

CLAP: A NEW ALGORITHM FOR PROMISE CSPS*

LORENZO CIARDO[†] AND STANISLAV ŽIVNÝ[†]

Abstract. We propose a new algorithm for Promise Constraint Satisfaction Problems (PCSPs). It is a combination of the Constraint Basic LP relaxation and the Affine IP relaxation (CLAP). We give a characterization of the power of CLAP in terms of a minion homomorphism. Using this characterization, we identify a certain weak notion of symmetry which, if satisfied by infinitely many polymorphisms of PCSPs, guarantees tractability. We demonstrate that there are PCSPs solved by CLAP that are not solved by any of the existing algorithms for PCSPs; in particular, not by the BLP + AIP algorithm of Brakensiek et al. [SIAM J. Comput., 49 (2020), pp. 1232–1248] and not by a reduction to tractable finite-domain CSPs.

Key words. promise constraint satisfaction, homomorphism problems, linear programming

MSC codes. 68Q25, 68R01, 90C05

DOI. 10.1137/22M1476435

1. Introduction.

Constraint satisfaction. Constraint Satisfaction Problems (CSPs) have driven some of the most influential developments in theoretical computer science, from NP-completeness to the PCP theorem [1, 2, 39] to semidefinite programming algorithms [63] to the Unique Games Conjecture [53].

A CSP over domain A is specified by a finite collection \mathbf{A} of relations over A , and is denoted by $\text{CSP}(\mathbf{A})$. Given on input a set of variables and a set of constraints, each of which uses relations from \mathbf{A} , the task is to decide the existence of an assignment of values from A to the variables that satisfies all the constraints. Classic examples of CSPs include 2-SAT, graph 3-coloring, and linear equations of fixed width over finite groups.

For Boolean CSPs, which are CSPs with $|A| = 2$, Schaefer proved that every such CSP is either solvable in polynomial time or is NP-complete [65]. Feder and Vardi famously conjectured that the same holds true for CSPs over arbitrary finite domains [41]. Furthermore, they realized the importance of considering closure properties of solution spaces of CSPs [41], which initiated the algebraic approach [26, 49, 50]. The key notion in the algebraic approach is that of *polymorphisms*, which are operations that take solutions to a CSP and are guaranteed to return, by a coordinatewise application, a solution to the same CSP. All CSPs admit projections (also known as dictators) as polymorphisms. However, the presence of less trivial polymorphisms, satisfying some notion of symmetry, is necessary for tractability. For instance, the set of solutions to 2-SAT is closed under the ternary majority operation $\text{maj} : \{0, 1\}^3 \rightarrow \{0, 1\}$ that satisfies the following notion of symmetry: $\text{maj}(a, a, b) = \text{maj}(a, b, a) = \text{maj}(b, a, a) = a$ for

*Received by the editors February 7, 2022; accepted for publication (in revised form) October 11, 2022; published electronically January 25, 2023. An extended abstract of this work appeared in the Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22) [32].

<https://doi.org/10.1137/22M1476435>

Funding: The authors' research was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 714532) and by UKRI EP/X024431/1. The second author's research was also supported by a Royal Society University Research Fellowship.

[†]Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, OX1 3QD Oxford, UK (lorenzo.ciarDO@cs.ox.ac.uk, standa.zivny@cs.ox.ac.uk).

any $a, b \in \{0, 1\}$. Similarly, the set of solutions to Horn-SAT is closed under the binary minimum operation $\min : \{0, 1\}^2 \rightarrow \{0, 1\}$ that satisfies a different notion of symmetry: $\min(a, a) = a$, $\min(a, b) = \min(b, a)$, and $\min(a, \min(b, c)) = \min(\min(a, b), c)$ for any $a, b, c \in \{0, 1\}$. The binary max operation—which is a polymorphism of dual Horn-SAT—has the same notion of symmetry, called semilattice [11]. Together with the ternary minority polymorphism, which captures linear equations on $\{0, 1\}$, this gives all nontrivial tractable cases from Schaefer’s dichotomy result.¹

The polymorphisms of any CSP form a clone, in that they include all projections and are closed under composition. For instance, since Horn-SAT has \min as a polymorphism, it also has the 4-ary minimum operation

$$\min_4(a, b, c, d) = \min(a, \min(b, \min(c, d)))$$

as a polymorphism. Building on the connection to universal algebra, the algebraic approach has been tremendously successful beyond decision CSPs, e.g., for robust satisfiability of CSPs [10, 35, 36], for exact optimization of CSPs [54, 57, 66], and for characterizing the power of algorithms [9, 14, 55, 56, 59, 67, 68]. The culmination of the algebraic approach is the positive resolution of the dichotomy conjecture by Bulatov [28] and Zhuk [71]. We refer the reader to [11] for a survey on the algebraic approach.

Promise constraint satisfaction. In this paper, we study Promise Constraint Satisfaction Problems (PCSPs), whose systematic study was initiated by Austrin, Guruswami, and Håstad [5], and Brakensiek and Guruswami [20]. PCSPs form a vast generalization of CSPs. In $\text{PCSP}(\mathbf{A}, \mathbf{B})$, each constraint comes in two forms, a strict one in \mathbf{A} and a weak one in \mathbf{B} . The goal is to distinguish between (i) the case in which (the strong form of) the constraints can be simultaneously satisfied in \mathbf{A} and (ii) the case in which (even the weak form of) the constraints cannot be simultaneously satisfied in \mathbf{B} . The promise is that it is never the case that the PCSP is not satisfiable in the strict sense but is satisfiable in the weak sense. If the strict and weak forms coincide in every constraint (i.e., if $\mathbf{A} = \mathbf{B}$) we get the (nonpromise) CSPs. However, PCSPs include many fundamental problems that are inexpressible as CSPs.

The simplest example of strict versus weak constraints is when the weak constraints are supersets of the strict constraints on the same domain (the first two examples below) or on a larger domain (the third example below); the notion of homomorphism from \mathbf{A} to \mathbf{B} formalizes this for any PCSP.

First, can we distinguish a g -satisfiable k -SAT instance (in the sense that there is an assignment that satisfies at least g literals in each clause) from an instance that is not even 1-satisfiable? This problem was studied in [5], where it was shown to be solvable in polynomial time if $\frac{g}{k} \geq \frac{1}{2}$ and NP-complete otherwise. Recently, this result has been generalized to arbitrary finite domains [23].

Second, can we distinguish a 3-SAT formula that admits an assignment satisfying exactly 1 literal in each clause (i.e., a satisfiable instance of **1-in-3-SAT**) from one that does not admit an assignment satisfying 1 or 2 literals in each clause (i.e., a nonsatisfiable instance of Not-All-Equal-3-SAT)? Remarkably, while both **1-in-3** and **NAE** are NP-hard, this promise version is solvable in polynomial time [19, 20].

Third, can we distinguish a k -colorable graph from a graph that is not even ℓ -colorable, where $k \leq \ell$? This is the *approximate graph coloring* problem, which is

¹The trivial cases, called 0- and 1-valid, are captured by the constant-0 and constant-1 polymorphisms, respectively.

believed to be NP-hard for any fixed $3 \leq k \leq \ell$, but has been elusive since the 1970s [43]. In particular, the larger the gap is between k and ℓ , the easier the problem could in principle be and, thus, the more challenging it is to prove NP-hardness. The current state of the art is NP-hardness for $k=3$ and $\ell=5$ [8], while already the case of $k=3$ and $\ell=6$ is open. For any $k \geq 4$ and $\ell = \ell(k) = \binom{k}{\lfloor k/2 \rfloor} - 1$, NP-hardness has been established in [70].

While a systematic study of PCSPs was initiated only recently [5, 20], concrete PCSPs have been considered for a while, e.g., approximate graph [16, 43, 44, 51, 52, 69] and hypergraph coloring [40]. A highlight result is the dichotomy of Boolean symmetric PCSPs [42] (in which all constraint relations are symmetric), following an earlier classification of Boolean symmetric PCSPs with disequalities [20]. Very recent works have investigated certain Boolean nonsymmetric PCSPs [24] and certain non-Boolean symmetric PCSPs [7]. Other recent results include, e.g., [4, 21, 45].

Most of the recent progress, including results on the approximate graph coloring problem [8, 70] and on the approximate graph homomorphism problem [58, 70], rely on the algebraic approach to PCSPs [8]. In particular, the breakthrough results in [8], building on [12], established that the complexity of PCSPs is captured by the polymorphism minions and certain types of symmetries these minions satisfy —these are nonnested identities on polymorphisms, such as the majority example but not the semilattice example. Crucially, minions are less structured than clones: A minion (of functions) is a set of operations closed under permuting coordinates, identifying coordinates, and introducing dummy coordinates, but not under composition.² Thus, unlike in our earlier CSP example (corresponding to Horn-SAT), a binary minimum polymorphism of a PCSP cannot, in general, be used to generate a 4-ary minimum polymorphism of the same PCSP.

Despite the momentous results in [8], there is a long way to go to classify all PCSPs, and it is not even clear whether a dichotomy for all PCSPs should be expected. When Feder and Vardi conjectured a CSP dichotomy [41], the Boolean case [65] and the graph case [46] had been fully classified. We seem quite far from these two cases being classified for PCSPs. Thus, further progress is needed on both the hardness and tractability part. This paper focuses on the latter.

Finite tractability. Although PCSPs are (much) more general than CSPs, some PCSPs can be reduced to tractable CSPs. This idea was introduced in [19] under the name of homomorphic sandwiching (cf. section 2 for a precise definition); PCSPs that are reducible to tractable (finite-domain) CSPs are called *finitely tractable*. Finite tractability is not sufficient to explain tractability of all tractable PCSPs. In particular, Barto et al. [8] showed that the above-mentioned example **1-in-3** versus **NAE** is not finitely tractable, despite being a tractable PCSP [20]. We remark that it is not inconceivable (and in fact was conjectured in [19]) that every tractable (finite-domain) PCSP could be reducible to a tractable CSP possibly over an infinite domain; this is the case for the **1-in-3** versus **NAE** problem [19]. However, while certain infinite-domain CSPs are amenable to algebraic methods, the complexity of infinite-domain CSPs is far from understood; cf. [13, 17, 18] for recent work.

Since finite tractability does not capture all tractable PCSPs, there is need for other algorithmic tools. One possibility is to attempt to extend algorithmic techniques developed for CSPs.

²In this work, we shall use the more abstract notion of minion introduced in [22]; cf. Definition 2.6.

There are two main algorithmic approaches to CSPs. On the one hand, there are local consistency methods [41], which have been studied in theoretical computer science but also in artificial intelligence, logic, and database theory. The power of local consistency for CSPs has been characterized in [9, 25], and it is known that the third level of consistency solves all so-called bounded-width CSPs [6]. On the other hand, there are CSPs solvable by algorithms based on generalizations of Gaussian elimination, most notably CSPs with a Mal'tsev polymorphism [29]. This method has been pushed to its limit, in a way, in [14, 48]. While the NP-hardness part of the CSP dichotomy has been known since [26], the challenge in proving the algorithmic part is the complicated interaction of these two very different algorithmic approaches. Although this interaction does not occur in Boolean CSPs, it occurs already in CSPs on three-element domains [27].

The characterization of the power of the first level of the consistency methods, 1-consistency (also known as arc-consistency [61]), has been lifted from CSPs [41] to PCSPs in [8]. Rather than establishing 1-consistency combinatorially, one can employ convex relaxations.

Relaxations. A canonical analogue of 1-consistency is the *basic linear programming relaxation* (BLP) [59], which in fact is stronger than 1-consistency [60]. The characterization of the power of BLP has been lifted from CSPs [59] to PCSPs in [8], both in terms of a minion and a property of polymorphisms. The power of BLP is captured by a minion consisting of rational stochastic vectors³ or, equivalently, by the presence of symmetric polymorphisms of all arities; these are polymorphisms invariant under any permutation of the coordinates. For example, we have seen that Horn-SAT, a classic CSP, has a binary symmetric polymorphism, namely min. We have also seen that min can generate a 4-ary operation \min_4 , which is symmetric. Similarly, min can generate (via composition) symmetric operations of all arities, and thus Horn-SAT is solved by BLP.

A different relaxation of PCSPs is the *basic affine integer programming relaxation* (AIP) [19]. The power of AIP has been characterized, both in terms of a minion and a property of polymorphisms, in [8]. The minion capturing AIP consists of integer affine vectors.⁴ Concerning polymorphisms, AIP is captured by polymorphisms of all odd arities that are invariant under permutations that only permute odd and even coordinates separately, and additionally satisfy that adjacent coordinates cancel each other out. The **1-in-3** versus **NAE** problem is solved by AIP (cf. Example 2.5).

Brakensiek et al. [22] proposed a combination of the two above-mentioned relaxations, called BLP + AIP. Their algorithm has many interesting features. First, it solves PCSPs that admit only infinitely many symmetric polymorphisms (i.e., not all arities are required as in the case of BLP). Second, it solves all tractable Boolean CSPs, thus demonstrating how research on PCSPs can shed new light on (nonpromise) CSPs. In fact, [22] established the power of BLP + AIP in terms of a minion and (a property of) polymorphisms. The minion capturing BLP + AIP is essentially a product of the BLP and AIP minions [22]. Concerning polymorphisms, BLP + AIP is captured by polymorphisms of all odd arities that are invariant under permutations that only permute odd and even coordinates.

It may be that BLP + AIP is the only algorithm needed to solve all tractable *Boolean* PCSPs. However, as already observed in [22], BLP + AIP does not solve some rather simple, tractable, *non-Boolean* PCSPs. Motivated by this, we investigate

³A vector is *stochastic* if its entries are nonnegative and sum up to one.

⁴An integer vector is *affine* if its entries sum up to one.

algorithms that are stronger than BLP + AIP. We note that all PCSPs hitherto known to be tractable are solved by BLP + AIP or by finite tractability (i.e., by a reduction to a tractable finite-domain CSP). In this work, we provide an example of a PCSP that is tractable (through our algorithm) but is not solved by either of those two algorithmic techniques.

Contributions. Building on the work of Brakensiek et al. [22], we study stronger relaxations for PCSPs and give three main contributions.

(1) CLAP Our first contribution is the introduction of CLAP to the study of PCSPs. Our goal was to design an algorithm that, unlike BLP + AIP, solves all CSPs of bounded width. While all bounded-width CSPs can be solved by 3-consistency [6], and thus also by the third level of the Sherali–Adams hierarchy for BLP (e.g., by [67]), Kozik showed that already (a special case of) the singleton arc-consistency (SAC) algorithm, introduced in [38] (cf. [15, 30]), solves all bounded-width CSPs [56]. Thus, we study the LP relaxation that we call the *singleton* BLP (SBLP), which is at least as strong as SAC. A special case of SBLP (without this name) implicitly appeared in the literature, e.g., in [5, 20] for Boolean PCSPs. The idea behind SBLP is essentially to run SAC but replace the arc-consistency check by the BLP; i.e., the algorithm repeatedly takes a variable-value pair (x, a) and tests the feasibility of the BLP with the requirement that x should be assigned the value a . If this LP is infeasible, then a is removed from the domain of x . This is repeated until convergence. If any variable ends up with an empty domain, then SBLP rejects, otherwise it accepts. Overall, the number of BLP calls occurring for an instance of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ with variable-set X is at most polynomial in the size of X . As mentioned above, this simple algorithm solves all bounded-width CSPs [56].

We adopt a modification of SBLP that turns out to be more naturally captured by a minion-oriented analysis: the *constraint* BLP (CBLP). This (possibly) stronger algorithm is a generalization of SBLP in which we do not consider only variable-value pairs (x, a) , but rather the constraint-assignment pairs (\mathbf{x}, \mathbf{a}) for every constraint in the instance. As in SBLP, if fixing a (local) assignment to a constraint yields an infeasible BLP, then the assignment is removed from the constraint relation. Upon convergence, which takes at most polynomially many BLP calls, if any constraint ends up with an empty relation, then CBLP rejects, otherwise it accepts.

Our algorithm CLAP first runs CBLP and then, upon termination, refines the solutions of CBLP by running (essentially) AIP. If one believes the suggestion in [22] that constantly many rounds of the Sherali–Adams hierarchy for BLP + AIP could solve all tractable (non-promise) CSPs, then it is not outrageous to believe that the same could be true for CLAP, and CLAP might be easier to analyze than such an algorithm.

(2) Characterization Our second contribution is a minion characterization of the power of CLAP, stated as Theorem 3.3. The objects in the minion are essentially matrices with a particular structure, which we call *skeletal* (cf. Definition 3.1). These matrices capture the CBLP part of CLAP and together with certain integer affine vectors form the minion (cf. Definition 3.2). Another, more conceptual contribution is the introduction of a minion of matrices to the study of PCSPs.

(3) H -symmetric polymorphisms The minion characterization is crucial to our third contribution: the identification of a sufficient condition for CLAP to work in terms of the symmetries of the polymorphisms. This is stated as Theorem 3.5, using the notion of H -symmetry. This condition can be more easily checked for concrete templates, thus allowing us to design a separating example that is not finitely tractable and is not solved by BLP + AIP (nor by local consistency methods; see [3]), but is

solved by CLAP. It follows that our new algorithm is strictly more powerful than BLP+AIP (and separated by an interesting PCSP that is not reducible to a tractable finite-domain CSP via “gadget reductions,” which capture the algebraic approach to PCSPs [8]).

For a matrix H , a polymorphism f is H -symmetric if f is invariant under permutations of the coordinates but only on a specific set of inputs determined by H (cf. Definition 3.4). For instance, if H is a row vector, then we obtain the requirement that f be symmetric on all inputs. If H is the identity matrix, then we require that f be symmetric only on inputs in which different entries occur with different multiplicities. In general, the intuition is that we capture “symmetry with exceptions that depend on multiplicities.” We refer the reader to the discussion in section 3 for details.

After necessary background material in section 2, our algorithm CLAP and the main results are presented in section 3; the proofs appear in sections 4 and 5.

2. Preliminaries. We let $\mathbb{N} = \{1, 2, \dots\}$ and $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. The cardinality of \mathbb{N} shall be denoted by \aleph_0 . For $k \in \mathbb{N}$, $[k]$ denotes the set $\{1, \dots, k\}$. For a set A , $\mathcal{P}(A)$ denotes the set of all subsets of A . We denote by \leq_p many-one polynomial-time reductions. We shall use standard notation for vectors and matrices. Vectors will be treated as column vectors and whenever convenient identified with the corresponding (row) tuples. Both tuples and vectors will be typed in bold font. We denote by \mathbf{e}_i the i th standard unit vector of the appropriate size (which will be clear from the context); i.e., \mathbf{e}_i is equal to 1 in the i th coordinate and 0 elsewhere. We denote by $\mathbf{0}_p$ and by $\mathbf{1}_p$ the all-zero and all-one vector, respectively, of size p ; if the size is clear, we occasionally drop the subscript. The *support* of a vector $\mathbf{v} = (v_i)$ of size p is the set $\text{supp}(\mathbf{v}) = \{i \in [p] : v_i \neq 0\}$. I_p denotes the identity matrix of order p , while O denotes an all-zero matrix of suitable size.

Promise CSPs. A *signature* σ is a finite set of relation symbols R , each with its arity $\text{ar}(R) \in \mathbb{N}$. A *relational structure* over a signature σ , or a σ -structure, is a finite universe A , called the *domain* of \mathbf{A} , and a relation $R^{\mathbf{A}} \subseteq A^{\text{ar}(R)}$ for each symbol $R \in \sigma$. For two σ -structures \mathbf{A} and \mathbf{B} , a mapping $f : A \rightarrow B$ is called a *homomorphism* from \mathbf{A} to \mathbf{B} , denoted by $f : \mathbf{A} \rightarrow \mathbf{B}$, if f preserves all relations; that is, for every $R \in \sigma$ and every tuple $\mathbf{a} \in R^{\mathbf{A}}$, we have $f(\mathbf{a}) \in R^{\mathbf{B}}$, where f is applied coordinatewise. The existence of a homomorphism from \mathbf{A} to \mathbf{B} is denoted by $\mathbf{A} \rightarrow \mathbf{B}$. A PCSP *template* is a pair (\mathbf{A}, \mathbf{B}) of relational structures over the same signature such that $\mathbf{A} \rightarrow \mathbf{B}$. Without loss of generality, we will often assume that A , the domain of \mathbf{A} , is $[n]$.

DEFINITION 2.1. *Let (\mathbf{A}, \mathbf{B}) be a PCSP template. Then, the decision version of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ is the following problem: Given as input a relational structure \mathbf{X} over the same signature as \mathbf{A} and \mathbf{B} , output Yes if $\mathbf{X} \rightarrow \mathbf{A}$ and No if $\mathbf{X} \not\rightarrow \mathbf{B}$. The search version of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ is the following problem: Given as input a relational structure \mathbf{X} over the same signature as \mathbf{A} and \mathbf{B} and such that $\mathbf{X} \rightarrow \mathbf{A}$, find a homomorphism from \mathbf{X} to \mathbf{B} .*

For a relational structure \mathbf{A} , the CSP with template \mathbf{A} [41], denoted by $\text{CSP}(\mathbf{A})$, is $\text{PCSP}(\mathbf{A}, \mathbf{A})$.

Example 2.2. For $k \geq 2$, let \mathbf{K}_k be the structure with domain $[k]$ and a binary relation $\{(i, j) \in [k]^2 \mid i \neq j\}$. Then, $\text{CSP}(\mathbf{K}_k)$ is the standard graph k -coloring problem. For $k \leq \ell$, $\text{PCSP}(\mathbf{K}_k, \mathbf{K}_\ell)$ is the *approximate graph coloring problem* [43]. In the decision version, the task is to decide whether a graph is k -colorable or not even ℓ -colorable. In the search version, given a k -colorable graph G , the task is to find an ℓ -coloring of G . It is widely believed that for any fixed $3 \leq k \leq \ell$, $\text{PCSP}(\mathbf{K}_k, \mathbf{K}_\ell)$

is NP-hard; i.e., constantly many colors do not help. The current most general NP-hardness result is known for $k = 3$ and $\ell = 5$ by Bulín, Krokhin, and Opršal [8] and for $k \geq 4$ and $\ell = \ell(k) = \binom{k}{\lfloor k/2 \rfloor} - 1$ by Wrochna and Živný [70].

We call a PCSP template (\mathbf{A}, \mathbf{B}) *tractable* if any instance of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ can be solved in polynomial time in the size of the input structure \mathbf{X} . It is easy to show that the decision version reduces to the search version [8] (but the converse is not known in general); for CSPs, the two versions are equivalent [26, 33]. Our results are for the decision version.

Example 2.3. Let **1-in-3** be the Boolean structure with domain $\{0, 1\}$ and a single ternary relation $\{(0, 0, 1), (0, 1, 0), (1, 0, 0)\}$. Let **NAE** be the structure with domain $\{0, 1\}$ and a single ternary relation $\{0, 1\}^3 \setminus \{(0, 0, 0), (1, 1, 1)\}$. Then, $\text{CSP}(\mathbf{1-in-3})$ is the (positive) 1-in-3-SAT problem and $\text{CSP}(\mathbf{NAE})$ is the (positive) Not-All-Equal-3-SAT problem. Since both of these problems are NP-hard [65], the PCSP templates $(\mathbf{1-in-3}, \mathbf{1-in-3})$ and $(\mathbf{NAE}, \mathbf{NAE})$ are both intractable. However, the PCSP template $(\mathbf{1-in-3}, \mathbf{NAE})$ is tractable, as shown by Brakensiek and Guruswami [20].

DEFINITION 2.4. Let (\mathbf{A}, \mathbf{B}) be a PCSP template with signature σ . An operation $f : A^L \rightarrow B$, where $L \in \mathbb{N}$, is a polymorphism of arity L of (\mathbf{A}, \mathbf{B}) if for every $R \in \sigma$ of arity $k = \text{ar}(R)$ and for any possible $L \times k$ matrix whose rows are tuples in $R^{\mathbf{A}}$, the application of f on the columns of the matrix gives a tuple in $R^{\mathbf{B}}$. We denote by $\text{Pol}(\mathbf{A}, \mathbf{B})$ the set of all polymorphisms of (\mathbf{A}, \mathbf{B}) .

