

Accepted Manuscript

Towards a characterization of constant-factor approximable finite-valued CSPs

Víctor Dalmau, Andrei Krokhin, Rajsekar Manokaran

PII: S0022-0000(18)30230-7
DOI: <https://doi.org/10.1016/j.jcss.2018.03.003>
Reference: YJCSS 3167

To appear in: *Journal of Computer and System Sciences*

Received date: 23 September 2016
Revised date: 17 January 2018
Accepted date: 15 March 2018

Please cite this article in press as: V. Dalmau et al., Towards a characterization of constant-factor approximable finite-valued CSPs, *J. Comput. Syst. Sci.* (2018), <https://doi.org/10.1016/j.jcss.2018.03.003>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Highlights

- New natural algebraic conditions for the finiteness of the integrality gap for the basic LP relaxation of VCSP are given.
- Efficient constant-factor approximation algorithms using the algebraic conditions are given.
- Algebraic conditions for NP-hardness of constant-factor approximation are given.

Towards a Characterization of Constant-Factor Approximable Finite-Valued CSPs^{☆,☆☆}

V́ctor Dalmau^a, Andrei Krokhin^b, Rajsekar Manokaran^c

^a*Department of Information and Communication Technologies, Universitat Pompeu Fabra, Spain*

^b*Department of Computer Science, Durham University, UK*

^c*Computer Science and Engineering department, IIT Madras, India*

Abstract

We study the approximability of (Finite-)Valued Constraint Satisfaction Problems (VCSPs) with a fixed finite constraint language Γ consisting of finitary functions on a fixed finite domain. Ene et al. have shown that, under a mild technical condition, the basic LP relaxation is optimal for constant-factor approximation for VCSP(Γ) unless the Unique Games Conjecture fails. Using the algebraic approach to the CSP, we give new natural algebraic conditions for the finiteness of the integrality gap for the basic LP relaxation of VCSP(Γ) and show how this leads to efficient constant-factor approximation algorithms for several examples that cover all previously known cases that are NP-hard to solve to optimality but admit constant-factor approximation. Finally, we show that the absence of another algebraic condition leads to NP-hardness of constant-factor approximation. Thus, our results indicate where the boundary of constant-factor approximability for VCSPs lies.

Keywords: constraint satisfaction problem, approximation algorithms, universal algebra

1. Introduction

The constraint satisfaction problem (CSP) provides a framework in which it is possible to express, in a natural way, many combinatorial problems encountered in computer science and AI [1, 2, 3]. Standard examples of CSPs include satisfiability of propositional formulas, graph colouring problems, and systems of linear equations. An instance of the CSP consists of a set of variables, a (not necessarily Boolean) domain of labels, and a set of constraints on combinations of values that can be taken by certain

[☆]This article is an extended version of a paper published in the proceedings of SODA'15.

^{☆☆}The first two authors were supported by UK EPSRC grant EP/J000078/01. The first author was also supported by MICCIN grant TIN2016-76573-C2-1P and Maria de Maeztu Units of Excellence Programme MDM-2015-0502. The third author was supported by ERC grant 226-203.

*Corresponding author

Email addresses: victor.dalmau@upf.edu (V́ctor Dalmau), Andrei.Krokhin@durham.ac.uk (Andrei Krokhin), rajsekar@gmail.com (Rajsekar Manokaran)

subsets of variables. The aim is then to find an assignment of labels to the variables that, in the decision version, satisfies all the constraints or, in the optimization version, maximizes (minimizes) the number of satisfied (unsatisfied, respectively) constraints.

Since the CSP is NP-hard in full generality, a major line of research in CSP tries to identify special cases that have desirable algorithmic properties (see, e.g. [1, 2, 3]), the primary motivation being the general picture rather than specific applications. The two main ingredients of a constraint are: (a) variables to which it is applied, and (b) relation specifying the allowed combinations of labels. Therefore, the main types of restrictions on CSP are: (a) *structural* where the hypergraph formed by sets of variables appearing in individual constraints is restricted [4, 5], and (b) *language-based* where the constraint language Γ , i.e. the set of relations that can appear in constraints, is fixed (see, e.g. [6, 1, 7, 3]); the corresponding decision/maximization/minimization problems are denoted by $\text{CSP}(\Gamma)$, $\text{Max CSP}(\Gamma)$, and $\text{Min CSP}(\Gamma)$, respectively. The ultimate sort of results in this direction are dichotomy results, pioneered by [8], which completely characterise the restrictions with a given desirable property modulo some complexity-theoretic assumptions. The language-based direction is considerably more active than the structural one, and there are many (partial and full) language-based complexity classification results, e.g. [9, 10, 11, 12, 13, 14, 15], but many questions are still open.

Problems Max CSP and Min CSP can be generalised by replacing relations (that specify allowed combinations of labels) with functions that specify a value in $[0, 1]$ (measuring the desirability or the cost, respectively) for each tuple of labels. The goal would then be to find an assignment of labels that maximizes the total desirability (minimizes the total cost, respectively). The maximization version was studied in [16, 17] under the name of Generalized CSP, or GCSP, (in fact, functions there can take values in $[-1, 1]$), while the minimization is known as (Finite-)Valued CSP [14]. In General-Valued CSP, functions can also take the infinite value to indicate infeasible tuples [18, 19, 13], but we will not consider this case in this paper. In this paper we write VCSP to mean *finite-valued* CSP. We note that [20] write Min CSP to mean what we call VCSP in this paper. Naturally, both GCSP and VCSP can be parameterized by constraint languages Γ , now consisting of functions instead of relations.

The CSP has always played an important role in mapping the landscape of approximability of NP-hard optimization problems, see e.g. surveys [21, 22]. For example, the famous PCP theorem has an equivalent reformulation in terms of inapproximability of a certain $\text{Max CSP}(\Gamma)$, see [23]; moreover, Dinur's combinatorial proof of this theorem [24] deals entirely with CSPs. The first optimal inapproximability results [25] by Håstad were about problems $\text{Max CSP}(\Gamma)$, and they led to the study of a new hardness notion called approximation resistance (see, e.g. [26, 27, 28]). The approximability of Boolean CSPs has been thoroughly investigated (see, e.g. [29, 1, 30, 31, 25, 27, 21, 32]). Much work around the Unique Games Conjecture (UGC) directly concerns CSPs [21]. This conjecture states that, for any $\epsilon > 0$, there is a large enough number $k = k(\epsilon)$ such that it NP-hard to tell ϵ -satisfiable from $(1 - \epsilon)$ -satisfiable instances of $\text{CSP}(\Gamma_k)$, where Γ_k consists of all graphs of bijections on a k -element set. Many approximation algorithms for classical optimization problems have been shown optimal assuming the UGC [21, 32]. Raghavendra proved [17] that one SDP-based algorithm provides optimal approximation for all problems $\text{GCSP}(\Gamma)$ assuming the UGC. In this paper, we investigate problems $\text{VCSP}(\Gamma)$ and $\text{Min CSP}(\Gamma)$ on an arbitrary finite domain

that belong to APX, i.e. admit a (polynomial-time) constant-factor approximation algorithm, proving some results that strongly indicate where the boundary of this property lies.

Related Work. Note that each problem Max CSP(Γ) trivially admits a constant-factor approximation algorithm because a random assignment of values to the variables is guaranteed to satisfy a constant fraction of constraints; this can be derandomized by the standard method of conditional probabilities. The same also holds for GCSP. Clearly, for Min CSP(Γ) to admit a constant-factor approximation algorithm, CSP(Γ) must be polynomial-time solvable.

The approximability of problems VCSP(Γ) has been studied, mostly for Min CSPs in the Boolean case (i.e., with domain $\{0, 1\}$, such CSPs are sometimes called “generalized satisfiability” problems), see [29, 1]. We need a few concepts from propositional logic. A clause is *Horn* if it contains at most one positive literal, and *negative* if it contains only negative literals. Let k -HORN be the constraint language over the Boolean domain that contains all Horn clauses with at most k variables. For $k \geq 2$, let k -IHBS be the subset of k -HORN that consists of all clauses that are negative or have at most 2 variables. It is known that, for each $k \geq 2$, Min CSP(k -IHBS) belongs to APX [1], and they (and the corresponding dual Horn problems) are essentially the only such Boolean Min CSPs unless the UGC fails [33]. For Min CSP(2-HORN), which is identical to Min CSP(2-IHBS), a 2-approximation (LP-based) algorithm is described in [31], which is optimal assuming the UGC, whereas it is NP-hard to constant-factor approximate Min CSP(3-HORN) [30]. If \neq_2 is the Boolean relation $\{(0, 1), (1, 0)\}$, then Min CSP($\{\neq_2\}$) is known as MINUNCUT. Min CSP(Γ) where Γ consists of 2-clauses is known as MIN 2CNF DELETION. The best currently known approximation algorithms for MINUNCUT and MIN 2CNF DELETION have approximation ratio $O(\sqrt{\log n})$ [29], and it follows from [32] that neither problem belongs to APX unless the UGC is false. The UGC is known to imply the optimality of the basic LP relaxation for any VCSP(Γ) such that Γ contains the (characteristic function of the) equality relation [20], extending the line of similar results for natural LP and SDP relaxations for various optimization CSPs [34, 35, 17].

An approximation algorithm for any VCSP(Γ) was also given in the 2013 conference version of [20] (that was claimed to match the LP integrality gap), but its analysis was later found to be faulty and this part was retracted in the 2015 update of [20]. The SDP rounding algorithm for GCSPs from [36] is discussed in detail in the book [37], where it is pointed out that the same algorithm does not work for VCSPs.

Constant-factor approximation algorithms for Min CSP are closely related to certain *robust algorithms* for CSP that attracted much attention recently [10, 33, 38, 39]. Call an algorithm for CSP(Γ) *robust* if, for every $\epsilon > 0$ and every $(1 - \epsilon)$ -satisfiable instance of CSP(Γ) (i.e. at most an ϵ -fraction of constraints can be removed to make the instance satisfiable), it outputs a $(1 - f(\epsilon))$ -satisfying assignment (i.e. that fails to satisfy at most a $f(\epsilon)$ -fraction of constraints) where f is a function such that $f(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$ and $f(0) = 0$. CSPs admitting a robust algorithm (with some function f) were completely characterised in [10]; when such an algorithm exists, one can always choose $f(\epsilon) = O(\log \log (1/\epsilon) / \log (1/\epsilon))$ for the randomized algorithm and $f(\epsilon) = O(\log \log (1/\epsilon) / \sqrt{\log (1/\epsilon)})$ for the derandomized version. A robust algorithm is said to have *linear loss* if the function f can be chosen so that $f(\epsilon) = O(\epsilon)$. The

problem of characterizing CSPs that admit a robust algorithm with linear loss was posed in [33]. It is easy to see that, for any Γ , $\text{CSP}(\Gamma)$ admits a robust algorithm with linear loss if and only if $\text{Min CSP}(\Gamma)$ has a constant-factor approximation algorithm. We will use this fact when referring to results in [33].

