Boolean symmetric vs. functional PCSP dichotomy

Tamio-Vesa Nakajima
University of Oxford
Oxford, UK
tamio-vesa.nakajima@cs.ox.ac.uk

Stanislav Živný
University of Oxford
Oxford, UK
standa.zivny@cs.ox.ac.uk

***Abstract*—As our first result, we establish a dichotomy for *promise constraint satisfaction problems* of the form $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$, where $\mathbf{A}$ is Boolean and symmetric and $\mathbf{B}$ is *functional* (on a domain of any size); i.e, all but one element of any tuple in a relation in $\mathbf{B}$ determine the last element. This includes PCSPs of the form $\mathrm{PCSP}(\mathbf{q\text{-}in\text{-}r},\mathbf{B})$, where $\mathbf{B}$ is functional, thus making progress towards a classification of $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3},\mathbf{B})$, which were studied by Barto, Battistelli, and Berg [STACS'21] for $\mathbf{B}$ on three-element domains.**

As our second result, we show that for $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$, where $\mathbf{A}$ contains a single symmetric relation and $\mathbf{B}$ is arbitrary (and thus not necessarily functional), the combined *basic linear programming* relaxation (BLP) and the *affine integer programming* relaxation (AIP) of Brakensiek et al. [SICOMP'20] is no more powerful than the (in general strictly weaker) AIP relaxation of Brakensiek and Guruswami [SICOMP'21].

***Index Terms*—algebraic approach, dichotomy, constraint satisfaction, promise CSP, polymorphisms, minions**

I. Introduction

Promise constraint satisfaction problems (PCSPs) are a generalisation of constraint satisfaction problems (CSPs) that allow for capturing many more computational problems [4], [8], [6].

A canonical example of a CSP is the 3-colouring problem: Given a graph $\mathbf{G}$, is it 3-colourable? This can be cast as a CSP. Let $\mathbf{K}_k$ denote a clique on k vertices. Then $\mathrm{CSP}(\mathbf{K}_3)$, the constraint satisfaction problem with the template $\mathbf{K}_3$, is the following computational problem (in the decision version): Given a graph $\mathbf{G}$, say YES if there is a homomorphism from $\mathbf{G}$ to $\mathbf{K}_3$ (indicated by $\mathbf{G}\to\mathbf{K}_3$) and say NO otherwise (indicated by $\mathbf{G}\not\to\mathbf{K}_3$). Here a graph homomorphism is an edge preserving map [23]. As graph homomorphisms from $\mathbf{G}$ to $\mathbf{K}_3$ are 3-colourings of $\mathbf{G}$, $\mathrm{CSP}(\mathbf{K}_3)$ is the 3-colouring problem.

Another example of a CSP is 1-in-3-SAT: Given a positive 3-CNF formula, is there an assignment that satisfies exactly one literal in each clause? This is $\mathrm{CSP}(\mathbf{1\text{-}in\text{-}3})$, where

$$\mathbf{1\text{-}in\text{-}3} = (\{0,1\}; \{(1,0,0),(0,1,0),(0,0,1)\}.$$

This work was supported by UKRI EP/X024431/1 and a Clarendon Fund Scholarship. All data is provided in full in the results section of this paper. The full version of the paper is available at arXiv:2210.03343.

979-8-3503-3587-3/23/$31.00 ©2023 IEEE

Yet another example is NAE-3-SAT: Given a positive 3-CNF formula, is there an assignment that satisfies one or two literals in each clause? This is $\mathrm{CSP}(\mathbf{NAE})$, where

$$\mathbf{NAE} = \{(0,1); \{0,1\}^3 \setminus \{(0,0,0),(1,1,1)\}).$$

A canonical example of a PCSP is the approximate graph colouring problem [21]: Fix $k\le\ell$. Given a graph $\mathbf{G}$, determine whether $\mathbf{G}$ is k-colourable or not even ℓ-colourable. (The case $k=\ell$ is the same as k-colouring.) This is the same as the PCSP over cliques; i.e., $\mathrm{PCSP}(\mathbf{K}_k,\mathbf{K}_\ell)$ is the following computational problem (in the decision version): Given a graph $\mathbf{G}$, say YES if $\mathbf{G}\to\mathbf{K}_k$ and say NO if $\mathbf{G}\not\to\mathbf{K}_\ell$. In the search version, one is given a k-colourable graph $\mathbf{G}$ and the task is to find an ℓ-colouring of $\mathbf{G}$ (which necessarily exists by the promise and the fact that $k\le\ell$).

Another example of a PCSP is $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3},\mathbf{NAE})$, identified in [8]: Given a satisfiable instance $\mathbf{X}$ of $\mathrm{CSP}(\mathbf{1\text{-}in\text{-}3})$, can one find a solution if $\mathbf{X}$ is seen as an instance of $\mathrm{CSP}(\mathbf{NAE})$? I.e., can one find a solution that satisfies one or two literals in each clause given the promise that a solution that satisfies exactly one literal in each clause exists? Although both $\mathrm{CSP}(\mathbf{1\text{-}in\text{-}3})$ and $\mathrm{CSP}(\mathbf{NAE})$ are NP-complete, Brakensiek and Guruswami showed in [8] that $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3},\mathbf{NAE})$ is solvable in polynomial time and in particular it is solved by the so-called affine integer programming relaxation (AIP), whose power was characterised in [6].

More generally, one fixes two relational structures $\mathbf{A}$ and $\mathbf{B}$ with $\mathbf{A}\to\mathbf{B}$. The $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$ is then, in the decision version, the computational problem of distinguishing (input) relational structures $\mathbf{X}$ with $\mathbf{X}\to\mathbf{A}$ from those with $\mathbf{X}\not\to\mathbf{B}$. In the search version, $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$ is the problem of finding a homomorphism from an input structure $\mathbf{X}$ to $\mathbf{B}$ given that one is promised that $\mathbf{X}\to\mathbf{A}$. One can think of $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$ as an approximation version of $\mathrm{CSP}(\mathbf{A})$ on satisfiable instances. Another way is to think of $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$ as $\mathrm{CSP}(\mathbf{B})$ with restricted inputs. We refer the reader to [25] for a very recent survey on PCSPs.

For CSPs, a dichotomy conjecture of Feder and Vardi [19] was resolved independently by Bulatov [15] and Zhuk [30] via the so-called algebraic approach [24], [14]: For every fixed finite $\mathbf{A}$, $\mathrm{CSP}(\mathbf{A})$ is either solvable in polynomial time or $\mathrm{CSP}(\mathbf{A})$ is NP-complete.

For PCSPs, even the case of graphs and structures on Boolean domains is widely open; these two were established for CSPs a long time ago [23], [28] and constituted important evidence for conjecturing a dichotomy. Following the important work of Barto et al. [6] on extending the algebraic framework from the realms of CSPs to the world of PCSPs, there have been several recent works on complexity classifications of fragments of PCSPs [20], [22], [8], [12], [5], [2], [9], [13], [27], [26], hardness conditions [6], [12], [7], [29], and power of algorithms [6], [10], [3], [17]. Nevertheless, a classification of more concrete fragments of PCSPs is needed for making progress with the general theory, such as finding hardness and tractability criteria, as well as with resolving longstanding open questions, such as approximate graph colouring.

Brakensiek and Guruswami classified $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ for all Boolean symmetric structures $\mathbf{A}$ and $\mathbf{B}$ with disequalities [8]. Ficak, Kozik, Olšák, and Stankiewicz generalised this result by classifying $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ for all Boolean symmetric structures $\mathbf{A}$ and $\mathbf{B}$ [20].

Barto, Battistelli, and Berg [5] studied symmetric PCSPs on non-Boolean domains and in particular PCSPs of the form $\operatorname{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$, where $\mathbf{B}$ contains a single ternary relation over the domain $\{0, 1, \dots, d-1\}$. For $d = 2$, a complete classification $\operatorname{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ is known [8], [20]. For $d = 3$, Barto et al. [5] managed to classify all but one structure $\mathbf{B}$. The remaining open case of "linearly ordered colouring" inspired further investigation in [27]. For $d = 4$, Barto et al. [5] obtained partial results. In particular, for certain structures $\mathbf{B}$ they managed to rule out the applicability of the $\mathrm{BLP} + \mathrm{AIP}$ algorithm from [10]. The significance of $\mathrm{BLP} + \mathrm{AIP}$ here is that it is the strongest known algorithm for PCSPs for which a characterisation of its power is known both in terms of a minion and also in terms of polymorphism identities. This suggests that those cases are NP-hard (or new algorithmic techniques are needed).

a) Contributions: We continue the work from [8], [20] and [5] and focus on promise constraint satisfaction problems of the form $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$, where $\mathbf{A}$ is symmetric, and $\mathbf{B}$ is over an arbitrary finite[1] domain. Our main motivation is the fact that more complexity classifications of PCSP fragments are needed to make progress with the general theory of PCSPs. Since the template $\mathbf{A}$ is symmetric, we can assume without loss of generality that $\mathbf{B}$ is symmetric, as observed in [5] and in [13].[2]

As our main result, we establish the following result. A structure $\mathbf{B}$ is called *functional* if, for any relation $R^{\mathbf{B}}$ in $\mathbf{B}$ of, say, arity r, and any tuple $\mathbf{x}$ in $R^{\mathbf{B}}$, any $r-1$ elements of $\mathbf{x}$ determine the last element. In detail, $(x_1, \dots, x_{r-1}, y), (x_1, \dots, x_{r-1}, z) \in R^{\mathbf{B}}$ implies $y = z$, and similarly for the other $r-1$ positions.[3] We say that a structure $\mathbf{A}$ has a *covering tuple* if and only if some tuple $(s_1, \dots, s_r) \in R^{\mathbf{A}}$ exists for some relation $R^{\mathbf{A}}$ of $\mathbf{A}$ such that A, the domain of $\mathbf{A}$, is equal to $\{s_1, \dots, s_r\}$. Finite tractability is defined in Section II. A tuple of arity different to 2 is called *non-binary*.

Theorem 1. *Let $\mathbf{A}$ be a symmetric structure with a non-binary covering tuple, and $\mathbf{B}$ a functional structure such that $\mathbf{A} \to \mathbf{B}$. Then, either $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ is solvable in polynomial time by AIP and is finitely tractable, or $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ is NP-hard.*

Theorem 1 implies the following result, advertised in the abstract.

Corollary 2. *Let $\mathbf{A}$ be a Boolean symmetric structure, and $\mathbf{B}$ a functional structure such that $\mathbf{A} \to \mathbf{B}$. Then, either $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ is solvable in polynomial time by AIP and is finitely tractable, or $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ is NP-hard.*

An example of a Boolean symmetric structure $\mathbf{A}$ is $\mathbf{1\text{-}in\text{-}3}$, and more generally $\mathbf{q\text{-}in\text{-}r}$.[4]

Several researchers have informally conjectured that PCSPs of the form $\operatorname{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ admit a dichotomy. The authors, as well as other researchers, believe that in fact not only is there a dichotomy but also all tractable cases are solved by AIP.