Example 2.5. The unary operation $\neg : \{0, 1\} \rightarrow \{0, 1\}$ defined by $\neg(a) = 1 - a$ is a polymorphism of $(\mathbf{NAE}, \mathbf{NAE})$ but not a polymorphism of $(\mathbf{1-in-3}, \mathbf{1-in-3})$. For any odd L , the L -ary operation $f : \{0, 1\}^L \rightarrow \{0, 1\}$ defined by $f(a_1, \dots, a_L) = 1$ if $a_1 - a_2 + a_3 - \dots + a_L > 0$ and $f(a_1, \dots, a_L) = 0$ otherwise is a polymorphism of $(\mathbf{1-in-3}, \mathbf{NAE})$.

Minions. Polymorphisms of CSPs form clones; i.e., $\text{Pol}(\mathbf{A}, \mathbf{A})$ contains all projections (also known as dictators) and is closed under composition [11]. Polymorphisms of the (more general) PCSPs form *minions*; i.e., they are closed under taking minors.⁵ Formally, given an L -ary function $f : A^L \rightarrow B$, its *minor* relative to a map $\pi : [L] \rightarrow [L']$ is the L' -ary function $f_{/\pi} : A^{L'} \rightarrow B$ defined by

$$(2.1) \quad f_{/\pi}(a_1, \dots, a_{L'}) = f(a_{\pi(1)}, \dots, a_{\pi(L)}).$$

Equivalently, a minor of f is a function obtained from f by identifying variables, permuting variables, and introducing dummy variables. Rather than focusing on minions of functions, we consider here abstract minions, as described and used in [22].

DEFINITION 2.6. A minion \mathcal{M} consists of the disjoint union of sets $\mathcal{M}^{(L)}$ for $L \in \mathbb{N}$ equipped with operations $(\cdot)_{/\pi} : \mathcal{M}^{(L)} \rightarrow \mathcal{M}^{(L')}$ for all functions $\pi : [L] \rightarrow [L']$, which satisfy

- $(M_{/\pi})_{/\tilde{\pi}} = M_{/\tilde{\pi} \circ \pi}$ for $\pi : [L] \rightarrow [L']$, $\tilde{\pi} : [L'] \rightarrow [L'']$ and
- $M_{/\text{id}} = M$

for all $M \in \mathcal{M}^{(L)}$.

DEFINITION 2.7. For two minions \mathcal{M} and \mathcal{N} , a minion homomorphism $\xi : \mathcal{M} \rightarrow \mathcal{N}$ is a map that preserves arities and minors: Given $M \in \mathcal{M}^{(L)}$ and $\pi : [L] \rightarrow [L']$, $\xi(M) \in \mathcal{N}^{(L')}$ and $\xi(M_{/\pi}) = \xi(M)_{/\pi}$.

⁵We remark that clones are also closed under taking minors.

For any PCSP template (\mathbf{A}, \mathbf{B}) , the set $\text{Pol}(\mathbf{A}, \mathbf{B})$ of its polymorphisms equipped with the operations described by (2.1) is a minion [8]. One of the results in [8] established that minion homomorphisms give rise to polynomial-time reductions: If there is a minion homomorphism from $\text{Pol}(\mathbf{A}, \mathbf{B})$ to $\text{Pol}(\mathbf{A}', \mathbf{B}')$, then $\text{PCSP}(\mathbf{A}', \mathbf{B}') \leq_p \text{PCSP}(\mathbf{A}, \mathbf{B})$. Minions are also useful for characterizing the power of algorithms, as we will discuss later.

Remark 2.8. Although we will not use this categorical view, we remark that a minion is nothing but a functor from the category of nonempty finite sets to the category of nonempty sets, and a minion homomorphism is a natural transformation.

Existing algorithms. One way to establish tractability of PCSPs is to reduce to CSPs. Let (\mathbf{A}, \mathbf{B}) be a PCSP template. A structure \mathbf{C} is called a (homomorphic) *sandwich* if $\mathbf{A} \rightarrow \mathbf{C} \rightarrow \mathbf{B}$. It is known that, in this case, $\text{PCSP}(\mathbf{A}, \mathbf{B}) \leq_p \text{CSP}(\mathbf{C})$.⁶ Thus, if \mathbf{C} is a tractable CSP template, then (\mathbf{A}, \mathbf{B}) is a tractable PCSP template. If \mathbf{C} has a finite domain, we say that (\mathbf{A}, \mathbf{B}) is *finitely tractable*.

Example 2.9. The PCSP template $(\mathbf{1-in-3}, \mathbf{NAE})$ from Example 2.3 is tractable, as shown in [20], but not finitely tractable unless $P = NP$, as shown in [8].

Another way to establish tractability for PCSPs is to leverage convex relaxations. In section 1, we mentioned three studied relaxations: BLP [59], AIP [20], and BLP + AIP [22]. Their powers have been characterized in [8, 22] in terms of certain minions and polymorphism identities. The details of these relaxations and the characterizations are provided in Appendix A.

All PCSPs hitherto known to be tractable are solved by finite tractability (i.e., by a reduction to a tractable finite-domain CSP) or by BLP + AIP. The next example identifies a simple PCSP template not captured by either of these two methods.

Example 2.10. Consider the relational structures $\mathbf{A} = (A; R_1^{\mathbf{A}}, R_2^{\mathbf{A}})$ and $\mathbf{B} = (B; R_1^{\mathbf{B}}, R_2^{\mathbf{B}})$ on the domain $A = B = \{0, \dots, 6\}$ with the following relations: $R_1^{\mathbf{A}} = \{(0, 0, 1), (0, 1, 0), (1, 0, 0)\}$ is **1-in-3** on $\{0, 1\}$, $R_1^{\mathbf{B}} = \{0, 1\}^3 \setminus \{(0, 0, 0), (1, 1, 1)\}$ is **NAE** on $\{0, 1\}$, and $R_2^{\mathbf{A}} = R_2^{\mathbf{B}} = \{(2, 3), (3, 2), (4, 5), (5, 6), (6, 4)\}$. The identity mapping is a homomorphism from \mathbf{A} to \mathbf{B} , so (\mathbf{A}, \mathbf{B}) is a PCSP template. Since the directed graph corresponding to $R_2^{\mathbf{A}} = R_2^{\mathbf{B}}$ is a disjoint union of a directed 2-cycle and a directed 3-cycle, [22, Example 6.1] shows that the BLP + AIP algorithm does not solve $\text{PCSP}(\mathbf{A}, \mathbf{B})$. We claim that the template (\mathbf{A}, \mathbf{B}) is not finitely tractable. For contradiction, assume that there is a finite relational structure $\mathbf{C} = (C; R_1^{\mathbf{C}}, R_2^{\mathbf{C}})$ such that $\mathbf{A} \rightarrow \mathbf{C} \rightarrow \mathbf{B}$ and $\text{CSP}(\mathbf{C})$ is tractable. We will argue that this would imply finite tractability of $(\mathbf{1-in-3}, \mathbf{NAE})$, which contradicts the result in [8] (unless $P = NP$); cf. Example 2.9. Indeed, the existence of such \mathbf{C} gives the following chain of homomorphisms:

(2.2)

$$\mathbf{1-in-3} = (\{0, 1\}; R_1^{\mathbf{A}}) \rightarrow (A; R_1^{\mathbf{A}}) \rightarrow (C; R_1^{\mathbf{C}}) \rightarrow (B; R_1^{\mathbf{B}}) \rightarrow (\{0, 1\}; R_1^{\mathbf{B}}) = \mathbf{NAE},$$

where the first map is the inclusion of $\{0, 1\}$ in A , the second and the third are the maps witnessing $\mathbf{A} \rightarrow \mathbf{C} \rightarrow \mathbf{B}$, and the fourth is any map $g : B \rightarrow \{0, 1\}$ such that $g(0) = 0$ and $g(1) = 1$. Let $\tilde{\mathbf{C}} = (C; R_1^{\mathbf{C}})$. Observe that $\tilde{\mathbf{C}}$ is tractable since the inclusion map gives a minion homomorphism $\text{Pol}(\mathbf{C}, \mathbf{C}) \rightarrow \text{Pol}(\tilde{\mathbf{C}}, \tilde{\mathbf{C}})$, and thus $\text{CSP}(\tilde{\mathbf{C}}) = \text{PCSP}(\tilde{\mathbf{C}}, \tilde{\mathbf{C}}) \leq_p \text{PCSP}(\mathbf{C}, \mathbf{C}) = \text{CSP}(\mathbf{C})$ by [8, Theorem 3.1]. This proves the claim, as (2.2) established $\mathbf{1-in-3} \rightarrow \tilde{\mathbf{C}} \rightarrow \mathbf{NAE}$.

⁶This is a special case of homomorphic relaxation [8], which we do not need here.

Notice that the assignment $f \mapsto g \circ f|_{\{0,1\}^L}$ (where f is a polymorphism of (\mathbf{A}, \mathbf{B}) of arity L and g is the map considered above) yields a minion homomorphism from $\text{Pol}(\mathbf{A}, \mathbf{B})$ to $\text{Pol}(\mathbf{1-in-3}, \mathbf{NAE})$. As established in [3, Corollary 4.2], the template $(\mathbf{1-in-3}, \mathbf{NAE})$ does not have bounded width—i.e., is not solved by local consistency methods. It follows from [8, Lemma 7.5] that (\mathbf{A}, \mathbf{B}) does not have bounded width either.

The template from Example 2.10 will be proved tractable later (in Example 3.6) using our new algorithm, which we will present next.

3. The CLAP algorithm. Let (\mathbf{A}, \mathbf{B}) be a PCSP template with signature σ , and let \mathbf{X} be an instance of $\text{PCSP}(\mathbf{A}, \mathbf{B})$. Without loss of generality, we assume that σ contains a unary symbol R_u such that $R_u^{\mathbf{X}} = X$, $R_u^{\mathbf{A}} = A$, and $R_u^{\mathbf{B}} = B$. If this is not the case, the signature and the instance can be extended without changing the set of solutions. Our algorithm—the combined **CBLP**+**AIP** algorithm (CLAP), presented in Algorithm 3.1 and discussed below—builds on BLP [8] and BLP + AIP [22].

CLAP works in two stages. In the first stage, it runs CBLP; i.e., a modified version of the singleton arc-consistency algorithm (cf. [38]) where (i) the “arc-consistency” part is replaced by BLP, and (ii) the “singleton” part is boosted by requiring that every constraint-assignment pair (as opposed to every variable-value pair) is fixed at each iteration. In the second stage, it refines CBLP by doing an additional sanity check: At least one of the solutions computed by CBLP should be compatible with a solution of AIP. As in [22], this second stage requires that the AIP solution should only use those variables from the CBLP solution that have nonzero weight. There are two equivalent ways to enforce this requirement: Either by storing the nonzero variables at each iteration of CBLP in the first stage of the algorithm, or by simply running BLP + AIP as a black box in the second stage of the algorithm. We adopt the latter option to achieve a simpler presentation. Concretely, the first stage of CLAP is performed by initializing the sets $S_{\mathbf{x}, R}$ of constraint-assignment pairs to the entire relation $R^{\mathbf{A}}$, and then progressively shrinking these sets by cycling over all constraint-assignment pairs and removing a pair whenever it yields an infeasible BLP. The second stage, that occurs if all sets $S_{\mathbf{x}, R}$ are nonempty, is performed by cycling over each feasible constraint-assignment pair and running BLP + AIP on it. As soon as one constraint-assignment pair is accepted by BLP + AIP, the algorithm terminates and outputs Yes. If no constraint-assignment pair is accepted, the algorithm outputs No.

As in Appendix A, where BLP, AIP, and BLP + AIP are presented in full detail for completeness, by $\lambda_{\mathbf{x}, R}(\mathbf{a})$ we denote the variable of $\text{BLP}(\mathbf{X}, \mathbf{A})$ associated with $\mathbf{x} \in R^{\mathbf{X}}$ and $\mathbf{a} \in R^{\mathbf{A}}$, where $R \in \sigma$. The algorithm has polynomial time complexity in the size of the input instance: Letting $g = \sum_{R \in \sigma} |R^{\mathbf{X}}| |R^{\mathbf{A}}|$, $\mathcal{O}(g^2)$ BLP calls and $\mathcal{O}(g)$ BLP + AIP calls occur. We say that CLAP *accepts* an instance \mathbf{X} of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ if Algorithm 3.1 returns Yes. We say that CLAP *solves* $\text{PCSP}(\mathbf{A}, \mathbf{B})$ if, for every instance \mathbf{X} of $\text{PCSP}(\mathbf{A}, \mathbf{B})$, we have (i) if $\mathbf{X} \rightarrow \mathbf{A}$, then CLAP accepts \mathbf{X} , and (ii) if \mathbf{X} is accepted by CLAP, then $\mathbf{X} \rightarrow \mathbf{B}$.

Characterization. Our first main result—Theorem 3.3—is a minion-theoretic characterization of the power of the CLAP algorithm. In particular, we will introduce in Definition 3.2 a minion \mathcal{C} such that, for any PCSP template (\mathbf{A}, \mathbf{B}) , the CLAP algorithm solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$ if and only if there is a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$. The two directions will be proved in Theorems 4.10 and 4.11, respectively, in section 4. Combining Theorem 4.10 with our second main result—

Algorithm 3.1 The CLAP algorithm.

Require: an instance \mathbf{X} of PCSP(\mathbf{A}, \mathbf{B}) of signature σ
Ensure: YES if $\mathbf{X} \rightarrow \mathbf{A}$ and NO if $\mathbf{X} \not\rightarrow \mathbf{B}$

```

1: for  $R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}$  do
2:   set  $S_{\mathbf{x}, R} := R^{\mathbf{A}}$ 
3: end for
4: repeat
5:   for  $R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in S_{\mathbf{x}, R}$  do
6:     if BLP( $\mathbf{X}, \mathbf{A}$ ) with  $\lambda_{\mathbf{x}, R}(\mathbf{a}) = 1$  and  $\lambda_{\mathbf{x}', R'}(\mathbf{a}') = 0$  for every  $R' \in \sigma$ ,
        $\mathbf{x}' \in R'^{\mathbf{X}}$ , and  $\mathbf{a}' \notin S_{\mathbf{x}', R'}$  is not feasible then
7:       remove  $\mathbf{a}$  from  $S_{\mathbf{x}, R}$ 
8:     end if
9:   end for
10:  until no set  $S_{\mathbf{x}, R}$  is changed
11:  if some  $S_{\mathbf{x}, R}$  is empty then
12:    return NO;
13:  else
14:    for  $R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in S_{\mathbf{x}, R}$  do
15:      if BLP + AIP( $\mathbf{X}, \mathbf{A}$ ) with  $\lambda_{\mathbf{x}, R}(\mathbf{a}) = 1$  and  $\lambda_{\mathbf{x}', R'}(\mathbf{a}') = 0$  for every  $R' \in \sigma$ ,
         $\mathbf{x}' \in R'^{\mathbf{X}}$ , and  $\mathbf{a}' \notin S_{\mathbf{x}', R'}$  is feasible then
16:        return YES
17:      else
18:        return NO
19:      end if
20:    end for
21:  end if

```

Theorem 3.5, proved in section 5—will then yield a sufficient condition for CLAP to solve a given PCSP template, in terms of a weak notion of symmetry for the polymorphisms of the template.

The L -ary objects of the minion \mathcal{C} are pairs $(M, \boldsymbol{\mu})$, where M is a matrix with L rows and infinitely many columns encoding the BLP computations of CLAP and $\boldsymbol{\mu}$ is an L -ary vector of integers encoding the AIP computation of CLAP. The matrices M in \mathcal{C} have a special structure, which we call “skeletal.”

DEFINITION 3.1. *Let M be a $p \times \aleph_0$ matrix with $p \in \mathbb{N}$. We say that M is skeletal if, for each $j \in [p]$, either $\mathbf{e}_j^T M = \mathbf{0}_{\aleph_0}^T$ or $M \mathbf{e}_i = \mathbf{e}_j$ for some $i \in \mathbb{N}$.*

In other words, either the j th row of M is the zero vector or some column of M is the j th standard unit vector. Equivalently, M is skeletal if there exist permutation matrices $P \in \mathbb{R}^{p, p}$ and $Q \in \mathbb{R}^{\aleph_0, \aleph_0}$ such that $PMQ = \begin{bmatrix} I_k & \tilde{M} \\ O & O \end{bmatrix}$ for some $k \leq p$ and some $\tilde{M} \in \mathbb{R}^{k, \aleph_0}$. The name indicates that the “body” of a skeletal matrix (the nonzero rows) is completely supported by a “skeleton” (the identity block).

We are now ready to define the minion \mathcal{C} . The L -ary objects of \mathcal{C} are pairs $(M, \boldsymbol{\mu})$, where M is a skeletal matrix of size $L \times \aleph_0$ and $\boldsymbol{\mu}$ is an affine vector (i.e., an integer vector whose entries sum up to one) of size L . We require that every column of M should be stochastic and M should have only finitely many different columns; the latter is formalized in (c_5) in Definition 3.2, which says that starting from some

point all the columns are equal. We also require a particular relationship between M and μ formalised in (c_4) .

DEFINITION 3.2. For $L \in \mathbb{N}$, let $\mathcal{C}^{(L)}$ be the set of pairs (M, μ) such that $M \in \mathbb{Q}^{L, \mathbb{N}_0}$, $\mu \in \mathbb{Z}^L$, and the following requirements are met:

- (c_1) M is entrywise nonnegative; (c_4) $\text{supp}(\mu) \subseteq \text{supp}(M\mathbf{e}_1)$;
- (c_2) $\mathbf{1}_L^T M = \mathbf{1}_{\mathbb{N}_0}^T$; (c_5) $\exists t \in \mathbb{N}$ such that $M\mathbf{e}_i = M\mathbf{e}_t \quad \forall i \geq t$;
- (c_3) $\mathbf{1}_L^T \mu = 1$; (c_6) M is skeletal.

We define \mathcal{C} as the disjoint union of L -ary parts, $\mathcal{C} := \bigcup_{L \geq 1} \mathcal{C}^{(L)}$.

We defined \mathcal{C} as a set. For \mathcal{C} to be a minion, we need to define the minor operation on \mathcal{C} and verify that it preserves the structure of \mathcal{C} . This is easy and done in section 4.1.

Our first result is the following characterization of the power of CLAP.

THEOREM 3.3. Let (\mathbf{A}, \mathbf{B}) be a PCSP template. Then, CLAP solves PCSP (\mathbf{A}, \mathbf{B}) if and only if there is a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$.

H-symmetry. Our second main result is a sufficient condition on a PCSP template (\mathbf{A}, \mathbf{B}) to guarantee that CLAP solves PCSP (\mathbf{A}, \mathbf{B}) . The condition is through symmetries satisfied by polymorphisms of the template. In particular, in Theorem 3.5 we will show that if $\text{Pol}(\mathbf{A}, \mathbf{B})$ contains infinitely many operations that are “ H -symmetric” for a suitable matrix H , then there is a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$, and thus CLAP solves PCSP (\mathbf{A}, \mathbf{B}) by Theorem 4.10.

In order to define the notion of H -symmetry, we need a few auxiliary definitions. A vector $\mathbf{w} = (w_i) \in \mathbb{R}^p$ is *tieless* if, for any two indices $i \neq i' \in [p]$, $w_i \neq 0 \Rightarrow w_i \neq w_{i'}$. A *tie matrix* is a matrix having integer nonnegative entries, each of whose columns is a tieless vector. Given an $m \times p$ tie matrix H , we say that a vector $\mathbf{v} \in \mathbb{R}^p$ is *H -tieless* if $H\mathbf{v}$ is tieless.

Let A be a finite set, let $L \in \mathbb{N}$, and take $\mathbf{a} = (a_1, \dots, a_L) \in A^L$. We define the (multiplicity) vector $\mathbf{a}^\#$ as the integer vector of size $|A|$ whose a th entry is $|\{i \in [L] : a_i = a\}|$ for each $a \in A$.

DEFINITION 3.4. Let A, B be finite sets, and consider a function $f : A^L \rightarrow B$ for some $L \in \mathbb{N}$. Given an $m \times |A|$ tie matrix H , we say that f is *H -symmetric* if

$$f_{/\pi}(\mathbf{a}) = f(\mathbf{a}) \quad \forall \pi : [L] \rightarrow [L] \text{ permutation, } \quad \forall \mathbf{a} \in A^L \text{ such that } \mathbf{a}^\# \text{ is } H\text{-tieless.}$$

Our second result is the following sufficient condition for tractability of PCSPs.

THEOREM 3.5. Let (\mathbf{A}, \mathbf{B}) be a PCSP template and suppose $\text{Pol}(\mathbf{A}, \mathbf{B})$ contains H -symmetric operations of arbitrarily large arity for some $m \times |A|$ tie matrix H , $m \in \mathbb{N}$. Then there exists a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$.

Recall from Definition 3.1 the notion of a skeletal matrix. As will be clear from the rest of this paper, the “skeleton” represents the link between CLAP and the above-defined notion of H -symmetry. Indeed, on the one hand the presence of the identity block in a skeletal matrix captures the fact that each BLP solution computed by CLAP gives probability 1 to some constraint-assignment pair and probability 0 to all other constraint-assignment pairs for the same constraint (cf. line 6 of Algorithm 3.1). On the other hand, Lemma 5.2 (stated and proved in section 5) shows that finitely many skeletal matrices can always be simultaneously reduced to H -tieless probability

distributions—which are exactly the distributions on which H -symmetric functions are symmetric (cf. Definition 3.4).

We now mention some consequences of Theorem 3.5. First, observe that a vector of size 1 is always tieless. Hence, if we take any $1 \times |A|$ integer nonnegative matrix as H , we have that H is a tie matrix and $\mathbf{a}^\#$ is H -tieless for each tuple \mathbf{a} in the domain of f ; therefore, for such an H , f being H -symmetric reduces to f being symmetric. On the other hand, having Definition 3.4 in mind, adding rows to H increases the chance that $H\mathbf{a}^\#$ has some ties, in which case f is released from the requirement of being symmetric on \mathbf{a} . In this sense, H encodes the “exceptions to symmetry” that f is allowed to have: The more rows H has, the stronger Theorem 3.5 becomes. If, for instance, H is the identity matrix of order $|A|$, then an H -symmetric operation needs to be symmetric only on those tuples where each entry occurs with a different multiplicity. A very special example of such an $I_{|A|}$ -symmetric operation is a function f that returns (the homomorphic image of) the *most-frequent* entry in the input tuple whenever it is unique, and, in any other case, f is say (the homomorphic image of) a projection. Other more creative choices for H allow capturing operations having more complex exceptions to symmetry, as shown in Example 3.6.

Theorems 3.3 and 3.5 together establish that the CLAP algorithm solves any PCSP template admitting arbitrarily large polymorphisms having some exceptions to symmetry that can be encoded via a tie matrix.

The importance of the next example lies in the fact that it provably separates CLAP from finite tractability and BLP + AIP; i.e., there are PCSP templates solvable by CLAP that are not finitely tractable and not solvable by BLP + AIP.

Example 3.6. Recall the PCSP template (\mathbf{A}, \mathbf{B}) from Example 2.10, where it was shown that $\text{PCSP}(\mathbf{A}, \mathbf{B})$ is not finitely tractable and not solved by the BLP + AIP algorithm from [22]. We will show that $\text{PCSP}(\mathbf{A}, \mathbf{B})$ is solved by CLAP.

