

The power of Sherali-Adams relaxations for general-valued CSPs*

Johan Thapper
Université Paris-Est, Marne-la-Vallée, France
thapper@u-pem.fr

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

Abstract

We give a precise algebraic characterisation of the power of Sherali-Adams relaxations for solvability of valued constraint satisfaction problems to optimality. The condition is that of bounded width which has already been shown to capture the power of local consistency methods for decision CSPs and the power of semidefinite programming for robust approximation of CSPs.

Our characterisation has several algorithmic and complexity consequences. On the algorithmic side, we show that several novel and many known valued constraint languages are tractable via the third level of the Sherali-Adams relaxation. For the known languages, this is a significantly simpler algorithm than the previously obtained ones. On the complexity side, we obtain a dichotomy theorem for valued constraint languages that can express an injective unary function. This implies a simple proof of the dichotomy theorem for conservative valued constraint languages established by Kolmogorov and Živný [JACM'13], and also a dichotomy theorem for the exact solvability of Minimum-Solution problems. These are generalisations of Minimum-Ones problems to *arbitrary* finite domains. Our result improves on several previous classifications by Khanna et al. [SICOMP'00], Jonsson et al. [SICOMP'08], and Uppman [ICALP'13].

1 Introduction

Convex relaxations are one of the most powerful techniques for designing polynomial-time exact and approximation algorithms [3, 15]. The idea is to formulate the problem at hand as an integer program and relax it to a convex program which can be solved in polynomial time, such as a linear program (LP) or a semidefinite program (SDP). A solution to the problem is then obtained by designing a (possibly randomised) polynomial-time algorithm that converts the solution to such a relaxation into an integer solution to the original problem.

Convex relaxations can be strengthened by including additional constraints which are satisfied by an integer solution. This process of generating stronger relaxations by adding larger (but still local) constraints is captured by various hierarchies of convex relaxations, including the hierarchy of linear programming relaxations proposed by Sherali and Adams [58], that by Lovász and Schrijver [52], and their semidefinite programming versions, including the hierarchy of Lasserre [50] (see also [51] for a nice comparison of these hierarchies). For an integer program with n variables taking values in $\{0, 1\}$, the convex program obtained by n levels of any of the above-mentioned hierarchies has integrality gap 1, that is, it gives an exact solution (but the program may take exponential time to solve). Since the size of a program obtained by k levels of these hierarchies is $n^{O(k)}$, for a constant k , the program can be solved in polynomial time.

In this paper we study *constant level* Sherali-Adams relaxations for *exact* solvability of discrete optimisation problems. We do this within the framework of constraint satisfaction

*An extended abstract of part of this work appeared in the *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP'15)* [65].

[†]Stanislav Živný was supported by a Royal Society University Research Fellowship. Part of this work was done when the second author was visiting the Simons Institute for the Theory of Computing at UC Berkeley.

problems, which captures a large family of both theoretical and practical problems. An instance of the *valued constraint satisfaction problem* (VCSP) is given by a collection of variables that is assigned labels from a given finite domain with the goal to *minimise* an objective function given by a sum of weighted relations (cost functions), each depending on some subset of the variables [19]. The weighted relations can take on finite rational values and positive infinity.

By varying the codomain of the weighted relations, we get a variety of interesting problems. When the codomain is $\{0, \infty\}$, we get the class of decision problems known as *constraint satisfaction problems* [28] with the goal to determine whether or not there is a labelling for all variables that evaluates the objective function to zero. When the codomain is $\{0, 1\}$, we get the class of optimisation problems known as *minimum constraint satisfaction problems* [21, 22, 40]. When the codomain is \mathbb{Q} , we get the class of optimisation problems known as *finite-valued* (or generalised [55]) constraint satisfaction problems [63]. The special case of having a domain of size two has been studied extensively under the name of pseudo-Boolean optimisation [8, 20]. Finally, by allowing a codomain to be both \mathbb{Q} and positive infinity, we get the large class of problems known as *valued constraint satisfaction problems* [19, 36]. Intuitively, the infinite value deems certain labellings forbidden and thus all constraints are required to be satisfied, whereas the rational values model the optimisation aspect of the problem.

We remark that this framework is more general than that of *mixed* CSPs with hard and soft constraints used in the approximation community [46], where each constraint is either hard or soft; hard constraints correspond to $\{0, \infty\}$ -valued weighted relations in our framework, and a soft constraint corresponds to a $\{0, w\}$ -valued weighted relation, where w is the weight of the constraint. Thus, all constraints in mixed CSPs are 2-valued.

Valued CSPs are sometimes also called *general-valued* CSPs to emphasise the fact that (decision) CSPs are a special case of valued CSPs.

For constraint satisfaction problems, an important algorithmic technique is *local consistency methods*, i.e. considering a bounded number of variables at a time and propagating infeasible partial assignments. Problems for which such techniques suffice to decide satisfiability are said to have *bounded width*. In an important series of papers, [6, 10, 49, 53], the property of having bounded width has been shown to be equivalent to a universal-algebraic condition, now known as the “bounded width condition”. There is a clear relation between the local propagation in consistency methods for decision CSPs and the consistent marginals-condition of Sherali-Adams relaxations. In this paper, we demonstrate the applicability of powerful universal-algebraic techniques, developed for decisions CSPs, in the study of linear programming hierarchies for valued CSPs.

Contributions

A set Γ of weighted relations on some fixed finite domain is called a *valued constraint language*. We denote by $\text{VCSP}(\Gamma)$ the class of VCSP instances with all weighted relations from Γ .

In our first result, we give an algebraic [11, 16] characterisation of the power of Sherali-Adams relaxations for VCSPs. Theorem 3, presented in Section 3, shows that for a valued constraint language Γ of finite size the following three statements are equivalent: (i) Γ is tractable via a constant level Sherali-Adams relaxation; (ii) Γ is tractable via the third level Sherali-Adams relaxation; (iii) the support clone of Γ contains (not necessarily idempotent) m -ary weak near-unanimity operations for every $m \geq 3$.¹ The condition (iii) is precisely that of “bounded width” for constraint languages with codomain $\{0, \infty\}$ (such languages are known as crisp) [6, 10, 49, 53]. Note that the implication “(ii) \implies (i)” is trivial.

The implication “(iii) \implies (ii)”, proved in Section 4, is shown via linear programming duality and fundamentally relies on [6] and [4]. This result simplifies and generalises several previously obtained tractability results for valued constraint languages, as discussed in Section 3.3. For

¹The precise definition of weak near-unanimity operations can be found in Section 2.

example, valued constraint languages with a tournament pair multimorphism were previously known to be tractable using ingenious application of various consistency techniques, advanced analysis of constraint networks using modular decompositions, and submodular function minimisation [18]. Here, we show that an even less restrictive condition (having a binary conservative commutative operation in some fractional polymorphism) ensures that the third level of the Sherali-Adams relaxation solves all instances to optimum.

The implication “(i) \implies (iii)”, proved in Section 5, is shown by proving that, given a language Γ that violates (iii), Γ can simulate linear equations in some Abelian group. This result is known for $\{0, \infty\}$ -valued constraint languages [6]. It suffices to show that linear equations can fool constant level Sherali-Adams relaxations, which is proved in Section 7, and that the “simulation” preserves bounded level of Sherali-Adams relaxations for valued constraint languages, which is proved in Section 6. Previously, it was only known that this “simulation” preserves polynomial-time reducibility. One immediate corollary of our result is a classification of conservative valued constraint languages [43] *without* relying on [59]. In fact we give an alternative and still simple proof of the complexity classification of conservative valued constraint languages [43], which implies that tractable conservative valued constraint languages are captured by admitting a majority operation in the support clone, which was not previously known.

Overall, we give a precise characterisation of the power of Sherali-Adams relaxations for exact solvability of VCSPs. This rather surprising result demonstrates how robust the concept of bounded width is, capturing not only the power of local consistency methods for decision CSPs [6, 10] and the class of decision CSPs that can be robustly approximated [5], but also the power of Sherali-Adams relaxations for exact solvability of VCSPs.

Minimum-Solution [38] problems are special types of VCSPs that involve $\{0, \infty\}$ -valued weighted relations together with a single unary \mathbb{Q} -valued weighted relation that is required to be injective. (The natural encoding of Vertex Cover as a VCSP instance is of this kind.) Minimum-Solution problems include integer programming over bounded domains and can be viewed as a generalisation of Min-Ones problems [22, 40] to larger domains. Compare this to the result [17, 54] that any VCSP instance is equivalent to a VCSP instance with only binary relations and unary (not necessarily injective) finite-valued weighted relations. Hence, unless we settle the CSP dichotomy conjecture [28], some additional requirement on the unary weighted relations (such as injectivity) is necessary.

As a corollary of our characterisation, we give, in Section 3.4, a complete complexity classification of exact solvability of Minimum-Solution problems over *arbitrary finite* domains, thus improving on previous partial classifications for domains of size two [40] and three [67], homogeneous and maximal (under a certain algebraic conjecture) languages [37] and on graphs with few vertices [39]. Theorem 6 shows that the Minimum-Solution problem is NP-hard unless it satisfies the bounded width condition. Previous partial results included ad-hoc algorithms for various special cases. Our result shows that one algorithm, the third level of the Sherali-Adams relaxation, solves all tractable cases and is thus universal. As a matter of fact, we actually prove a complexity classification for a larger class of problems that includes Minimum-Solutions problems as a special case, as described in detail in Section 3.4.