Many complexity classification results for CSP have been made possible by the introduction of the universal-algebraic approach (see, e.g., survey [40]), which extracts algebraic structure from a given constraint language Γ (via operations called *polymorphisms* of Γ) and uses it to analyze problem instances. This approach was extended to VCSP (see, e.g., survey [41]), where polymorphisms are replaced by certain probability distributions on operations called *fractional polymorphisms*. The universal-algebraic framework to study robust algorithms with a given loss was presented in [33], this approach was also used in [10, 39]. In this paper, we apply this framework with some old and some new algebraic conditions to study problems $\text{VCSP}(\Gamma)$ and $\text{Min CSP}(\Gamma)$. Our algebraic conditions use symmetric operations, which appear naturally when LP-based algorithms are used for (V)CSPs; other recent examples are [13, 42, 39, 14].

Contributions. Some of our results assume that Γ contains the equality relation. We characterise problems $\text{VCSP}(\Gamma)$ for which the basic LP relaxation has finite integrality gap. The characterisation is in terms of appropriately modified fractional polymorphisms. We then show how that a description of constant-factor approximable VCSPs can be reduced to that for Min CSPs. For Min CSPs, we give another algebraic condition that characterizes the property of being constant-factor approximable. This characterization uses the algebraic approach to CSP that has been extremely fruitful in proving complexity classification results for CSPs. The characterizing condition is in terms of Lipschitz probability distributions on symmetric polymorphisms of Γ . This condition can in principle be used to design efficient constant-factor approximation algorithms, provided one can efficiently sample from these distributions. We show that this is possible for some examples that cover all cases where such algorithms (but not algorithms finding an optimal solution) were previously known to exist.

It follows from the [20] that every Min CSP for which the basic LP relaxation does not have finite integrality gap is not constant-factor approximable, unless the UGC fails. For a class of Min CSPs we strengthen the UG-hardness to NP-hardness. A near-unanimity polymorphism is a type of polymorphism well known in the algebraic theory of CSP [43, 40, 7], and its presence follows from the existence of those Lipschitz distributions. We show $\text{Min CSP}(\Gamma)$ is NP-hard to constant-factor approximate if Γ has no near-unanimity polymorphism.

Thus, our results strongly indicate where the boundary of constant-factor approximability for VCSPs lies.

2. Preliminaries

Let A be a finite set. A k -tuple $\mathbf{a} = (a_1, \dots, a_k)$ on A is any element of A^k . A k -ary relation on A is a subset of A^k . We shall use $\text{arity}(R)$ to denote the arity of relation R . We shall denote by eq_A the binary relation $\{(a, a) \mid a \in A\}$.

An *instance* of the CSP is a triple $I = (V, A, \mathcal{C})$ with V a finite set of *variables*, A a finite set called *domain*, and \mathcal{C} a finite list of *constraints*. Each constraint in \mathcal{C}

is a pair $C = (\mathbf{v}, R)$, also denoted $R(\mathbf{v})$, where $\mathbf{v} = (v_1, \dots, v_k)$ is a tuple of variables of length k , called the *scope* of C , and R an k -ary relation on A , called the *constraint relation* of C . The *arity* of a constraint C , $\text{arity}(C)$, is the arity of its constraint relation. When considering optimization problems, we will assume that each constraint has a weight $w_C \in \mathbb{Q}_{>0}$. It is known (see, e.g. Lemma 7.2 in [1]) that allowing weights in Min CSP(Γ) does not affect membership in APX.

Very often we will say that a constraint C belongs to instance I when, strictly speaking, we should be saying that appears in the constraint list \mathcal{C} of I . Also, we might sometimes write (v_1, \dots, v_k, R) instead of $((v_1, \dots, v_k), R)$. A *constraint language* is any *finite* set Γ of relations on A . The problem CSP(Γ) consists of all instances of the CSP where all the constraint relations are from Γ . An assignment for I is a mapping $s : V \rightarrow A$. We say that s *satisfies* a constraint (\mathbf{v}, R) if $s(\mathbf{v}) \in R$ (where s is applied component-wise).

The *decision problem* for CSP(Γ) asks whether an input instance I of CSP(Γ) has a solution, i.e., an assignment satisfying all constraints. The natural *optimization problems* for CSP(Γ), Max CSP(Γ) and Min CSP(Γ), ask to find an assignment that maximizes the total weight of satisfied constraints or minimizes the total weight of unsatisfied constraints, respectively.

VCSP is the generalization of Min CSP obtained by allowing a richer set of constraints. Formally, a constraint in a VCSP instance is a pair $C = (\mathbf{v}, \varrho)$, also denoted $\varrho(\mathbf{v})$, where $\mathbf{v} = (v_1, \dots, v_k)$ is, as before, a tuple of variables, and $\varrho : A^k \rightarrow [0, 1]$ is a mapping from A^k to $[0, 1]$. Given an instance I of the VCSP, the goal is to find an assignment $s : V \rightarrow A$ that minimizes

$$\sum_{C=\varrho(\mathbf{v}) \in \mathcal{C}} w_C \cdot \varrho(s(\mathbf{v})).$$

Note that, to express Min CSP as VCSP, one needs to replace each relation R in a Min CSP instance by a function ϱ_R such that $\varrho_R(\mathbf{a}) = 0$ if $\mathbf{a} \in R$ and $\varrho_R(\mathbf{a}) = 1$ otherwise.

2.1. Basic linear program

Many approximation algorithms for optimization CSPs use the basic (aka standard) linear programming (LP) relaxation [33, 34, 39].

For any instance $I = (V, A, \mathcal{C})$ of VCSP(Γ), there is an equivalent canonical 0-1 integer program. It has variables $p_v(a)$ for every $v \in V$, $a \in A$, as well as variables $p_C(\mathbf{a})$ for every constraint $C = \varrho(\mathbf{v})$ and every tuple $\mathbf{a} \in A^{\text{arity}(\varrho)}$. The interpretation of $p_v(a) = 1$ is that variable v is assigned value a ; the interpretation of $p_C(\mathbf{a}) = 1$ is that \mathbf{v} is assigned (component-wise) tuple \mathbf{a} . More formally, the program ILP is the following:

$$\begin{aligned}
& \text{minimize:} && \sum_{C=\varrho(\mathbf{v}) \in \mathcal{C}, \mathbf{a} \in A^{\text{arity}(C)}} w_C \cdot p_C(\mathbf{a}) \cdot \varrho(\mathbf{a}) \\
& \text{subject to:} && \\
& p_v(A) = 1 && \text{for } v \in V; \quad (1) \\
& p_C(A^{j-1} \times \{a\} \times A^{\text{arity}(C)-j}) = p_{v_j}(a) && \text{for } C = (\mathbf{v}, \mathbf{R}) \in \mathcal{C}, \quad (2) \\
& && 1 \leq j \leq \text{arity}(C), a \in A.
\end{aligned}$$

Here, for every $v \in V$ and $S \subseteq A$, $p_v(S)$ is a shorthand for $\sum_{a \in S} p_v(a)$ and for every C and every $T \subseteq A^{\text{arity}(C)}$, $p_C(T)$ is a shorthand for $\sum_{\mathbf{a} \in T} p_C(\mathbf{a})$.

If we relax this ILP by allowing the variables to take values in the range $[0, 1]$ instead of $\{0, 1\}$, we obtain the *basic linear programming relaxation* for I , which we denote by $\text{BLP}(I)$. As Γ is fixed, an optimal solution to $\text{BLP}(I)$ can be computed in time polynomial in $|I|$.

For an instance I of $\text{VCSP}(\Gamma)$, we denote by $\text{Opt}(I)$ the value of an optimal solution to I , and by $\text{Opt}_{\text{LP}}(I)$ the value of an optimal solution to $\text{BLP}(I)$.

For any finite set X , we shall denote by $\Delta(X)$ the set of all probability distributions on X . Furthermore, for any $n \in \mathbb{N}$, we shall denote by $\Delta_n(X)$ the subset of $\Delta(X)$ consisting of all $q \in \Delta(X)$ such that $q(x) \cdot n$ is an integer for every $x \in X$. To simplify notation we shall write Δ_n and Δ as a shorthand of $\Delta_n(A)$ and $\Delta(A)$ respectively. If $p \in \Delta(A^r)$ and $p_1, \dots, p_r \in \Delta(A)$ will say that *the marginals of p are p_1, \dots, p_r* to indicate that for every $1 \leq i \leq r$, and every $a \in A$, $p(A^{i-1} \times \{a\} \times A^{r-i}) = p_i(a)$.

Restriction (1) of $\text{BLP}(I)$ expresses the fact that, for each $v \in V$, $p_v \in \Delta(A)$. Also, (1) and (2) together express the fact that, for each constraint $C = \varrho(\mathbf{v})$, of arity k , we have $p_C \in \Delta(A^k)$ and that the marginals of the p_C distribution are consistent with the p_v distributions.

Recall that the *integrality gap* of BLP for $\text{VCSP}(\Gamma)$ is defined as

$$\sup_I \frac{\text{Opt}(I)}{\text{Opt}_{\text{LP}}(I)}$$

where the supremum is taken over all instances I of $\text{VCSP}(\Gamma)$. In particular, if the integrality gap is finite, then $\text{Opt}(I) = 0$ whenever $\text{Opt}_{\text{LP}}(I) = 0$.

Recall that eq_A denotes the binary equality relation. In the following theorem, eq_A will also denote the binary function on A such that $\text{eq}_A(x, y)$ is equal to 0 if $x = y$ and equal to 1 otherwise. This will not cause any confusion.

Theorem 1 ([20]). *Let Γ be a constraint language such that $\text{eq}_A \in \Gamma$ and let α_{gap} be the integrality gap of BLP for $\text{VCSP}(\Gamma)$. For every real number $\beta < \alpha_{\text{gap}}$, it is NP-hard to approximate $\text{VCSP}(\Gamma)$ to within a factor β unless the UGC is false. In particular, if the integrality gap is infinite then there is no constant-factor approximation algorithm for $\text{VCSP}(\Gamma)$ unless the UGC is false.*

The setting in [20] assumes that each variable in an instance comes with its own list of allowed images (i.e. a subset of A), but this assumption is not essential in their reduction from the UGC.