Conjecture 3. *For every structure $\mathbf{B}$, either $\operatorname{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ is solvable in polynomial time by AIP, or $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ is NP-hard.*

Theorem 1 establishes the special case of Conjecture 3 for functional $\mathbf{B}$. We make further progress towards Conjecture 3 by proving that for any structure $\mathbf{A}$ with a single *(not necessarily Boolean)* symmetric relation, and any *(not necessarily functional)* structure $\mathbf{B}$ for which $\mathbf{A} \to \mathbf{B}$, $\mathrm{BLP} + \mathrm{AIP}$ from [10] is no more powerful for $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ than AIP from [8], although in general $\mathrm{BLP} + \mathrm{AIP}$ is strictly stronger than AIP [10], already for (non-promise) CSPs with *two* Boolean symmetric relations, cf. Remark 33. In fact, we establish a more general result. We say that a relation R is *balanced* if there exists a matrix M whose columns are tuples of R, where each tuple of R appears as a column (possibly a multiple times), and where the rows of M are permutations of each other. The matrix M below shows that the Boolean $\mathbf{1\text{-}in\text{-}3}$ relation is balanced:

$$M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Theorem 4. *Let $\mathbf{A}$ be any structure with a single relation. If the relation in $\mathbf{A}$ is balanced then, for any $\mathbf{B}$ such that $\mathbf{A} \to \mathbf{B}$, $\mathrm{BLP} + \mathrm{AIP}$ solves $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ if and only if AIP solves it.*

If the (only) relation in $\mathbf{A}$ is *binary* (i.e., a digraph), the condition of balancedness has a natural combinatorial

[1] All structures in this article can be assumed to be finite unless they are explicitly stated to be infinite.

[2] In detail, for any symmetric $\mathbf{A}$ and (not necessarily symmetric) $\mathbf{B}$ with $\mathbf{A} \to \mathbf{B}$, there is a symmetric $\mathbf{B}'$ with $\mathbf{A} \to \mathbf{B}'$ such that $\operatorname{PCSP}(\mathbf{A}, \mathbf{B})$ and $\operatorname{PCSP}(\mathbf{A}, \mathbf{B}')$ are polynomial-time equivalent [5], [13]. This $\mathbf{B}'$ is the largest symmetric substructure of $\mathbf{B}$. Observe that a functional structure has functional substructures, so if $\mathbf{B}$ is functional then $\mathbf{B}'$ remains functional.

[3] Note that for symmetric $\mathbf{B}$ the requirement "for the other $r-1$ positions" is satisfied automatically.

[4] $\mathbf{q\text{-}in\text{-}r}$ is the structure on $\{0, 1\}$ with a single (symmetric) relation of arity r containing all r-tuples with precisely q 1s (and $r - q$ 0s).

interpretation: A binary relation is balanced if and only if it is the disjoint union of strongly connected components.

Theorem 4 implies the following corollary. We say that a relation of arity r is *preserved* by a group of permutations of degree r if and only if permuting any tuple of the relation according to any permutation of the group gives another tuple of the relation.

Corollary 5. *Suppose that G is a transitive group of permutations. Further, suppose that $\mathbf{A}$ is a relational structure with one relation that is preserved by G. Then, for any $\mathbf{A} \to \mathbf{B}$, $\mathrm{BLP} + \mathrm{AIP}$ solves $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ if and only if AIP does.*

While Corollary 5 is more elegant than Theorem 4, it applies to fewer structures. Indeed, we will show in Remark 36 that there exist balanced relations that are not preserved by any transitive group. Examples of relations that are preserved by some transitive group of permutations G include symmetric relations (where G is the symmetric group) or cyclic relations (where G contains all cyclic shifts of appropriate degree).

II. Preliminaries

We let $[r] = \{1, \ldots, r\}$. We denote by 2^S the powerset of S.

a) Structures and PCSPs: Promise CSPs have been introduced in [4] and [8]. We follow the notation and terminology of [6].

A (relational) *structure* is a tuple $\mathbf{A} = (A; R_1^{\mathbf{A}}, \ldots, R_t^{\mathbf{A}})$, where $R_i^{\mathbf{A}} \subseteq A^{\mathrm{ar}(R_i)}$ is a relation of arity $\mathrm{ar}(R_i)$ on a set A, called the *domain*. A structure $\mathbf{A}$ is called *Boolean* if $A = \{0, 1\}$ and is called *symmetric* if $R_i^{\mathbf{A}}$ is a symmetric relation for each $i \in [t]$; i.e, if $(x_1, \ldots, x_{\mathrm{ar}(R_i)}) \in R_i^{\mathbf{A}}$ then for every permutation π on $[\mathrm{ar}(R_i)]$ we have $(x_{\pi(1)}, \ldots, x_{\pi(\mathrm{ar}(R_i))}) \in R_i^{\mathbf{A}}$. A structure $\mathbf{A}$ is called *functional* if

$$(x_1, \ldots, x_{\mathrm{ar}(R_i)-1}, y) \in R_i^{\mathbf{A}}$$

and

$$(x_1, \ldots, x_{\mathrm{ar}(R_i)-1}, z) \in R_i^{\mathbf{A}}$$

implies $y = z$ for any $i \in [t]$, and that the same hold for all other $r - 1$ positions in the tuple.

For two structures $\mathbf{A} = (A; R_1^{\mathbf{A}}, \ldots, R_t^{\mathbf{A}})$ and $\mathbf{B} = (B; R_1^{\mathbf{B}}, \ldots, R_t^{\mathbf{B}})$ with t relations with the same arities, a *homomorphism from* $\mathbf{A}$ *to* $\mathbf{B}$ is a function $h : A \to B$ such that, for any $i \in [t]$, for each $x = (x_1, \ldots, x_{\mathrm{ar}(R_i)}) \in R_i^{\mathbf{A}}$, we have $h(x) = (h(x_1), \ldots, h(x_{\mathrm{ar}(R_i)})) \in R_i^{\mathbf{B}}$. We denote the existence of a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ by $\mathbf{A} \to \mathbf{B}$.

Let $\mathbf{A}$ and $\mathbf{B}$ be two structures with $\mathbf{A} \to \mathbf{B}$; we call $(\mathbf{A}, \mathbf{B})$ a (PCSP) *template*. In the *search* version of the *promise constraint satisfaction problem* (PCSP) with the template $(\mathbf{A}, \mathbf{B})$, denoted by $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$, the task is: Given a structure $\mathbf{X}$ with the promise that $\mathbf{X} \to \mathbf{A}$, find a homomorphism from $\mathbf{X}$ to $\mathbf{B}$ (which necessarily exists as homomorphisms compose). In the *decision* version of $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$, the task is: Given a structure $\mathbf{X}$, output YES if $\mathbf{X} \to \mathbf{A}$, and output NO if $\mathbf{X} \not\to \mathbf{B}$. The decision version trivially reduces to the search version. We will use the decision version in this paper.

We will be interested in the complexity of $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$, in particular for symmetric $\mathbf{A}$, and functional $\mathbf{B}$. (As discussed in Section I, the symmetricity of $\mathbf{A}$ means that we can without loss of generality assume symmetricity of $\mathbf{B}$.)

b) Operations and polymorphisms: A function $h : A^n \to B$ is called an *operation* of arity n. A $(2n+1)$-ary operation $f : A^{2n+1} \to B$ is called *2-block-symmetric* if $f(a_1, \ldots, a_{2n+1}) = f(a_{\pi(1)}, \ldots, a_{\pi(2n+1)})$ for every $a_1, \ldots, a_{2n+1} \in A$ and every permutation π on $[2n+1]$ that preserves parity; i.e, π maps odd values to odd values and even values to even values.

A $(2n+1)$-ary operation $f : A^{2n+1} \to B$ is called *alternating* if it is 2-block-symmetric, and furthermore

$$f(a_1, \ldots, a_{2n-1}, a, a) = f(a_1, \ldots, a_{2n-1}, a', a')$$

for every

$$a_1, \ldots, a_{2n-1}, a, a' \in A.$$

Consider structures $\mathbf{A}, \mathbf{B}$ with t relations with the same arities. We call $h : A^n \to B$ a *polymorphism* of $(\mathbf{A}, \mathbf{B})$ if the following holds for any relation $R = R_i$, $i \in [t]$, of arity $r = \mathrm{ar}(R)$. For any $x^1, \ldots, x^r \in A^n$, where $x^i = (x_1^i, \ldots, x_n^i)$, with $(x_i^1, \ldots, x_i^r) \in R^{\mathbf{A}}$ for every $1 \le i \le n$, we have $(h(x^1), \ldots, h(x^r)) \in R^{\mathbf{B}}$. One can visualise this as an $(r \times n)$ matrix whose rows are the tuples $x^1, \ldots, x^r$. The requirement is that if every column of the matrix is in $R^{\mathbf{A}}$ then the application of h on the rows of the matrix results in a tuple from $R^{\mathbf{B}}$. We denote by $\mathrm{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$ the set of n-ary polymorphisms of $(\mathbf{A}, \mathbf{B})$ and by $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ the set of all polymorphisms of $(\mathbf{A}, \mathbf{B})$.

c) Relaxations: There are two standard polynomial-time solvable relaxations for PCSPs, the *basic linear programming* relaxation (BLP) and the *affine integer programming* relaxation (AIP) [8]. There is also a combination of the two, called $\mathrm{BLP} + \mathrm{AIP}$ [10], that is provably stronger than both BLP and AIP. We will show that for certain PCSPs, this is not the case (cf. Theorem 4). The precise definitions of the relaxations are not important for us as we will only need the notion of solvability of PCSPs by these relaxations and characterisations of the power of the relaxations; we refer the reader to [8], [6], [10] for details. Let $\mathbf{X}$ be an instance of $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$. It follows from the definitions of the relaxations that if $\mathbf{X} \to \mathbf{A}$ then both AIP and $\mathrm{BLP} + \mathrm{AIP}$ accept [8], [6]. We say that AIP ($\mathrm{BLP} + \mathrm{AIP}$, respectively) *solves* $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ if, for every $\mathbf{X}$ with $\mathbf{X} \not\to \mathbf{B}$, AIP ($\mathrm{BLP} + \mathrm{AIP}$, respectively) rejects.

The power of AIP and $\mathrm{BLP} + \mathrm{AIP}$ for PCSPs is characterised by the following results.

Theorem 6 ([6])**.** *$\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is solved by AIP if and only if $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ contains alternating operations of all odd arities.*

Theorem 7 ([10])**.** *$\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is solved $\mathrm{BLP} + \mathrm{AIP}$ if and only if $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ contains 2-block-symmetric operations of all odd arities.*

In particular, this immediately implies the following result. For this result, we first define *finite tractability* [6], [2]. We say that $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is finitely tractable if and only if $\mathbf{A} \to$

$\mathbf{E} \to \mathbf{B}$ for some finite structure $\mathbf{E}$ and $\mathrm{CSP}(\mathbf{E})$ is tractable. For a group G, we use the standard notation $H \lhd G$ to indicate that H is a normal subgroup of G.