Related work

The first level of the Sherali-Adams hierarchy is known as the *basic linear programming* (BLP) relaxation [14]. In [62], the authors gave a precise algebraic characterisation of Γ for which any instance of VCSP(Γ) is solved to optimality by BLP, see also [42]. The characterisation proved important not only in the study of VCSPs [34] and other classes of problems [32], but also in the design of fixed-parameter algorithms [69]. In [63], it was then shown that for finite-valued CSPs, the BLP solves *all* tractable cases; i.e. if BLP fails to solve any instance of some finite-valued constraint language then this language is NP-hard. The BLP has been considered in the context of CSPs for robust approximability [25, 47] and constant-factor approximation [24,

27]. Higher levels of Sherali-Adams hierarchy have been considered for (in)approximability of CSPs [12, 26, 70]. Semidefinite programming relaxations have also been considered in the context of CSPs for approximability [55] and robust approximability [5]. Concrete lower bounds on Sherali-Adams and other relaxations include [1, 13, 30, 57]. Whilst the complexity of valued constraint languages is open, it has been shown that a dichotomy for constraint languages, conjectured in [28], implies a dichotomy for valued constraint languages [41]. Our results give a complete complexity classification for a large class of VCSPs without any dependence on the dichotomy conjecture [28]. Since the announcement of our results [65], the tractability results obtained in this paper were shown using different methods (preprocessing combined with an LP relaxation) [41].

One ingredient of our proof is the fact that constant level Sherali-Adams relaxations cannot solve exactly instances involving equations over a non-trivial Abelian group. This is known to follow, via [66], from a stronger result of Grigoriev [31], later rediscovered by Schoenebeck [56], that even $\Omega(n)$ levels of Lasserre SDP relaxations cannot approximate such instances. However, a formal proof would require the definition of SDP relaxations that are not in the scope of this article. Rather, we provide here a direct, elementary proof of this fact and observe that our proof actually gives a gap instance for Sherali-Adams relaxations of level $\Theta(\sqrt{n})$. This also has the advantage of our proof being self-contained.

2 Preliminaries

2.1 Valued CSPs

Let $\overline{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$ denote the set of rational numbers extended with positive infinity. Throughout the paper, let D be a fixed finite set of size at least two, also called a *domain*; we call the elements of D *labels*.

Definition 1. An r -ary weighted relation over D is a mapping $\phi : D^r \rightarrow \overline{\mathbb{Q}}$. We write $\text{ar}(\phi) = r$ for the arity of ϕ .

A weighted relation $\phi : D^r \rightarrow \overline{\mathbb{Q}}$ is called *finite-valued* if $\phi(\mathbf{x}) < \infty$ for all $\mathbf{x} \in D^r$. A weighted relation $\phi : D^r \rightarrow \{0, \infty\}$ can be seen as the (ordinary) relation $\{\mathbf{x} \in D^r \mid \phi(\mathbf{x}) = 0\}$. We will use both viewpoints interchangeably.

For any r -ary weighted relation ϕ , we denote by $\text{Feas}(\phi) = \{\mathbf{x} \in D^r \mid \phi(\mathbf{x}) < \infty\}$ the underlying r -ary *feasibility relation*, and by $\text{Opt}(\phi) = \{\mathbf{x} \in \text{Feas}(\phi) \mid \forall \mathbf{y} \in D^r : \phi(\mathbf{x}) \leq \phi(\mathbf{y})\}$ the r -ary *optimality relation*, which contains the tuples on which ϕ is minimised.

Definition 2. Let $V = \{x_1, \dots, x_n\}$ be a set of variables. A valued constraint over V is an expression of the form $\phi(\mathbf{x})$ where ϕ is a weighted relation and $\mathbf{x} \in V^{\text{ar}(\phi)}$. The tuple \mathbf{x} is called the *scope* of the constraint.

We will use the notational convention to denote by X_i the *set* of variables occurring in the scope \mathbf{x}_i .

Definition 3. An instance I of the valued constraint satisfaction problem (VCSP) is specified by a finite set $V = \{x_1, \dots, x_n\}$ of variables, a finite set D of labels, and an objective function ϕ_I expressed as follows:

$$\phi_I(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i), \quad (1)$$

where each $\phi_i(\mathbf{x}_i)$, $1 \leq i \leq q$, is a valued constraint. Each constraint may appear multiple times in I . An assignment to I is a map $\sigma : V \rightarrow D$. The goal is to find an assignment that minimises the objective function.

For a VCSP instance I , we write $\text{Val}_{\text{VCSP}}(I, \sigma)$ for $\phi_I(\sigma(x_1), \dots, \sigma(x_n))$, and $\text{Opt}_{\text{VCSP}}(I)$ for the minimum of $\text{Val}_{\text{VCSP}}(I, \sigma)$ over all assignments.

An assignment σ with $\text{Val}_{\text{VCSP}}(I, \sigma) < \infty$ is called *satisfying*. A VCSP instance I is called *satisfiable* if there is a satisfying assignment to I . CSPs are a special case of VCSPs with (unweighted) relations with the goal to determine the existence of a satisfying assignment.

A *valued constraint language*, or just a constraint language, over D is a set of weighted relations over D . We denote by $\text{VCSP}(\Gamma)$ the class of all VCSP instances in which the weighted relations are all contained in Γ . A constraint language Δ is called *crisp* if Δ contains only (unweighted) relations. For a crisp language Δ we denote by $\text{CSP}(\Delta)$ the class $\text{VCSP}(\Delta)$ to emphasise the fact that there is no optimisation involved.

A valued constraint language Γ is called *tractable* if $\text{VCSP}(\Gamma')$ can be solved (to optimality) in polynomial time for every finite subset $\Gamma' \subseteq \Gamma$, and Γ is called *NP-hard* if $\text{VCSP}(\Gamma')$ is NP-hard for some finite $\Gamma' \subseteq \Gamma$.

2.2 Fractional Polymorphisms

Given an r -tuple $\mathbf{x} \in D^r$, we denote its i th entry by $\mathbf{x}[i]$ for $1 \leq i \leq r$. A mapping $f: D^m \rightarrow D$ is called an m -ary operation on D ; f is *idempotent* if $f(x, \dots, x) = x$. We apply an m -ary operation f to m r -tuples $\mathbf{x}_1, \dots, \mathbf{x}_m \in D^r$ coordinatewise, that is,

$$f(\mathbf{x}_1, \dots, \mathbf{x}_m) = (f(\mathbf{x}_1[1], \dots, \mathbf{x}_m[1]), \dots, f(\mathbf{x}_1[r], \dots, \mathbf{x}_m[r])). \quad (2)$$

Definition 4. Let ϕ be a weighted relation on D and let f be an m -ary operation on D . We call f a polymorphism of ϕ if, for any $\mathbf{x}_1, \dots, \mathbf{x}_m \in \text{Feas}(\phi)$, we have that $f(\mathbf{x}_1, \dots, \mathbf{x}_m) \in \text{Feas}(\phi)$.

For a valued constraint language Γ , we denote by $\text{Pol}(\Gamma)$ the set of all operations which are polymorphisms of all $\phi \in \Gamma$. We write $\text{Pol}(\phi)$ for $\text{Pol}(\{\phi\})$.

A probability distribution ω over the set of m -ary operations on D is called an m -ary *fractional operation*. We define $\text{supp}(\omega)$ to be the set of operations assigned positive probability by ω .

The following two notions are known to capture the complexity of valued constraint languages [16, 45] and will also be important in this paper.

Definition 5. Let ϕ be a weighted relation on D and let ω be an m -ary fractional operation on D . We call ω a fractional polymorphism of ϕ if $\text{supp}(\omega) \subseteq \text{Pol}(\phi)$ and for any $\mathbf{x}_1, \dots, \mathbf{x}_m \in \text{Feas}(\phi)$, we have

$$\mathbb{E}_{f \sim \omega} [\phi(f(\mathbf{x}_1, \dots, \mathbf{x}_m))] \leq \text{avg}\{\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_m)\}. \quad (3)$$

For a valued constraint language Γ , we denote by $\text{fPol}(\Gamma)$ the set of all fractional operations which are fractional polymorphisms of all weighted relations $\phi \in \Gamma$. We write $\text{fPol}(\phi)$ for $\text{fPol}(\{\phi\})$.

Example 1. Consider the domain $D = \{0, 1\}$ and the two binary operations \min and \max on D that return the smaller and the larger its two arguments, respectively. A valued constraint language on D is called *submodular* if it has the fractional polymorphism ω defined by $\omega(\min) = \omega(\max) = \frac{1}{2}$.

Definition 6. Let Γ be a valued constraint language on D . We define

$$\text{supp}(\Gamma) = \bigcup_{\omega \in \text{fPol}(\Gamma)} \text{supp}(\omega). \quad (4)$$

An m -ary *projection* is an operation of the form $\pi_i^{(m)}(x_1, \dots, x_m) = x_i$ for some $1 \leq i \leq m$. Projections are polymorphisms of all valued constraint languages.

The *composition* of an m -ary operation $f : D^m \rightarrow D$ with m n -ary operations $g_i : D^n \rightarrow D$ for $1 \leq i \leq m$ is the n -ary function $f[g_1, \dots, g_m] : D^n \rightarrow D$ defined by

$$f[g_1, \dots, g_m](x_1, \dots, x_n) = f(g_1(x_1, \dots, x_n), \dots, g_m(x_1, \dots, x_n)). \quad (5)$$

A *clone* of operations is a set of operations on D that contains all projections and is closed under composition. $\text{Pol}(\Gamma)$ is a clone for any valued constraint language Γ .

Lemma 1. *For any valued constraint language Γ , $\text{supp}(\Gamma)$ is a clone.*

We note that Lemma 1 has also been observed in [29] and in [45]. For completeness, we give a proof here. (Our proof is slightly different from the proofs in [29, 45] as we have defined fractional polymorphisms as probability distributions.)