2.2. Algebra

Most of the terminology introduced in this section is standard. See [6, 40] for more detail about the algebraic approach to the CSP. An n -ary *operation* on A is a map $f : A^n \rightarrow A$. Let us now define several types of operations that will be used in this paper. We usually define operations by identities, i.e. by equations where all variables are assumed to be universally quantified.

- An operation f is *idempotent* if it satisfies the identity $f(x, \dots, x) = x$.
- An operation f is *symmetric* if $f(x_1, \dots, x_n) = f(x_{\pi(1)}, \dots, x_{\pi(n)})$ for each permutation π on $\{1, \dots, n\}$.

Thus, a symmetric operation is one that depends only on the multiset of its arguments. Since there is an obvious one-to-one correspondence between $\Delta_n(A)$ and multisets of size n , n -ary symmetric operations on A can be naturally identified with functions from $\Delta_n(A)$ to A .

- An n -ary operation f on A is *totally symmetric* if $f(x_1, \dots, x_n) = f(y_1, \dots, y_n)$ whenever $\{x_1, \dots, x_n\} = \{y_1, \dots, y_n\}$.
- An n -ary ($n \geq 3$) operation f on A is called an NU (near-unanimity) operation if it satisfies the identities

$$f(y, x, x \dots, x, x) = f(x, y, x \dots x, x) = \dots = f(x, x, x \dots x, y) = x.$$

An n -ary operation f on A *preserves* (or is a *polymorphism* of) a k -ary relation R on A if for every n (not necessarily distinct) tuples $(a_1^i, \dots, a_k^i) \in R$, $1 \leq i \leq n$, the tuple

$$(f(a_1^1, \dots, a_1^n), \dots, f(a_k^1, \dots, a_k^n))$$

belongs to R as well. Given a set Γ of relations on A , we denote by $\text{Pol}(\Gamma)$ the set of all operations f such that f preserves each relation in Γ . If $f \in \text{Pol}(\Gamma)$ then Γ is said to be *invariant* under f . If R is a relation we might freely write $\text{Pol}(R)$ to denote $\text{Pol}(\{R\})$.

Example 1. Let $A = \{0, 1\}$.

1. It is well known and easy to check that, for each $n \geq 1$, the n -ary (totally symmetric) operation $f_n(x_1, \dots, x_n) = \bigwedge_{i=1}^n x_i$ is a polymorphism of 3-HORN.
2. It is well known and easy to check that, for each $k \geq 2$, constraint language k -IHBS, as defined in Section 1, has polymorphism $x \wedge (y \vee z)$, but the operation $x \vee y$ is not a polymorphism of k -IHBS.

In general, it is well known that the set $\text{Pol}(\Gamma)$ of any constraint language Γ is a clone, which means that it contains all projections (i.e. operations of the form $p_n^i(x_1, \dots, x_n) = x_i$) and is closed under composition. The latter condition is spelled out as follows: if f, g_1, \dots, g_n are polymorphisms of Γ where f is n -ary and g_1, \dots, g_n are m -ary then the m -ary operation $h(x_1, \dots, x_m) = f(g_1(x_1, \dots, x_m), \dots, g_n(x_1, \dots, x_m))$ is also a polymorphism of Γ .

The complexity of constant-factor approximation for Min CSP(Γ) is completely determined by $\text{Pol}(\Gamma)$, as the next theorem indicates.

Theorem 2 ([33]). *Let Γ and Γ' be constraint languages on A such that $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma')$. Assume, in addition, that Γ contains the equality relation eq_A . Then, if $\text{Min CSP}(\Gamma)$ has a constant-factor approximation algorithm then so does $\text{Min CSP}(\Gamma')$.*

We say that BLP decides $\text{CSP}(\Gamma)$ if, for any instance I of $\text{CSP}(\Gamma)$, I is satisfiable whenever $\text{Opt}_{\text{LP}}(I) = 0$.

Theorem 3 ([39]). *For any Γ , the following are equivalent:*

1. BLP decides $\text{CSP}(\Gamma)$,
2. Γ has symmetric polymorphisms of all arities.

Note that symmetric polymorphisms provide a natural rounding for $\text{BLP}(I)$, as follows. Let s be an optimal solution to $\text{BLP}(I)$ in which all variables are assigned rational numbers such that, for some $n \in \mathbb{N}$, $p_v \in \Delta_n(A)$ for each variable v in I and $p_C \in \Delta_n(A^{\text{arity}(C)})$ for each constraint C in I . Then each v can be assigned the element $f(p_v)$ where f is any fixed n -ary symmetric polymorphism of Γ . It is not hard to check (or see [39]) that if $\text{Opt}_{\text{LP}}(I) = 0$ then this assignment will satisfy all constraints in I .

It was claimed in [39] that the conditions in Theorem 3 are also equivalent to the condition of having totally symmetric polymorphisms of all arities, but a flaw was later discovered in the proof of this claim, and indeed a counterexample (see Section 3.4) was found in [44] (Example 99).

3. Results

We will first formally state our main results and then go into detailed proofs.

3.1. Formal statements of main results

For any function $\varrho : A^k \rightarrow [0, 1]$, let R_ϱ denote the k -ary relation $R_\varrho = \{\mathbf{a} \mid \varrho(\mathbf{a}) = 0\}$.

Theorem 4. *Let Γ_1 be a valued constraint language and let $\Gamma_2 = \{R_\varrho \mid \varrho \in \Gamma_1\}$. Then $\text{VCSP}(\Gamma_1)$ is in APX if and only if $\text{Min CSP}(\Gamma_2)$ is in APX.*

Hence, for every valued constraint language there is an equivalent (relational) constraint language. Due to this reduction, we can freely focus on Min CSPs. Regarding Min CSPs, we will formulate most of our results for constraint languages Γ that contain the equality relation eq_A . We make this restriction because some of the reductions in this paper and some papers that we use are currently known to work only with this restriction. We conjecture that this restriction is not essential, that is, for any Γ , $\text{Min CSP}(\Gamma)$ admits a constant-factor approximation algorithm if and only if $\text{Min CSP}(\Gamma \cup \{\text{eq}_A\})$ does so (though the optimal constants may differ).

As mentioned before, for any Γ , $\text{CSP}(\Gamma)$ admits a robust algorithm with linear loss if and only if $\text{Min CSP}(\Gamma)$ has a constant-factor approximation algorithm. For constraint languages Γ containing the equality relation eq_A , it follows from results in Section 3 of [33] that a complete characterisation of constant-factor approximability of Min CSP reduces to the case when Γ contains all unary singletons, i.e., relations $\{a\}$, $a \in A$.

Hence, some statements in the paper will (explicitly) assume this condition. Note that this condition implies that all polymorphisms of Γ are idempotent.

Theorem 1 provides evidence, in terms of integrality gap, that the BLP is optimal to design constant-factor approximation algorithms for $\text{Min CSP}(\Gamma)$. Our main result is a characterization of problems $\text{Min CSP}(\Gamma)$ for which BLP has a finite integrality gap.

For $p, q \in \Delta$, let $\text{dist}(p, q) = \max_{a \in A} |p(a) - q(a)|$. For a tuple $\mathbf{a} \in A^n$, let $d_{\mathbf{a}} \in \Delta_n$ be such that each element $x \in A$ appears in \mathbf{a} exactly $n \cdot d_{\mathbf{a}}(x)$ times. For tuples $\mathbf{a}, \mathbf{b} \in A^n$, define $\text{dist}(\mathbf{a}, \mathbf{b}) = \text{dist}(d_{\mathbf{a}}, d_{\mathbf{b}})$. An n -ary *fractional operation* ϕ on A is any probability distribution on the set of n -ary operations on A . For every real number $c \geq 0$, call ϕ *c-Lipschitz*¹ if its support consists of symmetric operations and, for all $\mathbf{a}, \mathbf{b} \in A^n$, we have $\Pr_{g \sim \phi} \{g(\mathbf{a}) \neq g(\mathbf{b})\} \leq c \cdot \text{dist}(\mathbf{a}, \mathbf{b})$.

Theorem 5. *For any Γ containing eq_A , the following are equivalent:*

1. *The integrality gap of BLP for $\text{Min CSP}(\Gamma)$ is finite.*
2. *There is $c \geq 0$ such that, for each $n \in \mathbb{N}$, there is an n -ary c -Lipschitz fractional operation ϕ_n on A whose support consists of symmetric polymorphisms of Γ .*

We now give an example of how Theorem 5 can be applied to prove negative results. Recall Example 1. It is known and not hard to check that the operation f_n is the only n -ary symmetric polymorphism of 3-HORN. It follows that there is only one fractional operation of arity n , ϕ_n , whose support consists of symmetric polymorphisms and that $\Pr_{g \sim \phi_n} \{g(\mathbf{a}) \neq g(\mathbf{b})\} = 1$ if we choose $\mathbf{a} = (1, 1, \dots, 1)$ and $\mathbf{b} = (0, 1, \dots, 1)$. Since $\text{dist}(\mathbf{a}, \mathbf{b}) = 1/n$, it follows that there is no constant $c \geq 0$ satisfying condition (2) of Theorem 5, and hence the integrality gap of BLP for $\text{Min CSP}(3\text{-HORN})$ is infinite (and constant-factor approximation for $\text{Min CSP}(3\text{-HORN})$ is UG-hard).

On the algorithmic side, any sequence ϕ_n , $n \in \mathbb{N}$, satisfying condition (2) of Theorem 5 can be used to obtain a (possibly efficient) randomized rounding procedure for BLP, as follows. As we explained after Theorem 3, if one has an optimal rational solution to $\text{BLP}(I)$, one can use a symmetric operation of appropriate arity n to round this solution to obtain a solution for I . If the symmetric operation is drawn from a c -Lipschitz distribution ϕ_n on n -ary symmetric polymorphisms (such as in Theorem 5) then this procedure is a randomized constant-factor approximation algorithm for $\text{Min CSP}(\Gamma)$ (this follows from the proof of direction (2) \Rightarrow (1) of Theorem 5). However it is not entirely clear how to efficiently sample from ϕ_n . In Subsection 3.4, we give two examples - a class of constraint languages and one specific language - with sequences of Lipschitz distributions that are nice enough to admit efficient sampling. The first of these examples, given in Theorem 6, covers (in a specific sense - see discussion in Subsection 3.4) all problems $\text{Min CSP}(\Gamma)$ that were previously known to belong to APX, but are not efficiently solvable to optimality.