Take $L \in \mathbb{N}$ and consider the function $f : A^L \rightarrow B$ defined as follows: For $\mathbf{a} = (a_1, \dots, a_L) \in A^L$,

- if $\mathbf{a} \in \{0, 1\}^L$, look at $\mathbf{a}_1^\#$, i.e., the multiplicity of $1 \in A$ in the tuple \mathbf{a} ;
 - * if $\mathbf{a}_1^\# < \frac{L}{3}$, set $f(\mathbf{a}) = 0$;
 - * if $\mathbf{a}_1^\# > \frac{L}{3}$, set $f(\mathbf{a}) = 1$;
 - * if $\mathbf{a}_1^\# = \frac{L}{3}$, set $f(\mathbf{a}) = a_1$;
- if $\mathbf{a} \in \{2, 3, 4, 5, 6\}^L$,
 - * if there is a unique element $a \in A$ having maximum multiplicity in \mathbf{a} , set $f(\mathbf{a}) = a$;
 - * if there is more than one element of A having maximum multiplicity in \mathbf{a} , set $f(\mathbf{a}) = a_1$;
- otherwise, set $f(\mathbf{a}) = 0$.⁷

We claim that $f \in \text{Pol}(\mathbf{A}, \mathbf{B})$. To see that f preserves R_1 , consider a tuple $\boldsymbol{\rho} = (\mathbf{r}_1, \dots, \mathbf{r}_L)$ of elements of $R_1^\mathbf{A}$, where $\mathbf{r}_i = (a_i, b_i, c_i)$ for $i \in [L]$. We shall let $\mathbf{a} = (a_1, \dots, a_L)$, $\mathbf{b} = (b_1, \dots, b_L)$, and $\mathbf{c} = (c_1, \dots, c_L)$. Notice that

$$(3.1) \quad \mathbf{a}_1^\# + \mathbf{b}_1^\# + \mathbf{c}_1^\# = L.$$

If $f(\mathbf{a}) = f(\mathbf{b}) = f(\mathbf{c}) = 0$, then $\mathbf{a}_1^\# \leq \frac{L}{3}$, $\mathbf{b}_1^\# \leq \frac{L}{3}$, and $\mathbf{c}_1^\# \leq \frac{L}{3}$; by (3.1), this implies that $\mathbf{a}_1^\# = \mathbf{b}_1^\# = \mathbf{c}_1^\# = \frac{L}{3}$. Hence, $(0, 0, 0) = (f(\mathbf{a}), f(\mathbf{b}), f(\mathbf{c})) = (a_1, b_1, c_1) = \mathbf{r}_1 \in R_1^\mathbf{A}$,

⁷Assigning any value in $\{0, \dots, 6\}$ to $f(\mathbf{a})$ would work here.

a contradiction. Similarly, $f(\mathbf{a}) = f(\mathbf{b}) = f(\mathbf{c}) = 1$ would yield $\mathbf{a}_1^\# \geq \frac{L}{3}$, $\mathbf{b}_1^\# \geq \frac{L}{3}$, and $\mathbf{c}_1^\# \geq \frac{L}{3}$; again by (3.1), this implies that $\mathbf{a}_1^\# = \mathbf{b}_1^\# = \mathbf{c}_1^\# = \frac{L}{3}$, hence $(1, 1, 1) = (f(\mathbf{a}), f(\mathbf{b}), f(\mathbf{c})) = (a_1, b_1, c_1) = \mathbf{r}_1 \in R_1^\mathbf{A}$, also a contradiction. We conclude that $f(\boldsymbol{\rho}) = (f(\mathbf{a}), f(\mathbf{b}), f(\mathbf{c})) \in R_1^\mathbf{B}$, thus showing that f preserves R_1 .

As for R_2 , let $\boldsymbol{\rho} = (\mathbf{r}_1, \dots, \mathbf{r}_L)$ be a tuple of elements of $R_2^\mathbf{A}$, where $\mathbf{r}_i = (a_i, b_i)$ for $i \in [L]$, and let $\mathbf{a} = (a_1, \dots, a_L)$ and $\mathbf{b} = (b_1, \dots, b_L)$. The directed graph having vertex set $\{2, 3, 4, 5, 6\}$ and edge set $R_2^\mathbf{A} = R_2^\mathbf{B}$ consists of the disjoint union of a directed 2-cycle and a directed 3-cycle and, hence, all of its vertices have in-degree and out-degree one. As a consequence, the multiplicity of a directed edge (a, b) in the tuple $\boldsymbol{\rho}$ equals both the multiplicity of a in \mathbf{a} and the multiplicity of b in \mathbf{b} . Therefore, if the tuple $\boldsymbol{\rho}$ has a unique element $\mathbf{r} = (a, b)$ with maximum multiplicity, then $f(\boldsymbol{\rho}) = (f(\mathbf{a}), f(\mathbf{b})) = (a, b) = \mathbf{r} \in R_2^\mathbf{B}$. Otherwise, $f(\boldsymbol{\rho}) = (a_1, b_1) = \mathbf{r}_1 \in R_2^\mathbf{B}$. This shows that f preserves R_2 , too, and is thus a polymorphism of (\mathbf{A}, \mathbf{B}) .

Consider the matrix $H = \text{diag}(1, 2, 1, 1, 1, 1)$, and observe that H is a tie matrix. We claim that f is H -symmetric. Let $\pi : [L] \rightarrow [L]$ be a permutation, and take a tuple $\mathbf{a} = (a_1, \dots, a_L) \in A^L$ such that $\mathbf{a}^\#$ is H -tieless; i.e., the vector $H\mathbf{a}^\# = (\mathbf{a}_0^\#, 2\mathbf{a}_1^\#, \mathbf{a}_2^\#, \mathbf{a}_3^\#, \mathbf{a}_4^\#, \mathbf{a}_5^\#, \mathbf{a}_6^\#)$ is tieless. Write $\tilde{\mathbf{a}} = (a_{\pi(1)}, \dots, a_{\pi(L)})$, and observe that $\tilde{\mathbf{a}}^\# = \mathbf{a}^\#$.

- If $\mathbf{a} \in \{0, 1\}^L$, we get $\mathbf{a}_0^\# \neq 2\mathbf{a}_1^\#$; since $\mathbf{a}_0^\# + \mathbf{a}_1^\# = L$, this gives $2\mathbf{a}_1^\# \neq L - \mathbf{a}_1^\#$ so that $\mathbf{a}_1^\# \neq \frac{L}{3}$. As a consequence, $f(\mathbf{a}) = f(\tilde{\mathbf{a}})$.
- If $\mathbf{a} \in \{2, 3, 4, 5, 6\}^L$, the condition above implies that the tuple

$$(\mathbf{a}_2^\#, \mathbf{a}_3^\#, \mathbf{a}_4^\#, \mathbf{a}_5^\#, \mathbf{a}_6^\#)$$

has a unique maximum element and, hence, there is a unique element a of A having maximum multiplicity in \mathbf{a} (and in $\tilde{\mathbf{a}}$). Therefore, $f(\mathbf{a}) = a = f(\tilde{\mathbf{a}})$.

- If $\mathbf{a} \notin \{0, 1\}^L \cup \{2, 3, 4, 5, 6\}^L$, then $f(\mathbf{a}) = 0 = f(\tilde{\mathbf{a}})$.

We conclude that, in each case, $f(\mathbf{a}) = f(\tilde{\mathbf{a}}) = f_{/\pi}(\mathbf{a})$, which means that f is H -symmetric. By Theorems 3.3 and 3.5, CLAP solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$.

Remark 3.7. Consider the minion $\mathcal{M}_{\text{BLP}+\text{AIP}}$ from [22] (cf. Appendix A.3). A direct consequence of Example 3.6, Theorem 3.3, and [22, Lemma 5.4] is that there is no minion homomorphism from $\mathcal{M}_{\text{BLP}+\text{AIP}}$ to \mathcal{C} . On the other hand, the function

$$\begin{aligned} \vartheta : \mathcal{C} &\rightarrow \mathcal{M}_{\text{BLP}+\text{AIP}} \\ (M, \boldsymbol{\mu}) &\mapsto (M\mathbf{e}_1, \boldsymbol{\mu}) \end{aligned}$$

is readily seen to be a minion homomorphism. It follows that CLAP solves any PCSP template solved by $\text{BLP} + \text{AIP}$ (as is also clear from the description of the two algorithms).

Remark 3.8. Similar to [22], the assumption in Theorem 3.5 can be weakened as follows: Instead of requiring H -symmetric polymorphisms of arbitrarily large arity, it turns out to be enough requiring H -block-symmetric polymorphisms of arbitrarily large width, where the definition of an H -block-symmetric operation mirrors that of a block-symmetric operation in [22]. The proof of this possibly stronger result is very similar to that of Theorem 3.5. For completeness, we include it in Appendix C. We point out that we do not know whether the condition in Theorem 3.5 (or the possibly weaker condition based on H -block-symmetric polymorphisms) is necessary for tractability via CLAP, but we suspect it is not.

Remark 3.9. A possibly stronger version of the CLAP algorithm consists of running $\text{BLP} + \text{AIP}$ (instead of just BLP) at each iteration in the **for** loop in lines 5–9 of

Algorithm 3.1, and then removing the additional **for** loop in lines 14–18. This algorithm can be called C(BLP+AIP). An analysis entirely analogous to the one presented in this paper shows that the power of C(BLP + AIP) is captured by the minion $\tilde{\mathcal{C}}$ defined like \mathcal{C} with the following difference: The L -ary elements of $\tilde{\mathcal{C}}$ are pairs (M, N) , where M is as in \mathcal{C} while N is an integer matrix of the same size as M taking the role of μ (in particular, N satisfies the “refinement condition” $\text{supp}(Ne_i) \subseteq \text{supp}(Me_i) \forall i \in \mathbb{N}$, analogous to (c_4) in Definition 3.2). A possible direction for future research is to investigate whether the richer structure of $\tilde{\mathcal{C}}$ can be exploited to obtain a stronger version of Theorem 3.5.

Remark 3.10. For CSPs, the characterization of bounded width [9, 25] and its collapse [6] was preceded by a characterization of width-1 CSPs [37, 41] and the collapse of width 2 to width 1 [34]. Thus the difference between width-1 CSPs and bounded-width CSPs is well understood. BLP and SBLP are the (convex relaxation) analogues of width 1 and SAC, respectively, and SAC solves all bounded-width CSPs [56]. Therefore, a natural question is whether a similar analysis can cast light on the difference in power between BLP on one side, and SBLP (and thus perhaps also of CBLP and CLAP) on the other side. We remark on two obstacles: First, BLP is strictly more powerful than width 1 for CSPs [60]. Second, a good characterization of the power of SBLP (and stronger algorithms studied in the present paper) would imply that these algorithms solve, in the special case of CSPs, all bounded-width CSPs—a nontrivial result implied by [56].

4. The power of the CLAP algorithm. The goal of this section is to prove Theorem 3.3. In section 4.1, we will verify that \mathcal{C} , which appears in the statement of Theorem 3.3, is indeed a minion. In sections 4.2 and 4.3, we will establish a compactness argument and present a condition that captures CLAP, respectively; both will be needed in the proof of Theorem 3.3. The two directions of Theorem 3.3 will be then proved in section 4.4.

Minions are not only useful for capturing the complexity of PCSPs but also for characterizing the power of algorithms. This will be done by using the concept of the free structure generated, for a given minion, by a relational structure [22] (cf. [8, Definition 4.1] for the definition in the special case of minions of functions).

DEFINITION 4.1. Let \mathcal{M} be a minion, and let \mathbf{A} be a (finite) relational structure with signature σ . The free structure $\mathbb{F}_{\mathcal{M}}(\mathbf{A})$ is a relational structure with domain $\mathcal{M}^{(|A|)}$ (potentially infinite) and signature σ . Given a relation $R \in \sigma$ of arity k , a tuple (M_1, \dots, M_k) of elements of $\mathcal{M}^{(|A|)}$ belongs to $R^{\mathbb{F}_{\mathcal{M}}(\mathbf{A})}$ if and only if there is some $Q \in \mathcal{M}^{(|R^{\mathbf{A}}|)}$ such that $M_i = Q_{/\pi_i}$ for each $i \in [k]$, where $\pi_i : R^{\mathbf{A}} \rightarrow A$ maps $\mathbf{a} \in R^{\mathbf{A}}$ to its i th coordinate a_i .

The next result will be useful to establish the connection between our algorithm CLAP, presented in section 3, and the minion \mathcal{C} .

LEMMA 4.2. Let \mathcal{M} be a minion, and let (\mathbf{A}, \mathbf{B}) be a PCSP template. Then there is a minion homomorphism from \mathcal{M} to $\text{Pol}(\mathbf{A}, \mathbf{B})$ if and only if $\mathbb{F}_{\mathcal{M}}(\mathbf{A}) \rightarrow \mathbf{B}$.

The proof of Lemma 4.2 is based on that of [8, Lemma 4.4], which proves one-to-one correspondence but only for minions of functions. For completeness, we prove Lemma 4.2 in Appendix B.

4.1. \mathcal{C} is a minion. The minor operation on \mathcal{C} is naturally defined via a matrix multiplication with a matrix that encodes the minor map. For a function $\pi : [L] \rightarrow [L']$, let P_π be the $L' \times L$ matrix whose (i, j) th entry is 1 if $\pi(j) = i$, and 0 otherwise. Note that $P_\pi^T \mathbf{1}_{L'} = \mathbf{1}_L$ and, for each $i \in [L']$, $P_\pi^T \mathbf{e}_i = \sum_{j \in \pi^{-1}(i)} \mathbf{e}_j$.

DEFINITION 4.3. For $(M, \boldsymbol{\mu}) \in \mathcal{C}^{(L)}$, we define $M_{/\pi} = P_{\pi}M$ and $\boldsymbol{\mu}_{/\pi} = P_{\pi}\boldsymbol{\mu}$, and we let the minor of $(M, \boldsymbol{\mu})$ with respect to π be $(M, \boldsymbol{\mu})_{/\pi} := (M_{/\pi}, \boldsymbol{\mu}_{/\pi})$.

We remark that this definition is consistent with the minions $\mathcal{Q}_{\text{conv}}$ and \mathcal{X}_{aff} studied in [8], and the minion $\mathcal{M}_{\text{BLP}+\text{AIP}}$ studied in [22]; cf. Appendices A.1, A.2, and A.3.

PROPOSITION 4.4. \mathcal{C} is a minion.

Proof. Write $M = [m_{ij}]$ and $\boldsymbol{\mu} = (\mu_i)$. Observe that $M_{/\pi} \in \mathbb{Q}^{L', \aleph_0}$ and $\boldsymbol{\mu}_{/\pi} \in \mathbb{Z}^{L'}$. The requirements $(c_1), (c_2), (c_3)$, and (c_5) are trivially satisfied by $(M, \boldsymbol{\mu})_{/\pi}$. As for (c_4) , suppose that $\mathbf{e}_i^T P_{\pi} M \mathbf{e}_1 = 0$ but $\mathbf{e}_i^T P_{\pi} \boldsymbol{\mu} \neq 0$. It follows that $\mu_j \neq 0$ for some $j \in \pi^{-1}(i)$. Hence, $m_{j1} > 0$, and then,

$$\mathbf{e}_i^T P_{\pi} M \mathbf{e}_1 = \sum_{j' \in \pi^{-1}(i)} \mathbf{e}_{j'}^T M \mathbf{e}_1 \geq \mathbf{e}_j^T M \mathbf{e}_1 > 0,$$

which is a contradiction. We now show that $M_{/\pi}$ is skeletal. Choose $j \in [L']$, and suppose that $\mathbf{e}_j^T M_{/\pi} \neq \mathbf{0}_{\aleph_0}^T$. We obtain

$$\mathbf{0}_{\aleph_0} \neq M^T P_{\pi}^T \mathbf{e}_j = \sum_{\ell \in \pi^{-1}(j)} M^T \mathbf{e}_{\ell}$$

and, in particular, $\exists \ell \in \pi^{-1}(j)$ such that $\mathbf{e}_{\ell}^T M \neq \mathbf{0}_{\aleph_0}^T$. Since M is skeletal, this implies that $M \mathbf{e}_i = \mathbf{e}_{\ell}$ for some $i \in \mathbb{N}$. This yields

$$M_{/\pi} \mathbf{e}_i = P_{\pi} M \mathbf{e}_i = P_{\pi} \mathbf{e}_{\ell} = \mathbf{e}_{\pi(\ell)} = \mathbf{e}_j$$

as required. Hence, (c_6) is satisfied, too, and $(M, \boldsymbol{\mu})_{/\pi} \in \mathcal{C}^{(L')}$.

Finally, considering $\tilde{\pi}: [L'] \rightarrow [L'']$ and the identity map $\text{id}: [L] \rightarrow [L]$, one readily checks that $P_{\tilde{\pi} \circ \pi} = P_{\tilde{\pi}} P_{\pi}$ and $P_{\text{id}} = I_L$. Hence, the minor operations defined above satisfy the requirements of Definition 2.6. \square

4.2. A compactness argument for \mathcal{C} . The set $\mathcal{C}^{(L)}$ of the L -ary objects in \mathcal{C} is infinite unless $L = 1$. As a consequence, given a relational structure \mathbf{A} whose domain has size at least 2, the free structure $\mathbb{F}_{\mathcal{C}}(\mathbf{A})$ has an infinite domain. We now describe a standard compactness argument analogous to [8, Remark 7.13] that will circumvent this inconvenience.

For $D, L \in \mathbb{N}$, consider the set

$$\mathcal{C}_D^{(L)} = \{(M, \boldsymbol{\mu}) \in \mathcal{C}^{(L)} : DM \text{ is entrywise integer,} \\ M \mathbf{e}_i = M \mathbf{e}_D \quad \forall i \geq D, \text{ and } \mathbf{1}_L^T |\boldsymbol{\mu}| \leq D\},$$

where $|\boldsymbol{\mu}|$ denotes the vector whose entries are the absolute values of the entries of $\boldsymbol{\mu}$. Since $\mathcal{C}_D^{(L)}$ is unambiguously determined by $L \times (D+1)$ integer numbers belonging to the set $\{-D, \dots, D\}$, it is finite. Observe that the set $\mathcal{C}_D = \bigcup_{L \in \mathbb{N}} \mathcal{C}_D^{(L)}$ is closed under taking minors. Indeed, given $(M, \boldsymbol{\mu}) \in \mathcal{C}_D^{(L)}$ and $\pi: [L] \rightarrow [L']$, $DP_{\pi}M = P_{\pi}DM$ is entrywise integer, $P_{\pi}M \mathbf{e}_i = P_{\pi}M \mathbf{e}_D \quad \forall i \geq D$, and $\mathbf{1}_{L'}^T |P_{\pi}\boldsymbol{\mu}| \leq \mathbf{1}_L^T |P_{\pi}\boldsymbol{\mu}| = \mathbf{1}_L^T |\boldsymbol{\mu}| \leq D$. Hence, \mathcal{C}_D is a subminion of \mathcal{C} . Observe also that $\mathcal{C} = \bigcup_{D \in \mathbb{N}} \mathcal{C}_D$. To see this, take $(M, \boldsymbol{\mu}) \in \mathcal{C}^{(L)}$ and suppose that $M \mathbf{e}_i = M \mathbf{e}_t \quad \forall i \geq t$. Let \tilde{D} be a common denominator of the finite set of rational numbers $\{m_{ij} : i \in [L], j \in [t]\}$, so that $\tilde{D}M$ is entrywise integer. Also let $\tilde{D} = \mathbf{1}_L^T |\boldsymbol{\mu}|$. Then, $(M, \boldsymbol{\mu}) \in \mathcal{C}_{\tilde{D}}^{(L)}$.

PROPOSITION 4.5. *Let \mathcal{M} be a minion such that $\mathcal{M}^{(L)}$ is finite for each $L \in \mathbb{N}$, and suppose that there exist minion homomorphisms $\xi_D : \mathcal{C}_D \rightarrow \mathcal{M}$ for each $D \in \mathbb{N}$. Then there exists a minion homomorphism $\zeta : \mathcal{C} \rightarrow \mathcal{M}$.*

Proof. For $D \in \mathbb{N}$, let $\mathcal{C}_{\{D\}} = \bigcup_{L \leq D} \mathcal{C}_{D!}^{(L)}$. Observe that $\mathcal{C}_{\{D\}}$ is a finite set and $\mathcal{C}_{\{D\}} \subseteq \mathcal{C}_{\{D+1\}}$. Moreover, $\bigcup_{D \in \mathbb{N}} \mathcal{C}_{\{D\}} = \bigcup_{D \in \mathbb{N}} \mathcal{C}_D = \mathcal{C}$. Indeed, given $D' \in \mathbb{N}$, we have that $\mathcal{C}_{\{D'\}} \subseteq \mathcal{C}_{D'!} \subseteq \bigcup_{D \in \mathbb{N}} \mathcal{C}_D$, and given $L \in \mathbb{N}$, $\mathcal{C}_{D'}^{(L)} \subseteq \mathcal{C}_{(D'L)!}^{(L)} \subseteq \mathcal{C}_{\{D'L\}} \subseteq \bigcup_{D \in \mathbb{N}} \mathcal{C}_{\{D\}}$. Consider an infinite rooted tree whose vertices are all the restrictions of the homomorphisms ξ_D to some $\mathcal{C}_{\{D'\}}$, whose root is the empty mapping, and the parent of a vertex corresponding to a function $\mathcal{C}_{\{D'+1\}} \rightarrow \mathcal{M}$ is the vertex corresponding to the restriction of the function to $\mathcal{C}_{\{D'\}}$. This is an infinite connected tree. Moreover, since $\mathcal{M}^{(L)}$ is finite for each $L \in \mathbb{N}$ and since minion homomorphisms preserve the arities, there exist only finitely many distinct restrictions of minion homomorphisms to $\mathcal{C}_{\{D\}}$; hence, the tree is locally finite. By König's lemma, it contains an infinite path, which corresponds to an infinite chain of maps $\zeta_i : \mathcal{C}_{\{i\}} \rightarrow \mathcal{M}$ such that ζ_{i+1} extends $\zeta_i \forall i \in \mathbb{N}$. Their union $\zeta : \mathcal{C} \rightarrow \mathcal{M}$ is then a minion homomorphism. \square

4.3. The CLAP condition. Given a finite set C , consider the set $\mathbb{S}(C)$ of the rational stochastic vectors of size $|C|$. Let $U \subseteq C^k$. For $i \in [k]$, consider the $|C| \times |U|$ matrix $E^{(U,i)}$ such that, for $c \in C$ and $\mathbf{c} = (c_1, \dots, c_k) \in U$, the (c, \mathbf{c}) th entry of $E^{(U,i)}$ is 1 if $c_i = c$, and 0 otherwise. Given $\xi \in \mathbb{S}(U)$ and $i \in [k]$, we define the i th *marginal* of ξ as

$$\xi^{(i)} = E^{(U,i)} \xi.$$

Observe that

$$\xi^{(i)T} \mathbf{1}_{|C|} = \xi^T E^{(U,i)T} \mathbf{1}_{|C|} = \xi^T \mathbf{1}_{|U|} = 1,$$

so that $\xi^{(i)} \in \mathbb{S}(C)$. We also define the set $\mathbb{Z}(C)$ of the integer vectors of size $|C|$ whose entries sum up to 1. Given $U \subseteq C^k$, $\zeta \in \mathbb{Z}(U)$, and $i \in [k]$, we define

$$\zeta^{(i)} = E^{(U,i)} \zeta.$$

As before, observe that

$$\zeta^{(i)T} \mathbf{1}_{|C|} = \zeta^T E^{(U,i)T} \mathbf{1}_{|C|} = \zeta^T \mathbf{1}_{|U|} = 1,$$

so $\zeta^{(i)} \in \mathbb{Z}(C)$.

Let \mathbf{A} be a relational structure having domain A and signature σ . We define the relational structures $\mathbb{S}(\mathbf{A})$ and $\mathbb{Z}(\mathbf{A})$ as follows:

- $\mathbb{S}(\mathbf{A})$ has domain $\mathbb{S}(A)$ and for every symbol $R \in \sigma$ of arity k , $R^{\mathbb{S}(\mathbf{A})} = \{(\xi^{(1)}, \dots, \xi^{(k)}) : \xi \in \mathbb{S}(R^{\mathbf{A}})\}$;
- $\mathbb{Z}(\mathbf{A})$ has domain $\mathbb{Z}(A)$ and for every symbol $R \in \sigma$ of arity k , $R^{\mathbb{Z}(\mathbf{A})} = \{(\zeta^{(1)}, \dots, \zeta^{(k)}) : \zeta \in \mathbb{Z}(R^{\mathbf{A}})\}$.

Remark 4.6. $\mathbb{S}(\mathbf{A})$ and $\mathbb{Z}(\mathbf{A})$ are denoted by $\text{LP}(\mathbf{A})$ and $\text{IP}(\mathbf{A})$ in [8], respectively. As noted in [8, Remarks 7.11 and 7.21], $\mathbb{S}(\mathbf{A})$ coincides with the free structure of the minion $\mathcal{Q}_{\text{conv}}$ generated by \mathbf{A} , and similarly, $\mathbb{Z}(\mathbf{A})$ is the free structure of the minion \mathcal{Z}_{aff} generated by \mathbf{A} . (See Appendices A.1 and A.2 for the definitions of $\mathcal{Q}_{\text{conv}}$ and \mathcal{Z}_{aff} , respectively.) In particular, given a relational structure \mathbf{X} with signature σ ,

BLP accepts \mathbf{X} as an instance of $\text{CSP}(\mathbf{A})$ if and only if $\mathbf{X} \rightarrow \mathbb{S}(\mathbf{A})$; similarly, AIP accepts \mathbf{X} as an instance of $\text{CSP}(\mathbf{A})$ if and only if $\mathbf{X} \rightarrow \mathbb{Z}(\mathbf{A})$.