Proof. Observe that $\text{supp}(\Gamma)$ contains all projections as $\tau_m \in \text{fPol}(\Gamma)$ for every $m \geq 1$, where τ_m is the fractional operation defined by $\tau_m(\pi_i^{(m)}) = \frac{1}{m}$ for every $1 \leq i \leq m$. Thus we only need to show that $\text{supp}(\Gamma)$ is closed under composition.

Let $f \in \text{supp}(\Gamma)$ be an m -ary operation with $\omega(f) > 0$ for some $\omega \in \text{fPol}(\Gamma)$. Moreover, let $g_i \in \text{supp}(\Gamma)$ be n -ary operations with $\mu_i(g_i) > 0$ for some $\mu_i \in \text{fPol}(\Gamma)$, where $1 \leq i \leq m$. We define an n -ary fractional operation

$$\omega'(p) = \Pr_{\substack{t \sim \omega \\ h_i \sim \mu_i}} [t[h_1, \dots, h_m] = p]. \quad (6)$$

Since $\omega(f) > 0$ and $\mu_i(g_i) > 0$ for all $1 \leq i \leq m$, we have $\omega'(f[g_1, \dots, g_m]) > 0$. A straightforward verification shows that $\omega' \in \text{fPol}(\Gamma)$. Consequently, $f[g_1, \dots, g_m] \in \text{supp}(\Gamma)$. \square

The following lemma is a generalisation of [61, Lemma 2.9] from arity one to arbitrary arity and from finite-valued to valued constraint languages, but the proof is analogous. A special case has also been observed, in the context of Min-Sol problems [67], by Hannes Uppman [68].

Lemma 2. *Let Γ be a valued constraint language of finite size on a domain D and let $f \in \text{Pol}(\Gamma)$. Then, $f \in \text{supp}(\Gamma)$ if, and only if, $f \in \text{Pol}(\text{Opt}(I))$ for all instances I of $\text{VCSP}(\Gamma)$.*

Proof. The operation f is in $\text{supp}(\Gamma)$ if, and only if, there exists a fractional polymorphism ω with $f \in \text{supp}(\omega)$. This is the case if, and only if, the following system of linear inequalities in the variables $\omega(g)$ for $g \in \text{Pol}(\Gamma)$ is satisfiable:

$$\begin{aligned} \sum_{g \in \text{Pol}(\Gamma)} \omega(g) \phi(g(\mathbf{x}_1, \dots, \mathbf{x}_m)) &\leq \text{avg}\{\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_m)\}, \quad \forall \phi \in \Gamma, \mathbf{x}_i \in \text{Feas}(\phi), \\ \sum_{g \in \text{Pol}(\Gamma)} \omega(g) &= 1, \\ \omega(f) &> 0, \\ \omega(g) &\geq 0, \quad \forall g \in \text{Pol}(\Gamma). \end{aligned} \quad (7)$$

By Farkas' lemma (e.g. [61, Lemma 2.8]), the system (7) is unsatisfiable if, and only if, the following system in variables $z(\phi, \mathbf{x}_1, \dots, \mathbf{x}_m)$, for $\phi \in \Gamma, \mathbf{x}_i \in \text{Feas}(\phi)$, is satisfiable:

$$\begin{aligned} \sum_{\substack{\phi \in \Gamma \\ \mathbf{x}_i \in \text{Feas}(\phi)}} z(\phi, \mathbf{x}_1, \dots, \mathbf{x}_m) (\text{avg}\{\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_m)\} - \phi(g(\mathbf{x}_1, \dots, \mathbf{x}_m))) &\leq 0, \quad \forall g \in \text{Pol}(\Gamma), \\ \sum_{\substack{\phi \in \Gamma \\ \mathbf{x}_i \in \text{Feas}(\phi)}} z(\phi, \mathbf{x}_1, \dots, \mathbf{x}_m) (\text{avg}\{\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_m)\} - \phi(f(\mathbf{x}_1, \dots, \mathbf{x}_m))) &< 0, \end{aligned}$$

$$z(\phi, \mathbf{x}_1, \dots, \mathbf{x}_m) \geq 0, \quad \forall \phi \in \Gamma, \mathbf{x}_i \in \text{Feas}(\phi). \quad (8)$$

First, assume that $f \notin \text{supp}(\Gamma)$ so that (8) has a feasible solution z . Note that by scaling we may assume that z is integral. Let $V^{(m)} = \{v_{\mathbf{x}} \mid \mathbf{x} \in D^m\}$ and let $\mathbf{v} = (v_1, \dots, v_n)$ be an enumeration of $V^{(m)}$. Define $\iota: V^{(m)} \rightarrow D^m$ by $\iota(v_{\mathbf{x}}) = \mathbf{x}$ and let I be the instance of VCSP(Γ) with variables $V^{(m)}$ and objective function:

$$\phi_I(\mathbf{v}) = \sum_{\phi \in \Gamma, \mathbf{x}_i \in \text{Feas}(\phi)} z(\phi, \mathbf{x}_1, \dots, \mathbf{x}_m) \phi(\iota^{-1}(\mathbf{x}_1[1], \dots, \mathbf{x}_m[1]), \dots, \iota^{-1}(\mathbf{x}_1[\text{ar}(\phi)], \dots, \mathbf{x}_m[\text{ar}(\phi)])),$$

where the multiplication by z is simulated by taking the corresponding constraint with multiplicity z . According to (8), every projection $\pi_i^{(m)}$ induces an optimal assignment $\pi_i^{(m)} \circ \iota$ to I . Interpreted as D^m -tuples, we therefore have $\pi_i^{(m)} \in \text{Opt}(\phi_I)$ for $1 \leq i \leq m$. On the other hand, (8) states that $f \circ \iota$ is not an optimal assignment, so $f(\pi_1^{(m)}, \dots, \pi_m^{(m)}) \notin \text{Opt}(\phi_I)$. In other words, $f \notin \text{Pol}(\text{Opt}(\phi_I))$, and ϕ_I is expressible in Γ .

For the opposite direction, assume that $f \in \text{supp}(\Gamma)$, so that (8) is unsatisfiable. Let ϕ be an arbitrary weighted relation expressible over Γ , and let I be an instance of VCSP(Γ) with objective function $\phi_I(x_1, \dots, x_m, y_1, \dots, y_p) = \sum_i \phi_i(\mathbf{x}_i)$ such that $\phi(x_1, \dots, x_m) = \min_{y_1, \dots, y_p} \phi_I$. Let $\sigma_1, \dots, \sigma_m \in \text{Opt}(\phi)$. Note that every assignment in $\text{Opt}(\phi)$ can be extended to an assignment in $\text{Opt}(\phi_I)$, and let $\tau_1, \dots, \tau_m \in \text{Opt}(\phi_I)$ be such that $\tau_i|_{\{v_1, \dots, v_m\}} = \sigma_i$. Construct an instance Z of VCSP(Γ) with variables $V^{(m)}$ by replacing each valued constraint $\phi_i(\mathbf{x}_i)$ in I by:

$$\phi_i(\iota^{-1}(\tau_1(\mathbf{x}_i[1]), \dots, \tau_1(\mathbf{x}_i[1])), \dots, \iota^{-1}(\tau_m(\mathbf{x}_i[\text{ar}(\phi_i)], \dots, \tau_m(\mathbf{x}_i[\text{ar}(\phi_i)]))).$$

Now, if $f \circ \iota$ were not an optimal solution to Z , then Z would be a solution to (8), a contradiction. Hence $f(\tau_1, \dots, \tau_m) \in \text{Opt}(\phi_I)$, so $f(\sigma_1, \dots, \sigma_m) \in \text{Opt}(\phi)$. Since the σ_i were chosen arbitrarily, $f \in \text{Pol}(\text{Opt}(\phi))$. This establishes the lemma. \square

2.3 Cores and Constants

Definition 7. Let Γ be a valued constraint language with domain D and let $S \subseteq D$. The sub-language $\Gamma[S]$ of Γ induced by S is the valued constraint language defined on domain S and containing the restriction of every weighted relation $\phi \in \Gamma$ onto S .

Definition 8. A valued constraint language Γ is a core if all unary operations in $\text{supp}(\Gamma)$ are bijections. A valued constraint language Γ' is a core of Γ if Γ' is a core and $\Gamma' = \Gamma[f(D)]$ for some $f \in \text{supp}(\Gamma)$.

The following lemma implies that when studying the computational complexity of a valued constraint language Γ , we may assume that Γ is a core.

Lemma 3. Let Γ be a valued constraint language and Γ' a core of Γ . Then, for all instances I of VCSP(Γ) and I' of VCSP(Γ'), where I' is obtained from I by substituting each weighted relation in Γ for its restriction in Γ' , the optimum of I and I' coincide.

A special case of Lemma 3 for finite-valued constraint languages was proved by the authors in [61]. Lemma 3, proved below using Lemma 2, has also been observed in [45] and in [64], where it was proved in a different way (and without the use of Lemma 2).

Proof. By definition, $\Gamma' = \Gamma[f(D)]$, where D is the domain of Γ and $f \in \text{supp}(\omega)$ for some unary fractional polymorphism ω . Assume that I is satisfiable, and let σ be an optimal assignment to I . Now $f \circ \sigma$ is a satisfying assignment to I' , and by Lemma 2, $f \circ \sigma$ is also an optimal assignment to I . Conversely, any satisfying assignment to I' is a satisfying assignment to I of the same value. \square

Let $\mathcal{C}_D = \{\{(a)\} \mid a \in D\}$ be the set of constant unary relations on the set D . It is known (cf. [45, Proposition 20]), that for a valued constraint language Γ on D and a core Γ' of Γ on $D' \subseteq D$, the problem $\text{VCSP}(\Gamma' \cup \mathcal{C}_{D'})$ polynomial-time reduces to $\text{VCSP}(\Gamma)$. In Theorem 16(5) in Section 5, we present a stronger form of this reduction.

