454–491, 1999.

23 Lorenzo Ciardo and Stanislav Živný. CLAP: A new algorithm for promise CSPs. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9–12, 2022*, pages 1057–1068. SIAM, 2022. doi:10.1137/1.9781611977073.46.

24 Lorenzo Ciardo and Stanislav Zivný. Hierarchies of minion tests for PCSPs through tensors. *CoRR*, abs/2207.02277, 2022. doi:10.48550/arXiv.2207.02277.

25 Lorenzo Ciardo and Stanislav Zivný. Approximate graph colouring and crystals. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 2256–2267. SIAM, 2023. doi:10.1137/1.9781611977554.CH86.

26 Lorenzo Ciardo and Stanislav Zivný. Approximate graph colouring and the hollow shadow. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 623–631. ACM, 2023. doi:10.1145/3564246.3585112.

27 Adam Ó Conghaile. Cohomology in constraint satisfaction and structure isomorphism. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022, August 22-26, 2022, Vienna, Austria*, volume 241 of *LIPIcs*, pages 75:1–75:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.MFCS.2022.75.

28 Víctor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems. *CoRR*, abs/2301.05084, 2023. `doi:10.48550/arXiv.2301.05084`.

29 Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1998. `doi:10.1137/S0097539794266766`.

30 Pierre Gillibert, Julius Jonusas, Michael Kompatscher, Antoine Mottet, and Michael Pinsker. Hrushovski's encoding and ω-categorical CSP monsters. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 131:1–131:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. `doi:10.4230/LIPIcs.ICALP.2020.131`.

31 Pierre Gillibert, Julius Jonusas, Michael Kompatscher, Antoine Mottet, and Michael Pinsker. When symmetries are not enough: A hierarchy of hard constraint satisfaction problems. *SIAM J. Comput.*, 51(2):175–213, 2022. `doi:10.1137/20m1383471`.

32 Jan Hubička and Jaroslav Nešetřil. All those Ramsey classes (Ramsey classes with closures and forbidden homomorphisms). *CoRR*, abs/1606.07979, 2016. `arXiv:1606.07979`.

33 Andrei A. Krokhin and Jakub Opršal. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News*, 9(3):30–59, 2022. `doi:10.1145/3559736.3559740`.

34 Andrei A. Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Zivný. Topology and adjunction in promise constraint satisfaction. *SIAM J. Comput.*, 52(1):38–79, 2023. `doi:10.1137/20M1378223`.

35 Gábor Kun, Ryan O'Donnell, Suguru Tamaki, Yuichi Yoshida, and Yuan Zhou. Linear programming, width-1 CSPs, and robust satisfaction. In Shafi Goldwasser, editor, *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 484–495. ACM, 2012. `doi:10.1145/2090236.2090274`.

36 Benoit Larose, Cynthia Loten, and Claude Tardif. A characterisation of first-order constraint satisfaction problems. *Log. Methods Comput. Sci.*, 3(4), 2007. `doi:10.2168/LMCS-3(4:6)2007`.

37 Antoine Mottet and Michael Pinsker. Smooth approximations and CSPs over finitely bounded homogeneous structures. In Christel Baier and Dana Fisman, editors, *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2–5, 2022*, pages 36:1–36:13. ACM, 2022. `doi:10.1145/3531130.3533353`.

38 Jaroslav Nešetřil and Claude Tardif. Duality theorems for finite structures (characterising gaps and good characterisations). *Journal of Combinatorial Theory, Series B*, 80(1):80–97, 2000. `doi:10.1006/jctb.2000.1970`.

39 Michael Pinsker. Current challenges in infinite-domain constraint satisfaction: Dilemmas of the infinite sheep. In *52nd IEEE International Symposium on Multiple-Valued Logic, ISMVL 2022, Dallas, TX, USA, May 18-20, 2022*, pages 80–87. IEEE, 2022. `doi:10.1109/ISMVL52857.2022.00019`.

40 Michael Pinsker and Manuel Bodirsky. Canonical functions: a proof via topological dynamics. *Contributions Discret. Math.*, 16(2):36–45, 2021. URL: `https://cdm.ucalgary.ca/article/view/71724`.

41 Benjamin Rossman. Homomorphism preservation theorems. *J. ACM*, 55(3):15:1–15:53, 2008. `doi:10.1145/1379759.1379763`.

42 Marcin Wrochna and Stanislav Živný. Improved hardness for H-colourings of G-colourable graphs. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1426–1435. SIAM, 2020. `doi:10.1137/1.9781611975994.86`.

43 Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 331–342. IEEE Computer Society, 2017. `doi:10.1109/FOCS.2017.38`.

44 Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. `doi:10.1145/3402029`.