Remark 4.7. The assignment $f : a \mapsto \mathbf{e}_a$ for each $a \in A$ yields both a canonical homomorphism from \mathbf{A} to $\mathbb{S}(\mathbf{A})$ and a canonical homomorphism from \mathbf{A} to $\mathbb{Z}(\mathbf{A})$. Indeed, for $R \in \sigma$ of arity k and $\mathbf{a} = (a_1, \dots, a_k) \in R^{\mathbf{A}}$,

$$f(\mathbf{a}) = (\mathbf{e}_{a_1}, \dots, \mathbf{e}_{a_k}) = (E^{(R^{\mathbf{A}}, 1)} \mathbf{e}_{\mathbf{a}}, \dots, E^{(R^{\mathbf{A}}, k)} \mathbf{e}_{\mathbf{a}})$$

which belongs to both $R^{\mathbb{S}(\mathbf{A})}$ and $R^{\mathbb{Z}(\mathbf{A})}$ since $\mathbf{e}_{\mathbf{a}} \in \mathbb{S}(R^{\mathbf{A}}) \cap \mathbb{Z}(R^{\mathbf{A}})$.

In Proposition 4.9, we characterize the instances of a given PCSP template for which the CLAP algorithm returns Yes in terms of the condition described in the following definition.

DEFINITION 4.8. Let (\mathbf{A}, \mathbf{B}) be a PCSP template, where \mathbf{A} and \mathbf{B} have signature σ . Given an instance \mathbf{X} of $\text{PCSP}(\mathbf{A}, \mathbf{B})$, we say that \mathbf{X} has the CLAP condition if the following holds: $\forall R \in \sigma$ of arity k $\exists s^R : R^{\mathbf{X}} \rightarrow \mathcal{P}(R^{\mathbf{A}}) \setminus \{\emptyset\}$ such that

(I) $\forall \mathbf{x} = (x_1, \dots, x_k) \in R^{\mathbf{X}}, \forall \mathbf{a} = (a_1, \dots, a_k) \in s^R(\mathbf{x})$ there is a homomorphism $h_{\mathbf{x}, \mathbf{a}} : \mathbf{X} \rightarrow \mathbb{S}(\mathbf{A})$ that satisfies the following:

1. $h_{\mathbf{x}, \mathbf{a}}(x_i) = \mathbf{e}_{a_i} \quad \forall i \in [k];$
2. $\forall \tilde{R} \in \sigma$ of arity \tilde{k} , $\forall \tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_{\tilde{k}}) \in \tilde{R}^{\mathbf{X}} \exists \xi \in \mathbb{S}(\tilde{R}^{\mathbf{A}})$ such that
 - * $h_{\mathbf{x}, \mathbf{a}}(\tilde{x}_i) = E^{(\tilde{R}^{\mathbf{A}}, i)} \xi \quad \forall i \in [\tilde{k}];$
 - * $\text{supp}(\xi) \subseteq s^{\tilde{R}}(\tilde{\mathbf{x}}).$

(II) $\exists \bar{R} \in \sigma, \bar{\mathbf{x}} \in \bar{R}^{\mathbf{X}}, \bar{\mathbf{a}} \in s^{\bar{R}}(\bar{\mathbf{x}})$ such that there is a homomorphism $g : \mathbf{X} \rightarrow \mathbb{Z}(\mathbf{A})$ that satisfies the following:

- 1'. $\forall \tilde{R} \in \sigma$ of arity \tilde{k} , $\forall \tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_{\tilde{k}}) \in \tilde{R}^{\mathbf{X}} \exists \xi \in \mathbb{S}(\tilde{R}^{\mathbf{A}}), \exists \zeta \in \mathbb{Z}(\tilde{R}^{\mathbf{A}})$ such that
 - * $h_{\bar{\mathbf{x}}, \bar{\mathbf{a}}}(\tilde{x}_i) = E^{(\tilde{R}^{\mathbf{A}}, i)} \xi \quad \forall i \in [\tilde{k}];$
 - * $g(\tilde{x}_i) = E^{(\tilde{R}^{\mathbf{A}}, i)} \zeta \quad \forall i \in [\tilde{k}];$
 - * $\text{supp}(\zeta) \subseteq \text{supp}(\xi) \subseteq s^{\tilde{R}}(\tilde{\mathbf{x}}).$

PROPOSITION 4.9. Given an instance \mathbf{X} of $\text{PCSP}(\mathbf{A}, \mathbf{B})$, CLAP accepts \mathbf{X} if and only if \mathbf{X} has the CLAP condition.

Proof. Suppose that CLAP accepts \mathbf{X} , and let $\{S_{\mathbf{x}, R} : R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}\}$ be the family of sets generated by the algorithm at termination. For each $R \in \sigma$, consider the map $s^R : R^{\mathbf{X}} \rightarrow \mathcal{P}(R^{\mathbf{A}}) \setminus \{\emptyset\}$ defined by $s^R(\mathbf{x}) = S_{\mathbf{x}, R}$. For each $\mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in s^R(\mathbf{x})$, consider the corresponding solution to $\text{BLP}(\mathbf{X}, \mathbf{A})$ generated by the algorithm. Letting w_x be the probability distribution on A associated with $x \in X$ in the linear program, we observe that the assignment $x \mapsto w_x$ yields a homomorphism (call it $h_{\mathbf{x}, \mathbf{a}}$) from \mathbf{X} to $\mathbb{S}(\mathbf{A})$ that satisfies the requirement 1. Moreover, letting ξ be the probability distribution associated with a constraint $\tilde{\mathbf{x}} \in \tilde{R}^{\mathbf{X}}$ for some $\tilde{R} \in \sigma$, observe that $h_{\mathbf{x}, \mathbf{a}}$ also satisfies the requirement 2. Finally, let $\bar{R} \in \sigma, \bar{\mathbf{x}} \in \bar{R}^{\mathbf{X}}, \bar{\mathbf{a}} \in S_{\bar{\mathbf{x}}, \bar{R}}$ be such that the condition in the **if** statement of line 15 of Algorithm 3.1 is met. Then 1' follows from the description of $\text{BLP} + \text{AIP}$.

The converse implication follows almost analogously, except for the following subtlety. The $\text{BLP} + \text{AIP}$ algorithm requires that the BLP solution should be picked from the *relative interior* of the polytope of the feasible solutions (cf. Algorithm A.1 in Appendix A.3). However, the homomorphism $h_{\bar{\mathbf{x}}, \bar{\mathbf{a}}}$, whose existence witnesses part (II) of the CLAP condition may correspond to a BLP solution that is not in the

relative interior of the polytope P of the feasible solutions of $\text{BLP}(\mathbf{X}, \mathbf{A})$ satisfying $\lambda_{\bar{\mathbf{x}}, \bar{\mathbf{a}}}(\bar{\mathbf{a}}) = 1$ and $\lambda_{\mathbf{x}', R'}(\mathbf{a}') = 0$ for every $R' \in \sigma$, $\mathbf{x}' \in R'^{\mathbf{X}}$, and $\mathbf{a}' \notin S_{\mathbf{x}', R'}$. If that is the case, the algorithm would not consider $(h_{\bar{\mathbf{x}}, \bar{\mathbf{a}}}, g)$ as a solution for $\text{BLP} + \text{AIP}$. However, letting h' be a solution in the relative interior of P , the conditions (I) and (II) of CLAP are still satisfied if we let h' replace $h_{\bar{\mathbf{x}}, \bar{\mathbf{a}}}$; and, in this case, the homomorphisms witnessing the CLAP condition do correspond to solutions found by the CLAP algorithm.⁸ Hence, CLAP accepts \mathbf{X} . \square

4.4. Proof of Theorem 3.3. Our first goal is to prove the following.

THEOREM 4.10. *If there is a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$, then CLAP solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$.*

Proof. Let \mathbf{X} be an instance of $\text{PCSP}(\mathbf{A}, \mathbf{B})$.

First, we show that if $\mathbf{X} \rightarrow \mathbf{A}$, then CLAP accepts \mathbf{X} , which is the easy direction. Consider a homomorphism $f: \mathbf{X} \rightarrow \mathbf{A}$. Given $R \in \sigma$ of arity k and $\mathbf{x} = (x_1, \dots, x_k) \in R^{\mathbf{X}}$, let $s^R(\mathbf{x}) = \{f(\mathbf{x})\}$. For $\mathbf{x} \in R^{\mathbf{X}}$ and $\mathbf{a} = (a_1, \dots, a_k) = f(\mathbf{x}) \in s^R(\mathbf{x})$, let $h_{\mathbf{x}, \mathbf{a}}: \mathbf{X} \rightarrow \mathbb{S}(\mathbf{A})$ be the homomorphism obtained by composing f with the canonical homomorphism from \mathbf{A} to $\mathbb{S}(\mathbf{A})$ of Remark 4.7—i.e., $h_{\mathbf{x}, \mathbf{a}}(x) = \mathbf{e}_{f(x)} \forall x \in X$. Observe that $h_{\mathbf{x}, \mathbf{a}}(x_i) = \mathbf{e}_{f(x_i)} = \mathbf{e}_{a_i}$ for any $i \in [k]$ and, given $\tilde{R} \in \sigma$ of arity \tilde{k} and $\tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_{\tilde{k}}) \in \tilde{R}^{\mathbf{X}}$, setting $\xi = \mathbf{e}_{f(\tilde{\mathbf{x}})}$ yields $h_{\mathbf{x}, \mathbf{a}}(\tilde{x}_i) = \mathbf{e}_{f(\tilde{x}_i)} = E^{(\tilde{R}^{\mathbf{A}}, i)} \mathbf{e}_{f(\tilde{\mathbf{x}})} = E^{(\tilde{R}^{\mathbf{A}}, i)} \xi$ for any $i \in [\tilde{k}]$, and $\text{supp}(\xi) = \text{supp}(\mathbf{e}_{f(\tilde{\mathbf{x}})}) = \{f(\tilde{\mathbf{x}})\} = s^{\tilde{R}}(\tilde{\mathbf{x}})$. This shows that part (I) of Definition 4.8 is satisfied. As for part (II), choose any $\bar{R} \in \sigma$ and $\bar{\mathbf{x}} \in \bar{R}^{\mathbf{X}}$, let $\bar{\mathbf{a}} = f(\bar{\mathbf{x}})$, and consider the homomorphism $g: \mathbf{X} \rightarrow \mathbb{Z}(\mathbf{A})$ obtained by composing f with the canonical homomorphism from \mathbf{A} to $\mathbb{Z}(\mathbf{A})$ of Remark 4.7—i.e., $g(x) = \mathbf{e}_{f(x)} \forall x \in X$. Given $\tilde{R} \in \sigma$ of arity \tilde{k} and $\tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_{\tilde{k}}) \in \tilde{R}^{\mathbf{X}}$, setting $\xi = \zeta = \mathbf{e}_{f(\tilde{\mathbf{x}})}$ yields $g(\tilde{x}_i) = h_{\bar{\mathbf{x}}, \bar{\mathbf{a}}}(\tilde{x}_i) = \mathbf{e}_{f(\tilde{x}_i)} = E^{(\tilde{R}^{\mathbf{A}}, i)} \mathbf{e}_{f(\tilde{\mathbf{x}})} = E^{(\tilde{R}^{\mathbf{A}}, i)} \xi = E^{(\tilde{R}^{\mathbf{A}}, i)} \zeta$ for any $i \in [\tilde{k}]$, and $\text{supp}(\zeta) = \text{supp}(\xi) = \{f(\tilde{\mathbf{x}})\} = s^{\tilde{R}}(\tilde{\mathbf{x}})$. It follows that \mathbf{X} has the CLAP condition. By Proposition 4.9, CLAP accepts \mathbf{X} .

Second, we show that if \mathbf{X} is accepted by CLAP, then $\mathbf{X} \rightarrow \mathbf{B}$. So, suppose that \mathbf{X} is accepted by CLAP. By Proposition 4.9, \mathbf{X} has the CLAP condition. Using the terminology of Definition 4.8, consider the set $\{h_1, \dots, h_t\} = \{h_{\mathbf{x}, \mathbf{a}}: R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in s^R(\mathbf{x})\}$, where each $h_{\mathbf{x}, \mathbf{a}}$ is a homomorphism from \mathbf{X} to $\mathbb{S}(\mathbf{A})$ described in part (I) of Definition 4.8. We also consider the homomorphism $g: \mathbf{X} \rightarrow \mathbb{Z}(\mathbf{A})$ of part (II) of Definition 4.8, corresponding to $\bar{R} \in \sigma, \bar{\mathbf{x}} \in \bar{R}^{\mathbf{X}}, \bar{\mathbf{a}} \in s^{\bar{R}}(\bar{\mathbf{x}})$. Without loss of generality, we set $h_1 = h_{\bar{\mathbf{x}}, \bar{\mathbf{a}}}$.

Let $n = |\mathbf{A}|$. Given $x \in X$, consider the matrix $M_x \in \mathbb{Q}^{n, n_0}$ and the vector $\mu_x \in \mathbb{Z}^n$ defined by

$$\begin{aligned} M_x \mathbf{e}_i &= h_i(x) \quad \forall i \in [t], \\ M_x \mathbf{e}_i &= h_t(x) \quad \forall i \in \mathbb{N} \setminus [t], \\ \mu_x &= g(x). \end{aligned}$$

We claim that $(M_x, \mu_x) \in \mathcal{C}^{(n)}$. The requirements $(c_1), (c_2), (c_3)$, and (c_5) in Definition 3.2 are easily seen to be satisfied. To check that M_x is skeletal, take $a \in \mathbf{A}$ and suppose that $\mathbf{e}_a^T M_x \neq \mathbf{0}_{n_0}^T$. This means that $\mathbf{e}_a^T M_x \mathbf{e}_d \neq 0$ for some $d \in [t]$.

⁸Another way to phrase this is by saying that the existence of a pair (h, g) of homomorphisms such that each variable for g is zero whenever the corresponding variable for h is zero is equivalent to the existence, for *any* h' in the nonempty relative interior of the polytope of solutions of the BLP, of a solution g' of AIP that sets to zero any variable that is zero in h' . This is implicit in the analysis in [22].

Hence, $a \in \text{supp}(M_x \mathbf{e}_d) = \text{supp}(h_d(x))$. Recall that we are assuming (with no loss of generality) that the signature σ of \mathbf{X} , \mathbf{A} , and \mathbf{B} contains a unary symbol R_u such that $R_u^{\mathbf{X}} = X$, $R_u^{\mathbf{A}} = A$, and $R_u^{\mathbf{B}} = B$. Notice that $E^{(R_u^{\mathbf{A}}, 1)} = I_n$. From part (I) of Definition 4.8, we deduce that $\text{supp}(h_d(x)) \subseteq s^{R_u}(x)$, and hence, $a \in s^{R_u}(x)$. We can then take the homomorphism $h_i = h_{x,a}$, which satisfies $h_i(x) = \mathbf{e}_a$, that is, $M_x \mathbf{e}_i = \mathbf{e}_a$. So, M_x is skeletal and (c_6) is satisfied. Finally, to check (c_4) , choose $a \in A$ and suppose that $\mathbf{e}_a^T M_x \mathbf{e}_1 = 0$. Since $M_x \mathbf{e}_1 = h_1(x) = h_{\tilde{x}, \tilde{a}}(x)$, this implies that $a \notin \text{supp}(h_{\tilde{x}, \tilde{a}}(x))$. Choosing R_u as \tilde{R} and x as \tilde{x} in 1' of Definition 4.8, and using again the fact that $E^{(R_u^{\mathbf{A}}, 1)} = I_n$, we see that $\text{supp}(g(x)) \subseteq \text{supp}(h_{\tilde{x}, \tilde{a}}(x))$. Therefore, $a \notin \text{supp}(g(x)) = \text{supp}(\boldsymbol{\mu}_x)$. Hence, (c_4) is satisfied, too, and the claim is proved.

Consider the map $\gamma : X \rightarrow \mathcal{C}^{(n)}$ defined by $x \mapsto (M_x, \boldsymbol{\mu}_x)$. We claim that γ is a homomorphism from \mathbf{X} to $\mathbb{F}_{\mathcal{C}}(\mathbf{A})$. With this claim, we can finish the proof. By assumption, there is a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$. By Lemma 4.2 applied to \mathcal{C} , we have $\mathbb{F}_{\mathcal{C}}(\mathbf{A}) \rightarrow \mathbf{B}$. Composing γ with this homomorphism yields $\mathbf{X} \rightarrow \mathbf{B}$, as required. It remains to establish the claim.

Claim. γ is a homomorphism from \mathbf{X} to $\mathbb{F}_{\mathcal{C}}(\mathbf{A})$.

Take $R \in \sigma$ of arity k , and let $\mathbf{x} = (x_1, \dots, x_k) \in R^{\mathbf{X}}$. We need to show that $((M_{x_1}, \boldsymbol{\mu}_{x_1}), \dots, (M_{x_k}, \boldsymbol{\mu}_{x_k})) \in R^{\mathbb{F}_{\mathcal{C}}(\mathbf{A})}$. For each $i \in [t] \setminus \{1\}$, consider a probability distribution $\boldsymbol{\xi}_i \in \mathbb{S}(R^{\mathbf{A}})$ corresponding to the homomorphism h_i and witnessing part 2 in Definition 4.8. Also, consider the probability distribution $\boldsymbol{\xi}_1 \in \mathbb{S}(R^{\mathbf{A}})$ and the integer distribution $\boldsymbol{\zeta} \in \mathbb{Z}(R^{\mathbf{A}})$ corresponding to h_1 and g , respectively, and witnessing 1'. We introduce the matrix $Q \in \mathbb{Q}^{|R^{\mathbf{A}}|, \aleph_0}$ and the vector $\boldsymbol{\delta} \in \mathbb{Z}^{|R^{\mathbf{A}}|}$ defined by

$$\begin{aligned} Q\mathbf{e}_i &= \boldsymbol{\xi}_i & \forall i \in [t], \\ Q\mathbf{e}_i &= \boldsymbol{\xi}_t & \forall i \in \mathbb{N} \setminus [t], \\ \boldsymbol{\delta} &= \boldsymbol{\zeta}. \end{aligned}$$

We claim that $(Q, \boldsymbol{\delta}) \in \mathcal{C}^{(|R^{\mathbf{A}}|)}$. The requirements (c_1) , (c_2) , (c_3) , and (c_5) in Definition 3.2 are easily seen to be satisfied. Suppose $\mathbf{e}_{\mathbf{a}}^T Q \neq \mathbf{0}_{\aleph_0}^T$ for some $\mathbf{a} = (a_1, \dots, a_k) \in R^{\mathbf{A}}$, so that there exists $d \in [t]$ such that $\mathbf{e}_{\mathbf{a}}^T Q \mathbf{e}_d \neq 0$. Hence, $\mathbf{a} \in \text{supp}(Q\mathbf{e}_d) = \text{supp}(\boldsymbol{\xi}_d) \subseteq s^{R^{\mathbf{A}}}(\mathbf{x})$. Pick $h_j = h_{\mathbf{x}, \mathbf{a}}$. We have that

$$E^{(R^{\mathbf{A}}, p)} \boldsymbol{\xi}_j = h_j(x_p) = h_{\mathbf{x}, \mathbf{a}}(x_p) = \mathbf{e}_{a_p} \quad \forall p \in [k].$$

Suppose that $\boldsymbol{\xi}_j \neq \mathbf{e}_{\mathbf{a}}$. Then, $\exists \mathbf{a}' = (a'_1, \dots, a'_k) \in R^{\mathbf{A}}$ such that $\mathbf{a}' \neq \mathbf{a}$ and $\mathbf{e}_{\mathbf{a}'}^T \boldsymbol{\xi}_j > 0$. Choose $q \in [k]$ such that $a'_q \neq a_q$, and observe that

$$0 = \mathbf{e}_{a'_q}^T \mathbf{e}_{a_q} = \mathbf{e}_{a'_q}^T E^{(R^{\mathbf{A}}, q)} \boldsymbol{\xi}_j \geq \mathbf{e}_{\mathbf{a}'}^T \boldsymbol{\xi}_j > 0,$$

which is a contradiction. Hence, $Q\mathbf{e}_j = \boldsymbol{\xi}_j = \mathbf{e}_{\mathbf{a}}$. We conclude that Q is skeletal and, therefore, (c_6) is satisfied. Finally, suppose that $\mathbf{a} \notin \text{supp}(Q\mathbf{e}_1) = \text{supp}(\boldsymbol{\xi}_1)$ for some $\mathbf{a} \in R^{\mathbf{A}}$. Recalling that $\boldsymbol{\xi}_1 \in \mathbb{S}(R^{\mathbf{A}})$ corresponds to the homomorphism $h_1 = h_{\tilde{x}, \tilde{a}}$, it follows from 1' that $\text{supp}(\boldsymbol{\zeta}) \subseteq \text{supp}(\boldsymbol{\xi}_1)$. Hence, $\mathbf{a} \notin \text{supp}(\boldsymbol{\zeta}) = \text{supp}(\boldsymbol{\delta})$, so that (c_4) is satisfied, too. As a consequence, $(Q, \boldsymbol{\delta}) \in \mathcal{C}^{(|R^{\mathbf{A}}|)}$, as claimed.

Now, we need to show that $(M_{x_\alpha}, \boldsymbol{\mu}_{x_\alpha}) = (Q, \boldsymbol{\delta}) / \pi_\alpha$ for each $\alpha \in [k]$, where $\pi_\alpha : R^{\mathbf{A}} \rightarrow A$ maps $\mathbf{a} \in R^{\mathbf{A}}$ to its α th coordinate. Observe first that, by definition, $P_{\pi_\alpha} = E^{(R^{\mathbf{A}}, \alpha)}$ for each $\alpha \in [k]$. We see that

$$\begin{aligned} Q_{/\pi_\alpha} \mathbf{e}_i &= P_{\pi_\alpha} Q \mathbf{e}_i = E^{(R^\mathbf{A}, \alpha)} Q \mathbf{e}_i = E^{(R^\mathbf{A}, \alpha)} \boldsymbol{\xi}_i = h_i(x_\alpha) = M_{x_\alpha} \mathbf{e}_i & \text{for } i \in [t], \\ Q_{/\pi_\alpha} \mathbf{e}_i &= P_{\pi_\alpha} Q \mathbf{e}_i = P_{\pi_\alpha} Q \mathbf{e}_t = M_{x_\alpha} \mathbf{e}_t = M_{x_\alpha} \mathbf{e}_i & \text{for } i \in \mathbb{N} \setminus [t], \end{aligned}$$

which yields $Q_{/\pi_\alpha} = M_{x_\alpha}$. Moreover,

$$\boldsymbol{\delta}_{/\pi_\alpha} = P_{\pi_\alpha} \boldsymbol{\delta} = E^{(R^\mathbf{A}, \alpha)} \boldsymbol{\delta} = E^{(R^\mathbf{A}, \alpha)} \boldsymbol{\zeta} = g(x_\alpha) = \boldsymbol{\mu}_{x_\alpha}.$$

It follows that $(M_{x_\alpha}, \boldsymbol{\mu}_{x_\alpha}) = (Q_{/\pi_\alpha}, \boldsymbol{\delta}_{/\pi_\alpha}) = (Q, \boldsymbol{\delta})_{/\pi_\alpha}$. By Definition 4.1,

$$((M_{x_1}, \boldsymbol{\mu}_{x_1}), \dots, (M_{x_k}, \boldsymbol{\mu}_{x_k})) \in R^{\mathbb{F}_{\mathcal{C}}(\mathbf{A})},$$

so $\gamma: \mathbf{X} \rightarrow \mathbb{F}_{\mathcal{C}}(\mathbf{A})$ is a homomorphism. \square

Our second goal is to prove the following.

THEOREM 4.11. *If CLAP solves PCSP(\mathbf{A}, \mathbf{B}), then there is a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$.*

Remark 4.12. The proof of Theorem 4.11 proceeds essentially by establishing that the free structure $\mathbb{F}_{\mathcal{C}}(\mathbf{A})$ has the CLAP condition as an instance of PCSP(\mathbf{A}, \mathbf{B}). However, some care is needed when handling Proposition 4.9, which only applies to *finite* structures, while $\mathbb{F}_{\mathcal{C}}(\mathbf{A})$ is not finite, in general. To overcome this problem, we use a compactness argument tailored to our minion \mathcal{C} discussed in section 4.2, which follows the ideas of [8].