Let Γ be a valued constraint language on D with $\mathcal{C}_D \subseteq \Gamma$. It is well known and easy to show that any $f \in \text{Pol}(\Gamma)$ is idempotent [11].

2.4 Relational Width

We define *relational width* which is the basis for our notion of valued relational width.

Let J be an instance of the CSP with $\phi_J(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i)$, $X_i \subseteq V = \{x_1, \dots, x_n\}$ and $\phi_i: D^{\text{ar}(\phi_i)} \rightarrow \{0, \infty\}$.

For a tuple $\mathbf{t} \in D^X$, we denote by $\pi_{X'}(\mathbf{t})$ its projection onto $X' \subseteq X$. For a constraint $\phi_i(\mathbf{x}_i)$, we define $\pi_{X'}(\phi_i) = \{\pi_{X'}(\mathbf{t}) \mid \mathbf{t} \in \text{Feas}(\phi_i)\}$ where $X' \subseteq X_i$.

Let $1 \leq k \leq \ell$ be integers. The following definition is equivalent² to the definition of (k, ℓ) -minimality [9] for CSP instances given in [4].

Definition 9. *Let J be an instance of the CSP with $\phi_J(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i)$, $X_i \subseteq V = \{x_1, \dots, x_n\}$ and $\phi_i: D^{\text{ar}(\phi_i)} \rightarrow \{0, \infty\}$. Then J is said to be (k, ℓ) -minimal if:*

- *For every $X \subseteq V$, $|X| \leq \ell$, there exists $1 \leq i \leq q$ such that $X = X_i$.*
- *For every $i, j \in [q]$ such that $|X_j| \leq k$ and $X_j \subseteq X_i$, $\phi_j = \pi_{X_j}(\phi_i)$.*

There is a straightforward polynomial-time algorithm for finding an equivalent (k, ℓ) -minimal instance [4]. This leads to the notion of *relational width*:

Definition 10. *A constraint language Δ has relational width (k, ℓ) if, for every instance J of $\text{CSP}(\Delta)$, an equivalent (k, ℓ) -minimal instance is non-empty if, and only if, J has a solution.*

An m -ary idempotent operation $f: D^m \rightarrow D$ is called a *weak near-unanimity* (WNU) operation if, for all $x, y \in D$,

$$f(y, x, x, \dots, x) = f(x, y, x, x, \dots, x) = f(x, x, \dots, x, y). \quad (9)$$

Definition 11. *We say that a clone of operations satisfies the bounded width condition (BWC) if it contains a (not necessarily idempotent) m -ary operation satisfying the identities (9), for every $m \geq 3$.*

The following result is known as the “bounded width theorem” as it characterises constraint languages of bounded relational width, that is, constraint languages that are tractable via the (k, ℓ) -minimality algorithm for some $k \leq \ell$.

Theorem 1 ([6, 10, 49]). *Let Δ be a constraint language of finite size containing all constant unary relations. Then, Δ has bounded relational width if, and only if, $\text{Pol}(\Delta)$ satisfies the BWC.*

Moreover, a collapse of relational width is known.

Theorem 2 ([4, 10]). *Let Δ be a constraint language of finite size containing all constant unary relations. If Δ has bounded relational width, then it has relational width $(2, 3)$.*

Remark 1. We remark that most of the papers cited above use a different bounded width condition, namely that of having WNU operations of all but finitely many arities [53, Theorem 1.2]. By [44, Theorem 1.6 (4)], this is equivalent to Definition 11. Also note that our definition of the BWC does not require idempotency of the operations. The reason is that we prove our main result, Theorem 3 below, without the requirement of including the constant unary relations, which is often assumed in the algebraic papers on the CSP.

²The two requirements in [4] are: for every $X \subseteq V$ with $|X| \leq \ell$ we have $X \subseteq X_i$ for some $1 \leq i \leq q$; and for every set $X \subseteq V$ with $|X| \leq k$ and every $1 \leq i, j \leq q$ with $X \subseteq X_i$ and $X \subseteq X_j$ we have $\pi_X(\phi_i) = \pi_X(\phi_j)$.

3 The Power of Sherali-Adams Relaxations

In this section, we state our main result on the power of the Sherali-Adams linear programming relaxation [58] to VCSPs. We also give a number of applications of this result. The Sherali-Adams linear programming relaxation is defined in Section 3.1 and the characterisation of its power is stated in Section 3.2. In Section 3.3, we give a number of algorithmic consequences of our result and Section 3.4 show how it can be used to derive complete complexity classifications for large families of valued constraint languages. In Section 3.5, we compare our result to the characterisation of valued relational width 1 which we obtained in previous work. Finally, in Section 3.6, we address the problem of finding an actual solution and of determining whether or not a valued constraint language has bounded valued relational width.

3.1 Valued Relation Width

Let I be an instance of the VCSP with $\phi_I(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i)$, $X_i \subseteq V = \{x_1, \dots, x_n\}$ and $\phi_i: D^{\text{ar}(\phi_i)} \rightarrow \overline{\mathbb{Q}}$. A *null constraint* is a constraint that has a weighted relation identical to 0. Make sure that for every non-empty $X \subseteq V$ with $|X| \leq \ell$ there is some constraint $\phi_i(\mathbf{x}_i)$ with $X_i = X$, possibly by adding null constraints.

The Sherali-Adams relaxation with parameters (k, ℓ) , henceforth called the $\text{SA}(k, \ell)$ -relaxation of I , is given by the following linear program. The variables are $\lambda_i(\sigma)$ for every $i \in [q]$ and assignment $\sigma: X_i \rightarrow D$. We slightly abuse notation by writing $\sigma \in \text{Feas}(\phi_i)$ for $\sigma: X_i \rightarrow D$ such that $\sigma(\mathbf{x}_i) \in \text{Feas}(\phi_i)$.

$$\begin{aligned} \min \sum_{i=1}^q \sum_{\sigma \in \text{Feas}(\phi_i)} \lambda_i(\sigma) \phi_i(\sigma(\mathbf{x}_i)) \\ \lambda_j(\tau) = \sum_{\substack{\sigma: X_i \rightarrow D \\ \sigma|_{X_j} = \tau}} \lambda_i(\sigma) \quad \forall i, j \in [q] : X_j \subseteq X_i, |X_j| \leq k, \tau: X_j \rightarrow D \end{aligned} \quad (10)$$

$$\sum_{\sigma: X_i \rightarrow D} \lambda_i(\sigma) = 1 \quad \forall i \in [q] \quad (11)$$

$$\lambda_i(\sigma) = 0 \quad \forall i \in [q], \sigma: X_i \rightarrow D, \sigma(\mathbf{x}_i) \notin \text{Feas}(\phi_i) \quad (12)$$

$$\lambda_i(\sigma) \geq 0 \quad \forall i \in [q], \sigma: X_i \rightarrow D \quad (13)$$

The relaxation $\text{SA}(k, k)$ is often referred to as “ k rounds of Sherali-Adams”.

We write $\text{Val}_{\text{LP}}(I, \lambda)$ for the value of the LP-solution λ to the $\text{SA}(k, \ell)$ -relaxation of I , and $\text{Opt}_{\text{LP}}(I)$ for its optimal value.

Definition 12. We say that a valued constraint language Γ has valued relational width (k, ℓ) if, for every instance I of $\text{VCSP}(\Gamma)$, $\text{Opt}_{\text{VCSP}}(I) = \text{Opt}_{\text{LP}}(I)$ (i.e. the optimum of I coincides with the optimum of the $\text{SA}(k, \ell)$ -relaxation of I).

When Γ has valued relational width (k, k) we also say that Γ has valued relational width k . When Γ has valued relational width k for some fixed $k \geq 1$, then we say that Γ has *bounded valued relational width*.

We say that an instance I of $\text{VCSP}(\Gamma)$ is a *gap instance* for $\text{SA}(k, \ell)$, if its $\text{SA}(k, \ell)$ optimum is strictly smaller than its VCSP optimum. Then, Γ having bounded valued relational width is equivalent to saying that there is some constant level of the Sherali-Adams hierarchy for which there are no gap instances in $\text{VCSP}(\Gamma)$.

Definition 13. Let Γ and Δ be two valued constraint languages. We write $\Delta \leq_{\text{SA}} \Gamma$ if there is a polynomial-time reduction from $\text{VCSP}(\Delta)$ to $\text{VCSP}(\Gamma)$ that preserves bounded valued relational width.

It is clear that the \leq_{SA} reductions compose.

3.2 A Characterisation of Bounded Valued Relational Width

The following characterisation of bounded valued relational width is our main result. It precisely captures the power of Sherali-Adams relaxations for exact optimisation of VCSPs.

Theorem 3 (Main). *Let Γ be a valued constraint language of finite size. The following are equivalent:*

- (i) Γ has bounded valued relational width.
- (ii) Γ has valued relational width $(2, 3)$.
- (iii) $\text{supp}(\Gamma)$ satisfies the BWC.

The proof of Theorem 3 is based on the following two theorems which show that the BWC is a sufficient and necessary condition, respectively, for a constraint language to have bounded valued relational width.

Theorem 4. *Let Γ be a valued constraint language of finite size containing all constant unary relations. If $\text{supp}(\Gamma)$ satisfies the BWC, then Γ has valued relational width $(2, 3)$.*

Theorem 5. *Let Γ be a valued constraint language of finite size containing all constant unary relations. If Γ has bounded valued relational width, then $\text{supp}(\Gamma)$ satisfies the BWC.*

We prove Theorems 4 and 5 in Section 4 and 5, respectively. In order to finish the proof of Theorem 3, we must reduce to the case when the language Γ is assumed to contain all constants. This is done by taking a core Γ' of Γ on a domain $D' \subseteq D$ and adding $\mathcal{C}_{D'}$ to Γ' . We need the following two lemmas to ensure that this can be carried out. Lemma 4 is proved in Section 6 (as Lemma 15). Lemma 5 is proved in Section 8.