Theorem 6. *Let A consist of subsets of a set and suppose that A is closed under intersection \cap and union \cup . If a constraint language Γ on A has polymorphism $x \cap (y \cup z)$ then $\text{Min CSP}(\Gamma)$ has a constant-factor approximation algorithm.*

¹In the conference version we used the terminology *stable* instead of *Lipschitz*.

Our last result strengthens UG-hardness to NP-hardness for a class of Min CSPs. A near-unanimity polymorphism (see definition in Subsection 2.2) is a type of polymorphism well known in the algebraic theory of CSP [43, 40, 7], and its presence is implied by the existence of those Lipschitz distributions (see Subsection 3.5). We show $\text{Min CSP}(\Gamma)$ is NP-hard to constant-factor approximate if Γ has no near-unanimity polymorphism.

Theorem 7. *Let Γ be a constraint language containing eq_A and all unary singleton relations. If $\text{Min CSP}(\Gamma)$ admits a constant-factor approximation algorithm then Γ has an NU polymorphism, unless $\text{P} = \text{NP}$.*

It is well known and easy to check that 3-HORN has no NU polymorphism of any arity, so, by Theorem 7, $\text{Min CSP}(3\text{-HORN})$ is NP-hard to constant-factor approximate.

3.2. Reduction from VCSP to Min CSP

Proof. (of Theorem 4). We can assume that some function in Γ_1 takes a positive value, since otherwise both problems are trivial. Let $m > 0$ denote the minimal positive value taken by any function in Γ_1 . Note, that there is a natural one-to-one correspondence between instances in $\text{VCSP}(\Gamma_1)$ and instances in $\text{Min CSP}(\Gamma_2)$, namely, every instance $I_1 = (V, A, C)$ in $\text{VCSP}(\Gamma_1)$ is associated to the instance I_2 in $\text{Min CSP}(\Gamma_2)$, obtained by replacing every constraint (\mathbf{v}, R_ϱ) in I_1 by $\varrho(\mathbf{v})$. The theorem follows from the observation that the values every assignment $s : V \rightarrow A$ in I_1 and I_2 are within a multiplicative factor of each other. More precisely, if v_1 and v_2 are the values of assignment s for instances I_1 and I_2 respectively, then

$$v_1 \leq v_2 \leq \frac{v_1}{m}.$$

□

3.3. Finite integrality gaps

In this subsection we prove Theorem 5. We need a few definitions and intermediate results.

Let I be any weighted instance in $\text{Min CSP}(\Gamma)$ with variable set V . A *fractional assignment* for I is any probability distribution, ϕ , on the set of assignments for I . For a real number $c \geq 1$, we say that a fractional assignment ϕ for I is *c-bounded* if, for every constraint $C = (v_1, \dots, v_r, R)$ in I ,

$$\Pr_{g \sim \phi} \{g(v_1), \dots, g(v_r) \notin R\} \leq c \cdot (1 - w_C)$$

where w_C is the weight in I of constraint C . We will apply it only to instances where $w_C \in [0, 1]$.

For every relation $R \in \Gamma$ of arity, say r , and every $p_1, \dots, p_r \in \Delta$ we define $\text{loss}(p_1, \dots, p_r, R) \in [0, 1]$ to be $\min_p (1 - p(R))$ where p ranges over all the probability distributions on A^r with marginals p_1, \dots, p_r .

In a technical sense, the function loss 'encodes' the contribution of each constraint in optimal solutions of BLP. This is formalized in the following observation.

Observation 1. *Let I be any instance of Min CSP(Γ) and let $C = (v_1, \dots, v_r, R)$ be any of its constraints. Then $1 - p_C(R) = \text{loss}(p_{v_1}, \dots, p_{v_r}, R)$ holds in any optimal solution of BLP(I).*

For every $n \in \mathbb{N}$, the n -th *universal instance* for Γ , $U_n(\Gamma)$, is the instance with variable set Δ_n containing for every relation R of arity, say r , in Γ , and every $p_1, \dots, p_r \in \Delta_n$, constraint (p_1, \dots, p_r, R) with weight $1 - \text{loss}(p_1, \dots, p_r, R)$. We write simply U_n if Γ is clear from the context.

The following is a variant of Farkas' lemma that we will use in our proofs.

Lemma 1. (*Farkas' Lemma*) *Let M be a $m \times n$ matrix, $b \in \mathbb{R}^m$. Then exactly one of the following two statements is true:*

1. *There is an $x \in (\mathbb{R}_{\geq 0})^n$ with $\|x\|_1 = 1$ ($\|x\|_1$ denotes the 1-norm of x) such that $Mx \leq b$.*
2. *There is a $y \in (\mathbb{R}_{\geq 0})^m$ with $\|y\|_1 = 1$ such that $y^T b < y^T M$ (i.e. each coordinate of $y^T M$ is strictly greater than $y^T b$).*

Proof. Condition (1) is equivalent to the existence of $x \in (\mathbb{R}_{\geq 0})^n$ such that $M'x \leq b'$ where M' and b' are obtained by adding two extra rows to M and b expressing that $\sum_{1 \leq i \leq n} x_i \leq 1$ and $\sum_{1 \leq i \leq n} -x_i \leq -1$. It then follows from Corollary 7.1f in [45] that the negation of condition (1) is equivalent to the existence of a vector $z \in (\mathbb{R}_{\geq 0})^{(m+2)}$ satisfying $z^T M' \geq 0$ and $z^T b' < 0$. It is easy to see that this is equivalent to condition (2). \square

Theorem 5 follows directly from Lemmas 2 and 3 below.

Lemma 2. *For every constraint language Γ and $c \geq 1$, the following are equivalent:*

1. *The integrality gap of BLP for Min CSP(Γ) is at most c .*
2. *For each $n \in \mathbb{N}$, there is a c -bounded fractional assignment for U_n .*

Proof. This proof is an adaptation of the proof of Theorem 1 in [42], and it also works for valued CSPs.

(2 \Rightarrow 1) Let $I = (V, A, \mathcal{C})$ be any instance of Min CSP(Γ), and let $p_v (v \in V)$, $p_C (C \in \mathcal{C})$ be any optimal solution of BLP(I). We can assume that there exists $n \in \mathbb{N}$ such that $p_v \in \Delta_n$ for every $v \in V$. For every assignment g for U_n , let s_g be the assignment for I defined as $s_g(v) = g(p_v), v \in V$.

Since (2) holds, it follows from Observation 1 and the definition of c -boundedness that, for every constraint $C = (v_1, \dots, v_r, R)$ in I , we have

$$\Pr_{g \sim \phi} \{(s_g(v_1), \dots, s_g(v_r)) \notin R\} \leq c \cdot (1 - p_C(R))$$

It follows that the expected value of s_g is at most $c \cdot \text{Opt}_{\text{LP}}(I)$. Consequently, there exists some s_g with value at most $c \cdot \text{Opt}_{\text{LP}}(I)$.

(1 \Rightarrow 2) We shall prove the contrapositive. Assume that for some $n \in \mathbb{N}$, there is no c -bounded fractional assignment for U_n . We shall write a system of linear inequalities

that expresses the existence of a c -bounded fractional assignment for U_n and then apply Lemma 1 to this system. To this end, we introduce a variable x_g for every assignment g for U_n . The system contains, for every constraint $C = (p_1, \dots, p_r, R)$ in U_n , the inequality:

$$\sum_{g \in G_n} x_g \cdot \mathbf{1}[(g(p_1), \dots, g(p_r)) \notin R] \leq c \cdot \text{loss}(C)$$

where G_n is the set of all assignments for U_n and $\mathbf{1}[(g(p_1), \dots, g(p_r)) \notin R]$ is 1 if $(g(p_1), \dots, g(p_r)) \notin R$ and 0 otherwise. Note that the system does not include equations for $x_g \geq 0$ and $\sum_{g \in G_n} x_g = 1$ since this is already built-in in the version of Farkas' lemma that we use.

Since there is no c -bounded fractional assignment for U_n it follows from Farkas' Lemma that the system containing for every $g \in G_n$ inequality

$$\sum_{C=(p_1, \dots, p_r, R) \in U_n} y_C \cdot c \cdot \text{loss}(C) < \sum_{C=R(p_1, \dots, p_r, R) \in U_n} y_C \cdot \mathbf{1}[(g(p_1), \dots, g(p_r)) \notin R] \quad (3)$$

has a solution where every variable y_C takes non-negative values and it holds that $\sum_C y_C = 1$. We can also assume the value of every variable in the solution is rational, since so are all the coefficients in the system.

Now consider instance $I = (V, A, \mathcal{C})$ where $V = \Delta_n$ and \mathcal{C} contains, for every relation $R \in \Gamma$ of arity, say r , and every $p_1, \dots, p_r \in \Delta_n$, constraint $C = (v_1, \dots, v_r, R)$ with weight y_C .

We shall construct a solution $p_v (v \in \Delta_n), p_C (C \in \mathcal{C})$ of $\text{BLP}(I)$. For every $v \in \Delta_n$, set p_v to v (note that v is a distribution on A). For every $C \in \mathcal{C}$ set p_C to the distribution q with $1 - q(C) = \text{loss}(C)$. Hence, the objective value of the solution of $\text{BLP}(I)$ thus constructed is $\sum_{C \in U_n} y_C \cdot \text{loss}(C)$, which is c times smaller than the left side of inequality (3). Furthermore, the total weight of falsified constraints by any assignment g for I is precisely the right side of inequality (3). It follows that the gap of instance I is larger than c . \square

For every set X , one can associate to every $p \in \Delta_n(X)$ the multiset p' such every element $x \in X$ occurs in p' exactly $p(x) \cdot n$ times. In a similar way, one obtains a one-to-one correspondence between the assignments (resp. fractional assignments) for U_n and the n -ary symmetric operations (resp. fractional operations with support consisting of n -ary symmetric operations).

Lemma 3. *For every constraint language Γ containing the equality relation eq_A , the following are equivalent:*

1. *There is $c \geq 1$, such that for each $n \in \mathbb{N}$, there is a c -bounded fractional assignment for $U_n(\Gamma)$.*
2. *There is $d \geq 0$ such that, for each $n \in \mathbb{N}$, there is an n -ary d -Lipschitz fractional operation on A whose support consists of symmetric polymorphisms of Γ .*

Proof. In this proof it is convenient to distinguish formally between a distribution y (resp. assignment, fractional assignment) and its associated multiset y (resp. operation,

fractional operation) that, whenever X and n are clear from the context, we shall denote by y' . The following observation will be useful.