Lemma 8. *Suppose $\mathbf{A} \to \mathbf{E} \to \mathbf{B}$, where $E = G$ for some Abelian group $(G, +)$, and each relation of $\mathbf{E}$ is either of the form (i) $c + H$ for some $r \in \mathbb{N}$, $c \in G^r$ and $H \lhd G^r$, or (ii) empty. Then $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is finitely tractable and solved by* AIP.

Proof. The following alternating operation is a polymorphism of $\mathbf{E}$

$$f(x_1, y_1, \ldots, y_k, x_{k+1}) = \sum_{i=1}^{k+1} x_i - \sum_{i=1}^{k} y_i.$$

Consider a relation $R^{\mathbf{E}}$ of $\mathbf{E}$, of the form $c + H$. Consider a matrix of inputs whose columns are $x_1, y_1, \ldots, y_k, x_{k+1} \in R^{\mathbf{E}}$. In other words, $x_i \in c + H$ and $y_i \in c + H$ for each x_i, y_i. Note that the column that results from applying f to the rows of this matrix is just

$$\begin{aligned} x_1 - y_1 + \cdots - y_k + x_{k+1} \in (c+H) - (c+H) + \cdots \\ - (c+H) + (c+H) \subseteq c + H \end{aligned}$$

Thus f is an alternating polymorphism of $\mathbf{E}$. It follows that $\mathrm{CSP}(\mathbf{E})$ is solved by AIP, from whence it follows that $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is finitely tractable and solved by AIP. ☐

d) Minions and hardness: We will use the theory of minions from [6].

Let $\mathcal{M}$ be a set, where each element $f \in \mathcal{M}$ is assigned an arity $\mathrm{ar}(f)$. We write $\mathcal{M}^{(n)} = \{f \in \mathcal{M} \mid \mathrm{ar}(f) = n\}$. Further, let $\mathcal{M}$ be endowed with, for each $\pi : [n] \to [m]$, a (so-called minor) map $f \mapsto f^\pi : \mathcal{M}^{(n)} \to \mathcal{M}^{(m)}$ such that, for $\pi : [m] \to [k]$ and $\sigma : [n] \to [m]$, and any $f \in \mathcal{M}^{(n)}$ we have $(f^\pi)^\sigma = f^{\sigma \circ \pi}$, and $f^{\mathrm{id}} = f$. Then, $\mathcal{M}$ is called a *minion*.[5] We often write $f \xrightarrow{\pi} g$ instead of $g = f^\pi$.

Consider two minions $\mathcal{M}, \mathcal{N}$; a *minion homomorphism* is a map $\xi : \mathcal{M} \to \mathcal{N}$ such that, for any $f \in \mathcal{M}^{(n)}$ and $\pi : [n] \to [m]$, we have that $\xi(f)^\pi = \xi(f^\pi)$.[6] If such a minion homomorphism exists, we write $\mathcal{M} \to \mathcal{N}$.

Given an n-ary operation $f : A^n \to B$ and a map $\pi : [n] \to [m]$, an m-ary operation $g : A^m \to B$ is called a *minor* of f given by the map π if

$$g(x_1, \ldots, x_m) = f(x_{\pi(1)}, \ldots, x_{\pi(n)}).$$

The polymorphisms $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ thus form a minion, where f^π is given by the minor of f at π.

The main hardness theorem that we will use is the following.

Theorem 9 ([6])**.** *Fix constants m and C. Take any template $(\mathbf{A}, \mathbf{B})$ such that $\mathrm{Pol}(\mathbf{A}, \mathbf{B}) = \bigcup_{i=1}^m \mathcal{M}_i$ is the union of m parts. Suppose that for each $i \in [m]$ there exists a map I_i that takes $f \in \mathcal{M}_i$ to a subset of $[\mathrm{ar}(f)]$ of size at most C such that the following holds: for any $f, g \in \mathcal{M}_i$ such that $g = f^\pi$ we have that $I_i(g) \cap \pi(I_i(f)) \neq \emptyset$. Then, $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is* NP*-hard.*

[5]We can see a minion as a functor from the skeleton of the category of finite sets to the category of sets.

[6]Minions are just functors and minion homomorphisms are just natural transformations.

III. Dichotomy

In this section we will show a dichotomy for $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ where $\mathbf{A}, \mathbf{B}$ are functional and symmetric, and $\mathbf{A}$ has a non-binary covering tuple.

Theorem 1. *Let $\mathbf{A}$ be a symmetric structure with a non-binary covering tuple, and $\mathbf{B}$ a functional structure such that $\mathbf{A} \to \mathbf{B}$. Then, either $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is solvable in polynomial time by* AIP *and is finitely tractable, or $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is* NP*-hard.*

Before proving this theorem, we will show how this immediately implies a stronger dichotomy when $\mathbf{A}$ is Boolean.

Corollary 2. *Let $\mathbf{A}$ be a Boolean symmetric structure, and $\mathbf{B}$ a functional structure such that $\mathbf{A} \to \mathbf{B}$. Then, either $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is solvable in polynomial time by* AIP *and is finitely tractable, or $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is* NP*-hard.*

Lemma 10. *Suppose $\mathbf{A}$ is Boolean, $\mathbf{B}$ is functional, and every relation of $\mathbf{A}$ is either binary or contains only constant tuples. Then $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is solved by* AIP *and is finitely tractable.*

Proof. Consider any $h : \mathbf{A} \to \mathbf{B}$. Suppose $h(0) = h(1)$. Then every relation in $\mathbf{B}$ contains a constant tuple of the form $(h(0), \ldots, h(0))$; in this case, $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is trivially solved by AIP and is finitely tractable. Thus suppose $h(0) \neq h(1)$. Any empty relation in $\mathbf{A}$ can be removed (together with the corresponding relation in $\mathbf{B}$) as it does not affect $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ and the complexity of $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$.

Since $\mathbf{B}$ is functional, the binary relations of $\mathbf{A}$ that do not contain only constant tuples must be the binary disequality. To see why, consider any relation $R^{\mathbf{A}}$ in $\mathbf{A}$ that contains the tuple $(0,1)$. $R^{\mathbf{A}}$ cannot contain $(0,0)$ or $(1,1)$, since the corresponding relation $R^{\mathbf{B}}$ in $\mathbf{B}$ contains $(h(0), h(1))$ and if it contained $(h(0), h(0))$ or $(h(1), h(1))$ it would not be functional. Since $\mathbf{A}$ is symmetric, $R^{\mathbf{A}}$ also contains the tuple $(1, 0)$. Thus $R^{\mathbf{A}} = \{(0,1), (1,0)\}$ is the disequality relation. It follows that every relation in $\mathbf{A}$ is either a binary disequality, or consists only of constant tuples. In this case, $\mathrm{CSP}(\mathbf{A})$ is solved by AIP and thus $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is solved by AIP and is finitely tractable. ☐

Proof of Corollary 2. If $\mathbf{A}$ has a covering non-binary tuple then we can apply Theorem 1 to get the desired conclusion. Otherwise, $\mathbf{A}$ contains only binary relations or relations that contain only constant tuples (since any non-constant tuple would be a covering tuple). In this case we can apply Lemma 10. ☐

We will now move on to a proof of Theorem 1.

Let $(s_1, \ldots, s_r) \in R^{\mathbf{A}}$ be the covering tuple of $\mathbf{A}$. Recall that $r \neq 2$ by assumption. If $r \leq 1$ then AIP immediately solves $\mathrm{CSP}(\mathbf{A})$ and thus $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$, so assume $r \geq 3$. Recall that we can assume without loss of generality that $\mathbf{B}$ is symmetric, since $\mathbf{A}$ is [5], [13]. Assume also that $\mathbf{A}$ and

$\mathbf{B}$ have domains of size greater than 1 (or else $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$ is obviously finitely tractable and solvable by AIP). We start with some important definitions.

Suppose throughout, without loss of generality, that $A = \{s_1,\ldots,s_r\} = [a]$. Note that a function $f : [a]^n \to B$ can be seen as a function from an ordered partition of $[n]$ into a sets to B: the first subset is the set of coordinates in the input set to 1 and so on. We let $a^{[n]}$ denote the set of ordered partitions of $[n]$ into a sets. We thus view any polymorphism $f \in \mathrm{Pol}^{(n)}(\mathbf{A},\mathbf{B})$ as a function $f : a^{[n]} \to B$ i.e. for a partition $S_1,\ldots,S_a \subseteq [n]$ we can evaluate $f(S_1,\ldots,S_a)$. We will use $\overrightarrow{S}$ to denote the partition $(S_1,\ldots,S_a) \in a^{[n]}$.

Definition 11. Fix some polymorphism $f : \mathrm{Pol}^{(n)}(\mathbf{A},\mathbf{B})$. For $i,j \in A$, let $f_{ij} : 2^{[n]} \to B$ be a function given by

$$f_{ij}(S) = f(x_1,\ldots,x_n),$$

where $x_k = j$ if $k \in S$ and $x_k = i$ otherwise.

Let $f^p : 2^{[n]} \to B^{a\times a}$ be a function given by

$$(f^p(S))_{ij} = f_{ij}(S).$$

(The p stands for "pairs".) Let $f^* : a^{[n]} \to (B^{a\times a})^a$ be a function given by

$$f^*(S_1,\ldots,S_a) = (f^p(S_1),\ldots,f^p(S_a)).$$

Definition 12. Consider a polymorphism $f \in \mathrm{Pol}^{(n)}(\mathbf{A},\mathbf{B})$. We call it *$k$-degenerate* if there exist $x_1,\ldots,x_k \in \mathrm{range}(f^p)$ such that for any $S_1,\ldots,S_k \subseteq [n]$ for which $f^p(S_i) = x_i$ we have that not all S_i are disjoint. Note that no polymorphism can be 1-degenerate as a single set is a disjoint family.

For any polymorphism $f \in \mathrm{Pol}^{(n)}(\mathbf{A},\mathbf{B})$, we call a set $S \subseteq [n]$ a *hard set* if, for any $T \supseteq S$, we have $f^p(T) \neq f^p(\emptyset)$.[7]

We will prove Theorem 1 using the following two cases. For the following, define $N_d = 1 + |B|^{a^2} a^{2r_{\max}}$ and $N_h = |B|^{a^2}$, where $r_{\max}$ is the maximum arity of any relation in $\mathbf{A}$.

Theorem 13. *If $\mathrm{Pol}(\mathbf{A},\mathbf{B})$ contains a polymorphism that is not k-degenerate, for any $k \leq N_d$, and that has no hard sets of size at most N_h, then $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$ is solved by AIP and is finitely tractable.*

Theorem 14. *If every polymorphism within $\mathrm{Pol}(\mathbf{A},\mathbf{B})$ is k-degenerate for some $k \leq N_d$, or has a hard set of size at most N_h, then $\mathrm{PCSP}(\mathbf{A},\mathbf{B})$ is NP-hard.*

These two theorems will be proved in their own sections later. We will first prove some results of common interest to both of them.