Lemma 4. *Let Γ be a valued constraint language of finite size on domain D . If Γ' is a core of Γ on domain $D' \subseteq D$, then $\Gamma' \cup \mathcal{C}_{D'} \leq_{\text{SA}} \Gamma$.*

Lemma 5. *Let Γ be a valued constraint language of finite size on domain D and Γ' a core of Γ on domain $D' \subseteq D$. Then, $\text{supp}(\Gamma)$ satisfies the BWC if, and only if, $\text{supp}(\Gamma' \cup \mathcal{C}_{D'})$ satisfies the BWC.*

Proof of Theorem 3. The implication $(ii) \implies (i)$ is trivial. We first prove the implication $(iii) \implies (ii)$. Suppose that $\text{supp}(\Gamma)$ satisfies the BWC. We start by going to a core of Γ and adding constant unary relations with the goal of applying Theorem 4. Let Γ' be a core of Γ on domain $D' \subseteq D$ and let $\Gamma'_c = \Gamma' \cup \mathcal{C}_{D'}$. By Lemma 5, $\text{supp}(\Gamma'_c)$ also satisfies the BWC. By Theorem 4, Γ'_c has valued relational width $(2, 3)$, so clearly Γ' has valued relational width $(2, 3)$ as well. Every feasible solution to the SA(2, 3)-relaxation of an instance I' of VCSP(Γ') is also a feasible solution to the SA(2, 3)-relaxation of the corresponding instance I of VCSP(Γ). The result now follows from Lemma 3 as the optimum of I' and I coincide.

It remains to prove the implication $(i) \implies (iii)$. Suppose that $\text{supp}(\Gamma)$ does not satisfy the BWC. We start by going to a core of Γ and adding constant unary relations with the goal of applying Theorem 5. Let Γ' be a core of Γ on domain $D' \subseteq D$ and let $\Gamma'_c = \Gamma' \cup \mathcal{C}_{D'}$. By Lemma 5, $\text{supp}(\Gamma'_c)$ does not satisfy the BWC. By Theorem 5, Γ'_c does not have bounded valued relational width. Finally, by Lemma 4, Γ does not have bounded valued relational width either. \square

3.3 Algorithmic Consequences

We now give examples of previously studied valued constraint languages and show that, as a corollary of Theorem 3, they all have, as well as their generalisations, valued relational width $(2, 3)$.

Example 2. Let ω be a ternary fractional operation defined by $\omega(f) = \omega(g) = \omega(h) = \frac{1}{3}$ for some (not necessarily distinct) majority operations f , g , and h . Cohen et al. proved the tractability of any language improved by ω by a reduction to CSPs with a majority polymorphism [19].

Example 3. Let ω be a ternary fractional operation defined by $\omega(f) = \frac{2}{3}$ and $\omega(g) = \frac{1}{3}$, where $f : \{0, 1\}^3 \rightarrow \{0, 1\}$ is the Boolean majority operation and $g : \{0, 1\}^3 \rightarrow \{0, 1\}$ is the Boolean minority operation. Cohen et al. proved the tractability of any language improved by ω by a simple propagation algorithm [19].

Example 4. Generalising Example 3 from Boolean to arbitrary domains, let ω be a ternary fractional operation such that $\omega(f) = \frac{1}{3}$, $\omega(g) = \frac{1}{3}$, and $\omega(h) = \frac{1}{3}$ for some (not necessarily distinct) conservative majority operations f and g , and a conservative minority operation h ; such an ω is called an MJN. Kolmogorov and Živný proved the tractability of any language improved by ω by a 3-consistency algorithm and a reduction, via Example 5, to submodular function minimisation [43].

The following corollary of Theorem 3 generalises Examples 2-4.

Corollary 1. *Let Γ be a valued constraint language of finite size such that $\text{supp}(\Gamma)$ contains a majority operation. Then, Γ has valued relational width $(2, 3)$.*

Proof. Let f be a majority operation in $\text{supp}(\Gamma)$. Then, for every $k \geq 3$, f generates a WNU g_k of arity k : $g_k(x_1, \dots, x_k) = f(x_1, x_2, x_3)$. By Lemma 1, $\text{supp}(\Gamma)$ is a clone, so $g_k \in \text{supp}(\Gamma)$ for all $k \geq 3$. Therefore, $\text{supp}(\Gamma)$ satisfies the BWC and the result follows from Theorem 3. \square

Example 5. Let ω be a binary fractional operation defined by $\omega(f) = \omega(g) = \frac{1}{2}$, where f and g are conservative and commutative operations and $f(x, y) \neq g(x, y)$ for every x and y ; such an ω is called a *symmetric tournament pair* (STP). Cohen et al. proved the tractability of any language improved by ω by a 3-consistency algorithm and an ingenious reduction to submodular function minimisation [18]. Such languages were shown to be the only tractable languages among conservative finite-valued constraint languages [43].

The following corollary of Theorem 3 generalises Example 5.

Corollary 2. *Let Γ be a valued constraint language of finite size such that $\text{supp}(\Gamma)$ contains two symmetric tournament operations (that is, binary operations f and g that are both conservative and commutative and $f(x, y) \neq g(x, y)$ for every x and y). Then, Γ has valued relational width $(2, 3)$.*

Proof. It is straightforward to verify that $h(x, y, z) = f(f(g(x, y), g(x, z)), g(y, z))$ is a majority operation, as observed in [18, Corollary 5.8]. The claim then follows from Corollary 1. \square

Example 6. Generalising Example 5, let ω be a binary fractional operation defined by $\omega(f) = \omega(g) = \frac{1}{2}$, where f and g are conservative and commutative operations; such an ω is called a *tournament pair*. Cohen et al. proved the tractability of any language improved by ω by a consistency-reduction relying on Bulatov's result [9], which in turn relies on 3-consistency, to the STP case from Example 5 [18].

The following corollary of Theorem 3 generalises Example 6.

Corollary 3. *Let Γ be a valued constraint language of finite size such that $\text{supp}(\Gamma)$ contains a tournament operation (that is, a binary conservative and commutative operation). Then, Γ has valued relational width $(2, 3)$.*

Proof. Let f be a tournament operation from $\text{supp}(\Gamma)$. We claim that f is a 2-semilattice; that is, f is idempotent, commutative, and satisfies the restricted associativity law $f(x, f(x, y)) = f(f(x, x), y)$. To see that, notice that $f(x, f(x, y)) = x$ if $f(x, y) = x$ and $f(x, f(x, y)) = y$ if $f(x, y) = y$; together, $f(x, f(x, y)) = f(x, y)$. On the other hand, trivially $f(f(x, x), y) = f(x, y)$.

Also note that $f(x, f(y, x)) = f(x, f(x, y)) = f(x, y)$, so f is a ternary WNU. For every $k \geq 3$, f generates a WNU g_k of arity k : $g_k(x_1, \dots, x_k) = f(f(\dots(f(x_1, x_2), x_3), \dots), x_k)$. By Lemma 1, $\text{supp}(\Gamma)$ is a clone, so $g_k \in \text{supp}(\Gamma)$ for all $k \geq 3$. Therefore, $\text{supp}(\Gamma)$ satisfies the BWC so the result follows from Theorem 3. \square

Example 7. In this example we denote by $\{\{\dots\}\}$ a multiset. Let ω be a binary fractional operation on D defined by $\omega(f) = \omega(g) = \frac{1}{2}$ and let μ be a ternary fractional operation on D defined by $\mu(h_1) = \mu(h_2) = \mu(h_3) = \frac{1}{3}$. Moreover, assume that $\{\{f(x, y), g(x, y)\}\} = \{\{x, y\}\}$ for every x and y and $\{\{h_1(x, y, z), h_2(x, y, z), h_3(x, y, z)\}\} = \{\{x, y, z\}\}$ for every x, y , and z . Let Γ be a language on D such that for every two-element subset $\{a, b\} \subseteq D$, either $\omega|_{\{a, b\}}$ is an STP or $\mu|_{\{a, b\}}$ is an MJN. Kolmogorov and Živný proved the tractability of Γ by a 3-consistency algorithm and a reduction, via Example 5, to submodular function minimisation [43]. Such languages were shown to be the only tractable languages among conservative valued constraint languages [43]. We will discuss conservative valued constraint languages in more detail in Section 3.4.

The following corollary of Theorem 3 covers Example 7.

Corollary 4. *Let Γ be a valued constraint language of finite size with fractional polymorphisms ω and μ as described in Example 7. Then, Γ has valued relational width $(2, 3)$.*

Proof. Let P be the set of 2-element subsets of D such that $\omega|_{\{a, b\}}$ is an STP for $\{a, b\} \in P$ and $\mu|_{\{a, b\}}$ is an MJN for $\{a, b\} \notin P$. Let $p(x, y, z) = f(f(g(y, x), g(x, z)), g(y, z))$. Observe that $p|_{\{a, b\}}$ is a majority for $\{a, b\} \in P$, and $p|_{\{a, b\}}$ is either $\pi_1^{(3)}$ or $\pi_2^{(3)}$ for $\{a, b\} \notin P$ (possibly different projections for different 2-element subsets from P). Now let $q(x, y, z) = p(h_1(x, y, z), h_2(x, y, z), h_3(x, y, z))$. For $x, y \in \{a, b\} \in P$, $q(x, x, y) = q(x, y, x) = q(y, x, x) = p(\{\{x, x, y\}\}) = x$. For $x, y \in \{a, b\} \notin P$, $q(x, x, y) = q(x, y, x) = q(y, x, x) = p(x, x, y) = x$ as p is either the first or the second projection. Thus, q is a majority operation. The claim then follows from Corollary 1. \square

3.4 Complexity Consequences

We now give some computational complexity consequences of Theorem 3. First, we obtain a new and simpler proof (in fact two proofs) of the complexity classification of conservative valued constraint languages [43]. Second, we obtain a complexity classification of (generalisation of) Minimum-Solution problems over arbitrary finite domains.