Observation 2. *For any assignment g for U_n and any distribution $p \in \Delta_n(A^r)$ we have that $(g(p_1), \dots, g(p_r)) = g'(p')$ where $p_1, \dots, p_r \in \Delta_n(A)$ are the marginals of p and $g'(p')$ denotes the r -ary tuple obtained by applying the symmetric n -ary operation g' (corresponding to g) to the n tuples in p' component-wise.*

(1) \Rightarrow (2) Assume that ϕ is a c -bounded fractional assignment for U_n . We claim that for every mapping g in the support of ϕ , g' is, in fact, a polymorphism of Γ . Indeed, let R be any relation of arity, say r , in Γ , let $t_1, \dots, t_n \in R$. We want to show that $g'(t_1, \dots, t_n) \in R$ where $g'(t_1, \dots, t_n)$ denotes the r -ary tuple obtained by applying g' to t_1, \dots, t_n component-wise.

Let $p \in \Delta_n(A^r)$ be the distribution associated to multiset $p' = [t_1, \dots, t_n]$ and consider constraint $C = (p_1, \dots, p_r, R)$ on U_n where p_1, \dots, p_r are the marginals of p . By the choice of p we have $p(R) = 1$. Since ϕ is c -bounded it follows that $\Pr_{g \sim \phi} \{(g(p_1), \dots, g(p_r)) \notin R\} \leq c \cdot \text{loss}(C) \leq 1 - p(R) = 0$. Hence, $(g(p_1), \dots, g(p_r)) \in R$ for every g in the support of ϕ . It follows from Observation 2 that $g'(t_1, \dots, t_n) = (g(p_1), \dots, g(p_r))$ and we are done.

We have just seen that the support of the fractional n -ary operation, ϕ' , associated to ϕ consists of polymorphisms of Γ . Since, by definition, the support of ϕ' only contains symmetric operations, in order to complete the proof it suffices to show that ϕ' is $(c \cdot |A|)$ -Lipschitz.

Let $p'_1, p'_2 \in A^n$ and consider constraint $C = (p_1, p_2, \text{eq}_A)$ in U_n where $p_1, p_2 \in \Delta_n$ are the distributions associated to p'_1 and p'_2 respectively and eq_A is the equality relation on A . It is not too difficult to find a distribution p on A^2 with marginals p_1 and p_2 such that $1 - p(\text{eq}_A) \leq |A| \cdot \text{dist}(p'_1, p'_2)$. A concrete example can be obtained as follows. For every $a \in A$, let $a_1 = \max\{p_1(a) - p_2(a), 0\}$, and $a_2 = \max\{p_2(a) - p_1(a), 0\}$. Also, let $s = \sum_a a_1 = \sum_a a_2$. Then we define p as follows:

$$p(a, b) = \begin{cases} \min\{p_1(a), p_2(b)\} & \text{if } a = b \\ \frac{a_1 \cdot b_2}{s} & \text{if } a \neq b \end{cases}$$

It is easy to verify that p satisfies the desired conditions. Finally, we have

$$\Pr_{g' \sim \phi'} \{g'(p'_1) \neq g'(p'_2)\} = \Pr_{g \sim \phi} \{(g(p_1), g(p_2)) \notin \text{eq}_A\} \leq c \cdot \text{loss}(C) \leq c \cdot |A| \cdot \text{dist}(p'_1, p'_2).$$

We note that this is the only part where the condition $\text{eq}_A \in \Gamma$ is required.

(2) \Rightarrow (1). For every $n \in \mathbb{N}$, let n' be a multiple of n to be fixed later, let ϕ' be a d -Lipschitz fractional operation of arity n' whose support consists of symmetric polymorphisms of Γ , and let ϕ be its associated fractional assignment for $U_{n'}$. We can without loss of generality assume that $d \geq 1$. We shall prove later that, for every constraint $C = (p_1, \dots, p_r, R)$ in U_n (note, not in $U_{n'}$), we have

$$\Pr_{g \sim \phi} \{(g(p_1), \dots, g(p_r)) \notin R\} \leq 2 \cdot r \cdot d \cdot \text{loss}(C) \quad (4)$$

Consider now the fractional assignment ϕ^* on U_n where for every assignment f on U_n , $\phi^*(f) = \sum_g \phi(g)$ where g ranges over all assignments for $U_{n'}$ that *extend* f (that is, such that $f(p) = g(p)$ for every $p \in \Delta_n$). It follows from the definition ϕ^* that

$$\Pr_{f \sim \phi^*} \{(f(p_1), \dots, f(p_r)) \notin R\} = \Pr_{g \sim \phi} \{(g(p_1), \dots, g(p_r)) \notin R\}$$

for every constraint (p_1, \dots, p_r, R) in U_n . This gives a way to construct, for every $n \in \mathbb{N}$, a $(2 \cdot K \cdot d)$ -bounded fractional assignment for U_n where K is the maximum arity of a relation in Γ .

To finish the proof it only remains to prove inequality (4) for any constraint $C = (p_1, \dots, p_r, R)$ in U_n . Let p be a distribution on A^r such that $1 - p(R) = \text{loss}(C)$ is achieved. We can assume that $\text{loss}(C) \leq 1/2$ since otherwise there is nothing to prove.

Note that we can assume that $p(t)$ is rational for every $t \in A^r$. Let q be the distribution on A^r defined as

$$q(t) = \begin{cases} p(t)/p(R) & t \in R \\ 0 & t \notin R \end{cases}$$

Consider constraint (q_1, \dots, q_r, R) where q_1, \dots, q_r are the marginals of q . Since the number of constraints in U_n is finite we can assume that n' has been picked such that $q \in \Delta_{n'}(A^r)$. We claim that $(g(q_1), \dots, g(q_r)) \in R$ for any g in the support of ϕ . Indeed, if $g' = [t_1, \dots, t_{n'}]$ is the multiset of tuples in A^r associated to q then by Observation 2 $(g(q_1), \dots, g(q_r)) = g'(t_1, \dots, t_{n'})$ and the latter tuple belongs to R because g' is a polymorphism of Γ .

We claim that $\text{dist}(p_i, q_i) \leq 2 \cdot \text{loss}(C)$ for every $1 \leq i \leq r$. By definition, for every $a \in A$, $q_i(a) = \sum_{t=(t^1, \dots, t^r) \in R, t^i=a} p(t)/p(R)$. Since $p(R) = 1 - \text{loss}(C)$, we have

$$\frac{p_i(a) - \text{loss}(C)}{1 - \text{loss}(C)} = \sum_{t^i=a} \frac{p(t)}{p(R)} - \sum_{t \notin R} \frac{p(t)}{p(R)} \leq q_i(a) \leq \sum_{t^i=a} \frac{p(t)}{p(R)} = \frac{p_i(a)}{1 - \text{loss}(C)}$$

for every $a \in A$. Moreover, we have $\frac{p_i(a)}{1 - \text{loss}(C)} \leq p_i(a) + 2 \cdot \text{loss}(C)$, which immediately follows from $\frac{p_i(a)}{1 - \text{loss}(C)} = p_i(a) + \frac{\text{loss}(C)p_i(a)}{1 - \text{loss}(C)}$, $\text{loss}(C) \leq 1/2$ and $p_i(a) \leq 1$. Hence, $q_i(a) \in [p_i(a) - \text{loss}(C), p_i(a) + 2 \cdot \text{loss}(C)]$ for every $a \in A$. We conclude that

$$\Pr_{g \sim \phi} \{(g(p_1), \dots, g(p_r)) \notin R\} \leq \Pr_{g \sim \phi} \{\exists i \text{ such that } g(p_i) \neq g(q_i)\} \leq 2 \cdot r \cdot d \cdot \text{loss}(C).$$

□

3.4. Algorithms

We now prove Theorem 6 and then describe another constraint language Γ which admits nicely structured Lipschitz distributions on symmetric polymorphisms (so that its Min CSP is constant-factor approximable).

Two classes of CSPs were introduced and studied in [46], one is a subclass of the other. We need two notions to define these classes. A *distributive lattice* (L, \cap, \cup) is a (lattice representable by a) family L of subsets of a set closed under intersection \cap and union \cup . We say that two constraint languages $\Gamma_1 = \{R_1^{(1)}, \dots, R_m^{(1)}\}$ on domain A and $\Gamma_2 = \{R_1^{(2)}, \dots, R_m^{(2)}\}$ on domain B , where the arities of corresponding relations

match, are *homomorphically equivalent* if there are two mappings $f : A \rightarrow B$, $g : B \rightarrow A$ such that for all $1 \leq i \leq m$, $f(t_1) \in R_i^{(2)}$ for every $t_1 \in R_i^{(1)}$ and $g(t_2) \in R_i^{(1)}$ for every $t_2 \in R_i^{(2)}$. The smaller class, which we shall call \mathcal{C} (from 'languages with caterpillar duality'), consists of constraint languages Γ such that Γ is homomorphically equivalent to a constraint language Γ' on some family L of subsets of a finite sets that has polymorphisms \cap and \cup , where \cap and \cup are the usual set-theoretic union and intersection (i.e. (L, \cap, \cup) is a finite distributive lattice). The larger class, which we shall call \mathcal{J} (from 'languages with jellfish duality') is defined similarly, but we require Γ' to have polymorphism $x \cap (y \cup z)$. Constraint languages k -IHBS (defined in Section 1) belong to the class \mathcal{J} , but not to \mathcal{C} (see Example 1). See [46] for other specific examples of CSPs contained in these classes. For every Γ in \mathcal{C} , $\text{Min CSP}(\Gamma)$ was shown to belong to APX in [39]. This result was extended to \mathcal{J} in [33] (see Theorems 5.8 and 4.8 there).

We will now show how Lipschitz distributions on symmetric polymorphisms can be used to provide a constant-factor approximation algorithm for $\text{Min CSP}(\Gamma)$ for every Γ in this class. Observe that if Γ and Γ' are homomorphically equivalent then $\text{Min CSP}(\Gamma)$ and $\text{Min CSP}(\Gamma')$ are essentially the same problem because there is an obvious one-to-one correspondence between instances of $\text{Min CSP}(\Gamma_1)$ and $\text{Min CSP}(\Gamma_2)$ (swapping $R_i^{(1)}$ and $R_i^{(2)}$ in all constraints) and the maps f and g allow one to move between solutions to corresponding instances without any loss of quality. So, we can assume that A consists of subsets of some set, and Γ has polymorphism $x \cap (y \cup z)$ where (A, \cap, \cup) is a distributive lattice.