Proof of Theorem 1. A result of Theorem 13 and Theorem 14. □

[7]These two notions are similar to those of *unbounded antichains* and *fixing sets* in [20]. The notion of hard-set is similar to the notion of an $f(\emptyset)$-avoiding set from [18], [29].

A. Common results

Lemma 15. *There exists a partial operation $+$ on $|B|^{a\times a}$ such that, for any $f \in \mathrm{Pol}^{(n)}(\mathbf{A},\mathbf{B})$, for disjoint $S,T \subseteq [n]$, we have that $f^p(S\cup T) = f^p(S) + f^p(T)$. In particular, $f^p(S) + f^p(T)$ must always be defined for such f,S,T.*

Proof. To show this, we show that the value of $f^p(S\cup T)$ is uniquely determined by $f^p(S)$ and $f^p(T)$ i.e. the values $f_{ij}(S)$, $f_{ij}(T)$ for any $i,j \in A$ determine the values $f_{i'j'}(S\cup T)$ for any $i',j' \in A$. Since $f_{ii}(S) = f_{ii}(T) = f_{ii}(S\cup T)$, it only remains to show this for $i \neq j$. In particular, we show that $f_{12}(S\cup T)$ satisfies this property. Without loss of generality, suppose $s_1 = 1, s_2 = 2$: we can do this since $\{s_1,\ldots,s_r\} = A$.

We will first consider the case where $S = \{1\}, T = \{2\}, n = 3, r = 3$. We will show at the end of the proof that this particular example generalises completely.

We show how to deduce $f(2,2,1) = f_{12}(S\cup T)$ from $f^p(S), f^p(T)$. First suppose that $s_3 = 1$. Then, consider the following two matrices of inputs:

$$\begin{pmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix} \qquad \begin{pmatrix} 2 & 2 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 2 \end{pmatrix}$$

Apply f to the rows of these matrices. The first two rows of the first matrix are $f_{12}(S)$ and $f_{12}(T)$; by the functionality of $\mathbf{B}$ and since f is a polymorphism, we can determine $f(1,1,2)$. Now, in the second matrix, the image through f of the second row is $f_{11}(S)$, and the image of the last row is the already determined $f(1,1,2)$. By functionality again these uniquely determine $f(2,2,1) = f_{12}(S\cup T)$. Thus we have our result.

Consider now if $s_3 = 2$. Then, symmetrically to before, we see that $f_{21}(S\cup T)$ is a function of $f^p(S)$, $f^p(T)$. Now, consider the matrix of values

$$\begin{pmatrix} 2 & 2 & 1 \\ 1 & 1 & 2 \\ 2 & 2 & 2 \end{pmatrix}$$

Since the image of the second and third row through f are a function of $f^p(S)$, $f^p(T)$, it follows, by functionality of $\mathbf{B}$, that the first row i.e. $f(2,2,1) = f_{12}(S\cup T)$ is as well.

Finally, suppose $s_3 \notin \{1,2\}$; without loss of generality, let $s_3 = 3$. Then consider the following two matrices of inputs

$$\begin{pmatrix} 3 & 1 & 2 \\ 2 & 3 & 3 \\ 1 & 2 & 1 \end{pmatrix} \qquad \begin{pmatrix} 3 & 1 & 2 \\ 1 & 3 & 3 \\ 2 & 2 & 1 \end{pmatrix}$$

Similarly to before, we see that $f(3,1,2)$ is uniquely determined by $f(2,3,3) = f_{32}(S)$ and $f(1,2,1) = f_{12}(T)$, and $f(3,1,2)$ together with $f(1,3,3) = f_{31}(S)$ uniquely determines $f(2,2,1) = f_{12}(S\cup T)$.

Now, note that this proof generalises to any values of n,S,T and r. To see why, note that if r is changed, we merely pad the matrix with rows containing constant tuples equal to the new values of s_i. Thus $f^p(S\cup T)$ remains a function of $f^p(S)$, $f^p(T)$ (recalling that the value of f applied to any constant tuple appears in $f^p(S)$ and $f^p(T)$, as $f(i,\ldots,i) =$

$f_{ii}(S) = f_{ii}(T)$). Furthermore, if S or T are changed, all that is necessary is to duplicate and permute the columns in the matrices presented above to fit the new values of S and T (and $[n]\setminus(S\cup T)$; crucially, our example above has $[n]\setminus(S\cup T) \neq \emptyset$). In other words, if π is a function that takes S to 1, T to 2 and $[n] \setminus (S \cup T)$ to 3, we apply the reasoning from above to f^π. □

For the next two results, let $S^c = [n] \setminus S$ i.e. S^c is the complement of S. Beware that the superscript T below denotes the matrix transpose.

Lemma 16. *For any n-ary polymorphism f and $S \subseteq [n]$, we have $f^p(S)^T = f^p(S^c)$.*

Proof. $(f^p(S)^T)_{ij} = f^p(S)_{ji} = f_{ji}(S) = f_{ij}(S^c) = f^p(S^c)_{ij}$. □

Lemma 17. *There exists a partial operation $-$ on $|B|^{a\times a}$ such that for any n-ary polymorphism f, for $S \subseteq T \subseteq [n]$, we have that $f^p(T \setminus S) = f^p(T) - f^p(S)$.*

Proof. Define $x - y = (x^T + y)^T$. For $S \subseteq T \subseteq [n]$, since $T^c \cap S = \emptyset$, we have

$$f^p(T) - f^p(S) = (f^p(T)^T + f^p(S))^T = f^p((T^c \cup S)^c) = f^p(T \setminus S).$$

Thus $f^p(T) - f^p(S) = f^p(T \setminus S)$ as required. □

Lemma 18. *Fix a polymorphism $f \in \mathrm{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$. Consider any family of disjoint sets $\mathcal{A} \subseteq 2^{[n]}$, containing at least $|B|^{a^2}$ sets. Then some nonempty subset $\mathcal{B} \subseteq \mathcal{A}$ exists such that $f^p(\bigcup \mathcal{B}) = f^p(\emptyset)$.*

The approach used to prove this is analogous to the following well known exercise (first set out by Vázsonyi and Sved, according to Erdös [1]): Prove that any sequence of n integers has a subsequence whose sum is divisible by n.

Proof. Note that $\mathcal{A}$ contains at least $|B|^{a^2} \geq |\operatorname{range}(f^p)|$ different sets. Let $A_1, \ldots, A_{|\operatorname{range}(f^p)|}$ be some of these sets. Define $B_i = \bigcup_{j\leq i} A_j$ for $0 \leq i \leq |\operatorname{range}(f^p)|$; note that $B_0 = \emptyset$. By the pigeonhole principle there exists $0 \leq i < j \leq |\operatorname{range}(f^p)|$ such that $f^p(B_i) = f^p(B_j)$. Then $f^p(B_j \setminus B_i) = f^p(B_j) - f^p(B_i) = f^p(B_i) - f^p(B_i) = f^p(B_i \setminus B_i) = f^p(\emptyset)$. Thus $\mathcal{B} = \{A_{i+1}, \ldots, A_j\}$ is the required family of sets. □

Lemma 19. *Fix a polymorphism $f \in \mathrm{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$. Consider any $S \subseteq [n]$. There exists $T \subseteq S$ of size at most $|B|^{a^2}$ such that $f^p(S) = f^p(T)$.*

Proof. Suppose this is not the case, and suppose that S is the minimal counterexample to this claim. Clearly $|S| > |B|^{a^2}$, or else taking $T = S$ shows that S is no counterexample at all. Thus, apply Lemma 18 to the set family $\{\{x\} \mid x \in S\}$ to find that some nonempty subset $U \subseteq S$ exists such that $f^p(U) = f^p(\emptyset)$. But now, take $S' = S \setminus U \subseteq S$, and note that $f^p(S') = f^p(S \setminus U) = f^p(S) - f^p(U) = f^p(S) - f^p(\emptyset) = f^p(S \setminus \emptyset) = f^p(S)$. By the minimality of S, S' has a subset T of size at most $|B|^{a^2}$ such that $f^p(T) = f^p(S') = f^p(S)$, which contradicts the fact that S is a counterexample. □

The following lemma elucidates the relation between f^* and f.

Lemma 20. *There exists a partial function $h : B^{a\times a} \to B$ such that, for any polymorphism $f \in \mathrm{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$, we have $f = h \circ f^*$.*

Proof. Without loss of generality suppose $s_1 = 1, s_2 = 2, \ldots, s_a = a$. To show our result, we will show that for any $\overrightarrow{S} = S_1, \ldots, S_a$, the value of $f(\overrightarrow{S})$ is uniquely determined by $f^*(\overrightarrow{S})$ i.e. the values $f_{ij}(S_k)$. In particular, we will show that $f(\overrightarrow{S})$ is uniquely determined by $f_{21}(S_2), \ldots, f_{a1}(S_a)$ and $f_{ii}(S_1) = f(i, i, \ldots, i)$ for $i \in [a]$.

To do this, create a matrix with r rows and n columns where the first row is distributed according to $\overrightarrow{S}$, and row i for $2 \leq i \leq a$ contains a 1 on positions j for which $j \in S_i$, and contains an i otherwise. Rows $a+1, \ldots, r$ will be constant and will contain the values $s_{a+1}, \ldots, s_r$ respectively. Apply f to the rows, and note that the image of the first row is $f(\overrightarrow{S})$, the image of the following $a-1$ rows is $f_{21}(S_2), \ldots, f_{a1}(S_a)$, and the last rows are constant and thus their images are equal to $f_{ii}(S_1)$ for some $i \in [a]$. Furthermore, each column contains a tuple of $R^{\mathbf{A}}$, namely a tuple that is a swap away from $(s_1, \ldots, s_r)$. By the functionality of $R^{\mathbf{B}}$ it follows that $f(\overrightarrow{S})$ is indeed uniquely determined by $f^*(\overrightarrow{S})$, as required. □

B. Proof of Theorem 13

In this section we assume that $\mathrm{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$ has a polymorphism f of arity n that is not k-degenerate for k at most N_d, and has no hard sets of size at most N_h. Given this, we will prove that $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is solved by AIP and is finitely tractable.

Definition 21. Define $0 = f^p(\emptyset)$ and $1 = f^p([n])$.