Minimum-Solution (Min-Sol) problems [38], studied under the name of Min-Ones on Boolean domains [22, 40], constitute a large and interesting subclass of VCSPs including, for instance, integer linear programming over bounded domains.

Definition 14. *A valued constraint language Γ on finite domain D is called a Min-Sol language if $\Gamma = \Delta \cup \{\nu\}$, where Δ is a crisp constraint language on D and $\nu : D \rightarrow \mathbb{Q}$ is an injective finite-valued weighted relation.*

In other words, in Min-Sol problems the optimisation part of the objective function is a sum of unary terms involving an injective finite-valued weighted relation.

As our main result in this section, we give a complexity classification of *all* Min-Sol languages on *arbitrary* finite domains, thus improving on previous classifications obtained for Min-Sol languages on domains with two elements [40], three elements [67], and other special cases [37–39].

By Lemma 4, we can, without loss of generality, restrict our attention to languages that include constants.

Theorem 6. *Let D be an arbitrary finite domain and let $\Gamma = \Delta \cup \{\nu\}$ be an arbitrary Min-Sol language of finite size on D with $\mathcal{C}_D \subseteq \Gamma$. Then, either $\text{supp}(\Gamma)$ satisfies the BWC, in which case Γ has valued relational width $(2, 3)$, or $\text{VCSP}(\Gamma)$ is NP-hard.*

In order to prove Theorem 6, we prove a more general result classifying valued constraint languages that can express an injective unary finite-valued weighted relation. Theorem 6 is then a simple corollary of the following result.

Theorem 7. *Let D be an arbitrary finite domain and let Γ be an arbitrary valued constraint language of finite size on D with $\mathcal{C}_D \subseteq \Gamma$. Assume that Γ expresses a unary finite-valued weighted relation ν that is injective on D . Then, either $\text{supp}(\Gamma)$ satisfies the BWC, in which case Γ has valued relational width $(2, 3)$, or $\text{VCSP}(\Gamma)$ is NP-hard.*

A valued constraint language Γ on D is called *conservative* if Γ contains all $\{0, 1\}$ -valued unary weighted relations. Note that any conservative language Γ is a core and by Lemma 4 we can assume that $\mathcal{C}_D \subseteq \Gamma$.

Theorem 7 implies the following dichotomy theorem, first established in [43] with the help of [59].

Theorem 8. *Let D be an arbitrary finite domain and let Γ be an arbitrary conservative valued constraint language on D . Then, either $\text{supp}(\Gamma)$ satisfies the BWC, in which case Γ has valued relational width $(2, 3)$, or $\text{VCSP}(\Gamma)$ is NP-hard.*

We now give a different proof classifying conservative valued constraint languages that relies on [59] but has the advantage of giving a more specific tractability criterion than the BWC that is different from the STP/MJN criterion established in [43] and discussed in Example 7.

The following theorem was proved by Takhanov [59] with a small strengthening in [43].

Theorem 9 ([43, 59]). *Let Γ be a conservative valued constraint language. If $\text{Pol}(\Gamma)$ does not contain a majority polymorphism, then $\text{VCSP}(\Gamma)$ is NP-hard.*

We can show the converse of Theorem 9. Consequently, the tractability criterion for conservative valued constraint languages is that of admitting a majority operation in the support clone of the language, which is simpler compared to the original criterion [43] that involved a binary STP multimorphism and a ternary MJN multimorphism (cf. Example 7).

Theorem 10. *Let Γ be a conservative valued constraint language. Either $\text{VCSP}(\Gamma)$ is NP-hard, or $\text{supp}(\Gamma)$ contains a majority operation and hence Γ has valued relational width $(2, 3)$.*

3.5 Related Work on BLP and Relational Width

The $\text{SA}(1, 1)$ relaxation is also known as the *basic linear programming relaxation* (BLP). The following result capturing the power of BLP has been established.³

An m -ary operation $f : D^m \rightarrow D$ is called *symmetric* if $f(x_1, \dots, x_m) = f(x_{\pi(1)}, \dots, x_{\pi(m)})$ for every permutation π of $\{1, \dots, m\}$.

Definition 15. *We say that a clone of operation satisfies the SYM condition if it contains an m -ary symmetric operation, for every $m \geq 2$.*

Theorem 11 ([42]). *Let Γ be a valued constraint language of finite size. Then the following are equivalent:*

1. Γ has valued relational width 1.
2. $\text{supp}(\Gamma)$ satisfies the SYM.

By definition, the $\text{SA}(1, \ell)$ -relaxation is at least as tight as the $\text{SA}(1, 1)$ -relaxation. Hence any language with valued relational width 1 has valued relational width $(1, \ell)$. We now show that for any fixed ℓ , $\text{SA}(1, 1)$ and $\text{SA}(1, \ell)$ have the same power.

³Theorem 11 as stated here follows from [42, Corollary 3] using Lemma 1.

Proposition 1. *Let Γ be a valued constraint language of finite size and let $\ell > 1$ be fixed. If Γ has valued relational width $(1, \ell)$ then Γ has valued relational width 1.*

Proof. Let I be an instance of $\text{VCSP}(\Gamma)$. Assume that $\text{Opt}_{\text{VCSP}}(I) = \text{Opt}_{\text{LP}}(I)$ for the $\text{SA}(1, \ell)$ -relaxation of I . For the sake of contradiction, suppose that $\text{Opt}_{\text{VCSP}}(I) > \text{Opt}_{\text{LP}}(I)$ for the $\text{SA}(1, 1)$ -relaxation of I and let λ^* be an optimal solution to $\text{SA}(1, 1)$ of value OPT^* . Define λ' as follows. If $\lambda_i(\sigma)$ is a variable of $\text{SA}(1, 1)$ then $\lambda'_i(\sigma) = \lambda_i^*(\sigma)$. Otherwise, let $\lambda_i(\sigma)$ correspond to the i th valued constraint $\phi_i(\mathbf{x}_i)$ with variables $\{x_1, \dots, x_r\}$. We define $\lambda'_i(\sigma)$ as the product of the λ^* 's corresponding to $\sigma(x_j)$, $1 \leq j \leq r$. More formally, if $\phi_j(x_j)$ are the valued constraints with the scope x_j , for $1 \leq j \leq r$, then we define $\lambda'_i(\sigma) = \prod_{j=1}^r \lambda_j^*(\sigma(x_j))$. By the definition of λ' , λ' is a feasible solution to $\text{SA}(1, \ell)$. By the definition of the SA relaxations, the extra valued constraints present in $\text{SA}(1, \ell)$ but missing in $\text{SA}(1, 1)$ are null and thus $\text{Val}_{\text{LP}}(I, \lambda') = \text{OPT}^* < \text{Opt}_{\text{VCSP}}(I)$. But this contradicts Γ having valued relational width $(1, \ell)$. \square

Corollary 5. *Let Γ be a valued constraint language of finite size. Then, the valued relational width of Γ is either 1, or 2, or $(2, 3)$, or unbounded.*

Proof. If the valued relational width of Γ is bounded then it is $(2, 3)$, by Theorem 3. If the valued relational width of Γ is $(1, \ell)$ for some $\ell > 1$ then it is 1, by Proposition 1. \square

There are valued constraint languages that have valued relational width $(2, 3)$ but not 1. For example, languages admitting a tournament pair fractional polymorphism [18], discussed in detail in Example 6 in Section 3.3, have valued relational width $(2, 3)$ by the results in this paper, but do not have valued relational width 1 as shown [42, Example 5] using Theorem 11.

It could be that either $\text{SA}(1)$ and $\text{SA}(2)$, or $\text{SA}(2)$ and $\text{SA}(2, 3)$ have the same power. The former happens in case of relational width. Dalmau proved that if a crisp language has relational width 2 then it has relational width 1 [23]. Together with Theorem 2 and the analogue of Proposition 1 for relational width established in [28], this gives a trichotomy for relational width.

Theorem 12 ([4, 23, 28]). *Let Δ be a crisp constraint language of finite size. Then precisely one of the following is true:*

1. Δ has relational width 1.
2. Δ has relational width $(2, 3)$ and does not have relational width 2, nor $(1, \ell)$ for any $\ell \geq 1$.
3. Δ does not have bounded relational width.

Remark 2. It follows from the definitions that if a crisp constraint language Δ has relational width (k, ℓ) then Δ also has valued relational width (k, ℓ) . However, the converse does not hold. There exists a constraint language on a three-element domain with two relations that has valued relational width 1 but not relational width 1 [48, Example 99].

3.6 Obtaining a Solution and the Meta Problem

We now address two questions related to our main result.

Firstly, we show that for any VCSP instance over a language of valued relational width $(2, 3)$ we can not only compute the value of an optimal solution but we can also find an optimal assignment in polynomial time.

Proposition 2. *Let Γ be a valued constraint language of finite size and I an instance of $\text{VCSP}(\Gamma)$. If $\text{supp}(\Gamma)$ satisfies the BWC, then an optimal assignment to I can be found in polynomial time.*

Proof. Let Γ' be a core of Γ on domain D' , and let $\Gamma_c = \Gamma' \cup \{\mathcal{C}_{D'}\}$. By Lemma 5, $\text{supp}(\Gamma_c)$ satisfies the BWC, so by Theorem 4 we can obtain the optimum of I by solving a linear programming relaxation. Now, we can use self-reduction to obtain an optimal assignment. It suffices to modify the instance I to successively force each variable to take on each value of D' . Whenever the optimum of the modified instance matches that of the original instance, we can move on to assign the next variable. This means that we need to solve at most $1 + |V| |D'|$ linear programming relaxations before finding an optimal assignment, where V is the set of variables of I . \square

Secondly, we show that testing for the BWC is a decidable problem. We rely on the following result that was proved in [44], and also follows from results in [4].