Throughout the section, K will denote the maximum arity of a relation in such Γ . For every $1 \leq h \leq n$, let $g_{h,n}(x_1, \dots, x_n)$ be the n -ary symmetric operation on A defined as

$$\bigcup_{I \subseteq \{1, \dots, n\}, |I|=h} \left(\bigcap_{i \in I} x_i \right)$$

Lemma 4. *For all $h, n \in \mathbb{N}$ with $(1 - \frac{1}{|A|^K})n < h \leq n$, we have $g_{h,n} \in \text{Pol}(\Gamma)$.*

Proof. It is not difficult to see that $x \cap y$ is also a polymorphism of Γ . Indeed, for every relation R and every pair of tuples $t, t' \in R$, we have that $t \cap t' = t \cap (t' \cup t')$ and hence it belongs to R . Using composition, we shall show that R has polymorphism $f_{h,n}$ where $f_{h,n}(x_0, x_1, \dots, x_n)$ is the $(1+n)$ -ary operation defined as

$$x_0 \cap g_{h,n}(x_1, \dots, x_n) = x_0 \cap \left(\bigcup_{I \subseteq \{1, \dots, n\}, |I|=h} \left(\bigcap_{i \in I} x_i \right) \right)$$

First, we observe that for every $m \geq 2$, the m -ary operation $x_1 \cap \dots \cap x_m$ preserves R as it can be obtained by composition from $x \cap y$ by $x_1 \cap (x_2 \cap (x_3 \cap \dots \cap (x_{m-1} \cap x_m) \dots))$. In a bit more complicated fashion we can show that $x_0 \cap (x_1 \cup \dots \cup x_m)$ preserves R for every $m \geq 3$. If $m = 3$ it follows that $x_0 \cap ((x_0 \cap (x_1 \cup x_2)) \cup x_3)$ is equal to $x_0 \cap (x_1 \cup x_2 \cup x_3)$ (recall that \cup and \cap are the set union and intersection respectively). The pattern generalizes easily to arbitrary values for m . Finally, one obtains $f_{h,n}$ by suitably composing $x_0 \cap (x_1 \cup \dots \cup x_{\binom{n}{h}})$ and $x_1 \cap \dots \cap x_h$.

Let R be a relation in Γ of arity, say, r and let t_1, \dots, t_n be a list of (not necessarily distinct) tuples in R . By the pigeon-hole principle, there exists a tuple t appearing at

least $\lceil n/|A|^r \rceil$ times in t_1, \dots, t_n . It follows from the choice of h and t , that for every $I \subseteq \{1, \dots, n\}$, with $|I| = h$, there exists $i \in I$ such that $t = t_i$. It then follows that $f_{h,n}(t, t_1, \dots, t_n)$, which necessarily belongs to R , is precisely $g_{h,n}(t_1, \dots, t_n)$ \square

For every natural number $n \in \mathbb{N}$, consider the n -ary fractional operation ϕ_n with support $\{g_{h,n} \mid (1 - \frac{1}{|A|^K})n < h \leq n\}$ that distributes uniformly among the operations of its support.

Lemma 5. *There exists some $c \geq 0$ such that ϕ_n is c -Lipschitz for every $n \in \mathbb{N}$.*

Proof. Let $\mathbf{a} = (a_1, \dots, a_n), \mathbf{b} = (b_1, \dots, b_n) \in A^n$. Recall that from distributivity we assume that every element $a \in A$ is a subset of some set that we call S . Note that, according to the definition of $g_{h,n}$, an element $j \in S$ belongs to $g_{h,n}(\mathbf{a})$ iff $|\mathbf{a}|_j \geq h$ where $|\mathbf{a}|_j$ is defined to be $|\{1 \leq i \leq n \mid j \in a_i\}|$. Consequently, $g_{h,n}(\mathbf{a}) \neq g_{h,n}(\mathbf{b})$ iff there exists some $j \in S$ such that $|\mathbf{a}|_j < h \leq |\mathbf{b}|_j$ or $|\mathbf{b}|_j < h \leq |\mathbf{a}|_j$. It follows that

$$\begin{aligned} \Pr_{g \sim \phi_n} \{g(\mathbf{a}) \neq g(\mathbf{b})\} &= \frac{|\{h \mid \exists j (|\mathbf{a}|_j < h \leq |\mathbf{b}|_j \vee |\mathbf{b}|_j < h \leq |\mathbf{a}|_j)\}|}{n/|A|^K} \\ &\leq \sum_{j \in S} \frac{|\{h \mid |\mathbf{a}|_j < h \leq |\mathbf{b}|_j \vee |\mathbf{b}|_j < h \leq |\mathbf{a}|_j\}|}{n/|A|^K} \\ &= \frac{1}{n/|A|^K} \sum_{j \in S} \left| |\mathbf{a}|_j - |\mathbf{b}|_j \right| \leq |A|^K \cdot |S| \cdot \text{dist}(\mathbf{a}, \mathbf{b}). \end{aligned}$$

\square

With the help of the sequence ϕ_n , we can prove Theorem 6, i.e. obtain a constant-factor approximation algorithm for Min CSP(Γ). A different proof of this result was given in [33].

Proof. (of Theorem 6). Let $I = (V, A, \mathcal{C})$ be any instance of Min CSP(Γ) and let $p_v(v \in V), p_C(C \in \mathcal{C})$ be an optimal solution of BLP(I) with objective value $\text{Opt}_{\text{LP}}(I)$. We can assume that there exists some $n \in \mathbb{N}$ such that all the probabilities in the solution are of the form n'/n where n' is a non-negative integer. Also we can assume that $\log(n)$ is polynomial in the size of instance I .

Consider an assignment s for I obtained in the following way: draw $g_{h,n}$ according to ϕ_n (i.e. select $(1 - \frac{1}{|A|^K})n < h \leq n$ uniformly at random) and assign $s(v) = g_{h,n}(p'_v)$ where p'_v is any tuple such that every $a \in A$ appears exactly $p_v(a) \cdot n$ times in p'_v . It can be shown (this is basically the proof of direction $(2 \Rightarrow 1)$ of Theorem 5) that there exists some $c' \geq 1$ such that expected value of assignment s is $c' \cdot \text{Opt}_{\text{LP}}(I)$. In particular, c' can be taken to be $2Kc$ where c is the Lipschitz constant of ϕ_n .

We shall prove that there is a randomized polynomial-time algorithm that constructs s . Select $(1 - \frac{1}{|A|^K})n < h \leq n$ uniformly at random. Recall that we assume that every element $a \in A$ is a subset of some set that we call S . Hence, in order to compute $g_{h,n}(p'_v)$, it is only necessary to give an efficient procedure that decides, for every $j \in S$, whether $j \in g_{h,n}(p'_v)$. Note that, according to the definition of $g_{h,n}$, $j \in g_{h,n}(p'_v)$ iff the

number, $|p'_v|_j$, of entries in tuple p'_v that contain j is at least h . This number can be easily computed from p_v as $|p'_v|_j = n \cdot \sum_{\{a \in A \mid j \in a\}} p_v(a)$. \square

We finish this subsection by introducing another constraint language Γ such that $\text{Min CSP}(\Gamma)$ admits a constant-factor approximation algorithm. The interest of this result is in the fact that it is the first known example of a constraint language where $\text{Min CSP}(\Gamma)$ has a constant-factor approximation algorithm but is not invariant under totally symmetric polymorphisms of all arities (i.e. Γ does not have the so-called width 1 property [7]). This constraint language has domain $A = \{-1, 0, +1\}$ and contains relations $R_+ = \{(a_1, a_2, a_3) \in A^3 \mid a_1 + a_2 + a_3 \geq 1\}$ and $R_- = \{(a_1, a_2, a_3) \in A^3 \mid a_1 + a_2 + a_3 \leq -1\}$. This is the example in [44] that we mentioned after Theorem 3. It is easy to show that this constraint language has no totally symmetric polymorphism of arity 3.

However $\{R_+, R_-\}$ have many symmetric polymorphisms. In particular, it is not difficult to see that, for all $h, n \in \mathbb{N}$ with $h < \lfloor n/3 \rfloor$, operation

$$s_{h,n}(x_1, \dots, x_n) = \begin{cases} 1 & \text{if } h < \sum_i x_i \\ 0 & \text{if } -h \leq \sum_i x_i \leq h \\ -1 & \text{if } \sum_i x_i < -h \end{cases}$$

preserves Γ . It is also easy to show that the n -ary fractional operation with support $\{s_{h,n} \mid h < \lfloor n/3 \rfloor\}$ that distributes uniformly among the operations of its support is 3-Lipschitz and that can be efficiently sampled. Consequently, $\text{Min CSP}(\{R_+, R_-\})$ has a constant-factor approximation algorithm.

3.5. NP-hardness result

In this subsection we prove Theorem 7, i.e. show that, modulo $P \neq NP$, if $\text{Min CSP}(\Gamma)$ admits a constant-factor-approximation algorithm then Γ must have a near-unanimity (NU) polymorphism (recall the definition of an NU operation from Section 2.2). NU polymorphisms have been well studied in universal algebra [47] and have been applied in CSP [43, 40, 7, 38]. For example, every relation invariant under an n -ary NU operation is uniquely determined by its $(n-1)$ -ary projections [47], and NU polymorphisms characterize CSPs of “bounded strict width” [7].

We can assume (proved in Lemma 3.7 of [33]) that Γ contains all unary singleton relations $\{a\}$, $a \in A$. This implies that polymorphisms of Γ are idempotent. It can be easily derived from Theorem 5 that, modulo UGC, Γ must have a near-unanimity polymorphism of some (large enough) arity. Indeed, for any n -ary fractional operation ϕ_n with support on symmetric polymorphisms of Γ and every pair $a, b \in A$, the mass of operations g in the support of ϕ_n such that $g(b, a, \dots, a) \neq g(a, a, \dots, a) (= a)$ is at most $\frac{c}{n}$. Since c is constant, if we choose n large enough, some g in the support of ϕ_n will satisfy the near-unanimity identity.