We remark that the compactness argument for relational structures in the form stated in [22, Lemma A.6] does not entirely fit our proof structure, as the element $(\mathbf{e}_1 \mathbf{1}_{\mathbb{N}_0}^T, \mathbf{e}_1)$ having the role of $\bar{\mathbf{x}}$ in Definition 4.8 does not belong to every induced substructure of $\mathbb{F}_{\mathcal{C}}(\mathbf{A})$. A different option would have been to use the general compactness argument known as the (uncountable version of the) *compactness theorem of logic* [62], which applies to all *minion tests*⁹ as derived in [31, Proposition 6] through [64].

Proof of Theorem 4.11. Let $n = |\mathbf{A}|$. For $D \in \mathbb{N}$, denote $\mathbb{F}_{\mathcal{C}_D}(\mathbf{A})$ by \mathbf{F} (where \mathcal{C}_D is the subminion of \mathcal{C} introduced in section 4.2). Hence, the domain of \mathbf{F} is $\mathcal{C}_D^{(n)}$, which is finite. We claim that \mathbf{F} has the CLAP condition as an instance of PCSP(\mathbf{A}, \mathbf{B}).

For each $R \in \sigma$ of arity k and for each $\boldsymbol{\tau} = ((M_1, \boldsymbol{\mu}_1), \dots, (M_k, \boldsymbol{\mu}_k)) \in R^{\mathbf{F}}$, take $(Q_{\boldsymbol{\tau}}, \boldsymbol{\delta}_{\boldsymbol{\tau}}) \in \mathcal{C}_D^{(|R^\mathbf{A}|)}$ satisfying $(M_j, \boldsymbol{\mu}_j) = (Q_{\boldsymbol{\tau}}, \boldsymbol{\delta}_{\boldsymbol{\tau}})_{/\pi_j} \forall j \in [k]$, where $\pi_j: R^\mathbf{A} \rightarrow A$ maps $\mathbf{a} \in R^\mathbf{A}$ to its j th coordinate; i.e., $M_j = E^{(R^\mathbf{A}, j)} Q_{\boldsymbol{\tau}}$ and $\boldsymbol{\mu}_j = E^{(R^\mathbf{A}, j)} \boldsymbol{\delta}_{\boldsymbol{\tau}} \forall j \in [k]$. Given $R \in \sigma$ of arity k , consider the map

$$\begin{aligned} s^R: R^{\mathbf{F}} &\rightarrow \mathcal{P}(R^\mathbf{A}) \setminus \{\emptyset\} \\ \boldsymbol{\tau} &\mapsto \bigcup_{i \in \mathbb{N}} \text{supp}(Q_{\boldsymbol{\tau}} \mathbf{e}_i). \end{aligned}$$

Let us first check part (I) of Definition 4.8. Pick $\boldsymbol{\tau} = ((M_1, \boldsymbol{\mu}_1), \dots, (M_k, \boldsymbol{\mu}_k)) \in R^{\mathbf{F}}$ and $\mathbf{a} = (a_1, \dots, a_k) \in s^R(\boldsymbol{\tau})$. We have that $\mathbf{a} \in \text{supp}(Q_{\boldsymbol{\tau}} \mathbf{e}_\alpha)$ for some $\alpha \in \mathbb{N}$, i.e., $\mathbf{e}_\alpha^T Q_{\boldsymbol{\tau}} \mathbf{e}_\alpha \neq 0$. Since $Q_{\boldsymbol{\tau}}$ is skeletal, the set $L_{\boldsymbol{\tau}, \mathbf{a}} = \{\ell \in \mathbb{N} : Q_{\boldsymbol{\tau}} \mathbf{e}_\ell = \mathbf{e}_\mathbf{a}\}$ is nonempty; let $\ell(\boldsymbol{\tau}, \mathbf{a}) := \min(L_{\boldsymbol{\tau}, \mathbf{a}})$. Consider the map

$$\begin{aligned} h_{\boldsymbol{\tau}, \mathbf{a}}: \mathcal{C}_D^{(n)} &\rightarrow \mathbb{S}(A) \\ (\hat{M}, \hat{\boldsymbol{\mu}}) &\mapsto \hat{M} \mathbf{e}_{\ell(\boldsymbol{\tau}, \mathbf{a})}. \end{aligned}$$

⁹Compare Remark 4.13.

We claim that $h_{\tau, \mathbf{a}}$ is a homomorphism from \mathbf{F} to $\mathbb{S}(\mathbf{A})$. Take $\tilde{R} \in \sigma$ of arity \tilde{k} , and let $\tilde{\tau} = ((\tilde{M}_1, \tilde{\mu}_1), \dots, (\tilde{M}_{\tilde{k}}, \tilde{\mu}_{\tilde{k}})) \in \tilde{R}^{\mathbf{F}}$. Consider the pair $(Q_{\tilde{\tau}}, \delta_{\tilde{\tau}}) \in \mathcal{C}_D^{(|\tilde{R}^{\mathbf{A}}|)}$. We have that

$$h_{\tau, \mathbf{a}}(\tilde{\tau}) = (\tilde{M}_1 \mathbf{e}_{\ell(\tau, \mathbf{a})}, \dots, \tilde{M}_{\tilde{k}} \mathbf{e}_{\ell(\tau, \mathbf{a})}) = \left(E^{(\tilde{R}^{\mathbf{A}}, 1)} Q_{\tilde{\tau}} \mathbf{e}_{\ell(\tau, \mathbf{a})}, \dots, E^{(\tilde{R}^{\mathbf{A}}, \tilde{k})} Q_{\tilde{\tau}} \mathbf{e}_{\ell(\tau, \mathbf{a})} \right).$$

Since $Q_{\tilde{\tau}} \mathbf{e}_{\ell(\tau, \mathbf{a})} \in \mathbb{S}(\tilde{R}^{\mathbf{A}})$, we deduce that $h_{\tau, \mathbf{a}}(\tilde{\tau}) \in \tilde{R}^{\mathbb{S}(\mathbf{A})}$, as wanted. Therefore, $h_{\tau, \mathbf{a}}$ is a homomorphism from \mathbf{F} to $\mathbb{S}(\mathbf{A})$. We now check that the requirements 1 and 2 in Definition 4.8 are met. The former follows from

$$h_{\tau, \mathbf{a}}((M_i, \mu_i)) = M_i \mathbf{e}_{\ell(\tau, \mathbf{a})} = E^{(R^{\mathbf{A}}, i)} Q_{\tau} \mathbf{e}_{\ell(\tau, \mathbf{a})} = E^{(R^{\mathbf{A}}, i)} \mathbf{e}_{\mathbf{a}} = \mathbf{e}_{a_i} \quad \forall i \in [k].$$

To check the latter requirement, take $\tilde{R} \in \sigma$ of arity \tilde{k} and

$$\tilde{\tau} = ((\tilde{M}_1, \tilde{\mu}_1), \dots, (\tilde{M}_{\tilde{k}}, \tilde{\mu}_{\tilde{k}})) \in \tilde{R}^{\mathbf{F}},$$

and consider $\xi := Q_{\tilde{\tau}} \mathbf{e}_{\ell(\tau, \mathbf{a})}$. Observe that

- $h_{\tau, \mathbf{a}}((\tilde{M}_i, \tilde{\mu}_i)) = \tilde{M}_i \mathbf{e}_{\ell(\tau, \mathbf{a})} = E^{(\tilde{R}^{\mathbf{A}}, i)} Q_{\tilde{\tau}} \mathbf{e}_{\ell(\tau, \mathbf{a})} = E^{(\tilde{R}^{\mathbf{A}}, i)} \xi \quad \forall i \in [\tilde{k}],$
- $\text{supp}(\xi) = \text{supp}(Q_{\tilde{\tau}} \mathbf{e}_{\ell(\tau, \mathbf{a})}) \subseteq \bigcup_{i \in \mathbb{N}} \text{supp}(Q_{\tilde{\tau}} \mathbf{e}_i) = s^{\tilde{R}}(\tilde{\tau}).$

We now check part (II) of Definition 4.8. Take R_u as \bar{R} , and observe that

$$R_u^{\mathbf{F}} = \{(M, \mu) \in \mathcal{C}_D^{(n)} : \exists (Q, \delta) \in \mathcal{C}_D^{(n)} \text{ such that } M = E^{(R_u^{\mathbf{A}}, 1)} Q, \\ \mu = E^{(R_u^{\mathbf{A}}, 1)} \delta\} = \mathcal{C}_D^{(n)},$$

where we have used that $E^{(R_u^{\mathbf{A}}, 1)} = I_n$. Consider the element $\bar{\tau} = (\mathbf{e}_1 \mathbf{1}_{\mathbb{N}_0}^T, \mathbf{e}_1) \in \mathcal{C}_D^{(n)} = R_u^{\mathbf{F}}$. Using again that $E^{(R_u^{\mathbf{A}}, 1)} = I_n$, we see that $(Q_{\bar{\tau}}, \delta_{\bar{\tau}}) = \bar{\tau}$. We obtain

$$s^{R_u}(\bar{\tau}) = \bigcup_{i \in \mathbb{N}} \text{supp}(\mathbf{e}_1 \mathbf{1}_{\mathbb{N}_0}^T \mathbf{e}_i) = \bigcup_{i \in \mathbb{N}} \text{supp}(\mathbf{e}_1) = \{1\}.$$

Hence, we pick $\bar{a} = 1$. Notice that

$$\ell(\bar{\tau}, \bar{a}) = \min\{\ell \in \mathbb{N} : \mathbf{e}_1 \mathbf{1}_{\mathbb{N}_0}^T \mathbf{e}_\ell = \mathbf{e}_1\} = \min\{\ell \in \mathbb{N} : \mathbf{e}_1 = \mathbf{e}_1\} = \min \mathbb{N} = 1.$$

Consider the function

$$g : \mathcal{C}_D^{(n)} \rightarrow \mathbb{Z}(A) \\ (\hat{M}, \hat{\mu}) \mapsto \hat{\mu}.$$

Following the same procedure as for $h_{\tau, \mathbf{a}}$, we easily check that g is a homomorphism from \mathbf{F} to $\mathbb{Z}(\mathbf{A})$. We now verify that condition 1' of Definition 4.8 is satisfied. Given $\tilde{R} \in \sigma$ of arity \tilde{k} and $\tilde{\tau} = ((\tilde{M}_1, \tilde{\mu}_1), \dots, (\tilde{M}_{\tilde{k}}, \tilde{\mu}_{\tilde{k}})) \in \tilde{R}^{\mathbf{F}}$, let $\xi := Q_{\tilde{\tau}} \mathbf{e}_1 \in \mathbb{S}(\tilde{R}^{\mathbf{A}})$ and $\zeta := \delta_{\tilde{\tau}} \in \mathbb{Z}(\tilde{R}^{\mathbf{A}})$. Then, given $i \in [k]$,

- * $h_{\bar{\tau}, \bar{a}}((\tilde{M}_i, \tilde{\mu}_i)) = \tilde{M}_i \mathbf{e}_{\ell(\bar{\tau}, \bar{a})} = \tilde{M}_i \mathbf{e}_1 = E^{(\tilde{R}^{\mathbf{A}}, i)} Q_{\tilde{\tau}} \mathbf{e}_1 = E^{(\tilde{R}^{\mathbf{A}}, i)} \xi;$
- * $g((\tilde{M}_i, \tilde{\mu}_i)) = \tilde{\mu}_i = E^{(\tilde{R}^{\mathbf{A}}, i)} \delta_{\tilde{\tau}} = E^{(\tilde{R}^{\mathbf{A}}, i)} \zeta;$
- * $\text{supp}(\zeta) = \text{supp}(\delta_{\tilde{\tau}}) \subseteq \text{supp}(Q_{\tilde{\tau}} \mathbf{e}_1) = \text{supp}(\xi) \subseteq \bigcup_{i \in \mathbb{N}} \text{supp}(Q_{\tilde{\tau}} \mathbf{e}_i) = s^{\tilde{R}}(\tilde{\tau}),$

where, for the first inclusion in the third line, we have used (c_4) in Definition 3.2.

It follows that \mathbf{F} has the CLAP condition as an instance of PCSP(\mathbf{A}, \mathbf{B}), as claimed. Then, Proposition 4.9 implies that CLAP accepts \mathbf{F} . Since, by hypothesis,

CLAP solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$, we deduce that $\mathbb{F}_{\mathcal{C}_D}(\mathbf{A}) = \mathbf{F} \rightarrow \mathbf{B}$. By Lemma 4.2, there exists a minion homomorphism from \mathcal{C}_D to $\text{Pol}(\mathbf{A}, \mathbf{B})$. Finally, since the set of polymorphisms of (\mathbf{A}, \mathbf{B}) of arity L is finite for every $L \in \mathbb{N}$, Proposition 4.5 allows us to conclude that there exists a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$. \square

Remark 4.13. It follows from the proofs of Theorems 4.10 and 4.11 that CLAP fits within the framework of *minion tests* recently introduced in [31]. More precisely, $\text{CLAP} = \text{Test}_{\mathcal{C}}$, which means that, for two σ -structures \mathbf{X} and \mathbf{A} , $\text{CLAP}(\mathbf{X}, \mathbf{A})$ accepts if and only if $\mathbf{X} \rightarrow \mathbb{F}_{\mathcal{C}}(\mathbf{A})$. Additionally, it follows from [31] that CLAP is a *conic* minion test, which essentially means that one can build a progressively tighter hierarchy of relaxations based on CLAP whose k th level correctly classifies all instances of size k .

5. H-symmetric polymorphisms. This section contains the proof of Theorem 3.5. We remark that the machinery developed here can be extended to the more general setting of H -block-symmetric polymorphisms, at the only cost of dealing with a more cumbersome notation. This is done in Appendix C and results in Theorem C.3—a slightly stronger version of Theorem 3.5.

We shall need two helpful lemmas. The first lemma shows a property of H -symmetric functions that will be useful in the proof of Theorem 3.5. Throughout this section, without loss of generality, we consider $A = [n]$.

LEMMA 5.1. *Let $f : A^L \rightarrow B$ be H -symmetric for some $m \times n$ tie matrix H , with $m \in \mathbb{N}$. Consider two maps $\pi, \tilde{\pi} : [L] \rightarrow [n]$ such that $P_{\pi} \mathbf{1}_L = P_{\tilde{\pi}} \mathbf{1}_L$ and the vector $P_{\pi} \mathbf{1}_L$ is H -tieless. Then*

$$f_{/\pi}(1, \dots, n) = f_{/\tilde{\pi}}(1, \dots, n).$$

Proof. For $a \in [n]$, we have

$$|\pi^{-1}(a)| = \sum_{i \in [L]} (P_{\pi})_{ai} = \sum_{i \in [L]} \mathbf{e}_a^T P_{\pi} \mathbf{e}_i = \mathbf{e}_a^T P_{\pi} \mathbf{1}_L = \mathbf{e}_a^T P_{\tilde{\pi}} \mathbf{1}_L = |\tilde{\pi}^{-1}(a)|.$$

Hence, we can consider bijections $\varphi_a : \pi^{-1}(a) \rightarrow \tilde{\pi}^{-1}(a)$ for each $a \in [n]$. Clearly, their union

$$\varphi = \bigcup_{a \in [n]} \varphi_a : [L] \rightarrow [L]$$

is also a bijection. For each $i \in [L]$, we have

$$(\tilde{\pi} \circ \varphi)(i) = \tilde{\pi}(\varphi(i)) = \tilde{\pi}(\varphi_{\pi(i)}(i)) = \pi(i)$$

and, hence, $\tilde{\pi} \circ \varphi = \pi$. Let $\tilde{\mathbf{a}} = (\tilde{\pi}(1), \dots, \tilde{\pi}(L))$. Notice that, for each $a \in [n]$,

$$\mathbf{e}_a^T \tilde{\mathbf{a}}^{\#} = |\{i \in [L] : \tilde{\pi}(i) = a\}| = \mathbf{e}_a^T P_{\tilde{\pi}} \mathbf{1}_L$$

and, therefore, $\tilde{\mathbf{a}}^{\#} = P_{\tilde{\pi}} \mathbf{1}_L = P_{\pi} \mathbf{1}_L$, which is H -tieless. Using that f is H -symmetric, we find

$$f_{/\tilde{\pi}}(1, \dots, n) = f(\tilde{\mathbf{a}}) = f_{/\varphi}(\tilde{\mathbf{a}}) = (f_{/\varphi})_{/\tilde{\pi}}(1, \dots, n) = f_{/\tilde{\pi} \circ \varphi}(1, \dots, n) = f_{/\pi}(1, \dots, n),$$

as required. \square

One intriguing property of skeletal matrices is that they can be simultaneously reduced to H -tieless vectors, in the sense of the next lemma. We say that a vector is *finitely supported* if it only has a finite number of nonzero entries.

LEMMA 5.2 (tiebreak lemma). *For $k, p, m \in \mathbb{N}$, let $M_1, \dots, M_k \in \mathbb{Q}^{p, \aleph_0}$ be skeletal matrices, and let H be an $m \times p$ tie matrix. Then there exists a stochastic finitely supported vector $\mathbf{v} \in \mathbb{Q}^{\aleph_0}$ with $\mathbf{e}_1^T \mathbf{v} > 0$ such that $M_j \mathbf{v}$ is H -tieless for any $j \in [k]$.*

Proof. Let Ω be the set of rational stochastic finitely supported vectors of size \aleph_0 whose first entry is nonzero, and consider the map

$$f : \Omega \rightarrow \mathbb{N}_0$$

$$\hat{\mathbf{v}} \mapsto \sum_{j \in [k]} |\{(i, i') \in [m]^2 : i \neq i' \text{ and } \mathbf{e}_i^T H M_j \hat{\mathbf{v}} = \mathbf{e}_{i'}^T H M_j \hat{\mathbf{v}} \neq 0\}|.$$

In other words, $f(\hat{\mathbf{v}})$ counts the total number of ties in the set of vectors $\{H M_j \hat{\mathbf{v}} : j \in [k]\}$. Let \mathbf{v} attain the minimum of f over Ω . If $f(\mathbf{v}) = 0$, we are done. Otherwise, let $j \in [k]$, $i, i' \in [m]$ be such that $i \neq i'$ and $\mathbf{e}_i^T H M_j \mathbf{v} = \mathbf{e}_{i'}^T H M_j \mathbf{v} \neq 0$. From $\mathbf{e}_i^T H M_j \mathbf{v} \neq 0$, we see that $\exists \beta \in [p]$ such that $\mathbf{e}_i^T H \mathbf{e}_\beta \neq 0$ and $\mathbf{e}_\beta^T M_j \mathbf{v} \neq 0$. In particular, we have $\mathbf{e}_\beta^T M_j \neq \mathbf{0}_{\aleph_0}^T$; since M_j is skeletal, this implies that $M_j \mathbf{e}_\alpha = \mathbf{e}_\beta$ for some $\alpha \in \mathbb{N}$. For $\epsilon \in \mathbb{Q}$, $0 < \epsilon < 1$, consider the vector $\mathbf{v}_\epsilon = (1 - \epsilon)\mathbf{v} + \epsilon \mathbf{e}_\alpha$. Observe that $\mathbf{v}_\epsilon \in \Omega$. For $g \in [k]$, we have $H M_g \mathbf{v}_\epsilon = (1 - \epsilon)H M_g \mathbf{v} + \epsilon H M_g \mathbf{e}_\alpha$. By choosing ϵ sufficiently small, we can assume that, for each $g \in [k]$, $H M_g \mathbf{v}_\epsilon$ does not have new ties other than those in $H M_g \mathbf{v}$. Moreover,

$$H M_j \mathbf{v}_\epsilon = (1 - \epsilon)H M_j \mathbf{v} + \epsilon H M_j \mathbf{e}_\alpha = (1 - \epsilon)H M_j \mathbf{v} + \epsilon H \mathbf{e}_\beta,$$

and hence,

$$\begin{aligned} \mathbf{e}_i^T H M_j \mathbf{v}_\epsilon &= (1 - \epsilon)\mathbf{e}_i^T H M_j \mathbf{v} + \epsilon \mathbf{e}_i^T H \mathbf{e}_\beta = (1 - \epsilon)\mathbf{e}_{i'}^T H M_j \mathbf{v} + \epsilon \mathbf{e}_i^T H \mathbf{e}_\beta \\ &\neq (1 - \epsilon)\mathbf{e}_{i'}^T H M_j \mathbf{v} + \epsilon \mathbf{e}_{i'}^T H \mathbf{e}_\beta = \mathbf{e}_{i'}^T H M_j \mathbf{v}_\epsilon, \end{aligned}$$

where the disequality follows from $\mathbf{e}_i^T H \mathbf{e}_\beta \neq 0$ and from the fact that $H \mathbf{e}_\beta$ is a tieless vector by the definition of tie matrix. We conclude that $f(\mathbf{v}_\epsilon) < f(\mathbf{v})$, which contradicts our assumption. \square

THEOREM 5.3 (Theorem 3.5 restated). *Let (\mathbf{A}, \mathbf{B}) be a PCSP template, and suppose $\text{Pol}(\mathbf{A}, \mathbf{B})$ contains H -symmetric operations of arbitrarily large arity for some $m \times |A|$ tie matrix H , $m \in \mathbb{N}$. Then there exists a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$.*

Remark 5.4. Before proving Theorem 3.5, we provide some intuition on the construction of the minion homomorphism whose existence shall establish the result. First, one fixes an H -symmetric polymorphism f . Then, the image of an L -ary element $(M, \boldsymbol{\mu})$ of \mathcal{C} under the homomorphism is the function that (i) takes a tuple (a_1, \dots, a_L) of variables in A as input, (ii) *deforms* the tuple by changing the frequency of each variable according to the information carried by M and $\boldsymbol{\mu}$, and (iii) returns as output the evaluation of f on the deformed tuple. The deformation in step (ii) is encoded by the map φ defined in (5.1). Essentially, φ decides what frequency to assign to a variable a_i on the basis of the weight of i in the probability distribution $M\mathbf{v}$, where \mathbf{v} is the *tie-breaking* vector from Lemma 5.2. The integer distribution $\boldsymbol{\mu}$ is also taken into account by φ , and its role is essentially to fill the gap between the size of the deformed tuple obtained above and the arity of f . If $\text{Pol}(\mathbf{A}, \mathbf{B})$ is rich enough to provide H -symmetric polymorphisms of whichever arity we need, $\boldsymbol{\mu}$ is inessential (cf. Remark 5.5).

Proof of Theorem 3.5. For $D \in \mathbb{N}$, consider the subminion \mathcal{C}_D of \mathcal{C} described in section 4.2. Observe that $S = \{M : (M, \boldsymbol{\mu}) \in \mathcal{C}_D^{(n)}\}$ is a finite set of skeletal matrices. Therefore, we can apply the tiebreak lemma, Lemma 5.2, to find a stochastic finitely supported vector $\mathbf{v} \in \mathbb{Q}^{\mathbb{N}_0}$ with $\mathbf{e}_1^T \mathbf{v} > 0$ such that $M\mathbf{v}$ is H -tieless for any $M \in S$. Since \mathbf{v} is finitely supported, we can find $N' \in \mathbb{N}$ such that $N'\mathbf{v}$ has integer entries. Let σ_1^H denote the largest singular value of H , i.e., the square root of the largest eigenvalue of $H^T H$. Set $N = 2\lceil \sigma_1^H + 1 \rceil D^2 N'$, and let f be an H -symmetric polymorphism of arity $c \geq N^2$. Write $c = N\alpha + \beta$ with $\alpha, \beta \in \mathbb{N}_0$, $\beta \leq N - 1$. Note that $N^2 \leq c = N\alpha + \beta \leq N\alpha + N - 1 < N(\alpha + 1)$, so $N < \alpha + 1$, and hence, $\beta < \alpha$.