Theorem 13 ([44]). *An idempotent clone of operations satisfies the BWC if, and only if, it contains a ternary WNU f and a quaternary WNU g with $f(y, x, x) = g(y, x, x, x)$ for all x and y .*

Proposition 3. *Testing whether a valued constraint language of finite size satisfies the BWC is decidable.*

Proof. Let Γ be a valued constraint language of finite size on domain D . Let Γ' be a core of Γ defined on domain $D' \subseteq D$. Finding D' and Γ' can be done via linear programming [61, Section 4]. By Lemma 5, $\text{supp}(\Gamma)$ satisfies the BWC if, and only if, $\text{supp}(\Gamma' \cup \mathcal{C}_{D'})$ satisfies the BWC. As constant unary relations enforce idempotency, by Theorem 13, $\text{supp}(\Gamma' \cup \mathcal{C}_{D'})$ satisfies the BWC if, and only if, $\text{supp}(\Gamma' \cup \mathcal{C}_{D'})$ contains a ternary WNU f and a 4-ary WNU g with $f(y, x, x) = g(y, x, x, x)$ for all x and y . It is easy to write a linear program that checks for this condition, as it has been done in the context of finite-valued constraint languages [61, Section 4]. \square

4 Sufficiency: Proof of Theorem 4

In this section, we prove that the BWC is a sufficient condition for a valued constraint language to have valued relational width $(2, 3)$.

We start with a technical lemma. For a feasible solution λ of $\text{SA}(k, \ell)$, let $\text{supp}(\lambda_i) = \{\sigma : X_i \rightarrow D \mid \lambda_i(\sigma) > 0\}$.

Lemma 6. *Let I be an instance of $\text{VCSP}(\Gamma)$. Assume that $\text{SA}(k, \ell)$ for I is feasible. Then, there exists an optimal solution λ^* to $\text{SA}(k, \ell)$ such that, for every i , $\text{supp}(\lambda_i^*)$ is closed under every operation in $\text{supp}(\Gamma)$.*

Proof. Let ω be an arbitrary m -ary fractional polymorphism of Γ , and let λ be any feasible solution λ to $\text{SA}(k, \ell)$. Define λ^ω by

$$\lambda_i^\omega(\sigma) = \Pr_{\substack{f \sim \omega \\ \sigma_1, \dots, \sigma_m \sim \lambda_i}} [f \circ (\sigma_1, \dots, \sigma_m) = \sigma].$$

We show that λ^ω is a feasible solution to $\text{SA}(k, \ell)$, and that if λ is optimal, then so is λ^ω .

Clearly λ_i^ω is a probability distribution for each $i \in [q]$, so (11) and (13) hold. Since ω is a fractional polymorphism of Γ , we have $\sigma \in \text{Feas}(\phi_i)$ for any choice of $f \in \text{supp}(\omega)$ and $\sigma_1, \dots, \sigma_m \in \text{supp}(\lambda_i)$. Hence, $\lambda_i^\omega(\sigma) = 0$ for $\sigma \notin \text{Feas}(\phi_i)$, so (12) holds.

Finally, let $j \in [q]$ be such that $X_j \subseteq X_i$, $|X_j| \leq k$, and let $\tau: X_j \rightarrow D$. Then,

$$\begin{aligned}
\sum_{\sigma: X_i \rightarrow D, \sigma|_{X_j} = \tau} \lambda_i^\omega(\sigma) &= \sum_{\sigma: X_i \rightarrow D, \sigma|_{X_j} = \tau} \Pr_{f \sim \omega, \sigma_1, \dots, \sigma_m \sim \lambda_i} [f \circ (\sigma_1, \dots, \sigma_m) = \sigma] \\
&= \Pr_{f \sim \omega, \sigma_1, \dots, \sigma_m \sim \lambda_i} [(f \circ (\sigma_1, \dots, \sigma_m))|_{X_j} = \tau] \\
&= \Pr_{f \sim \omega, \sigma_1, \dots, \sigma_m \sim \lambda_i} [f \circ (\sigma_1|_{X_j}, \dots, \sigma_m|_{X_j}) = \tau] \\
&= \sum_{\tau_1, \dots, \tau_m: X_j \rightarrow D} \Pr_{f \sim \omega, \sigma_1, \dots, \sigma_m \sim \lambda_i} [\sigma_1|_{X_j} = \tau_1, \dots, \sigma_m|_{X_j} = \tau_m, f \circ (\tau_1, \dots, \tau_m) = \tau] \\
&= \sum_{\tau_1, \dots, \tau_m: X_j \rightarrow D} \lambda_j(\tau_1) \cdots \lambda_j(\tau_m) \Pr_{f \sim \omega} [f \circ (\tau_1, \dots, \tau_m) = \tau] \\
&= \Pr_{f \sim \omega, \tau_1, \dots, \tau_m \sim \lambda_j} [f \circ (\tau_1, \dots, \tau_m) = \tau] \\
&= \lambda_j^\omega(\tau),
\end{aligned}$$

where, we have used the fact that (10) can be read as $\lambda_j(\tau) = \Pr_{\sigma \sim \lambda_i} [\sigma|_{X_j} = \tau]$. It follows that (10) also holds for λ^ω , so λ^ω is feasible.

For each $i \in [q]$, we have:

$$\begin{aligned}
\sum_{\sigma \in \text{Feas}(\phi_i)} \lambda_i(\sigma) \phi_i(\sigma(\mathbf{x}_i)) &= \mathbb{E}_{\sigma \sim \lambda_i} \phi_i(\sigma) = \mathbb{E}_{\sigma_1, \dots, \sigma_m \sim \lambda_i} \frac{1}{m} \sum_{j=1}^m \phi_i(\sigma_j(\mathbf{x}_i)) \\
&\geq \mathbb{E}_{f \sim \omega, \sigma_1, \dots, \sigma_m \sim \lambda_i} \phi_i(f(\sigma_1(\mathbf{x}_i), \dots, \sigma_m(\mathbf{x}_i))) \\
&= \sum_{\sigma \in \text{Feas}(\phi_i)} \left(\Pr_{f \sim \omega, \sigma_1, \dots, \sigma_m \sim \lambda_i} [f \circ (\sigma_1, \dots, \sigma_m) = \sigma] \right) \phi_i(\sigma(\mathbf{x}_i)) \\
&= \sum_{\sigma \in \text{Feas}(\phi_i)} \lambda_i^\omega(\sigma) \phi_i(\sigma(\mathbf{x}_i)).
\end{aligned}$$

Therefore, if λ is optimal, then λ^ω must also be optimal.

Now assume that λ is an optimal solution and that $\text{supp}(\lambda)$ is not closed under some operation $f \in \text{supp}(\omega)$ for $\omega \in \text{fPol}(\Gamma)$, i.e. for some $\sigma_1, \dots, \sigma_m \in \text{supp}(\lambda)$, we have $f(\sigma_1, \dots, \sigma_m) \notin \text{supp}(\lambda)$. But note that $f(\sigma_1, \dots, \sigma_m) \in \text{supp}(\lambda_i^\omega)$. Therefore, $\lambda' = \frac{1}{2}(\lambda + \lambda^\omega)$ is an optimal solution such that $\text{supp}(\lambda_i) \subsetneq \text{supp}(\lambda'_i) \subseteq D^{X_i}$. For each $i \in [q]$, D^{X_i} is finite. Hence, by repeating this procedure, we obtain a sequence of optimal solutions with strictly increasing support until, after a finite number of steps, we obtain a λ^* that is closed under every operation in $\text{supp}(\Gamma)$. \square

We now have everything that is needed to prove Theorem 4.

Proof of Theorem 4. Let I be an instance of VCSP(Γ) with $\phi_I(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i)$, $X_i \subseteq V = \{x_1, \dots, x_n\}$ and $\phi_i: D^{\text{ar}(\phi_i)} \rightarrow \overline{\mathbb{Q}}$.

The dual of the SA(k, ℓ) relaxation can be written in the following form, with variables z_i for $i \in [q]$ and $y_{j, \tau, i}$ for $i, j \in [q]$ such that $X_j \subseteq X_i$, $|X_j| \leq k$, and $\tau: X_j \rightarrow D$. The dual variables corresponding to $\lambda_i(\sigma) = 0$ are eliminated together with the dual inequalities for $i, \sigma \notin \text{Feas}(\phi_i)$.

$$\max \sum_{i=1}^q z_i$$

$$z_i \leq \phi_i(\sigma) + \sum_{j \in [q], X_j \subseteq X_i} y_{j, \sigma|_{X_j}, i} - \sum_{j \in [q], X_i \subseteq X_j} y_{i, \sigma, j} \quad \forall i \in [q], |X_i| \leq k, \sigma \in \text{Feas}(\phi_i) \quad (14)$$

$$z_i \leq \phi_i(\sigma) + \sum_{\substack{j \in [q], X_j \subseteq X_i \\ |X_j| \leq k}} y_{j, \sigma|_{X_j}, i} \quad \forall i \in [q], |X_i| > k, \sigma \in \text{Feas}(\phi_i) \quad (15)$$

It is clear that if I has a feasible solution, then so does the $\text{SA}(k, \ell)$ primal. Assume that the $\text{SA}(2, 3)$ -relaxation has a feasible solution.

By Lemma 6, there exists an optimal primal solution λ^* such that, for every $i \in [q]$, $\text{supp}(\lambda_i^*)$ is closed under $\text{supp}(\Gamma)$. Let y^*, z^* be an optimal dual solution.