Lemma 22. *$G = (\operatorname{range}(f^p), +, 0)$ forms a group.*[8]

Proof. We prove this in a few parts.
Closure, well-definedness: Consider $x, y \in \operatorname{range}(f^p)$. As f is not 2-degenerate, there exist disjoint S, T such that $f^p(S) = x, f^p(T) = y$. Thus

$$x + y = f^p(S) + f^p(T) = f^p(S \cup T) \in \operatorname{range}(f^p),$$

so $+$ is closed and well-defined.
Associativity: Consider any $x, y, z \in \operatorname{range}(f^p)$. Since f is not 3-degenerate, there exist disjoint $S, T, U \subseteq [n]$ such that $f^p(S) = x, f^p(T) = y, f^p(U) = z$. Thus,

$$\begin{aligned} x + (y + z) &= x + f(T \cup U) = f(S \cup (T \cup U)) \\ &= f((S \cup T) \cup U) = f(S \cup T) + z = (x + y) + z. \end{aligned}$$

Identity element: Consider any $x \in \operatorname{range}(f^p)$. Suppose $f^p(S) = x$ for some $S \subseteq [n]$. Thus,

$$x + 0 = f^p(S \cup \emptyset) = f^p(S) = x.$$

[8] This group happens to be Abelian, but this is not needed for the proof.

Inverses: Consider any $x \in \mathrm{range}(f^p)$. Suppose that $f(S) = x$; by Lemma 19, some $T \subseteq S$ exists with size at most $|B|^{a^2}$ such that $f^p(T) = f^p(S) = x$. Since f has no hard sets of size at most $N_h = |B|^{a^2}$, T is not a hard set, and thus some $U \supseteq T$ exists such that $f^p(U) = 0$. Thus $x + f^p(U \setminus T) = f^p(T) + f^p(U \setminus T) = f^p(U) = 0$, so x has an inverse.

Thus we conclude that $(\mathrm{range}(f^p), +, 0)$ is a group. $\square$

Definition 23. Let G be the Abelian subgroup of

$$(\mathrm{range}(f^p), +, 0)$$

generated by $1 = f^p([n])$. Let m be the order of 1. Thus $G \cong \mathbb{Z}_m$. Note that $m \leq |\mathrm{range}(f^p)| \leq |B|^{a^2}$. We will identify $\mathbb{Z}_m$ with G.

Define the Abelian group $(H, +) = G^a$. We will identify H with $\mathbb{Z}_m^a$. We will also define 0 to be the 0 element in H as well as G.

For any $i \in [a]$, define $\overline{i} \in H$ as the unit vector that has a 1 at position i. For some tuple $(x_1, \dots, x_s) \in [a]^s$, define $\overline{(x_1, \dots, x_s)} = (\overline{x_1}, \dots, \overline{x_s}) \in H^s$. Define 0 to be the zero vector in H^s as well.[9]

For any relation $Q^{\mathbf{A}}$ of $\mathbf{A}$ of arity s, define $M(Q^{\mathbf{A}})$ to be the subgroup of H^s generated by $\overline{p} - \overline{q}$ for $p, q \in Q^{\mathbf{A}}$. Since H^s is Abelian, $M(Q^{\mathbf{A}})$ is a normal subgroup.

Lemma 24. *Fix some relation $Q^{\mathbf{A}}$ of $\mathbf{A}$; suppose it has arity s. Let t be some tuple of $Q^{\mathbf{A}}$. Define $M = M(Q^{\mathbf{A}})$. Consider any $(a_1, \dots, a_s) \in H^s$ such that $(a_1, \dots, a_s) \equiv \overline{t} \bmod M$. There exists a matrix (x_{ij}) with $N \leq N_d$ columns and s rows, where $N \equiv 1 \bmod m$, with elements in $[a]$, such that each column is a tuple of $Q^{\mathbf{A}}$, and, for each row i, we have*

$$\sum_{j=1}^{N} \overline{x_{ij}} = a_i.$$

Proof. Note that every element in H^s has order that divides m (since $H^s \cong (G^a)^s \cong (\mathbb{Z}_m^a)^s$). Thus, since $(a_1, \dots, a_s) \equiv \overline{t} \bmod M$, and since M is generated by $\overline{p} - \overline{q}$ for $p, q \in Q^{\mathbf{A}}$, it follows that there exist coefficients $c_{pq} \in \{0, \dots, m-1\}$ for $p, q \in Q^{\mathbf{A}}$ such that

$$(a_1, \dots, a_s) = \overline{t} + \sum_{p,q \in Q^{\mathbf{A}}} c_{pq}(\overline{p} - \overline{q}) = \overline{t} + \sum_{p,q \in Q^{\mathbf{A}}} c_{pq}\overline{p} + (m - c_{pq})\overline{q}. \quad (1)$$

This indicates the matrix we will use: let (x_{ij}) be a matrix whose first column is t, and, for each $p, q \in Q^{\mathbf{A}}$, has c_{pq} columns equal to p and $m - c_{pq}$ columns equal to q. Clearly we use $N \leq 1 + m|Q^{\mathbf{A}}|^2 = 1 + |B|^{a^2} a^{2r_{\max}} = N_d$ columns, and $N \equiv 1 \bmod m$. To see why $\sum_{j=1}^{N} \overline{x_{ij}} = a_i$ for each i, note that this condition is equivalent to $(a_1, \dots, a_s) = \sum_{j=1}^{N} \overline{c_j}$, where $c_1, \dots, c_N$ are the columns of the matrix. But this is precisely Equation (1). Thus we have created the required matrix. $\square$

[9]We can see the elements of H as frequency vectors modulo m. Indeed, for $x_1, \dots, x_n \in [a]$, $\overline{x_1} + \cdots + \overline{x_n}$ counts the number of appearances of $1, 2, \dots, a$ modulo m among $x_1, \dots, x_n$. In line with this, the elements of H^s can be seen as tuples of s frequency vectors. Under this view, for $t_1, \dots, t_n \in [a]^s$, the sum $\overline{t_1} + \cdots + \overline{t_n}$ is a tuple of s frequency vectors, where the i-th frequency vector counts the frequencies of the elements of $[a]$ among the i-th elements of the tuples $t_1, \dots, t_n$, modulo m.

We can now prove the main theorem in this subsection.

Proof of Theorem 13. We will show that $(\mathbf{A}, \mathbf{B})$ admits a homomorphic sandwich $\mathbf{A} \to \mathbf{E} \to \mathbf{B}$, where $\mathbf{E}$ is a relational structure whose domain is H, and where each relation will be of the form (i) $c + M$ for some $c \in H^s$ and $M \lhd H^s$, or (ii) empty. By Lemma 8 this implies our desired conclusion. The homomorphism $\mathbf{A} \to \mathbf{E}$ will be given by the map $g(x) = \overline{x}$. The homomorphism $\mathbf{E} \to \mathbf{B}$ will be given by any function h for which $f = h \circ f^*$. (Recall that such a function exists by Lemma 20.)[10] We will construct $\mathbf{E}$ relation by relation, showing along the way that g and h are in fact homomorphisms.

Consider some relation $Q^{\mathbf{A}}$ of $\mathbf{A}$, of arity s, that corresponds to a relation $Q^{\mathbf{B}}$ of $\mathbf{B}$, and $Q^{\mathbf{E}}$ in $\mathbf{E}$. If $Q^{\mathbf{A}}$ is empty, then we can simply set $Q^{\mathbf{E}}$ to be empty, and then g and h map tuples of $Q^{\mathbf{A}}$ to tuples of $Q^{\mathbf{E}}$, and then to tuples of $Q^{\mathbf{B}}$ vacuously. Thus, suppose $t = (t_1, \dots, t_s)$ is some tuple of $Q^{\mathbf{A}}$, and let $M = M(Q^{\mathbf{A}})$. Then we set $Q^{\mathbf{E}} = \overline{t} + M$; in other words, a tuple $x \in H^s$ will belong to this relation if and only if $x \equiv \overline{t} \bmod M$.

We first show that g maps $Q^{\mathbf{A}}$ to $\overline{t} + M$. Indeed, consider any tuple $x \in Q^{\mathbf{A}}$. We know that $g(x) = \overline{x}$ by definition. Thus, $g(x) = \overline{x} = \overline{t} + (\overline{x} - \overline{t}) \in \overline{t} + M$. Thus g maps $Q^{\mathbf{A}}$ to $\overline{t} + M$.

We now show that h maps $\overline{t} + M$ to $Q^{\mathbf{B}}$. Consider any tuple $(a_1, \dots, a_s) \in \overline{t} + M$. By Lemma 24 there exists some matrix (x_{ij}) with $N \leq N_d$ columns and s rows, where $N \equiv 1 \bmod m$, such that each column is an element of $Q^{\mathbf{A}}$, and for each $i \in [s]$ we have

$$\sum_{j=1}^{N} \overline{x_{ij}} = a_i.$$

Let $I_j^i = \{k \mid k \in [N], x_{ik} = j\}$; in other words, I_j^i is the set of columns k such that the (i, k)-th entry of the matrix is equal to j. Clearly $\overrightarrow{I}^i = I_1^i, \dots, I_a^i$ is a partition of $[N]$.

By assumption, f is not N-degenerate. Thus there exist disjoint subsets $S_1, \dots, S_N$ of $[n]$ where $f^p(S_1) = \cdots = f^p(S_N) = f^p([n]) = 1$. Let $T = [n] \setminus (S_1 \cup \ldots \cup S_N)$. Note that $S_1, \dots, S_N, T$ form a partition of $[n]$. Furthermore,

$$1 = f^p([n]) = f^p(S_1) + \cdots + f^p(S_N) + f^p(T) = N + f^p(T) = 1 + f^p(T).$$

The last equation holds as $N \equiv 1 \bmod m$, and addition is done in $G \cong \mathbb{Z}_m$. Thus $f^p(T) = 0$.

We will now create partitions $\overrightarrow{U}^1, \dots, \overrightarrow{U}^s \in a^{[n]}$ such that (the vectors that correspond to) these partitions constitute valid inputs for the polymorphism f (i.e. they are the rows of a matrix whose columns are tuples of $Q^{\mathbf{A}}$), and $f^*(\overrightarrow{U}^i) = a_i$.

[10]The h given by Lemma 20 is a partial function; take h to be any extension of it to the domain of $\mathbf{E}$.

First, let T_j^i be the empty set if $t_i \neq j$ and T otherwise. In either case, $f^p(T_j^i) = 0$. Define $\overrightarrow{U}^i = U_1^i, \dots, U_a^i$ by

$$U_j^i = T_j^i \cup \bigcup_{k \in I_j^i} S_k.$$

Note that $U_1^i, \dots, U_a^i$ form a partition of $[n]$, since each of $S_1, \dots, S_N, T$ will appear in exactly one of these sets (T in the set indicated by t_i, and S_k according to the partition $\overrightarrow{I}^i$), and $S_1, \dots, S_N, T$ form a partition of $[n]$. To see why $\overrightarrow{U}^1, \dots, \overrightarrow{U}^s$ form a valid input to the polymorphism f, suppose we create a matrix whose rows are distributed according to $\overrightarrow{U}^1, \dots, \overrightarrow{U}^s$. Consider column $j \in [n]$: if $j \in T$ then this column is equal to $t \in Q^{\mathbf{A}}$, and if $j \in S_k$ then this column is equal to column k of the matrix (x_{ij}), which is a tuple of $Q^{\mathbf{A}}$.