Consider the function

$$\xi_D : \mathcal{C}_D \rightarrow \text{Pol}(\mathbf{A}, \mathbf{B})$$

defined as follows. Given $L \in \mathbb{N}$ and $(M, \boldsymbol{\mu}) \in \mathcal{C}_D^{(L)}$, take the map $\varphi : [c] \rightarrow [L]$ such that the corresponding $L \times c$ matrix P_φ is

$$(5.1) \quad P_\varphi = \begin{pmatrix} \mathbf{1}_{\mathbf{e}_1^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu})}^T & \mathbf{0}^T & \cdots & \mathbf{0}^T \\ \mathbf{0}^T & \mathbf{1}_{\mathbf{e}_2^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu})}^T & \cdots & \mathbf{0}^T \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}^T & \mathbf{0}^T & \cdots & \mathbf{1}_{\mathbf{e}_L^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu})}^T \end{pmatrix}.$$

To verify that (5.1) is well defined, observe first that

$$\begin{aligned} \sum_{i=1}^L \mathbf{e}_i^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu}) &= \mathbf{1}_L^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu}) \\ &= \alpha N \mathbf{1}_L^T M \mathbf{v} + \beta \mathbf{1}_L^T \boldsymbol{\mu} = \alpha N \mathbf{1}_{\mathbb{N}_0}^T \mathbf{v} + \beta = \alpha N + \beta = c. \end{aligned}$$

Moreover, for each $i \in [L]$, $\mathbf{e}_i^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu}) = \mathbf{e}_i^T(2\alpha\lceil \sigma_1^H + 1 \rceil D(DM)(N'\mathbf{v}) + \beta \boldsymbol{\mu})$ is an integer. If $\mathbf{e}_i^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu})$ was negative, then $\mathbf{e}_i^T \boldsymbol{\mu} < 0$. By the requirement (c_4) in Definition 3.2, this would imply that $\mathbf{e}_i^T M \mathbf{e}_1 > 0$, and hence, $0 < \mathbf{e}_i^T M \mathbf{e}_1 \mathbf{e}_1^T \mathbf{v} \leq \mathbf{e}_i^T M \mathbf{v}$. As a consequence, $\mathbf{e}_i^T(DM)(N'\mathbf{v}) \geq 1$ so that

$$\begin{aligned} \mathbf{e}_i^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu}) &= 2\alpha\lceil \sigma_1^H + 1 \rceil D \mathbf{e}_i^T(DM)(N'\mathbf{v}) + \beta \mathbf{e}_i^T \boldsymbol{\mu} \\ &\geq 2\alpha\lceil \sigma_1^H + 1 \rceil D + \beta \mathbf{e}_i^T \boldsymbol{\mu} \geq \alpha D - \beta D > 0, \end{aligned}$$

which is a contradiction. In conclusion, the numbers $\mathbf{e}_i^T(\alpha N M \mathbf{v} + \beta \boldsymbol{\mu})$ are nonnegative integers summing up to c , so (5.1) is well defined.

We define $\xi_D((M, \boldsymbol{\mu})) := f_{/\varphi}$. Clearly, $\xi_D((M, \boldsymbol{\mu})) \in \text{Pol}(\mathbf{A}, \mathbf{B})$. We claim that the map ξ_D is a minion homomorphism. It is straightforward to check that ξ_D preserves arities so, to conclude, we need to show that it also preserves minors. Take $L' \in \mathbb{N}$ and choose a map $\pi : [L] \rightarrow [L']$. Letting $\tilde{\varphi} : [c] \rightarrow [L']$ be the map corresponding to the matrix

$$P_{\tilde{\varphi}} = \begin{pmatrix} \mathbf{1}_{\mathbf{e}_1^T(\alpha N P_\pi M \mathbf{v} + \beta P_\pi \boldsymbol{\mu})}^T & \mathbf{0}^T & \cdots & \mathbf{0}^T \\ \mathbf{0}^T & \mathbf{1}_{\mathbf{e}_2^T(\alpha N P_\pi M \mathbf{v} + \beta P_\pi \boldsymbol{\mu})}^T & \cdots & \mathbf{0}^T \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}^T & \mathbf{0}^T & \cdots & \mathbf{1}_{\mathbf{e}_{L'}^T(\alpha N P_\pi M \mathbf{v} + \beta P_\pi \boldsymbol{\mu})}^T \end{pmatrix},$$

we see that $\xi_D((M, \boldsymbol{\mu})/\pi) = f/\tilde{\varphi}$. Moreover, $\xi_D((M, \boldsymbol{\mu})/\pi) = (f/\varphi)/\pi = f/\pi \circ \varphi$, where φ corresponds to the matrix P_φ in (5.1). Take $\mathbf{a} = (a_1, \dots, a_{L'}) \in A^{L'}$, and consider the map

$$\begin{aligned} \pi_{\mathbf{a}} : [L'] &\rightarrow [n] \\ i &\mapsto a_i. \end{aligned}$$

Observe that

$$(5.2) \quad \begin{aligned} f/\tilde{\varphi}(\mathbf{a}) &= (f/\tilde{\varphi})/\pi_{\mathbf{a}}(1, \dots, n) = f/\pi_{\mathbf{a}} \circ \tilde{\varphi}(1, \dots, n), \quad \text{and similarly,} \\ f/\pi \circ \varphi(\mathbf{a}) &= (f/\pi \circ \varphi)/\pi_{\mathbf{a}}(1, \dots, n) = f/\pi_{\mathbf{a}} \circ \pi \circ \varphi(1, \dots, n). \end{aligned}$$

Notice that

$$\begin{aligned} P_{\pi_{\mathbf{a}} \circ \tilde{\varphi}} \mathbf{1}_c &= P_{\pi_{\mathbf{a}}} P_{\tilde{\varphi}} \mathbf{1}_c = P_{\pi_{\mathbf{a}}} (\alpha N P_{\pi} M \mathbf{v} + \beta P_{\pi} \boldsymbol{\mu}) \\ &= P_{\pi_{\mathbf{a}}} P_{\pi} (\alpha N M \mathbf{v} + \beta \boldsymbol{\mu}) = P_{\pi_{\mathbf{a}}} P_{\pi} P_{\varphi} \mathbf{1}_c = P_{\pi_{\mathbf{a}} \circ \pi \circ \varphi} \mathbf{1}_c. \end{aligned}$$

We claim that the vector $P_{\pi_{\mathbf{a}} \circ \tilde{\varphi}} \mathbf{1}_c$ is H -tieless. Let $\mathbf{u} = (u_i) = H P_{\pi_{\mathbf{a}} \circ \tilde{\varphi}} \mathbf{1}_c$; the claim is equivalent to \mathbf{u} being tieless. Let $\mathbf{w} = (w_i) = \alpha N H P_{\pi_{\mathbf{a}} \circ \pi} M \mathbf{v}$ and $\mathbf{z} = (z_i) = \beta H P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}$, so that $\mathbf{u} = \mathbf{w} + \mathbf{z}$. Choose $i, i' \in [m]$ such that $i \neq i'$ and $u_i \neq 0$. We need to show that $u_i \neq u_{i'}$. Suppose $w_i = 0$. We can write $H^T \mathbf{e}_i = \sum_{g \in G} \lambda_g \mathbf{e}_g$ for $G = \text{supp}(H^T \mathbf{e}_i)$, where each λ_g is a positive integer (note that $G \neq \emptyset$ since, otherwise, $H^T \mathbf{e}_i = \mathbf{0}_n$, which would imply $u_i = 0$). Let $F = (\pi_{\mathbf{a}} \circ \pi)^{-1}(G)$. From $w_i = 0$, we obtain

$$\begin{aligned} 0 &= \mathbf{e}_i^T H P_{\pi_{\mathbf{a}} \circ \pi} M \mathbf{v} = (H^T \mathbf{e}_i)^T P_{\pi_{\mathbf{a}} \circ \pi} M \mathbf{v} \\ &= \sum_{g \in G} \lambda_g \mathbf{e}_g^T P_{\pi_{\mathbf{a}} \circ \pi} M \mathbf{v} = \sum_{g \in G} \lambda_g \sum_{j \in (\pi_{\mathbf{a}} \circ \pi)^{-1}(g)} \mathbf{e}_j^T M \mathbf{v}, \end{aligned}$$

and hence, the following chain of implications holds:

$$\begin{aligned} 0 &= \sum_{g \in G} \sum_{j \in (\pi_{\mathbf{a}} \circ \pi)^{-1}(g)} \mathbf{e}_j^T M \mathbf{v} = \sum_{j \in F} \mathbf{e}_j^T M \mathbf{v} \quad \Rightarrow \quad \mathbf{e}_j^T M \mathbf{v} = 0 \quad \forall j \in F \\ \Rightarrow \quad \mathbf{e}_j^T M \mathbf{e}_1 &= 0 \quad \forall j \in F \quad \Rightarrow \quad \mathbf{e}_j^T \boldsymbol{\mu} = 0 \quad \forall j \in F \end{aligned}$$

(where the second implication follows from $\mathbf{e}_1^T \mathbf{v} > 0$, and the third follows from (c₄) in Definition 3.2). Hence,

$$z_i = \beta \mathbf{e}_i^T H P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu} = \beta \sum_{g \in G} \lambda_g \mathbf{e}_g^T P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu} = \beta \sum_{g \in G} \lambda_g \sum_{j \in (\pi_{\mathbf{a}} \circ \pi)^{-1}(g)} \mathbf{e}_j^T \boldsymbol{\mu} = 0,$$

so that $u_i = w_i + z_i = 0$, a contradiction. Hence, $w_i > 0$. Observe that

$$(M/\pi_{\mathbf{a}} \circ \pi, \boldsymbol{\mu}/\pi_{\mathbf{a}} \circ \pi) \in \mathcal{C}_D^{(n)},$$

and hence, $M/\pi_{\mathbf{a}} \circ \pi \in S$. By the choice of \mathbf{v} , this implies that the vector $P_{\pi_{\mathbf{a}} \circ \pi} M \mathbf{v} = M/\pi_{\mathbf{a}} \circ \pi \mathbf{v}$ is H -tieless; i.e., $H P_{\pi_{\mathbf{a}} \circ \pi} M \mathbf{v}$ is tieless. It follows that the vector

$$H P_{\pi_{\mathbf{a}} \circ \pi} (D M) (N' \mathbf{v}) = \frac{1}{2\alpha[\sigma_1^H + 1]D} \mathbf{w}$$

is also tieless; being that it is entrywise integer, and since $\frac{1}{2\alpha[\sigma_1^H+1]D}w_i > 0$, we obtain

$$\left| \frac{1}{2\alpha[\sigma_1^H+1]D}w_i - \frac{1}{2\alpha[\sigma_1^H+1]D}w_{i'} \right| \geq 1, \quad \text{which yields} \quad |w_i - w_{i'}| \geq 2\alpha[\sigma_1^H+1]D.$$

Denote the ℓ_1 -norm and the ℓ_2 -norm of a vector by $\|\cdot\|_1$ and $\|\cdot\|_2$, respectively. Recall that the largest singular value of a matrix is its spectral operator norm, i.e., $\sigma_1^H = \max_{\mathbf{0} \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|H\mathbf{x}\|_2}{\|\mathbf{x}\|_2}$ (see [47]). In particular, $\|H\mathbf{x}\|_2 \leq \sigma_1^H \|\mathbf{x}\|_2$ for each vector \mathbf{x} of size n . Using the Cauchy-Schwarz inequality and the fact that the ℓ_1 -norm of a vector is greater than or equal to its ℓ_2 -norm, we find

$$\begin{aligned} |z_i - z_{i'}| &= \beta |(\mathbf{e}_i - \mathbf{e}_{i'})^T H P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}| \leq \beta \|\mathbf{e}_i - \mathbf{e}_{i'}\|_2 \|H P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}\|_2 \\ &\leq \beta \|\mathbf{e}_i - \mathbf{e}_{i'}\|_2 \sigma_1^H \|P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}\|_2 \\ &\leq \beta \|\mathbf{e}_i - \mathbf{e}_{i'}\|_1 [\sigma_1^H + 1] \|P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}\|_1 \\ &= 2\beta [\sigma_1^H + 1] \mathbf{1}_n^T P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu} \leq 2\beta [\sigma_1^H + 1] \mathbf{1}_n^T P_{\pi_{\mathbf{a}} \circ \pi} |\boldsymbol{\mu}| \\ &= 2\beta [\sigma_1^H + 1] \mathbf{1}_L^T |\boldsymbol{\mu}| \leq 2\beta [\sigma_1^H + 1] D < 2\alpha [\sigma_1^H + 1] D. \end{aligned}$$

We conclude the proof of the claim by noting that

$$\begin{aligned} |u_i - u_{i'}| &= |(w_i - w_{i'}) - (z_{i'} - z_i)| \geq |w_i - w_{i'}| - |z_i - z_{i'}| \\ &> 2\alpha [\sigma_1^H + 1] D - 2\alpha [\sigma_1^H + 1] D = 0, \end{aligned}$$

which implies $u_i \neq u_{i'}$. As a consequence, the vector $P_{\pi_{\mathbf{a}} \circ \tilde{\varphi}} \mathbf{1}_c$ is H -tieless. We can then apply Lemma 5.1 to conclude that $f_{/\pi_{\mathbf{a}} \circ \tilde{\varphi}}(1, \dots, n) = f_{/\pi_{\mathbf{a}} \circ \pi \circ \varphi}(1, \dots, n)$. Hence, by (5.2), $f_{/\tilde{\varphi}} = f_{/\pi \circ \varphi}$. Therefore, $\xi_D((M, \boldsymbol{\mu})_{/\pi}) = \xi_D((M, \boldsymbol{\mu}))_{/\pi}$, as required. It follows that ξ_D is a minion homomorphism.

Since the set of polymorphisms of (\mathbf{A}, \mathbf{B}) of arity L is finite for every $L \in \mathbb{N}$, we can apply Proposition 4.5 to conclude that there exists a minion homomorphism $\zeta : \mathcal{C} \rightarrow \text{Pol}(\mathbf{A}, \mathbf{B})$. \square

Remark 5.5. If $\text{Pol}(\mathbf{A}, \mathbf{B})$ contains H -symmetric operations of *all* arities—as it happens for the PCSP template (\mathbf{A}, \mathbf{B}) from Example 2.10, cf. Example 3.6—the AIP part of CLAP is not required. Indeed, in that case, we can choose f in the proof of Theorem 3.5 to be an H -symmetric polymorphism of arity $c = N^2$, which implies $\beta = 0$. Therefore, the affine vector $\boldsymbol{\mu}$ does not have any role in the definition of P_φ in (5.1), nor in the definition of the minion homomorphism ξ_D . It follows that, under this stronger hypothesis, $\text{Pol}(\mathbf{A}, \mathbf{B})$ admits a minion homomorphism from a minion \mathcal{C} whose L -ary elements are matrices in $\mathbb{Q}^{L, \mathbb{N}_0}$ satisfying the requirements $(c_1), (c_2), (c_5), (c_6)$ of Definition 3.2; notice that the projection $(M, \boldsymbol{\mu}) \mapsto M$ yields a natural minion homomorphism from \mathcal{C} to \mathcal{C} . The proofs of Theorems 4.10 and 4.11 can be straightforwardly modified to show that \mathcal{C} captures the power of the algorithm CBLP, i.e., the simplified version of CLAP that does not run BLP + AIP at the end (cf. the discussion in section 3).

Appendix A. Existing relaxations for PCSPs. Every CSP can be equivalently expressed as a 0–1 integer program in a standard way.

If the variables are allowed to take values in $[0, 1]$, we obtain the so-called basic linear programming relaxation (BLP) [59]. This naturally extends to PCSPs [8], as we describe in Appendix A.1.

$$\begin{aligned}
0 &\leq \lambda_{\mathbf{x},R}(\mathbf{a}) \leq 1 && \forall R \in \sigma, \forall \mathbf{x} \in R^{\mathbf{X}}, \forall \mathbf{a} \in R^{\mathbf{A}} \\
\sum_{\mathbf{a} \in R^{\mathbf{A}}} \lambda_{\mathbf{x},R}(\mathbf{a}) &= 1 && \forall R \in \sigma, \forall \mathbf{x} \in R^{\mathbf{X}} \\
\sum_{\mathbf{a} \in R^{\mathbf{A}}, a_i = a} \lambda_{\mathbf{x},R}(\mathbf{a}) &= \lambda_{x_i, R_u}(a) && \forall R \in \sigma, \forall \mathbf{x} \in R^{\mathbf{X}}, \forall a \in A, \forall i \in [\text{ar}(R)]
\end{aligned}$$

FIG. A.1. Definition of BLP(\mathbf{X}, \mathbf{A}).

$$\begin{aligned}
\tau_{\mathbf{x},R}(\mathbf{a}) &\in \mathbb{Z} && \forall R \in \sigma, \forall \mathbf{x} \in R^{\mathbf{X}}, \forall \mathbf{a} \in R^{\mathbf{A}} \\
\sum_{\mathbf{a} \in R^{\mathbf{A}}} \tau_{\mathbf{x},R}(\mathbf{a}) &= 1 && \forall R \in \sigma, \forall \mathbf{x} \in R^{\mathbf{X}} \\
\sum_{\mathbf{a} \in R^{\mathbf{A}}, a_i = a} \tau_{\mathbf{x},R}(\mathbf{a}) &= \tau_{x_i, R_u}(a) && \forall R \in \sigma, \forall \mathbf{x} \in R^{\mathbf{X}}, \forall a \in A, \forall i \in [\text{ar}(R)]
\end{aligned}$$

FIG. A.2. Definition of AIP(\mathbf{X}, \mathbf{A}).

If the variables are allowed to take integer values, we obtain the so-called basic affine integer programming relaxation (AIP) [20], studied in detail in [8], as we describe in Appendix A.2.

A combination of the two relaxations, called the BLP + AIP relaxation, was proposed in [22] and its power characterized in [22], as we describe in Appendix A.3.

Let (\mathbf{A}, \mathbf{B}) be a PCSP template with signature σ , and let \mathbf{X} be an instance of PCSP(\mathbf{A}, \mathbf{B}). In all three relaxations described below, we assume without loss of generality that σ contains a unary symbol R_u such that $R_u^{\mathbf{X}} = X$, $R_u^{\mathbf{A}} = A$, and $R_u^{\mathbf{B}} = B$. If this is not the case, the signature and the instance can be extended without changing the set of solutions.

A.1. BLP. The BLP of \mathbf{X} , denoted by $\text{BLP}(\mathbf{X}, \mathbf{A})$, is defined as follows.¹⁰ The variables are $\lambda_{\mathbf{x},R}(\mathbf{a})$ for every $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{a} \in R^{\mathbf{A}}$, and the constraints are given in Figure A.1.

We say that $\text{BLP}(\mathbf{X}, \mathbf{A})$ accepts if the LP in Figure A.1 is feasible, and rejects otherwise. By construction, if $\mathbf{X} \rightarrow \mathbf{A}$, then $\text{BLP}(\mathbf{X}, \mathbf{A})$ accepts. We say that BLP solves PCSP(\mathbf{A}, \mathbf{B}) if for every instance \mathbf{X} accepted by $\text{BLP}(\mathbf{X}, \mathbf{A})$ we have $\mathbf{X} \rightarrow \mathbf{B}$.

We denote by $\mathcal{Q}_{\text{conv}}$ the minion of stochastic vectors on \mathbb{Q} with the minor operation defined as in section 4.1; i.e., if $\mathbf{q} \in \mathcal{Q}_{\text{conv}}^{(L)}$ and $\pi : [L] \rightarrow [L']$, then $\mathbf{q}_{/\pi} = P_{\pi}\mathbf{q}$, where P_{π} is the $L' \times L$ matrix whose (i, j) th entry is 1 if $\pi(j) = i$, and 0 otherwise.

An L -ary operation $f : A^L \rightarrow B$ is called *symmetric* if

$$f(a_1, \dots, a_L) = f(a_{\pi(1)}, \dots, a_{\pi(L)})$$

for every $a_1, \dots, a_L \in A$ and every permutation $\pi : [L] \rightarrow [L]$.

The power of BLP for PCSPs is characterized in the following result.

THEOREM A.1 (see [8]). *Let (\mathbf{A}, \mathbf{B}) be a PCSP template. The following are equivalent:*

¹⁰The definition does not depend on \mathbf{B} and is the same as the BLP of an instance \mathbf{X} of CSP(\mathbf{A}); the same holds for AIP and BLP + AIP.

Algorithm A.1. The BLP + AIP algorithm.

Require: an instance \mathbf{X} of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ of signature σ

Ensure: YES if $\mathbf{X} \rightarrow \mathbf{A}$ and no if $\mathbf{X} \not\rightarrow \mathbf{B}$ find a relative interior point

$(\lambda_{\mathbf{x}, R}(\mathbf{a}))_{R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}}$ of $\text{BLP}(\mathbf{X}, \mathbf{A})$

- 1: **if** no relative interior point exists **then**
- 2: **return** NO
- 3: **end if**
- 4: refine $\text{AIP}(\mathbf{X}, \mathbf{A})$ by setting $\tau_{\mathbf{x}, R}(\mathbf{a}) = 0$ if $\lambda_{\mathbf{x}, R}(\mathbf{a}) = 0$
- 5: **if** the refined $\text{AIP}(\mathbf{X}, \mathbf{A})$ accepts **then**
- 6: **return** YES
- 7: **else**
- 8: **return** NO
- 9: **end if**

- (1) BLP solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$.
- (2) $\text{Pol}(\mathbf{A}, \mathbf{B})$ admits a minion homomorphism from $\mathcal{Q}_{\text{conv}}$.
- (3) $\text{Pol}(\mathbf{A}, \mathbf{B})$ contains symmetric operations of all arities.

A.2. AIP The AIP of \mathbf{X} , denoted by $\text{AIP}(\mathbf{X}, \mathbf{A})$, is defined as follows. The variables are $\tau_{\mathbf{x}, R}(\mathbf{a})$ for every $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{a} \in R^{\mathbf{A}}$, and the constraints are given in Figure A.2.

We say that $\text{AIP}(\mathbf{X}, \mathbf{A})$ accepts if the affine program in Figure A.2 is feasible, and rejects otherwise. By construction, if $\mathbf{X} \rightarrow \mathbf{A}$, then $\text{AIP}(\mathbf{X}, \mathbf{A})$ accepts. We say that AIP solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$ if for every instance \mathbf{X} accepted by $\text{AIP}(\mathbf{X}, \mathbf{A})$ we have $\mathbf{X} \rightarrow \mathbf{B}$.

We denote by \mathcal{Z}_{aff} the minion of affine vectors on \mathbb{Z} with the minor operation defined as in section 4.1; i.e., if $\mathbf{z} \in \mathcal{Z}_{\text{aff}}^{(L)}$ and $\pi : [L] \rightarrow [L']$, then $\mathbf{z}_{/\pi} = P_{\pi} \mathbf{z}$, where P_{π} is the $L' \times L$ matrix whose (i, j) th entry is 1 if $\pi(j) = i$, and 0 otherwise.

A $(2L+1)$ -ary operation $f : A^{2L+1} \rightarrow B$ is called *alternating* if $f(a_1, \dots, a_{2L+1}) = f(a_{\pi(1)}, \dots, a_{\pi(2L+1)})$ for every $a_1, \dots, a_{2L+1} \in A$ and every permutation $\pi : [2L+1] \rightarrow [2L+1]$ that preserves parity, and $f(a_1, \dots, a_{2L-1}, a, a) = f(a_1, \dots, a_{2L-1}, a', a')$ for every $a_1, \dots, a_{2L-1}, a, a' \in A$. Intuitively, an alternating operation is invariant under permutations of its odd and even coordinates and has the property that adjacent coordinates cancel each other out.

The power of AIP for PCSPs is characterized in the following result.

THEOREM A.2 (see [8]). *Let (\mathbf{A}, \mathbf{B}) be a PCSP template. The following are equivalent:*

- (1) AIP solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$.
- (2) $\text{Pol}(\mathbf{A}, \mathbf{B})$ admits a minion homomorphism from \mathcal{Z}_{aff} .
- (3) $\text{Pol}(\mathbf{A}, \mathbf{B})$ contains alternating operations of all odd arities.

A.3. BLP+AIP The combined basic LP and affine IP algorithm (BLP + AIP) is presented in Algorithm A.1.

If $\mathbf{X} \rightarrow \mathbf{A}$, then BLP + AIP accepts \mathbf{X} [22]. We say that BLP + AIP solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$ if for every instance \mathbf{X} accepted by BLP + AIP we have $\mathbf{X} \rightarrow \mathbf{B}$.