In this section we shall prove it assuming only $P \neq NP$. As an intermediate step, we consider the variant of $\text{CSP}(\Gamma)$ where some constraints in an instance can be designated as *hard*, meaning that they must be satisfied in any feasible solution, while the other constraints are soft and can be falsified. It makes sense to investigate approximation algorithms for this mixed version of CSP (see, e.g. [30]). In particular, the value of a

feasible assignment for a instance of mixed Min CSP(Γ) is defined to be the number (or total weight) of soft constraints it violates. It is not difficult to see, and was mentioned in [30] that mixed Min CSP(Γ) has a constant-factor approximation algorithm if and only if the ordinary, not mixed, Min CSP(Γ) has such an algorithm.

The proof of our NP-hardness result makes use of a result about hardness of approximation for the problem Max IS $_k$ in which the goal is to find a maximum independent set in a given k -uniform hypergraph. Recall that an independent set in a hypergraph is a subset of its vertices that does not include any of its hyperedges (entirely). For real numbers $0 \leq \alpha, \beta \leq 1$, say that an algorithm (α, β) -distinguishes Max IS $_k$ if, given a k -uniform hypergraph $H = (V, E)$, it correctly decides between the following two cases:

1. the size of the largest independent set of H is at least $\beta \cdot |V|$
2. the size of the largest independent set of H is at most $\alpha \cdot |V|$.

Note that it does not matter what the algorithm does for a hypergraph falling into neither of these cases.

Theorem 8 ([48]). *For any integer $k \geq 3$ and any real number $\epsilon > 0$, it is NP-hard to $(\epsilon, 1 - \frac{1}{k-1} - \epsilon)$ -distinguish Max IS $_k$.*

The key in proof of Theorem 7 is to show that, roughly, if Γ has no NU polymorphisms then Γ can simulate (pp-define, to be precise), for every $k \geq 3$, a k -ary relation R_k such that $R_k \cap \{a, b\}^k = \{a, b\}^k \setminus \{(a, \dots, a)\}$ for some distinct $a, b \in A$. This relation, used in hard constraints, can encode a k -uniform hypergraph, while soft unary constraints using relation $\{a\}$ simulate a choice of an independent set. To make this precise we will need a few definitions.

We say that R is pp-definable from Γ if there exists a (primitive positive) formula

$$\phi(x_1, \dots, x_k) \equiv \exists y_1, \dots, y_l \psi(x_1, \dots, x_k, y_1, \dots, y_l)$$

where ψ is a conjunction of atomic formulas with relations in Γ and eq_A such that for every $(a_1, \dots, a_k) \in A^k$

$$(a_1, \dots, a_k) \in R \text{ if and only if } \phi(a_1, \dots, a_k) \text{ holds.}$$

Note that in the definition of primitive positive formulas we are slightly abusing notation by identifying a relation with its relation symbol. It is shown in [33] that if Γ contains eq_A and R is pp-definable from Γ then the problems Min CSP(Γ) and Min CSP($\Gamma \cup \{R\}$) simultaneously belong or do not belong to APX.

An n -ary operation on A is called a *weak near-unanimity (WNU)* operation if it is idempotent and satisfies the identities

$$f(y, x, \dots, x) = f(x, y, \dots, x) = \dots = f(x, x, \dots, y).$$

Proof. (of Theorem 7) Assume, towards a contradiction, that Γ falsifies the statement of the theorem.

The following lemma can be derived from a combination of several known results. We give a (more or less) direct proof for completeness.

Lemma 6. For every $k \geq 1$, there is a k -ary relation, R , pp-definable from Γ , and $a, b \in A$ such that

$$R \cap \{a, b\}^k = \{a, b\}^k \setminus \{(a, \dots, a)\}$$

Proof. It follows easily from [47] that if $\text{Pol}(\Gamma)$ does not contain any NU operation, then for every $n \geq 3$ there is a relation $T \subseteq A^n$ which is pp-definable from Γ and a tuple $(a_1, \dots, a_n) \notin T$ such that for every $1 \leq i \leq n$ there exists $c_i \in A$ such that $(a_1, \dots, a_{i-1}, c_i, a_{i+1}, \dots, a_n) \in T$. Setting $n \geq (k+2)|A|^2$ it follows from the pigeon-hole principle that there exists $a, c \in A$ and $I = \{i_1, \dots, i_{k+2}\} \subseteq \{1, \dots, n\}$ of size $k+2$ such that $a_i = a$ and $c_i = c$ for every $i \in I$. Consider relation S defined as

$$S = \{(x_{i_1}, \dots, x_{i_{k+2}}) \mid (x_1, \dots, x_n) \in T, \forall i \notin I (x_i = a)\}$$

Clearly, S is pp-definable using T and the unary singletons. It follows that S is pp-definable from Γ as well. We have that $(a, a, \dots, a) \notin S$, $t_1 = (c, a, \dots, a) \in S$, $t_2 = (a, c, \dots, a) \in S$, \dots , and $t_{k+2} = (a, a, \dots, c) \in S$. We can also assume that, in addition to the previous property, S is symmetric, meaning that if (x_1, \dots, x_{k+2}) belongs to S then so does any tuple obtained by permuting its entries. This is because we can always replace S by the relation $\{(x_1, \dots, x_{k+2}) \mid (x_{\sigma(1)}, \dots, x_{\sigma(k+2)}) \in S \text{ for every permutation } \sigma\}$ which is pp-definable from S . Since, by assumption, we have that $\text{Min CSP}(\Gamma)$ admits a constant-factor approximation algorithm it follows from Theorem 9 of [33] that Γ has a certain property, called *bounded width* (or else $\text{P} = \text{NP}$). Theorem 2.8 in [49] states that this property implies that $\text{Pol}(\Gamma)$ contains WNU polymorphisms g_3, g_4 of arity 3 and 4, respectively, such that $g_3(y, x, x) = g_4(y, x, x, x)$ holds for every $x, y \in A$. The proof of Theorem 2.8 in [49] shows how to obtain g_n for $n = 3, 4$, but the proof generalizes immediately to show that, for each $n \geq 3$, Γ has an n -ary WNU polymorphism g_n , of arity n , and the identity $g_n(y, x, \dots, x) = g_{n'}(y, x, \dots, x)$ holds for all n, n' .

Let $b = g_n(c, a, \dots, a)$ and let j be minimum with the property that S contains every tuple $t \in \{a, b\}^{k+2}$ with at least j b 's. We claim that $1 \leq j \leq 3$. The lower bound follows from the fact that $(a, \dots, a) \notin S$. For the upper bound, it follows from the fact every g_n is a WNU (and so idempotent), that every tuple $t \in \{a, b\}^{k+2}$ with $j(\geq 3)$ b 's can be obtained by applying g_j component-wise to tuples t_{i_1}, \dots, t_{i_j} where i_1, \dots, i_j are the components in t that contain a b . Since S is symmetric then it does not contain any tuple in $\{a, b\}^{k+2}$ with exactly $j-1$ b 's.

Finally, consider relation R defined as

$$R = \{(x_1, \dots, x_k) \mid \underbrace{(b, \dots, b)}_{j-1}, \underbrace{(a, \dots, a)}_{3-j}, x_1, \dots, x_k \in S\}$$

As before we infer that R is pp-definable from Γ . It follows from the definition that R , a and b satisfy the statement of the lemma. \square

Lemma 7. For every $k \geq 1$, there is a linear algorithm that, for a given k -regular hypergraph $H = (V, E)$, returns an instance I of mixed $\text{Min CSP}(\Gamma)$ such that the value of optimal solution for I is $1 - m/|V|$ where m is the size of the maximum independent set in H .

Proof. Fix $k \geq 1$ and let R and a, b be as in Lemma 6. Let $\exists y_1, \dots, y_l \psi(x_1, \dots, x_k, y_1, \dots, y_l)$ be a primitive positive formula defining R from Γ . It is well known that ψ can be seen as an instance J of CSP(Γ). More precisely, define J to be the instance that has variables $x_1, \dots, x_k, y_1, \dots, y_l$ and contains for every atomic formula $S(v_1, \dots, v_r)$ in ψ , the constraint $((v_1, \dots, v_r), S)$. It follows that for any assignment $s : \{x_1, \dots, x_k, y_1, \dots, y_l\} \rightarrow A$, s is a solution of J if and only if $\psi(s(x_1), \dots, s(x_k), s(y_1), \dots, s(y_l))$ holds.

Consider the algorithm that, given a k -regular hypergraph, $H = (V, E)$, constructs an instance I of mixed Min CSP(Γ) as follows. The set of variables of I contains, in addition to all nodes in V , some other fresh variables to be introduced later. Then, for every hyperedge $E = \{v_1, \dots, v_k\}$, add a copy of J where the variables have been renamed so that $x_1 = v_1, \dots, x_k = v_k$ and y_1, \dots, y_l are different fresh variables (different for each hyperedge). All the constraints added so far are designated as hard. Finally, add for every $v \in V$ a soft constraint $(v, \{a\})$ requiring v to take value a .

Note that as k is fixed, this can be carried out in linear time. It follows from the construction of I that for every independent set, X , of H there is an assignment for I satisfying all hard constraints that maps every node in X to a and every node in $V \setminus X$ to b . This assignment violates exactly $|V| - |X|$ soft constraints. Conversely, for every assignment s in I , the set $X = \{v \in V \mid s(v) = a\}$ is an independent set of H . \square

We are finally in a position to obtain a contradiction. As discussed above, if Min CSP(Γ) admits a constant-factor approximation algorithm then so does its mixed variant. Let δ satisfy $0 < \delta/2 \leq 1 - c \cdot \delta$ where c is any constant larger than the approximation factor for the mixed Min CSP(Γ) algorithm. We shall prove that, for every k , there is a polynomial time algorithm that $(1 - c \cdot \delta, 1 - \delta)$ -distinguishes Max IS $_k$, obtaining a contradiction since this task is NP-hard, as follows by setting $\epsilon = \delta/2$ and $k = 1 + 2/\delta$ in Theorem 8. Let us prove our claim. In order to distinguish whether the size, m , of the maximal independent set of a k -regular hypergraph $H = (V, E)$ is at most $(1 - c \cdot \delta)|V|$ or at least $(1 - \delta)|V|$ we do the following: compute the instance I of mixed Min CSP(Γ) using the linear algorithm of Lemma 7 and run the constant approximation algorithm for Min CSP(Γ) with instance I . Then, we only need to compare the value of the assignment, s , returned by the approximation algorithm with $c \cdot \delta$ to safely distinguish between the two cases. Indeed, if $m \leq (1 - c \cdot \delta)|V|$ then the optimum, Opt, of instance I , has value at least $c \cdot \delta$ from which it follows that the value of s is necessarily at least $c \cdot \delta$. Otherwise, if $m \geq (1 - \delta)|V|$ then the value of Opt is at most δ , from which it follows that the value of s is less than $c \cdot \delta$. \square

4. Conclusion

We have reduced a classification of constant-factor approximable finite-valued CSPs to that for Min CSPs. Due to technical limitations, we proved most of our results for constraint languages Γ containing the equality relation eq_A . It is in open question whether adding eq_A can ever change constant-factor approximability of Min CSP(Γ). We provided (Theorem 5) an algebraic characterisation of constraint languages Γ such that (Γ contains eq_A and) the integrality gap of BLP for Min CSP(Γ) is finite. We conjecture that Min CSP(Γ) is constant-factor approximable for all such languages, even

without the assumption on eq_A . One way to prove this could be to strengthen the algebraic characterisation so that it features only fractional operations that can be efficiently sampled from. We showed that this works, in particular, for all constraint languages Γ for which $\text{Min CSP}(\Gamma)$ was previously known to be constant-factor approximable.