We must now show that $f^*(\overrightarrow{U}^i) = a_i$. Consider component $j \in [a]$ of the vector $f^*(\overrightarrow{U}^i)$ i.e. $f^p(U_j^i)$. Note that

$$f^p(U_j^i) = f^p\left(T_j^i \cup \bigcup_{k \in I_j^i} S_k\right) = f^p(T_j^i) + \sum_{k \in I_j^i} f^p(S_k) = 0 + |I_j^i| 1 = |I_j^i|.$$

Thus, since $f^*(\overrightarrow{U})$ is a vector that has $|I_j^i|$ as its j-th element (modulo m), and since $\overline{j}$ is the j-th unit vector, we can deduce that

$$f^*(\overrightarrow{U}^i) = \sum_{j=1}^{a} |I_j^i| \overline{j}.$$

Now, recall that $I_j^i = \{k \mid x_{ik} = j\}$. In other words, $|I_j^i|$ counts the multiplicity of $\overline{j}$ in the sum $\sum_{k=1}^N \overline{x_{ik}}$. But then clearly

$$f^*(\overrightarrow{U}^i) = \sum_{j=1}^{a} |I_j^i| \overline{j} = \sum_{k=1}^{N} \overline{x_{ik}} = a_i.$$

Thus $f^*(\overrightarrow{U}^i) = a_i$ as required.

Now, to see why h maps the tuple $(a_1, \dots, a_s) \in c + M$ to a tuple of $Q^{\mathbf{B}}$, note that

$$(h(a_1), \dots, h(a_s)) = (h(f^*(\overrightarrow{U}^1)), \dots, h(f^*(\overrightarrow{U}^s))) = (f(\overrightarrow{U}^1), \dots, f(\overrightarrow{U}^s)) \in Q^{\mathbf{B}}.$$

The last inclusion holds since f is a polymorphism, and $\overrightarrow{U}^1, \dots, \overrightarrow{U}^s$ are valid inputs to this polymorphism.

Thus we note that $\mathbf{A} \to \mathbf{E} \to \mathbf{B}$ for some structure $\mathbf{E}$ that satisfies the conditions in Lemma 8. In conclusion, $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is finitely tractable and solved by AIP. $\square$

C. *Proof of Theorem 14*

In this section we will prove that $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is NP-hard if each polymorphism $f \in \mathrm{Pol}(\mathbf{A}, \mathbf{B})$ is k-degenerate for some k at most N_d, or has a hard set of size at most N_h.

Lemma 25. *If $f \in \mathrm{Pol}(\mathbf{A}, \mathbf{B})$, then f cannot have more than $|B|^{a^2}$ disjoint hard sets.*

Proof. Consider any n-ary polymorphism $f \in \mathrm{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$ and suppose it has more than $|B|^{a^2}$ disjoint hard sets. Let $\mathcal{H}$ be a family of more than $|B|^{a^2}$ disjoint hard sets. Apply Lemma 18 to $\mathcal{H}$ to find a nonempty subfamily $\mathcal{G}$ of hard sets for which $f^p(\bigcup \mathcal{G}) = f^p(\emptyset)$. For any $G \in \mathcal{G}$, this contradicts the fact that G is a hard set. $\square$

Lemma 26. *Suppose $f \in \mathrm{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$ and $\pi : [n] \to [m]$. Then $(f^\pi)^p = f^p \circ \pi^{-1}$.*

Proof. Note that $f_{ij}(S) = f(T_1, \dots, T_a)$ where $T_j = S$, $T_i = [n] \setminus S$, and all the other inputs are $\emptyset$. Now, $(f^\pi)_{ij}(S) = f^\pi(T_1, \dots, T_a) = f(\pi^{-1}(T_1), \dots, \pi^{-1}(T_a)) = f^p(\pi^{-1}(S))$, and so $(f^\pi)_{ij} = f_{ij} \circ \pi^{-1}$. Our conclusion follows by applying this fact for each $i, j \in A$. $\square$

Lemma 27. *Suppose $f \in \mathrm{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$ and $\pi : [n] \to [m]$. If S is a hard set of f then $\pi(S)$ is a hard set of f^π.*

Proof. We prove this by contrapositive. Suppose $\pi(S)$ is not a hard set of f^π. Then some $T \supseteq \pi(S)$ exists such that $(f^\pi)^p(T) = (f^\pi)^p(\emptyset)$. So,

$$(f^p)(\pi^{-1}(T)) = (f^\pi)^p(T) = (f^\pi)^p(\emptyset) = (f^p)(\pi^{-1}(\emptyset)) = f^p(\emptyset).$$

Thus $f^p(\pi^{-1}(T)) = f^p(\emptyset)$, and S is not a hard set, as $S \subseteq \pi^{-1}(T)$. $\square$

Let $\mathcal{M}_h$ denote the subset of $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ whose polymorphisms have hard sets of size at most N_h. Let $\mathcal{M}_{x_1, \dots, x_k}$ denote the subset of $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ whose polymorphisms are k-degenerate, yet not $(k-1)$-degenerate, where $x_1, \dots, x_k \in B^{a \times a}$ are witnesses to this degeneracy. By assumption, and as no polymorphism is 1-degenerate,

$$\mathrm{Pol}(\mathbf{A}, \mathbf{B}) = \mathcal{M}_h \cup \bigcup_{k=2}^{N_d} \bigcup_{x_1, \dots, x_k \in B^{a \times a}} \mathcal{M}_{x_1, \dots, x_k}. \quad (2)$$

Lemma 28. *There exists some assignment I that takes $f \in \mathcal{M}_h^{(n)}$ to a subset of $[n]$ of size at most $|B|^{2a^2}$ such that, whenever $g \in \mathcal{M}_h^{(m)}$ and $g = f^\pi$ for some $\pi : [n] \to [m]$, we have that $\pi(I(f)) \cap I(g) \neq \emptyset$.*

Proof. To construct $I(f)$, let $S_1, \dots$ be a maximal sequence of disjoint hard sets of f of size at most $|B|^{a^2}$, constructed greedily, and then set $I(f)$ to be the union of these sets. Since there can be at most $|B|^{a^2}$ disjoint hard sets by Lemma 25, it follows that $|I(f)| \leq |B|^{2a^2}$.

Consider now any $f, g \in \mathcal{M}_h$ such that $g = f^\pi$. Note that $I(f)$ contains within it a hard set S of size at most $|B|^{a^2}$. Thus $\pi(I(f)) \supseteq \pi(S)$, which is a hard set of size at most $|B|^{a^2}$ by Lemma 27, and thus must intersect $I(g)$ by maximality. It follows that $\pi(I(f)) \cap I(g) \neq \emptyset$. $\square$

Lemma 29. *For $k \geq 2, x_1, \dots, x_k \in A$, there exists some assignment I that takes $f \in \mathcal{M}_{x_1, \dots, x_k}^{(n)}$ to a subset of $[n]$ of size at most $k|B|^{a^2}$ such that, whenever $g \in \mathcal{M}_{x_1, \dots, x_k}^{(m)}$ and $g = f^\pi$ for some $\pi : [n] \to [m]$ we have that $\pi(I(f)) \cap I(g) \neq \emptyset$.*

Proof. To construct $I(f)$, take $S_1, \ldots, S_{k-1}$ to be disjoint sets such that $f^p(S_i) = x_i$, and take T to be any set such that $f(T) = x_k$. Such sets exist since f is not $(k-1)$-degenerate, and we can take all of these sets to be of size at most $|B|^{a^2}$, by replacing them with the subsets given by Lemma 19. Let $I(f)$ be the union of $S_1, \ldots, S_{k-1}, T$. Note that $|I(f)| \leq k|B|^{a^2}$.

Consider now any $f, g \in \mathcal{M}_{x_1,\ldots,x_k}$ such that $g = f^\pi$. Note that $I(f)$ contains within it disjoint sets $S_1, \ldots, S_{k-1}$ such that $f^p(S_i) = x_i$, and $I(g)$ contains within it a set T such that $g^p(T) = x_k$. Now, $f^p(\pi^{-1}(T)) = (f^\pi)^p(T) = g^p(T) = x_k$, and thus by the k-degeneracy of f and the disjointness of $S_1, \ldots, S_{k-1}$ it follows that $\pi^{-1}(T)$ and $S_1, \ldots, S_{k-1}$ must intersect. It follows that $\pi(I(f)) \cap I(g) \neq \emptyset$, as required. □

Proof of Theorem 14. We see in (2) that $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ is the union of $m = 1 + \sum_{k=2}^{N_d} \left(|B|^{a^2}\right)^k$ sets, each of which has an assignment I that satisfies the condition of Theorem 9 for $C = \max(N_d|B|^{a^2}, |B|^{2a^2})$. Thus $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ is NP-hard. □

IV. BLP + AIP = AIP when A has one balanced relation

In this section we prove Theorem 4 and Corollary 5. Recall that we say that a relation R is *balanced* if and only if there exists a matrix whose columns are tuples of R, that contains every tuple of R as a column, and whose rows are permutations of each other.

Theorem 4. *Let* $\mathbf{A}$ *be any structure with a single relation. If the relation in* $\mathbf{A}$ *is balanced then, for any* $\mathbf{B}$ *such that* $\mathbf{A} \to \mathbf{B}$*,* $\mathrm{BLP} + \mathrm{AIP}$ *solves* $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ *if and only if* AIP *solves it.*

Suppose that $A = [a]$, and the relation of $\mathbf{A}$ is $R = R^{\mathbf{A}}$. Furthermore suppose that each element in $[a]$ appears in R (otherwise these elements can just be eliminated from A). Suppose $A \neq \emptyset, R \neq \emptyset$ (otherwise the conclusion is trivially true). Suppose also that the columns of the matrix that witness the balancedness of R are $t_1, \ldots, t_N \in R$.

For any $i \in [a]$, let $\overline{i}$ be a unit vector in $\mathbb{Z}^a$; i.e., it has a 1 at position i. For any tuple $(a_1, \ldots, a_r) \in A^r$, let $\overline{(a_1, \ldots, a_r)} = (\overline{a_1}, \ldots, \overline{a_r}) \in (\mathbb{Z}^a)^r$. Let $\overline{R} = \{\overline{t} \mid t \in R\} \subseteq (\mathbb{Z}^a)^r$. Endow all of these vectors with additive structure. For any $k \in \mathbb{Z}$, define $S_k \subseteq \mathbb{Z}^a$ to be the set of vectors that sum up to k, with non-negative coordinates.[11]

Lemma 30. $(k+1)\overline{R} - k\overline{R} + k\sum_i \overline{t_i} \subseteq (kN+1)\overline{R}$.

Proof. If $x \in (k+1)\overline{R} - k\overline{R} + k\sum_i \overline{t_i}$, it can be written as a sum of $k+1$ vectors from $\overline{R}$, minus k vectors from $\overline{R}$, plus k copies of each vector $\overline{t_i}$. Since each tuple of R appears among $t_1, \ldots, t_N$, the last kN vectors in the sum above include at least k copies of each vector in $\overline{R}$. By removing the k subtracted vectors from the k copies of each vector from $\overline{R}$, we find that x can be written as a sum of $k+1-k+kN = kN+1$ vectors from $\overline{R}$, i.e. $x \in (kN+1)\overline{R}$. □

[11] We see the elements of $\mathbb{Z}^a$ as frequency vectors, and the elements of $(\mathbb{Z}^a)^r$ as tuples of frequency vectors.