We denote by $\mathcal{M}_{\text{BLP+AIP}}$ the minion whose L -ary objects are pairs (\mathbf{q}, \mathbf{z}) , where $\mathbf{q} \in \mathbb{Q}^L$ is a stochastic vector and $\mathbf{z} \in \mathbb{Z}^L$ is an affine vector, with the property that, for every $i \in [L]$, $q_i = 0$ implies $z_i = 0$. As before, the minor operation is defined as in

section 4.1; i.e., if $(\mathbf{q}, \mathbf{z}) \in \mathcal{M}_{\text{BLP}+\text{AIP}}^{(L)}$ and $\pi: [L] \rightarrow [L']$, then $(\mathbf{q}, \mathbf{z})/\pi = (P_\pi \mathbf{q}, P_\pi \mathbf{z})$, where P_π is the $L' \times L$ matrix whose (i, j) th entry is 1 if $\pi(j) = i$, and 0 otherwise.

A $(2L+1)$ -ary operation $f: A^{2L+1} \rightarrow B$ is called *2-block symmetric* if

$$f(a_1, \dots, a_{2L+1}) = f(a_{\pi(1)}, \dots, a_{\pi(2L+1)})$$

for every $a_1, \dots, a_{2L+1} \in A$ and every permutation $\pi: [2L+1] \rightarrow [2L+1]$ that preserves parity.

The power of BLP + AIP for PCSPs is characterized in the following result.

THEOREM A.3 (see [22]). *Let (\mathbf{A}, \mathbf{B}) be a PCSP template. The following are equivalent:*

- (1) BLP + AIP solves PCSP (\mathbf{A}, \mathbf{B}) .
- (2) Pol (\mathbf{A}, \mathbf{B}) admits a minion homomorphism from $\mathcal{M}_{\text{BLP}+\text{AIP}}$.
- (3) Pol (\mathbf{A}, \mathbf{B}) contains 2-block-symmetric operations of all odd arities.

Appendix B. Proof of Lemma 4.2. In this section, we shall prove Lemma 4.2, which we restate below. The proof is based on that of [8, Lemma 4.4], which concerns minions of functions.

LEMMA B.1 (Lemma 4.2 restated). *Let \mathcal{M} be a minion, and let (\mathbf{A}, \mathbf{B}) be a PCSP template. Then there is a minion homomorphism from \mathcal{M} to Pol (\mathbf{A}, \mathbf{B}) if and only if $\mathbb{F}_{\mathcal{M}}(\mathbf{A}) \rightarrow \mathbf{B}$.*

Proof. Let $A = [n]$, and let σ be the signature of \mathbf{A} and \mathbf{B} . Suppose $\xi: \mathcal{M} \rightarrow \text{Pol}(\mathbf{A}, \mathbf{B})$ is a minion homomorphism, and consider the function

$$\begin{aligned} f: \mathcal{M}^{(n)} &\rightarrow B \\ M &\mapsto \xi(M)(1, \dots, n). \end{aligned}$$

For $R \in \sigma$ of arity k , consider a tuple $(M_1, \dots, M_k) \in R^{\mathbb{F}_{\mathcal{M}}(\mathbf{A})}$. List the elements of $R^{\mathbf{A}}$ as $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(m)}$. From Definition 4.1, $\exists Q \in \mathcal{M}^{(m)}$ such that $M_i = Q_{/\pi_i}$ for each $i \in [k]$, where $\pi_i: [m] \rightarrow A$ maps j to the i th coordinate of $\mathbf{a}^{(j)}$. It follows that, for each $i \in [k]$,

$$f(M_i) = f(Q_{/\pi_i}) = \xi(Q_{/\pi_i})(1, \dots, n) = \xi(Q)_{/\pi_i}(1, \dots, n) = \xi(Q)(\pi_i(1), \dots, \pi_i(m)).$$

Hence,

$$\begin{aligned} f(M_1, \dots, M_k) &= (\xi(Q)(\pi_1(1), \dots, \pi_1(m)), \dots, \xi(Q)(\pi_k(1), \dots, \pi_k(m))) \\ &= \xi(Q)(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(m)}) \in R^{\mathbf{B}} \end{aligned}$$

since $\xi(Q)$ is a polymorphism of (\mathbf{A}, \mathbf{B}) . Therefore, f is a homomorphism from $\mathbb{F}_{\mathcal{M}}(\mathbf{A})$ to \mathbf{B} .

Conversely, let $f: \mathbb{F}_{\mathcal{M}}(\mathbf{A}) \rightarrow \mathbf{B}$ be a homomorphism, and consider the function $\xi: \mathcal{M} \rightarrow \text{Pol}(\mathbf{A}, \mathbf{B})$ defined by $\xi(M)(a_1, \dots, a_L) = f(M_{/\rho})$ for each $L \in \mathbb{N}$, $M \in \mathcal{M}^{(L)}$, $(a_1, \dots, a_L) \in A^L$, where

$$\begin{aligned} \rho: [L] &\rightarrow [n] \\ i &\mapsto a_i. \end{aligned}$$

Let us first check that ξ is well defined, i.e., that $\xi(M) \in \text{Pol}(\mathbf{A}, \mathbf{B})$. For $R \in \sigma$ of arity k , consider a matrix $Z \in A^{L,k}$ such that each row of Z corresponds to a tuple in $R^{\mathbf{A}}$. We need to show that $\xi(M)(Z) \in R^{\mathbf{B}}$. Consider the maps

$$\begin{aligned} \tau : [L] &\rightarrow R^{\mathbf{A}}, & \rho_j : [L] &\rightarrow [n], & \pi_j : R^{\mathbf{A}} &\rightarrow [n], \\ i &\mapsto Z^T \mathbf{e}_i, & i &\mapsto \mathbf{e}_i^T Z \mathbf{e}_j, & \mathbf{a} &\mapsto \mathbf{e}_j^T \mathbf{a} \end{aligned}$$

for $j \in [k]$. Observe that $\rho_j = \pi_j \circ \tau$, and set $Q = M_{/\tau} \in \mathcal{M}(|R^{\mathbf{A}}|)$. We obtain

$$\xi(M)(Z) = f(M_{/\rho_1}, \dots, M_{/\rho_k}) = f(M_{/\pi_1 \circ \tau}, \dots, M_{/\pi_k \circ \tau}) = f(Q_{/\pi_1}, \dots, Q_{/\pi_k}) \in R^{\mathbf{B}}$$

since $(Q_{/\pi_1}, \dots, Q_{/\pi_k}) \in R^{\mathbb{F}\mathcal{M}(\mathbf{A})}$ and f is a homomorphism. Finally, we show that ξ is a minion homomorphism. Clearly, ξ preserves arities. To check that it preserves minors, let $M \in \mathcal{M}^{(L)}$ and take a map $\pi : [L] \rightarrow [L']$. Given $(a_1, \dots, a_{L'}) \in A^{L'}$, consider the maps

$$\begin{aligned} \rho' : [L'] &\rightarrow [n], & \rho'' : [L] &\rightarrow [n], \\ i &\mapsto a_i, & i &\mapsto a_{\pi(i)}, \end{aligned}$$

and observe that $\rho'' = \rho' \circ \pi$. We obtain

$$\begin{aligned} \xi(M_{/\pi})(a_1, \dots, a_{L'}) &= f((M_{/\pi})_{/\rho'}) = f(M_{/\rho' \circ \pi}) = f(M_{/\rho''}) = \xi(M)(a_{\pi(1)}, \dots, a_{\pi(L)}) \\ &= \xi(M)_{/\pi}(a_1, \dots, a_{L'}), \end{aligned}$$

which yields $\xi(M_{/\pi}) = \xi(M)_{/\pi}$, as desired. \square

Appendix C. H-block-symmetric polymorphisms. Let $\mathcal{C} = (\mathcal{C}_1, \dots, \mathcal{C}_\ell)$ be a partition of $c \in \mathbb{N}$; i.e., the sets \mathcal{C}_i are pairwise disjoint and their union is $[c]$. Let $c_i = |\mathcal{C}_i|$, so that $c = \sum_{i \in [\ell]} c_i$. For each $i \in [\ell]$, we consider the unique monotonically increasing function $\vartheta_i : [c_i] \rightarrow [c]$ such that $\vartheta_i([c_i]) = \mathcal{C}_i$. We also consider the function $\chi_i : \mathcal{C}_i \rightarrow [c_i]$ such that $\vartheta_i \circ \chi_i$ is the inclusion map of \mathcal{C}_i in $[c]$. Given $c' \in \mathbb{N}$ and a map $\pi : [c] \rightarrow [c']$, we let $\pi_{(i)} = \pi \circ \vartheta_i$.

DEFINITION C.1. Let A, B be finite sets, and consider a function $f : A^c \rightarrow B$ for some $c \in \mathbb{N}$. Given an $m \times |A|$ tie matrix H and a partition $\mathcal{C} = (\mathcal{C}_1, \dots, \mathcal{C}_\ell)$ of c , we say that f is H - \mathcal{C} -block-symmetric if

$$\begin{aligned} f_{/\pi}(\mathbf{a}) &= f(\mathbf{a}) & \forall \pi : [c] \rightarrow [c'] \text{ permutation such that } \pi(\mathcal{C}_i) = \mathcal{C}_i \quad \forall i \in [\ell], \\ & & \forall \mathbf{a} \in A^c \text{ such that } (P_{\vartheta_i}^T \mathbf{a})^\# \text{ is } H\text{-tieless } \forall i \in [\ell]. \end{aligned}$$

We say that f is H -block-symmetric with width W if W is the largest integer for which there is a partition \mathcal{C} of c such that each part of \mathcal{C} has size at least W and f is H - \mathcal{C} -block-symmetric.¹¹ Without loss of generality, we consider $A = [n]$.

LEMMA C.2. Let $f : A^c \rightarrow B$ be H - \mathcal{C} -block-symmetric for some $m \times n$ tie matrix H ($m \in \mathbb{N}$) and some partition $\mathcal{C} = (\mathcal{C}_1, \dots, \mathcal{C}_\ell)$ of c . Consider two maps $\pi, \tilde{\pi} : [c] \rightarrow [n]$ such that, for each $i \in [\ell]$, $P_{\pi(i)} \mathbf{1}_{c_i} = P_{\tilde{\pi}(i)} \mathbf{1}_{c_i}$ and the vector $P_{\pi(i)} \mathbf{1}_{c_i}$ is H -tieless. Then

$$f_{/\pi}(1, \dots, n) = f_{/\tilde{\pi}}(1, \dots, n).$$

¹¹The notion of H -block-symmetric operation is the H -analogue of that of block-symmetric operation in [22] (cf. Theorem A.3).

Proof. For $i \in [\ell]$ and $a \in [n]$, we have

$$|\pi_{(i)}^{-1}(a)| = \mathbf{e}_a^T P_{\pi_{(i)}} \mathbf{1}_{c_i} = \mathbf{e}_a^T P_{\tilde{\pi}_{(i)}} \mathbf{1}_{c_i} = |\tilde{\pi}_{(i)}^{-1}(a)|.$$

Hence, we can consider bijections $\varphi_{i,a} : \pi_{(i)}^{-1}(a) \rightarrow \tilde{\pi}_{(i)}^{-1}(a)$ for each $i \in [\ell], a \in [n]$. The union

$$\varphi_i = \bigcup_{a \in [n]} \varphi_{i,a} : [c_i] \rightarrow [c_i]$$

is also a bijection. Define $\varphi : [c] \rightarrow [c]$ by letting $\varphi|_{c_i} = \vartheta_i \circ \varphi_i \circ \chi_i$ for each $i \in [\ell]$. Notice that $\varphi(\mathcal{C}_i) = \mathcal{C}_i$ for each $i \in [\ell]$, so φ is a bijection. Take $j \in [c]$ and suppose that $j \in \mathcal{C}_i$. We have

$$\begin{aligned} (\tilde{\pi} \circ \varphi)(j) &= \tilde{\pi}(\varphi(j)) = \tilde{\pi}(\vartheta_i(\varphi_i(\chi_i(j)))) = \tilde{\pi}_{(i)}(\varphi_{i,\pi_{(i)}(\chi_i(j))}(\chi_i(j))) \\ &= \pi_{(i)}(\chi_i(j)) = (\pi \circ \vartheta_i \circ \chi_i)(j) = \pi(j), \end{aligned}$$

and, hence, $\tilde{\pi} \circ \varphi = \pi$. Let $\tilde{\mathbf{a}} = (\tilde{\pi}(1), \dots, \tilde{\pi}(c))$. Notice that, for each $i \in [\ell]$ and $a \in [n]$,

$$\begin{aligned} \mathbf{e}_a^T (P_{\vartheta_i}^T \tilde{\mathbf{a}})^\# &= |\{j \in [c_i] : \mathbf{e}_j^T P_{\vartheta_i}^T \tilde{\mathbf{a}} = a\}| = |\{j \in [c_i] : \mathbf{e}_{\vartheta_i(j)}^T \tilde{\mathbf{a}} = a\}| \\ &= |\{j \in [c_i] : \tilde{\pi}(\vartheta_i(j)) = a\}| \\ &= |\{j \in [c_i] : \tilde{\pi}_{(i)}(j) = a\}| = \mathbf{e}_a^T P_{\tilde{\pi}_{(i)}} \mathbf{1}_{c_i}, \end{aligned}$$

and therefore, $(P_{\vartheta_i}^T \tilde{\mathbf{a}})^\# = P_{\tilde{\pi}_{(i)}} \mathbf{1}_{c_i} = P_{\pi_{(i)}} \mathbf{1}_{c_i}$, which is H -tieless. Using that f is H - \mathcal{C} -block-symmetric, we find

$$f_{/\tilde{\pi}}(1, \dots, n) = f(\tilde{\mathbf{a}}) = f_{/\varphi}(\tilde{\mathbf{a}}) = (f_{/\varphi})_{/\tilde{\pi}}(1, \dots, n) = f_{/\tilde{\pi} \circ \varphi}(1, \dots, n) = f_{/\pi}(1, \dots, n),$$

as required. \square

THEOREM C.3. *Let (\mathbf{A}, \mathbf{B}) be a PCSP template, and suppose $\text{Pol}(\mathbf{A}, \mathbf{B})$ contains H -block-symmetric operations of arbitrarily large width for some $m \times |A|$ tie matrix H , $m \in \mathbb{N}$. Then there exists a minion homomorphism from \mathcal{C} to $\text{Pol}(\mathbf{A}, \mathbf{B})$.*

Proof. For $D \in \mathbb{N}$, consider the subminion \mathcal{C}_D of \mathcal{C} described in section 4.2. Observe that $S = \{M : (M, \boldsymbol{\mu}) \in \mathcal{C}_D^{(n)}\}$ is a finite set of skeletal matrices. Therefore, we can apply the tiebreak lemma, Lemma 5.2, to find a stochastic finitely supported vector $\mathbf{v} \in \mathbb{Q}^{\mathbb{N}_0}$ with $\mathbf{e}_1^T \mathbf{v} > 0$ such that $M\mathbf{v}$ is H -tieless for any $M \in S$. Since \mathbf{v} is finitely supported, we can find $N' \in \mathbb{N}$ such that $N'\mathbf{v}$ has integer entries. Let σ_1^H denote the largest singular value of H , i.e., the square root of the largest eigenvalue of $H^T H$. Set $N = 2\lceil \sigma_1^H + 1 \rceil D^2 N'$, and let f be an H -block-symmetric polymorphism of width $W \geq N^2$. Letting c be the arity of f , consider a partition $\mathcal{C} = (\mathcal{C}_1, \dots, \mathcal{C}_\ell)$ of c such that $c_i = |\mathcal{C}_i| \geq W$ for each $i \in [\ell]$ and f is H - \mathcal{C} -block-symmetric. Write $c_i = N\alpha_i + \beta_i$ with $\alpha_i, \beta_i \in \mathbb{N}_0$, $\beta_i \leq N - 1$. Note that $N^2 \leq W \leq c_i = N\alpha_i + \beta_i \leq N\alpha_i + N - 1 < N(\alpha_i + 1)$, so $N < \alpha_i + 1$ and, hence, $\beta_i < \alpha_i$.

Consider the function

$$\xi_D : \mathcal{C}_D \rightarrow \text{Pol}(\mathbf{A}, \mathbf{B})$$

defined as follows. Given $L \in \mathbb{N}$ and $(M, \boldsymbol{\mu}) \in \mathcal{C}_D^{(L)}$, for each $i \in [\ell]$ take the map $\varphi_i : [c_i] \rightarrow [L]$ such that the corresponding $L \times c_i$ matrix P_{φ_i} is

$$(C.1) \quad P_{\varphi_i} = \begin{pmatrix} \mathbf{1}_{\mathbf{e}_1^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu})}^T & \mathbf{0}^T & \cdots & \mathbf{0}^T \\ \mathbf{0}^T & \mathbf{1}_{\mathbf{e}_2^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu})}^T & \cdots & \mathbf{0}^T \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}^T & \mathbf{0}^T & \cdots & \mathbf{1}_{\mathbf{e}_L^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu})}^T \end{pmatrix}.$$

To verify that (C.1) is well defined, observe first that

$$\begin{aligned} \sum_{j=1}^L \mathbf{e}_j^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu}) &= \mathbf{1}_L^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu}) = \alpha_i N \mathbf{1}_L^T M \mathbf{v} + \beta_i \mathbf{1}_L^T \boldsymbol{\mu} \\ &= \alpha_i N \mathbf{1}_{\mathbb{N}_0}^T \mathbf{v} + \beta_i = \alpha_i N + \beta_i = c_i. \end{aligned}$$

Moreover, for each $j \in [L]$, $\mathbf{e}_j^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu}) = \mathbf{e}_j^T(2\alpha_i[\sigma_1^H + 1]D(DM)(N'\mathbf{v}) + \beta_i \boldsymbol{\mu})$ is an integer. If $\mathbf{e}_j^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu})$ was negative, then $\mathbf{e}_j^T \boldsymbol{\mu} < 0$. By the requirement (c_4) in Definition 3.2, this would imply that $\mathbf{e}_j^T M \mathbf{e}_1 > 0$, and hence, $0 < \mathbf{e}_j^T M \mathbf{e}_1 \mathbf{e}_1^T \mathbf{v} \leq \mathbf{e}_j^T M \mathbf{v}$. As a consequence, $\mathbf{e}_j^T(DM)(N'\mathbf{v}) \geq 1$ so that

$$\begin{aligned} \mathbf{e}_j^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu}) &= 2\alpha_i[\sigma_1^H + 1]D\mathbf{e}_j^T(DM)(N'\mathbf{v}) + \beta_i \mathbf{e}_j^T \boldsymbol{\mu} \\ &\geq 2\alpha_i[\sigma_1^H + 1]D + \beta_i \mathbf{e}_j^T \boldsymbol{\mu} \geq \alpha_i D - \beta_i D > 0, \end{aligned}$$

which is a contradiction. In conclusion, the numbers $\mathbf{e}_j^T(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu})$ are nonnegative integers summing up to c_i , so (C.1) is well defined.

Consider the function $\varphi : [c] \rightarrow [L]$ defined by $\varphi|_{c_i} = \varphi_i \circ \chi_i \ \forall i \in [\ell]$, and let $\xi_D((M, \boldsymbol{\mu})) := f_{/\varphi}$. Clearly, $\xi_D((M, \boldsymbol{\mu})) \in \text{Pol}(\mathbf{A}, \mathbf{B})$. We claim that the map ξ_D is a minion homomorphism. It is straightforward to check that ξ_D preserves arities so, to conclude, we need to show that it also preserves minors. Take $L' \in \mathbb{N}$ and choose a map $\pi : [L] \rightarrow [L']$. Letting $\tilde{\varphi}_i : [c_i] \rightarrow [L']$ be the map corresponding to the matrix

$$P_{\tilde{\varphi}_i} = \begin{pmatrix} \mathbf{1}_{\mathbf{e}_1^T(\alpha_i N P_\pi M \mathbf{v} + \beta_i P_\pi \boldsymbol{\mu})}^T & \mathbf{0}^T & \cdots & \mathbf{0}^T \\ \mathbf{0}^T & \mathbf{1}_{\mathbf{e}_2^T(\alpha_i N P_\pi M \mathbf{v} + \beta_i P_\pi \boldsymbol{\mu})}^T & \cdots & \mathbf{0}^T \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}^T & \mathbf{0}^T & \cdots & \mathbf{1}_{\mathbf{e}_{L'}^T(\alpha_i N P_\pi M \mathbf{v} + \beta_i P_\pi \boldsymbol{\mu})}^T \end{pmatrix}$$

for each $i \in [\ell]$, and considering $\tilde{\varphi} : [c] \rightarrow [L']$ such that $\tilde{\varphi}|_{c_i} = \tilde{\varphi}_i \circ \chi_i \ \forall i \in [\ell]$, we see that $\xi_D((M, \boldsymbol{\mu})/\pi) = f_{/\tilde{\varphi}}$. Moreover, $\xi_D((M, \boldsymbol{\mu}))/\pi = (f_{/\varphi})/\pi = f_{/\pi \circ \varphi}$, where φ is the map defined above. Take $\mathbf{a} = (a_1, \dots, a_{L'}) \in A^{L'}$, and consider the map

$$\begin{aligned} \pi_{\mathbf{a}} : [L'] &\rightarrow [n] \\ i &\mapsto a_i. \end{aligned}$$

Observe that

$$(C.2) \quad \begin{aligned} f_{/\tilde{\varphi}}(\mathbf{a}) &= (f_{/\tilde{\varphi}})_{/\pi_{\mathbf{a}}}(1, \dots, n) = f_{/\pi_{\mathbf{a}} \circ \tilde{\varphi}}(1, \dots, n), \quad \text{and similarly,} \\ f_{/\pi \circ \varphi}(\mathbf{a}) &= (f_{/\pi \circ \varphi})_{/\pi_{\mathbf{a}}}(1, \dots, n) = f_{/\pi_{\mathbf{a}} \circ \pi \circ \varphi}(1, \dots, n). \end{aligned}$$

Notice that, for each $i \in [\ell]$, $\varphi \circ \vartheta_i = \varphi_i$ and $\tilde{\varphi} \circ \vartheta_i = \tilde{\varphi}_i$. Hence,

$$\begin{aligned} P_{(\pi_{\mathbf{a}} \circ \tilde{\varphi})(i)} \mathbf{1}_{c_i} &= P_{\pi_{\mathbf{a}} \circ \tilde{\varphi} \circ \vartheta_i} \mathbf{1}_{c_i} = P_{\pi_{\mathbf{a}}} P_{\tilde{\varphi} \circ \vartheta_i} \mathbf{1}_{c_i} = P_{\pi_{\mathbf{a}}} P_{\tilde{\varphi}_i} \mathbf{1}_{c_i} = P_{\pi_{\mathbf{a}}}(\alpha_i N P_\pi M \mathbf{v} + \beta_i P_\pi \boldsymbol{\mu}) \\ &= P_{\pi_{\mathbf{a}}} P_\pi(\alpha_i N M \mathbf{v} + \beta_i \boldsymbol{\mu}) = P_{\pi_{\mathbf{a}}} P_\pi P_{\varphi_i} \mathbf{1}_{c_i} \\ &= P_{\pi_{\mathbf{a}} \circ \pi \circ \varphi \circ \vartheta_i} \mathbf{1}_{c_i} = P_{(\pi_{\mathbf{a}} \circ \pi \circ \varphi)(i)} \mathbf{1}_{c_i}. \end{aligned}$$