On the hardness side, infinite integrality gap is known [20] to imply UG-hardness of constant-factor approximation for all constraint languages Γ containing the equality relation eq_A . For a large subclass of such languages, we improved UG-hardness to NP-hardness. Proving this for all such languages is (probably) beyond current techniques, even for very special cases such as MinUncut , but a further extension of our subclass could be within reach.

It is an open question whether condition (2) in Theorem 5 is decidable. We remark that the decidability question for the related property of having symmetric polymorphisms of all arities (see Theorem 3) is also open - see [50, 51] for related results. However, another related property - of having so-called fractional symmetric polymorphisms of all arities - is decidable [42].

5. Acknowledgments

The authors would like to thank Per Austrin, Johan Håstad, and Venkatesan Guruswami for useful discussions. We would also like to thank the anonymous referees for useful comments.

- [1] N. Creignou, S. Khanna, M. Sudan, Complexity Classifications of Boolean Constraint Satisfaction Problems, Vol. 7 of SIAM Monographs on Discrete Mathematics and Applications, 2001.
- [2] N. Creignou, P. Kolaitis, H. Vollmer (Eds.), Complexity of Constraints, Vol. 5250 of LNCS, Springer, 2008.
- [3] A. Krokhin, S. Živný (Eds.), The Constraint Satisfaction Problem: Complexity and Approximability, Vol. 7 of Dagstuhl Follow-Ups, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017.
- [4] G. Gottlob, G. Greco, F. Scarcello, Tractable optimization problems through hypergraph-based structural restrictions, in: Proceedings of ICALP'09, 2009, pp. 16–30.
- [5] D. Marx, Tractable hypergraph properties for constraint satisfaction and conjunctive queries, J. ACM 60 (6) (2013) Article 42.
- [6] A. Bulatov, P. Jeavons, A. Krokhin, Classifying complexity of constraints using finite algebras, SIAM Journal on Computing 34 (3) (2005) 720–742.
- [7] T. Feder, M. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory, SIAM Journal on Computing 28 (1998) 57–104.
- [8] T. Schaefer, The complexity of satisfiability problems, in: STOC'78, 1978, pp. 216–226.

- [9] L. Barto, M. Kozik, Constraint satisfaction problems solvable by local consistency methods, *Journal of the ACM* 61 (1) (2014) Article 3.
- [10] L. Barto, M. Kozik, Robust satisfiability of constraint satisfaction problems, *SIAM Journal on Computing* 45 (4) (2016) 1646–1669.
- [11] A. Bulatov, A dichotomy theorem for nonuniform CSPs, in: *FOCS'17, 2017*, pp. 319–330.
- [12] V. Deineko, P. Jonsson, M. Klasson, A. Krokhin, The approximability of Max CSP with fixed-value constraints, *Journal of the ACM* 55 (4) (2008) Article 16.
- [13] V. Kolmogorov, A. Krokhin, M. Rolínek, The complexity of general-valued CSPs, *SIAM Journal on Computing* 46 (3) (2017) 1087–1110.
- [14] J. Thapper, S. Živný, The complexity of finite-valued CSPs, *JACM* 63 (4) (2016) 37:1–37:33.
- [15] D. Zhuk, The proof of CSP dichotomy conjecture, in: *FOCS'17, 2017*, pp. 331–342.
- [16] J. Brown-Cohen, P. Raghavendra, Combinatorial optimization algorithms via polymorphisms, Tech. rep., ArXiv:1501.01598 (2015).
- [17] P. Raghavendra, Optimal algorithms and inapproximability results for every CSP?, in: *STOC'08, 2008*, pp. 245–254.
- [18] D. Cohen, M. Cooper, P. Creed, P. Jeavons, S. Živný, An algebraic theory of complexity for discrete optimisation, *SIAM Journal on Computing* 42 (5) (2013) 1915–1939.
- [19] D. Cohen, M. Cooper, P. Jeavons, A. Krokhin, The complexity of soft constraint satisfaction, *Artificial Intelligence* 170 (11) (2006) 983–1016.
- [20] A. Ene, J. Vondrák, Y. Wu, Local distribution and the symmetry gap: Approximability of multiway partitioning problems, in: *SODA, 2013*, pp. 306–325, update at arXiv1503.03905.
- [21] S. Khot, On the unique games conjecture, in: *CCC'10, 2010*, pp. 99–121.
- [22] K. Makarychev, Y. Makarychev, Approximation algorithms for CSPs, in: *The Constraint Satisfaction Problem: Complexity and Approximability*, Vol. 7 of Dagstuhl Follow-Ups, 2017, pp. 287–325.
- [23] S. Arora, B. Barak, *Computational Complexity: A Modern Approach*, Cambridge University Press, 2009.
- [24] I. Dinur, The PCP theorem by gap amplification, *J. ACM* 54 (3).
- [25] J. Håstad, Some optimal inapproximability results, *J. ACM* 48 (2001) 798–859.

- [26] P. Austrin, Conditional inapproximability and limited independence, Ph.D. thesis, KTH (2008).
- [27] J. Håstad, Every 2-CSP allows nontrivial approximation, *Computational Complexity* 17 (4) (2008) 549–566.
- [28] S. Khot, M. Tulsiani, P. Worah, A characterization of strong approximation resistance, in: *Symposium on Theory of Computing, STOC 2014*, 2014, pp. 634–643.
- [29] A. Agarwal, M. Charikar, K. Makarychev, Y. Makarychev, $O(\sqrt{\log n})$ approximation algorithms for MinUnCut, Min 2CNF Deletion, and Directed Cut problems, in: *STOC’05*, 2005, pp. 573–581.
- [30] V. Guruswami, E. Lee, Complexity of approximating CSP with balance/hard constraints, *Theory Comput. Syst.* 59 (1) (2016) 76–98.
- [31] V. Guruswami, Y. Zhou, Tight bounds on the approximability of Almost-satisfiable Horn SAT and Exact Hitting Set, *Theory of Computing* 8 (2012) Article 11.
- [32] S. Khot, G. Kindler, E. Mossel, R. O’Donnell, Optimal inapproximability results for Max-Cut and other 2-variable CSPs?, *SIAM Journal on Computing* 37 (1) (2007) 319–357.
- [33] V. Dalmau, A. Krokhin, Robust satisfiability for CSPs: Hardness and algorithmic results, *ACM Transactions on Computation Theory* 5 (4) (2013) Article 15.
- [34] A. Kumar, R. Manokaran, M. Tulsiani, N. Vishnoi, On LP-based approximability for strict CSPs, in: *SODA*, 2011, pp. 1560–1573.
- [35] R. Manokaran, J. Naor, P. Raghavendra, R. Schwartz, SDP gaps and UGC hardness for Multiway Cut, 0-Extension, and Metric Labeling, in: *STOC*, 2008, pp. 11–20.
- [36] P. Raghavendra, D. Steurer, How to round any CSP, in: *FOCS’09*, 2009, pp. 586–594.
- [37] B. Gärtner, J. Matoušek, *Approximation algorithms and semidefinite programming*, Springer, 2012.
- [38] V. Dalmau, M. Kozik, A. Krokhin, K. Makarychev, Y. Makarychev, J. Opršal, Robust algorithms with polynomial loss for near-unanimity CSPs, in: *SODA’17*, 2017, pp. 340–357.
- [39] G. Kun, R. O’Donnell, T. Suguru, Y. Yoshida, Y. Zhou, Linear programming, width-1 CSPs, and robust satisfaction, in: *ITCS’12*, 2012, pp. 484–495.
- [40] L. Barto, A. Krokhin, R. Willard, Polymorphisms, and how to use them, in: *The Constraint Satisfaction Problem: Complexity and Approximability*, 2017, pp. 1–44.

- [41] A. Krokhin, S. Živný, The complexity of valued CSPs, in: *The Constraint Satisfaction Problem: Complexity and Approximability*, Vol. 7 of Dagstuhl Follow-Ups, 2017, pp. 233–266.
- [42] V. Kolmogorov, J. Thapper, S. Živný, The power of linear programming for general-valued CSPs, *SIAM Journal on Computing* 44 (1) (2015) 1–36.
- [43] L. Barto, M. Kozik, R. Willard, Near unanimity constraints have bounded path-width duality, in: *LICS*, 2012, pp. 125–134.
- [44] G. Kun, M. Szegedy, A new line of attack on the dichotomy conjecture, *Eur. J. Comb.* 52 (2016) 338–367.
- [45] A. Schrijver, *Theory of Linear and Integer Programming*, Wiley, 1986.
- [46] C. Carvalho, V. Dalmau, A. Krokhin, Two new homomorphism dualities and lattice operations, *Journal of Logic and Computation* 21 (6) (2011) 1065–1092.
- [47] K. Baker, A. Pixley, Polynomial interpolation and the chinese remainder theorem, *Mathematische Zeitschrift* 143 (1975) 165–174.
- [48] I. Dinur, V. Guruswami, S. Khot, O. Regev, A new multilayered PCP and the hardness of hypergraph vertex cover, *SIAM J. Comput.* 34 (5) (2005) 1129–1146.
- [49] M. Kozik, A. Krokhin, M. Valeriote, R. Willard, Characterizations of several Maltsev conditions, *Algebra Universalis* 73 (3) (2015) 205–224.
- [50] C. Carvalho, A. Krokhin, On algebras with many symmetric operations, *IJAC* 26 (5) (2016) 1019–1032.
- [51] H. Chen, B. Larose, Asking the metaquestions in constraint tractability, *TOCT* 9 (3) (2017) 11:1–11:27.