Lemma 31. *If* $(\mathbf{A}, \mathbf{B})$ *has a 2-block-symmetric polymorphism* f *of arity* $2k+1$ *then there exists a function* $g : S_k \times S_{k+1} \to B$ *such that* $(g(x_1, y_1), \ldots, g(x_r, y_r)) \in R^{\mathbf{B}}$ *for all* $(x_1, \ldots, x_r) \in k\overline{R}, (y_1, \ldots, y_r) \in (k+1)\overline{R}$.

Proof. To compute $g(x, y)$, create sequences of elements in $[a]$, of lengths k and $k+1$, whose frequencies correspond to x and y respectively (i.e. the sequence for $x = (x_1, \ldots, x_a)$ has x_i appearances of i), and interleave these to create a sequence $a_1, \ldots a_{2k+1}$. Then $g(x, y) = f(a_1, \ldots, a_{2k+1})$.

To see why this function satisfies the required condition, suppose $(x_1, \ldots, x_r) \in k\overline{R}$ and $(y_1, \ldots, y_r) \in (k+1)\overline{R}$. Thus we can, by definition, create matrices M and N, with k and $k+1$ columns respectively, and r rows, where each column is an element of R, and each row i has frequencies corresponding to x_i and y_i respectively. Interleave the columns of these matrices to create a matrix A. Apply f to the rows of A. We find that the image of row i of A is $g(x_i, y_i)$ by the symmetry of f; furthermore, the images of the rows of A must form a tuple of $R^{\mathbf{B}}$, since f is a polymorphism. This is the desired conclusion. □

Lemma 32. *Assume there exists a function* $f : (S_{k+1} - S_k) \to B$ *such that* $(f(x_1), \ldots, f(x_r)) \in R^{\mathbf{B}}$ *for any* $x_1, \ldots, x_r \in S_{k+1} - S_k$ *with* $(x_1, \ldots, x_r) \in (k+1)\overline{R} - k\overline{R}$*. Then,* $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ *has an alternating polymorphism of arity* $2k+1$.

Proof. If such a function exists, then

$$g(x_1, \ldots, x_{2k+1}) = f(\overline{x_1} + \overline{x_3} + \cdots + \overline{x_{2k+1}} - \overline{x_2} - \overline{x_4} - \cdots - \overline{x_{2k}})$$

is the required polymorphism. □

Proof of Theorem 4. By Theorem 6, AIP solves $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ if and only if $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ contains alternating operations of all odd arities. By Theorem 7, $\mathrm{BLP} + \mathrm{AIP}$ solves $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ if and only if $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ contains 2-block-symmetric operations of all odd arities. As any alternating operation is 2-block-symmetric, it follows that any PCSP solved by AIP is also solved by $\mathrm{BLP} + \mathrm{AIP}$.[12] It suffices to show that 2-block-symmetric operations in $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ imply alternating operations.

Fix some natural number k; we will now show that there exists an alternating operation in $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ of arity $2k+1$. Since $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ contains a 2-block-symmetric operation of arity $2kN+1$, let $f : S_{kN} \times S_{kN+1} \to B$ be the function given by Lemma 31. We will construct the function required by Lemma 32 in order to prove the existence of an alternating polymorphism.

Consider the vector $v = \sum_i \overline{t_i} \in (\mathbb{Z}^a)^r$. We claim that v is a constant vector. To see why this is the case, observe that one way to compute $\sum_i \overline{t_i}$ is to make $t_1, \ldots, t_N$ into the columns of a matrix, and then to compute the frequencies of each element of $[a]$ in each row. Element i of v is a tuple,

[12] This also directly follows from the definitions of the AIP and BLP + AIP algorithms [10].

where element j is the number of appearances of j in row i in this matrix. But, since $t_1, \dots, t_N$ witness the balancedness of R, these frequencies are equal for each row. Thus v is indeed a constant vector; suppose that $v = (c, \dots, c)$ for some $c \in S_N$. Note that each element in $[a]$ appears in some tuple of R by assumption, and each tuple of R appears in the sum $\sum_i \overline{t_i}$. Thus each coordinate in $c \in \mathbb{Z}^a$ is at least 1.

The function we are interested in is $g : (S_{k+1} - S_k) \to B$, where

$$g(x) = f(kc, x + kc).$$

First note that these inputs are legal inputs for the function f. To see why, note first that $c \in S_N$ and thus $kc \in S_{kN}$. Second, consider $x + kc$. As $x \in S_{k+1} - S_k$, the elements in x sum up to 1. Since the elements in kc sum up to kN, it follows that the elements in $x + kc$ sum up to $1 + kN$ as required. Furthermore, all the elements of $x + kc$ are non-negative: each element of x is at least $-k$, whereas each element of c is at least 1, and thus each element of kc is at least k. Thus $x + kc \in S_{kN+1}$.

Why does g satisfy the conditions from Lemma 32? Consider any $x_1, \dots, x_r \in S_{k+1} - S_k$ such that $(x_1, \dots, x_r) \in (k+1)\overline{R} - k\overline{R}$. Note that

$$(kc, \dots, kc) = k(c, \dots, c) = k \sum_i \overline{t_i} \in kN\overline{R},$$

and

$$\begin{aligned}(x_1 + kc, \dots, x_r + kc) &= (x_1, \dots, x_r) + k(c, \dots, c) \in \\ &(k+1)\overline{R} - k\overline{R} + k\sum_i \overline{t_i} \subseteq (kN+1)\overline{R},\end{aligned}$$

due to Lemma 30. Thus, since f satisfies the conditions in Lemma 31,

$$\begin{aligned}&(g(x_1), \dots, g(x_r)) = \\ &\quad (f(kc, x_1 + kc), \dots, f(kc, x_r + kc)) \in R^B.\end{aligned}$$

Thus $(g(x_1), \dots, g(x_r)) \in R^{\mathbf{B}}$, as required. ☐

Theorem 4 does not generalise to structures with multiple relations (even just two), as the following examples shows.

Remark 33. Consider a Boolean symmetric template $\mathbf{A}$ that has two balanced relations, namely $R^{\mathbf{A}} = \{(0)\}$ and $Q^{\mathbf{A}} = \{(0,1),(1,0),(1,1)\}$, which are unary and binary, respectively. Then $\mathrm{CSP}(\mathbf{A})$ is solved by $\mathrm{BLP} + \mathrm{AIP}$, and indeed by BLP, since the symmetric operation $\max(x_1, \dots, x_n)$ is a polymorphism for any n [6], but *not* by AIP. This is because $\mathbf{A}$ fails to have any alternating non-unary polymorphisms, even of arity 3: suppose $f(x,y,z)$ is such a polymorphism. Then $f(1,1,0) = f(0,0,0) = f(0,1,1)$ as f is alternating; and $f(0,0,0) = 0$ due to $R^{\mathbf{A}}$. However, due to $Q^{\mathbf{A}}$, $f(1,1,0)$ and $f(0,1,1)$ cannot both be 0. This contradiction implies our conclusion.

One cannot simply remove the balancedness condition from Theorem 4, as the following example shows.

Remark 34. Let $\mathbf{A}$ be a Boolean template with relation $S^{\mathbf{A}} = \{(0,0,1),(0,1,0),(0,1,1)\}$. Note that $S^{\mathbf{A}}$ is *not* balanced. Then $\mathrm{CSP}(\mathbf{A})$ is solved by $\mathrm{BLP} + \mathrm{AIP}$, and indeed by BLP, since the symmetric operation $\max(x_1, \dots, x_n)$ is a polymorphism for any n [6], but not by AIP. $\mathbf{A}$ fails to have any alternating polymorphism, even of arity 3, for exactly the same reason as the problem from Remark 33. (The identities that would result from $R^{\mathbf{A}}$ in that example now result from the first element in each tuple in $S^{\mathbf{A}}$, and the identities that would result from $Q^{\mathbf{A}}$ in that example now result from the last two elements in each tuple in $S^{\mathbf{A}}$.)

On the other hand, there are templates that are unbalanced for which AIP and $\mathrm{BLP} + \mathrm{AIP}$ have equivalent power, as the following example shows.

Remark 35. Consider a Boolean template $\mathbf{A}$ that has one relation $P^{\mathbf{A}} = \{(0,1)\}$. Then $\mathrm{CSP}(\mathbf{A})$ is solved by AIP and by $\mathrm{BLP} + \mathrm{AIP}$, since the alternating operation $x_1 + \cdots + x_n \bmod 2$ is a polymorphism of $\mathbf{A}$ for every odd n. This is in spite of the fact that $P^{\mathbf{A}}$ is unbalanced.

We now prove Corollary 5.

Corollary 5. *Suppose that G is a transitive group of permutations. Further, suppose that $\mathbf{A}$ is a relational structure with one relation that is preserved by G. Then, for any $\mathbf{A} \to \mathbf{B}$, $\mathrm{BLP} + \mathrm{AIP}$ solves $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ if and only if AIP does.*

Proof. Let R be the relation of $\mathbf{A}$, of arity r. It is sufficient to show that R is balanced. Let M be a matrix whose columns are the tuples of R. Suppose that the rows of M are $r_1, \dots, r_n$. We show that row i is a permutation of row j, for arbitrary $i, j \in [r]$.

Represent the elements of G as permutation matrices. Let $\pi \in G$ be a permutation (matrix) that sends i to j (it exists by transitivity). Consider πM. Note that no two columns of πM can be equal, since then two columns of $\pi^{-1}\pi M = M$ would be equal, which is false. Furthermore each column of πM is a tuple of R, and thus a column of M, since R is preserved by π. Thus we see that πM can be seen as M but with its columns permuted. In other words, for some permutation matrix σ, we have $\pi M = M\sigma^T$.

Now, let us look at row j in $\pi M = M\sigma^T$. In πM this is r_i (since π sends i to j). In $M\sigma^T$ this is $r_j\sigma^T$. Thus $r_i = r_j\sigma^T$, i.e. row i of M and row j of M are permutations of each other. We conclude that R is balanced, as required. ☐

This corollary applies to fewer structures than Theorem 4, as shown in the next example.

Remark 36. Consider any digraph $\mathbf{A}$ with edge relation $E^{\mathbf{A}}$ that is strongly connected but not symmetric. Then, it is easy to show that $E^{\mathbf{A}}$ is balanced. On the other hand, the unique transitive permutation group with degree 2 (i.e. the group containing the identity permutation and the permutation swapping two elements) does not preserve $E^{\mathbf{A}}$.

V. Conclusion

Our first result classifies problems $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ where $\mathbf{A}$ is symmetric and has a non-binary covering tuple and $\mathbf{B}$ is functional into being either NP-hard or solvable in polynomial time. This is the first step towards the following more general problem.