We claim that the vector $P_{(\pi_{\mathbf{a}} \circ \tilde{\varphi})_{(i)}} \mathbf{1}_{c_i} = P_{\pi_{\mathbf{a}} \circ \tilde{\varphi}_i} \mathbf{1}_{c_i}$ is H -tieless. Let $\mathbf{u} = (u_t) = HP_{\pi_{\mathbf{a}} \circ \tilde{\varphi}_i} \mathbf{1}_{c_i}$; the claim is equivalent to \mathbf{u} being tieless. Let

$$\mathbf{w} = (w_t) = \alpha_i NHP_{\pi_{\mathbf{a}} \circ \pi} M\mathbf{v}$$

and $\mathbf{z} = (z_t) = \beta_i HP_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}$, so that $\mathbf{u} = \mathbf{w} + \mathbf{z}$. Choose $t, t' \in [m]$ such that $t \neq t'$ and $u_t \neq 0$. We need to show that $u_t \neq u_{t'}$. Suppose $w_t = 0$. We can write $H^T \mathbf{e}_t = \sum_{g \in G} \lambda_g \mathbf{e}_g$ for $G = \text{supp}(H^T \mathbf{e}_t)$, where each λ_g is a positive integer (note that $G \neq \emptyset$ since, otherwise, $H^T \mathbf{e}_t = \mathbf{0}_n$, which would imply $u_t = 0$). Let $F = (\pi_{\mathbf{a}} \circ \pi)^{-1}(G)$. From $w_t = 0$, we obtain

$$\begin{aligned} 0 = \mathbf{e}_t^T HP_{\pi_{\mathbf{a}} \circ \pi} M\mathbf{v} &= (H^T \mathbf{e}_t)^T P_{\pi_{\mathbf{a}} \circ \pi} M\mathbf{v} = \sum_{g \in G} \lambda_g \mathbf{e}_g^T P_{\pi_{\mathbf{a}} \circ \pi} M\mathbf{v} \\ &= \sum_{g \in G} \lambda_g \sum_{j \in (\pi_{\mathbf{a}} \circ \pi)^{-1}(g)} \mathbf{e}_j^T M\mathbf{v}, \end{aligned}$$

and hence, the following chain of implications holds:

$$\begin{aligned} 0 &= \sum_{g \in G} \sum_{j \in (\pi_{\mathbf{a}} \circ \pi)^{-1}(g)} \mathbf{e}_j^T M\mathbf{v} = \sum_{j \in F} \mathbf{e}_j^T M\mathbf{v} \Rightarrow \mathbf{e}_j^T M\mathbf{v} = 0 \quad \forall j \in F \\ \Rightarrow \mathbf{e}_j^T M\mathbf{e}_1 &= 0 \quad \forall j \in F \Rightarrow \mathbf{e}_j^T \boldsymbol{\mu} = 0 \quad \forall j \in F \end{aligned}$$

(where the second implication follows from $\mathbf{e}_1^T \mathbf{v} > 0$, and the third follows from (c_4) in Definition 3.2). Hence,

$$z_t = \beta_i \mathbf{e}_t^T HP_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu} = \beta_i \sum_{g \in G} \lambda_g \mathbf{e}_g^T P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu} = \beta_i \sum_{g \in G} \lambda_g \sum_{j \in (\pi_{\mathbf{a}} \circ \pi)^{-1}(g)} \mathbf{e}_j^T \boldsymbol{\mu} = 0,$$

so that $u_t = w_t + z_t = 0$, a contradiction. Hence, $w_t > 0$. Observe that

$$(M/\pi_{\mathbf{a}} \circ \pi, \boldsymbol{\mu}/\pi_{\mathbf{a}} \circ \pi) \in \mathcal{C}_D^{(n)},$$

and hence, $M/\pi_{\mathbf{a}} \circ \pi \in S$. By the choice of \mathbf{v} , this implies that the vector $P_{\pi_{\mathbf{a}} \circ \pi} M\mathbf{v} = M/\pi_{\mathbf{a}} \circ \pi \mathbf{v}$ is H -tieless; i.e., $HP_{\pi_{\mathbf{a}} \circ \pi} M\mathbf{v}$ is tieless. It follows that the vector

$$HP_{\pi_{\mathbf{a}} \circ \pi} (DM)(N'\mathbf{v}) = \frac{1}{2\alpha_i \lceil \sigma_1^H + 1 \rceil D} \mathbf{w}$$

is also tieless; being it entrywise integer, and since $\frac{1}{2\alpha_i \lceil \sigma_1^H + 1 \rceil D} w_t > 0$, we obtain

$$\left| \frac{1}{2\alpha_i \lceil \sigma_1^H + 1 \rceil D} w_t - \frac{1}{2\alpha_i \lceil \sigma_1^H + 1 \rceil D} w_{t'} \right| \geq 1,$$

which yields

$$|w_t - w_{t'}| \geq 2\alpha_i \lceil \sigma_1^H + 1 \rceil D.$$

Denote the ℓ_1 -norm and the ℓ_2 -norm of a vector by $\|\cdot\|_1$ and $\|\cdot\|_2$, respectively. Recall that the largest singular value of a matrix is its spectral operator norm, i.e., $\sigma_1^H = \max_{\mathbf{0} \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|H\mathbf{x}\|_2}{\|\mathbf{x}\|_2}$ (see [47]). In particular, $\|H\mathbf{x}\|_2 \leq \sigma_1^H \|\mathbf{x}\|_2$ for each vector \mathbf{x}

of size n . Using the Cauchy–Schwarz inequality and the fact that the ℓ_1 -norm of a vector is greater than or equal to its ℓ_2 -norm, we find that

$$\begin{aligned} |z_t - z_{t'}| &= \beta_i |(\mathbf{e}_t - \mathbf{e}_{t'})^T H P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}| \leq \beta_i \|\mathbf{e}_t - \mathbf{e}_{t'}\|_2 \|H P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}\|_2 \\ &\leq \beta_i \|\mathbf{e}_t - \mathbf{e}_{t'}\|_2 \sigma_1^H \|P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}\|_2 \leq \beta_i \|\mathbf{e}_t - \mathbf{e}_{t'}\|_1 [\sigma_1^H + 1] \|P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu}\|_1 \\ &= 2\beta_i [\sigma_1^H + 1] \mathbf{1}_n^T P_{\pi_{\mathbf{a}} \circ \pi} \boldsymbol{\mu} \leq 2\beta_i [\sigma_1^H + 1] \mathbf{1}_n^T P_{\pi_{\mathbf{a}} \circ \pi} |\boldsymbol{\mu}| \\ &= 2\beta_i [\sigma_1^H + 1] \mathbf{1}_L^T |\boldsymbol{\mu}| \leq 2\beta_i [\sigma_1^H + 1] D < 2\alpha_i [\sigma_1^H + 1] D. \end{aligned}$$

We conclude the proof of the claim by noting that

$$\begin{aligned} |u_t - u_{t'}| &= |(w_t - w_{t'}) - (z_{t'} - z_t)| \geq |w_t - w_{t'}| - |z_t - z_{t'}| \\ &> 2\alpha_i [\sigma_1^H + 1] D - 2\alpha_i [\sigma_1^H + 1] D = 0, \end{aligned}$$

which implies $u_t \neq u_{t'}$. As a consequence, the vector $P_{(\pi_{\mathbf{a}} \circ \tilde{\varphi})(i)} \mathbf{1}_{c_i}$ is H -tieless. We can then apply Lemma C.2 to conclude that $f_{/\pi_{\mathbf{a}} \circ \tilde{\varphi}}(1, \dots, n) = f_{/\pi_{\mathbf{a}} \circ \pi \circ \varphi}(1, \dots, n)$. Hence, by (C.2), $f_{/\tilde{\varphi}} = f_{/\pi \circ \varphi}$. Therefore, $\xi_D((M, \boldsymbol{\mu})_{/\pi}) = \xi_D((M, \boldsymbol{\mu}))_{/\pi}$, as required. It follows that ξ_D is a minion homomorphism.

Since the set of polymorphisms of (\mathbf{A}, \mathbf{B}) of arity L is finite for every $L \in \mathbb{N}$, we can apply Proposition 4.5 to conclude that there exists a minion homomorphism $\zeta : \mathcal{C} \rightarrow \text{Pol}(\mathbf{A}, \mathbf{B})$. \square

Acknowledgment. We would like to thank the anonymous referees of both the conference [32] and this full version of the paper.

REFERENCES

- [1] S. ARORA, C. LUND, R. MOTWANI, M. SUDAN, AND M. SZEGEDY, *Proof verification and the hardness of approximation problems*, J. ACM, 45 (1998), pp. 501–555, <https://doi.org/10.1145/278298.278306>.
- [2] S. ARORA AND S. SAFRA, *Probabilistic checking of proofs: A new characterization of NP*, J. ACM, 45 (1998), pp. 70–122, <https://doi.org/10.1145/273865.273901>.
- [3] A. ATSERIAS AND V. DALMAU, *Promise constraint satisfaction and width*, in Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22), 2022, pp. 1129–1153, <https://doi.org/10.1137/1.9781611977073.48>.
- [4] P. AUSTRIN, A. BHANGALE, AND A. POTUKUCHI, *Improved inapproximability of rainbow coloring*, in Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA'20), 2020, pp. 1479–1495, <https://doi.org/10.1137/1.9781611975994.90>.
- [5] P. AUSTRIN, V. GURUSWAMI, AND J. HÅSTAD, *$(2+\epsilon)$ -sat is NP-hard*, SIAM J. Comput., 46 (2017), pp. 1554–1573, <https://doi.org/10.1137/15M1006507>.
- [6] L. BARTO, *The collapse of the bounded width hierarchy*, J. Log. Comput., 26 (2016), pp. 923–943, <https://doi.org/10.1093/logcom/exu070>.
- [7] L. BARTO, D. BATTISTELLI, AND K. M. BERG, *ymmetric promise constraint satisfaction problems: Beyond the Boolean case*, in Proceedings of the 38th International Symposium on Theoretical Aspects of Computer Science (STACS'21), LIPIcs. Leibniz Int. Proc. Inform. 187, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2021, 10, <https://doi.org/10.4230/LIPIcs.STACS.2021.10>.
- [8] L. BARTO, J. BULÍN, A. A. KROKHIN, AND J. OPRŠAL, *Algebraic approach to promise constraint satisfaction*, J. ACM, 68 (2021), 28, <https://doi.org/10.1145/3457606>.
- [9] L. BARTO AND M. KOZIK, *Constraint satisfaction problems solvable by local consistency methods*, J. ACM, 61 (2014), 3, <https://doi.org/10.1145/2556646>.
- [10] L. BARTO AND M. KOZIK, *Robustly solvable constraint satisfaction problems*, SIAM J. Comput., 45 (2016), pp. 1646–1669, <https://doi.org/10.1137/130915479>.
- [11] L. BARTO, A. KROKHIN, AND R. WILLARD, *Polymorphisms, and how to use them*, in The Constraint Satisfaction Problem: Complexity and Approximability, A. KROKHIN, AND S. ŽIVNÝ, Dagstuhl Follow-Ups 7, Schloss Dagstuhl. Leibniz-Zent. Inform., Dagstuhl, Germany, 2017, pp. 1–44, <https://doi.org/10.4230/DFU.Vol7.15301.1>.

- [12] L. BARTO, J. OPRŠAL, AND M. PINSKER, *The wonderland of reflections*, Isr. J. Math., 223 (2018), pp. 363–398, <https://doi.org/10.1007/s11856-017-1621-9>.
- [13] L. BARTO AND M. PINSKER, *Topology is irrelevant (in a dichotomy conjecture for infinite domain constraint satisfaction problems)*, SIAM J. Comput., 49 (2020), pp. 365–393, <https://doi.org/10.1137/18M1216213>.
- [14] J. BERMAN, P. IDZIAK, P. MARKOVIĆ, R. MCKENZIE, M. VALERIOTE, AND R. WILLARD, *Varieties with few subalgebras of powers*, Trans. Amer. Math. Soc., 362 (2010), pp. 1445–1473.
- [15] C. BESSIERE AND R. DEBRUYNE, *Theoretical analysis of singleton arc consistency and its extensions*, Artif. Intell., 172 (2008), pp. 29–41, <https://doi.org/10.1016/j.artint.2007.09.001>.
- [16] A. BLUM, *New approximation algorithms for graph coloring*, J. ACM, 41 (1994), pp. 470–516, <https://doi.org/10.1145/176584.176586>.
- [17] M. BODIRSKY, B. MARTIN, AND A. MOTTET, *Discrete temporal constraint satisfaction problems*, J. ACM, 65 (2018), 9, <https://doi.org/10.1145/3154832>.
- [18] M. BODIRSKY, A. MOTTET, M. OLŠÁK, J. OPRŠAL, M. PINSKER, AND R. WILLARD, *ω -categorical structures avoiding height 1 identities*, Trans. Amer. Math. Soc., 374 (2021), pp. 327–350, <https://doi.org/10.1090/tran/8179>.
- [19] J. BRAKENSIEK AND V. GURUSWAMI, *An algorithmic blend of LPs and ring equations for promise CSPs*, in Proceedings of the 2019 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'19), SIAM, Philadelphia, 2019, pp. 436–455, <https://doi.org/10.1137/1.9781611975482.28>.
- [20] J. BRAKENSIEK AND V. GURUSWAMI, *Promise constraint satisfaction: Algebraic structure and a symmetric boolean dichotomy*, SIAM J. Comput., 50 (2021), pp. 1663–1700, <https://doi.org/10.1137/19M128212X>.
- [21] J. BRAKENSIEK, V. GURUSWAMI, AND S. SANDEEP, *Conditional dichotomy of Boolean ordered promise CSPs*, in Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP'21), LIPIcs. Leibniz Int. Proc. Inform. 198, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2021, 37, <https://doi.org/10.4230/LIPIcs.ICALP.2021.37>.
- [22] J. BRAKENSIEK, V. GURUSWAMI, M. WROCHNA, AND S. ŽIVNÝ, *The power of the combined basic linear programming and affine relaxation for promise constraint satisfaction problems*, SIAM J. Comput., 49 (2020), pp. 1232–1248, <https://doi.org/10.1137/20M1312745>.
- [23] A. BRANDTS, M. WROCHNA, AND S. ŽIVNÝ, *The complexity of promise SAT on non-Boolean domains*, ACM Trans. Comput. Theory, 13 (2021), 26, <https://doi.org/10.1145/3470867>.
- [24] A. BRANDTS AND S. ŽIVNÝ, *Beyond PCSP(1-in-3,NAE)*, Inform. and Comput., 289 (2022), 104954, <https://doi.org/10.1016/j.ic.2022.104954>.
- [25] A. BULATOV, *Bounded relational width*, unpublished manuscript, 2009, <https://www2.cs.sfu.ca/~abulatov/papers/relwidth.pdf>.
- [26] A. BULATOV, P. JEAVONS, AND A. KROKHIN, *Classifying the complexity of constraints using finite algebras*, SIAM J. Comput., 34 (2005), pp. 720–742, <https://doi.org/10.1137/S0097539700376676>.
- [27] A. A. BULATOV, *A dichotomy theorem for constraint satisfaction problems on a 3-element set*, J. ACM, 53 (2006), pp. 66–120, <https://doi.org/10.1145/1120582.1120584>.
- [28] A. A. BULATOV, *A dichotomy theorem for nonuniform CSPs*, in Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17), 2017, pp. 319–330, <https://doi.org/10.1109/FOCS.2017.37>.
- [29] A. A. BULATOV AND V. DALMAU, *A simple algorithm for Mal'tsev constraints*, SIAM J. Comput., 36 (2006), pp. 16–27, <https://doi.org/10.1137/050628957>.
- [30] H. CHEN, V. DALMAU, AND B. GRUSSEN, *Arc consistency and friends*, J. Log. Comput., 23 (2013), pp. 87–108, <https://doi.org/10.1093/logcom/exr039>.
- [31] L. CIARDO AND S. ŽIVNÝ, *Hierarchies of minion tests for PCSPs through tensors*, in Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA'23), 2023, to appear, <https://arxiv.org/abs/2207.02277>.
- [32] L. CIARDO AND S. ŽIVNÝ, *CLAP: A new algorithm for promise CSPs*, in Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22), SIAM, Philadelphia, 2022, pp. 1057–1068, <https://doi.org/10.1137/1.9781611977073.46>.
- [33] D. A. COHEN, *Tractable decision for a constraint language implies tractable search*, Constraints, 9 (2004), pp. 219–229, <https://doi.org/10.1023/B:CONS.0000036045.82829.94>.
- [34] V. DALMAU, *There are no pure relational width 2 constraint satisfaction problems*, Inform. Process. Lett., 109 (2009), pp. 213–218, <https://doi.org/10.1016/j.ipl.2008.10.005>.
- [35] V. DALMAU, M. KOZIK, A. A. KROKHIN, K. MAKARYCHEV, Y. MAKARYCHEV, AND J. OPRŠAL, *Robust algorithms with polynomial loss for near-unanimity CSPs*, SIAM J. Comput., 48 (2019), pp. 1763–1795, <https://doi.org/10.1137/18M1163932>.

- [36] V. DALMAU AND A. A. KROKHIN, *Robust satisfiability for CSPs: Hardness and algorithmic results*, ACM Trans. Comput. Theory, 5 (2013), 15, <https://doi.org/10.1145/2540090>.
- [37] V. DALMAU AND J. PEARSON, *Closure functions and width 1 problems*, in Proceedings of the 4th International Conference on Principles and Practice of Constraint Programming (CP'99), Lecture Notes in Comput. Sci. 1713, Springer, 1999, pp. 159–173, https://doi.org/10.1007/978-3-540-48085-3_12.
- [38] R. DEBRUYNE AND C. BESSIÈRE, *Some practicable filtering techniques for the constraint satisfaction problem*, in Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI'97), Morgan Kaufmann, Burlington, MA1997, pp. 412–417.
- [39] I. DINUR, *The PCP theorem by gap amplification*, J. ACM, 54 (2007), 12, <https://doi.org/10.1145/1236457.1236459>.
- [40] I. DINUR, O. REGEV, AND C. SMYTH, *The hardness of 3-uniform hypergraph coloring*, Combinatorica, 25 (2005), pp. 519–535, <https://doi.org/10.1007/s00493-005-0032-4>.
- [41] 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., 28 (1998), pp. 57–104, <https://doi.org/10.1137/S0097539794266766>.
- [42] M. FICAK, M. KOZIK, M. OLŠÁK, AND S. STANKIEWICZ, *Dichotomy for symmetric Boolean PCSPs*, in Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP'19), LIPIcs. Leibniz Int. Proc. Inform. 132, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019, 57, <https://doi.org/10.4230/LIPIcs.ICALP.2019.57>.
- [43] M. R. GAREY AND D. S. JOHNSON, *The complexity of near-optimal graph coloring*, J. ACM, 23 (1976), pp. 43–49, <https://doi.org/10.1145/321921.321926>.
- [44] V. GURUSWAMI AND S. KHANNA, *On the hardness of 4-coloring a 3-colorable graph*, SIAM J. Discrete Math., 18 (2004), pp. 30–40, <https://doi.org/10.1137/S0895480100376794>.
- [45] V. GURUSWAMI AND S. SANDEEP, *d-To-1 Hardness of Coloring 3-Colorable Graphs with $O(1)$ Colors*, in Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP'20), LIPIcs. Leibniz Int. Proc. Inform. 168, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020, 62, <https://doi.org/10.4230/LIPIcs.ICALP.2020.62>.
- [46] P. HELL AND J. NEŠETŘIL, *On the complexity of H -coloring*, J. Combin. Theory, Ser. B, 48 (1990), pp. 92–110, [https://doi.org/10.1016/0095-8956\(90\)90132-J](https://doi.org/10.1016/0095-8956(90)90132-J).
- [47] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, 2012.
- [48] P. M. IDZIAK, P. MARKOVIC, R. MCKENZIE, M. VALERIOTE, AND R. WILLARD, *Tractability and learnability arising from algebras with few subpowers*, SIAM J. Comput., 39 (2010), pp. 3023–3037, <https://doi.org/10.1137/090775646>.
- [49] P. G. JEAVONS, *On the algebraic structure of combinatorial problems*, Theor. Comput. Sci., 200 (1998), pp. 185–204, [https://doi.org/10.1016/S0304-3975\(97\)00230-2](https://doi.org/10.1016/S0304-3975(97)00230-2).
- [50] P. G. JEAVONS, D. A. COHEN, AND M. GYSSENS, *Closure properties of constraints*, J. ACM, 44 (1997), pp. 527–548, <https://doi.org/10.1145/263867.263489>.
- [51] S. KHANNA, N. LINIAL, AND S. SAFRA, *On the hardness of approximating the chromatic number*, Combinatorica, 20 (2000), pp. 393–415, <https://doi.org/10.1007/s004930070013>.
- [52] S. KHOT, *Improved Inapproximability results for MaxClique, chromatic number and approximate graph coloring*, in Proceedings of 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS'01), IEEE Computer Society, 2001, pp. 600–609, <https://doi.org/10.1109/SFCS.2001.959936>.
- [53] S. KHOT, *On the power of unique 2-prover 1-round games*, in Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC'02), ACM, 2002, pp. 767–775, <https://doi.org/10.1145/509907.510017>.
- [54] V. KOLMOGOROV, A. A. KROKHIN, AND M. ROLÍNEK, *The complexity of general-valued CSPs*, SIAM J. Comput., 46 (2017), pp. 1087–1110, <https://doi.org/10.1137/16M1091836>.
- [55] V. KOLMOGOROV, J. THAPPER, AND S. ŽIVNÝ, *The power of linear programming for general-valued CSPs*, SIAM J. Comput., 44 (2015), pp. 1–36, <https://doi.org/10.1137/130945648>.
- [56] M. KOZIK, *Solving CSPs using weak local consistency*, SIAM J. Comput., 50 (2021), pp. 1263–1286, <https://doi.org/10.1137/18M117577X>.
- [57] M. KOZIK AND J. OCHREMIK, *Algebraic properties of valued constraint satisfaction problem*, in Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP'15), Lecture Notes in Comput. Sci. 9134, Springer, 2015, pp. 846–858, https://doi.org/10.1007/978-3-662-47672-7_69.
- [58] A. KROKHIN AND J. OPRŠAL, *The complexity of 3-colouring H -colourable graphs*, in Proceedings of the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS'19), 2019, pp. 1227–1239, <https://doi.org/10.1109/FOCS.2019.00076>.

- [59] G. KUN, R. O'DONNELL, S. TAMAKI, Y. YOSHIDA, AND Y. ZHOU, *Linear programming, width-1 CSPs, and robust satisfaction*, in Proceedings of the 3rd Innovations in Theoretical Computer Science (ITCS'12), ACM, 2012, pp. 484–495, <https://doi.org/10.1145/2090236.2090274>.
- [60] G. KUN AND M. SZEGEDY, *A new line of attack on the dichotomy conjecture*, European J. Combin., 52 (2016), pp. 338–367, <https://doi.org/10.1016/j.ejc.2015.07.011>.
- [61] A. K. MACKWORTH, *Consistency in networks of relations*, Artif. Intell., 8 (1977), pp. 99–118, [https://doi.org/10.1016/0004-3702\(77\)90007-8](https://doi.org/10.1016/0004-3702(77)90007-8).
- [62] A. MALCEV, *Untersuchungen aus dem gebiete der mathematischen logik*, J. Symbolic Logic, 2 (1937), 84.
- [63] P. RAGHAVENDRA, *Optimal algorithms and inapproximability results for every CSP?*, in Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC'08), 2008, pp. 245–254, <https://doi.org/10.1145/1374376.1374414>.
- [64] D. RORABAUGH, C. TARDIF, AND D. L. WEHLAU, *Logical compactness and constraint satisfaction problems*, Log. Methods Comput. Sci., 13 (2017), [https://doi.org/10.23638/LMCS-13\(1:1\)2017](https://doi.org/10.23638/LMCS-13(1:1)2017).
- [65] T. SCHAEFER, *The complexity of satisfiability problems*, in Proceedings of the 10th Annual ACM Symposium on the Theory of Computing (STOC'78), 1978, pp. 216–226, <https://doi.org/10.1145/800133.804350>.
- [66] J. THAPPER AND S. ŽIVNÝ, *The complexity of finite-valued CSPs*, J. ACM, 63 (2016), 37, <https://doi.org/10.1145/2974019>.
- [67] J. THAPPER AND S. ŽIVNÝ, *The power of Sherali–Adams relaxations for general-valued CSPs*, SIAM J. Comput., 46 (2017), pp. 1241–1279, <https://doi.org/10.1137/16M1079245>.
- [68] J. THAPPER AND S. ŽIVNÝ, *The limits of SDP relaxations for general-valued CSPs*, ACM Trans. Comput. Theory, 10 (2018), 12, <https://doi.org/10.1145/3201777>.
- [69] A. WIGDERSON, *Improving the performance guarantee for approximate graph coloring*, J. ACM, 30 (1983), pp. 729–735, <https://doi.org/10.1145/2157.2158>.
- [70] M. WROCHNA AND S. ŽIVNÝ, *Improved hardness for H -colourings of G -colourable graphs*, in Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA'20), SIAM, Philadelphia, 2020, pp. 1426–1435, <https://doi.org/10.1137/1.9781611975994.86>.
- [71] D. ZHUK, *A proof of the CSP dichotomy conjecture*, J. ACM, 67 (2020), 30, <https://doi.org/10.1145/3402029>.