Problem 37. Classify the complexity of $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ when $\mathbf{B}$ is functional.

Looking more specifically at the case $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$, we note that our proof of Theorem 1 implies that, for functional $\mathbf{B}$, we have that $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ is tractable if and only if $\mathbf{Eqn}_{m,1} \to \mathbf{B}$ for some $m \leq |B|$, where $\mathbf{Eqn}_{m,1}$ is a relational structure over $\{0, \ldots, m-1\}$ with one ternary relation defined by $x+y+z \equiv 1 \bmod m$. By using the Chinese remainder theorem, $\mathbf{Eqn}_{m,1} = \mathbf{Eqn}_{3^p,1} \times \mathbf{Eqn}_{q,1}$, where q is coprime to 3. Since this latter template contains a constant tuple (namely (x, x, x) where x is the inverse of 3 modulo q), we find that, for functional $\mathbf{B}$, $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ is tractable if and only if $\mathbf{Eqn}_{3^p,1} \to \mathbf{B}$.

Looking at non-functional templates $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ that are tractable, all the examples the authors are aware of are either tractable for the same reason as a functional template is (i.e. $\mathbf{Eqn}_{3^p,1} \to \mathbf{B}$), or because they include the not-all-equal predicate (i.e. $\mathbf{NAE} \to \mathbf{B}$). Thus, we pose the following problem.

Problem 38. Is $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ tractable if and only if $\mathbf{Eqn}_{3^p,1} \times \mathbf{NAE} \to \mathbf{B}$ for some p?

Problem 38 has a link with the problem of determining the complexity of $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{C}_k^+)$, where $\mathbf{C}_k^+$ is a ternary symmetric template on domain $[k]$ which contains tuples of the form $(1,1,2), \ldots, (k-1, k-1, k), (k, k, 1)$, as well as all tuples of three distinct elements (rainbow tuples). Such templates are called *cyclic*, with the cycle being $1 \to \ldots \to k \to 1$.

The link is the following: $\mathbf{Eqn}_{3^p,1} \times \mathbf{NAE}$ is a template containing one cycle of length 2×3^p, together with certain rainbow tuples — in other words, $\mathbf{Eqn}_{3^p,1} \times \mathbf{NAE} \to \mathbf{C}^+_{2\times 3^p}$. Likewise, $\mathbf{Eqn}_{3^p,1}$ has a cycle of length 3^p and some rainbow tuples, i.e. $\mathbf{Eqn}_{3^p,1} \to \mathbf{C}^+_{3^p}$. That $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{C}_k^+)$ is tractable whenever $k = 3^p$ or $k = 2 \times 3^p$ was first observed in [11]. If Problem 38 were answered in the affirmative then we would have that $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{C}_k^+)$ is tractable if and only if $k = 3^p$ or $k = 2 \times 3^p$. In particular, this would mean that $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{C}_4^+)$ is NP-hard, as conjectured in [5].[13]

Answering Problem 38 in the affirmative would resolve Conjecture 3, i.e., $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ would be tractable (via AIP) if and only if $\mathbf{Eqn}_{3^p,1} \times \mathbf{NAE} \to \mathbf{B}$. Perhaps determining whether this equivalence is true might be easier than resolving Conjecture 3; thus we pose the following problem.

Problem 39. Is $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ solved by AIP if and only if $\mathbf{Eqn}_{3^p,1} \times \mathbf{NAE} \to \mathbf{B}$ for some p?

There already exists such a characterisation for the power of AIP using an infinite structure [6]. In particular, if we let $\mathbf{Z}$ be an infinite structure whose domain is $\mathbb{Z}$, and with a tuple (x, y, z) in the relation if and only if $x + y + z = 1$, then $\mathrm{PCSP}(\mathbf{1\text{-}in\text{-}3}, \mathbf{B})$ is solved by AIP if and only if $\mathbf{Z} \to \mathbf{B}$. We are interested in a finite template of this kind.

Turning from problems to algorithms, our second result shows us that, for certain problems of the form $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ where $\mathbf{B}$ *need not be functional*, and $\mathbf{A}, \mathbf{B}$ have one relation, AIP and BLP + AIP have the same power. A natural question is for which other templates is it true?

Problem 40. For which templates do AIP and BLP + AIP have the same power?

We remark that the recent work [16] does not answer any problem from this section, and the results from [16] are consistent with positive answers to Problems 38 and 39.

[13] Our structure $\mathbf{C}_4^+$ is called $\check{\mathbf{C}}^+$ in [5].

Acknowledgment

We thank all three anonymous reviewers for their comments and suggestions for changes.

References

[1] M. Aigner and G. M. Ziegler, *Proofs from THE BOOK*, 4th ed. Germany: Springer Publishing Company, Incorporated, 2009.

[2] K. Asimi and L. Barto, "Finitely tractable promise constraint satisfaction problems," in *Proc. 46th International Symposium on Mathematical Foundations of Computer Science (MFCS'21)*, ser. LIPIcs, vol. 202. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, pp. 11:1–11:16.

[3] A. Atserias and V. Dalmau, "Promise Constraint Satisfaction and Width," in *Proc. 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22)*. USA: SIAM, 2022, pp. 1129–1153.

[4] P. Austrin, V. Guruswami, and J. Håstad, "(2+ϵ)-Sat is NP-hard," *SIAM J. Comput.*, vol. 46, no. 5, pp. 1554–1573, 2017. [Online]. Available: https://eccc.weizmann.ac.il/report/2013/159/

[5] L. Barto, D. Battistelli, and K. M. Berg, "Symmetric Promise Constraint Satisfaction Problems: Beyond the Boolean Case," in *Proc. 38th International Symposium on Theoretical Aspects of Computer Science (STACS'21)*, ser. LIPIcs, vol. 187. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, pp. 10:1–10:16.

[6] L. Barto, J. Bulín, A. A. Krokhin, and J. Opršal, "Algebraic approach to promise constraint satisfaction," *J. ACM*, vol. 68, no. 4, pp. 28:1–28:66, 2021.

[7] L. Barto and M. Kozik, "Combinatorial Gap Theorem and Reductions between Promise CSPs," in *Proc. 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22)*. USA: SIAM, 2022, pp. 1204–1220.

[8] J. Brakensiek and V. Guruswami, "Promise Constraint Satisfaction: Algebraic Structure and a Symmetric Boolean Dichotomy," *SIAM J. Comput.*, vol. 50, no. 6, pp. 1663–1700, 2021.

[9] ——, "The quest for strong inapproximability results with perfect completeness," *ACM Trans. Algorithms*, vol. 17, no. 3, pp. 27:1–27:35, 2021. [Online]. Available: https://eccc.weizmann.ac.il/report/2017/080/

[10] J. Brakensiek, V. Guruswami, M. Wrochna, and S. Živný, "The power of the combined basic LP and affine relaxation for promise CSPs," *SIAM J. Comput.*, vol. 49, pp. 1232–1248, 2020.

[11] A. Brandts, "Promise constraint satisfaction problems," Ph.D. dissertation, University of Oxford, 2022. [Online]. Available: https://ora.ouls.ox.ac.uk/objects/uuid:5efeff56-d5c8-42e8-9bd2-95243a1367ad

[12] A. Brandts, M. Wrochna, and S. Živný, "The complexity of promise SAT on non-Boolean domains," *ACM Trans. Comput. Theory*, vol. 13, no. 4, pp. 26:1–26:20, 2021.

[13] A. Brandts and S. Živný, "Beyond PCSP(1-in-3,NAE)," *Information and Computation*, vol. 289, no. Part A, p. 104954, 2022.

[14] A. Bulatov, P. Jeavons, and A. Krokhin, "Classifying the complexity of constraints using finite algebras," *SIAM J. Comput.*, vol. 34, no. 3, pp. 720–742, 2005.

[15] A. A. Bulatov, "A dichotomy theorem for nonuniform CSPs," in *Proc. 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*. USA: IEEE, 2017, pp. 319–330.

[16] L. Ciardo, M. Kozik, A. Krokhin, T.-V. Nakajima, and S. Živný, "On the complexity of the approximate hypergraph homomorphism problem," Tech. Rep., 2023.

[17] L. Ciardo and S. Živný, "CLAP: A New Algorithm for Promise CSPs," *SIAM J. Comput.*, vol. 52, no. 1, pp. 1–37, 2023.

[18] I. Dinur, O. Regev, and C. Smyth, "The hardness of 3-uniform hypergraph coloring," *Comb.*, vol. 25, no. 5, pp. 519–535, 2005.

[19] T. Feder and M. Y. Vardi, "The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory," *SIAM J. Comput.*, vol. 28, no. 1, pp. 57–104, 1998.

[20] M. Ficak, M. Kozik, M. Olšák, and S. Stankiewicz, "Dichotomy for Symmetric Boolean PCSPs," in *Proc. 46th International Colloquium on Automata, Languages, and Programming (ICALP'19)*, vol. 132. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019, pp. 57:1–57:12.

[21] M. R. Garey and D. S. Johnson, "The complexity of near-optimal graph coloring," *J. ACM*, vol. 23, no. 1, pp. 43–49, 1976.

[22] V. Guruswami and S. Sandeep, "d-To-1 Hardness of Coloring 3-Colorable Graphs with O(1) Colors," in *Proc. 47th International Colloquium on Automata, Languages, and Programming (ICALP'20)*, ser. LIPIcs, vol. 168. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020, pp. 62:1–62:12. [Online]. Available: https://eccc.weizmann.ac.il/report/2019/116/

[23] P. Hell and J. Nešetřil, "On the Complexity of H-coloring," *Journal of Combinatorial Theory, Series B*, vol. 48, no. 1, pp. 92–110, 1990.

[24] P. G. Jeavons, "On the Algebraic Structure of Combinatorial Problems," *Theor. Comput. Sci.*, vol. 200, no. 1-2, pp. 185–204, 1998.

[25] A. Krokhin and J. Opršal, "An invitation to the promise constraint satisfaction problem," *ACM SIGLOG News*, vol. 9, no. 3, pp. 30–59, 2022.

[26] A. A. Krokhin, J. Opršal, M. Wrochna, and S. Živný, "Topology and adjunction in promise constraint satisfaction," *SIAM J. Comput.*, vol. 52, no. 1, pp. 38–79, 2023.

[27] T.-V. Nakajima and S. Živný, "Linearly Ordered Colourings of Hypergraphs," *ACM Transactions on Computation Theory*, vol. 14, no. 3–4, pp. 12:1–12:19, 2022.

[28] T. Schaefer, "The complexity of satisfiability problems," in *Proc. 10th Annual ACM Symposium on the Theory of Computing (STOC'78)*. USA: ACM, 1978, pp. 216–226.

[29] M. Wrochna, "A note on hardness of promise hypergraph colouring," 2022.

[30] D. Zhuk, "A proof of the CSP dichotomy conjecture," *J. ACM*, vol. 67, no. 5, pp. 30:1–30:78, 2020.