Let $\Delta = \{\phi'_i\}_{i=1}^q \cup \{\mathcal{C}_D\}$, where $\phi'_i = \text{supp}(\lambda_i^*)$, i.e. $\phi'_i(\mathbf{x}) = 0$ if $\mathbf{x} \in \text{supp}(\lambda_i^*)$ and $\phi'_i(\mathbf{x}) = \infty$ otherwise. We consider the instance J of $\text{CSP}(\Delta)$ with $\phi_J(x_1, \dots, x_n) = \sum_{i=1}^q \phi'_i(\mathbf{x}_i)$.

We make the following observations:

1. By construction of λ^* , $\text{supp}(\Gamma) \subseteq \text{Pol}(\Delta)$, so Δ contains all constant unary relations and satisfies the BWC. By Theorems 1 and 2, the language Δ has relational width $(2, 3)$.
2. The first set of constraints in the primal say that if $i, j \in [q]$, $|X_j| \leq 2$ and $X_j \subseteq X_i$, then $\lambda_j^*(\tau) > 0$ (i.e. $\tau \in C_j$) if, and only if, $\sum_{\sigma: X_i \rightarrow D, \sigma|_{X_j} = \tau} \lambda_i^*(\sigma) > 0$ (i.e. τ satisfies $\pi_{X_j}(C_i)$).

In other words, J is $(2, 3)$ -minimal.

These two observations imply that J has a satisfying assignment $\alpha: V \rightarrow D$. Let $\alpha_i = \alpha|_{X_i}$. By complementary slackness, since $\lambda_i^*(\alpha_i) > 0$ for every $i \in [q]$, we must have equality in the corresponding rows in the dual indexed by i and α_i . We sum these rows over i :

$$\sum_{i=1}^q z_i^* = \sum_{i=1}^q \phi_i(\alpha(\mathbf{x}_i)) + \left(\sum_{i=1}^q \sum_{\substack{j \in [q], X_j \subseteq X_i \\ |X_j| \leq 2}} y_{j, \alpha_i|_{X_j}, i}^* - \sum_{\substack{i \in [q] \\ |X_i| \leq 2}} \sum_{\substack{j \in [q] \\ X_i \subseteq X_j}} y_{i, \alpha_i, j}^* \right). \quad (16)$$

By noting that $\alpha_i|_{X_j} = \alpha_j$ when $X_j \subseteq X_i$, we can rewrite the expression in parenthesis on the right-hand side of (16) as:

$$\sum_{\substack{i, j \in [q], X_j \subseteq X_i \\ |X_j| \leq 2}} y_{j, \alpha_j, i}^* - \sum_{\substack{i, j \in [q], X_i \subseteq X_j \\ |X_i| \leq 2}} y_{i, \alpha_i, j}^* = 0. \quad (17)$$

Therefore,

$$\sum_{i=1}^q \sum_{\sigma \in \text{Feas}(\phi_i)} \lambda_i^*(\sigma) \phi_i(\sigma(\mathbf{x}_i)) = \sum_{i=1}^q z_i^* = \sum_{i=1}^q \phi_i(\alpha(\mathbf{x}_i)),$$

where the first equality follows by strong LP-duality, and the second by (16) and (17). Since I was an arbitrary instance of $\text{VCSP}(\Gamma)$, the theorem follows. \square

5 Necessity: Proof of Theorem 5

In this section, we prove that the BWC is a necessary condition for a valued constraint language to have bounded valued relational width.

The main idea of the proof is to show that if $\text{supp}(\Gamma)$ does not satisfy the BWC, then Γ can, in a sense, simulate linear equations in some Abelian group. We show that such linear equations do not have bounded valued relational width, and that the simulation preserves bounded valued relational width. We first state the result on linear equations in an Abelian group and then discuss the precise meaning of “simulation”.

Let \mathcal{G} be an Abelian group over a finite set G and let $r \geq 1$ be an integer. Denote by $E_{\mathcal{G},r}$ the crisp constraint language over domain G with, for every $a \in G$, and $1 \leq m \leq r$, a relation $R_a^m = \{(x_1, \dots, x_m) \in G^m \mid x_1 + \dots + x_m = a\}$. In Section 7, we prove the following.

Theorem 14. *Let \mathcal{G} be a finite non-trivial Abelian group. Then, the constraint language $E_{\mathcal{G},3}$ does not have bounded valued relational width.*

Definition 16. *We say that an m -ary weighted relation ϕ is expressible over a valued constraint language Γ if there exists an instance I of $\text{VCSP}(\Gamma)$ with variables $x_1, \dots, x_m, v_1, \dots, v_p$ such that*

$$\phi(x_1, \dots, x_m) = \min_{v_1, \dots, v_p} \phi_I(x_1, \dots, x_m, v_1, \dots, v_p). \quad (18)$$

For a fixed set D , let ϕ_{\equiv}^D denote the binary equality relation $\{(x, x) \mid x \in D\}$. Denote by $\langle \Gamma \rangle$ all weighted relations expressible in $\Gamma \cup \{\phi_{\equiv}^D\}$, where D is the domain of Γ . A weighted relation being expressible over $\Gamma \cup \{\phi_{\equiv}^D\}$ is the analogue of a relation being definable by a *primitive positive* (pp) formula (using existential quantification and conjunction) over a relational structure with equality. Indeed, when Γ is crisp, the two notions coincide.

Definition 17. *Let Γ and Δ be valued constraint languages on domain D and D' , respectively. We say that Δ has an interpretation in Γ with parameters (d, S, h) if there exists a $d \in \mathbb{N}$, a set $S \subseteq D^d$, and a surjective map $h : S \rightarrow D'$ such that $\langle \Gamma \rangle$ contains the following weighted relations:*

- $\phi_S : D^d \rightarrow \overline{\mathbb{Q}}$ defined by $\phi_S(\mathbf{x}) = 0$ if $\mathbf{x} \in S$ and $\phi_S(\mathbf{x}) = \infty$ otherwise;
- $h^{-1}(\phi_{\equiv}^{D'})$; and
- $h^{-1}(\phi_i)$, for every weighted relation $\phi_i \in \Delta$,

where $h^{-1}(\phi_i)$, for an m -ary weighted relation ϕ_i , is the dm -ary weighted relation on D defined by $h^{-1}(\phi_i)(\mathbf{x}_1, \dots, \mathbf{x}_m) = \phi_i(h(\mathbf{x}_1), \dots, h(\mathbf{x}_m))$, for all $\mathbf{x}_1, \dots, \mathbf{x}_m \in S$.

When Γ is crisp, the notion of an interpretation coincides with the notion of a *pp-interpretation* for relational structures [7].

Theorem 15. *Let Δ be a crisp constraint language of finite size that contains all constant unary relations. If $\text{Pol}(\Delta)$ does not satisfy the BWC, then there exists a finite non-trivial Abelian group \mathcal{G} such that Δ interprets $E_{\mathcal{G},r}$, for every $r \geq 1$.*

Proof. It has been shown in [44, Theorem 1.6 (4)] that if the polymorphism algebra B of Δ does not satisfy the BWC, then the variety generated by B admits type **1** or **2** (the notion of admitting types comes from Tame Congruence Theory [33]). By [2, Lemmas 20 and 21], this implies that there exists a finite non-trivial Abelian group \mathcal{G} such that the variety generated by B contains a reduct A of the polymorphism algebra of $E_{\mathcal{G},r}$, for every $r \geq 1$. For finite algebras A and B , A is contained in the variety generated by B if, and only if, A is contained in the pseudovariety generated by B . In terms of pp-interpretations [7], this is equivalent to $E_{\mathcal{G},r}$ having a pp-interpretation in Δ (see also [11]). \square

Our notion of reduction will be the \leq_{SA} reduction from Definition 13.

The following theorem shows that we can augment a valued constraint language with various additional weighted relations. The transformations in these reductions have previously been used to prove polynomial-time reductions [11, 16, 45, 63, 65]. Here, we show that they all *additionally* preserve bounded valued relational width.

Theorem 16. *Let Γ be a valued constraint language of finite size on domain D . The following holds:*

1. *If ϕ is expressible in Γ , then $\Gamma \cup \{\phi\} \leq_{\text{SA}} \Gamma$.*
2. *$\Gamma \cup \{\phi_{=}^D\} \leq_{\text{SA}} \Gamma$.*
3. *If Γ interprets the valued constraint language Δ of finite size, then $\Delta \leq_{\text{SA}} \Gamma$.*
4. *If $\phi \in \Gamma$, then $\Gamma \cup \{\text{Opt}(\phi)\} \leq_{\text{SA}} \Gamma$ and $\Gamma \cup \{\text{Feas}(\phi)\} \leq_{\text{SA}} \Gamma$.*
5. *If Γ' is a core of Γ on domain $D' \subseteq D$, then $\Gamma' \cup \mathcal{C}_{D'} \leq_{\text{SA}} \Gamma$.*

Note that Theorem 16(5) is just a restatement of Lemma 4.

A formal proof of Theorem 16 is given in Section 6. Here is the main idea.

All of the reductions are based on replacing each constraint $\phi_i(\mathbf{x}_i)$ of an instance I of the left-hand side by some gadget, given as an instance J_i of the right-hand side. The instance J is then defined as the sum of all objective functions ϕ_{J_i} .

If the replacements satisfy certain conditions, then we show that, for any $1 \leq k' \leq \ell'$, there exist $1 \leq k \leq \ell$ such that if Δ' has valued relational width (k', ℓ') , then Δ has valued relational width (k, ℓ) , so the reductions preserve bounded valued relational width. The conditions are: (a) for every optimal solution α of J , there is a satisfying and optimal assignment σ^α of I such that $\text{Val}_{\text{VCSP}}(I, \sigma^\alpha) \leq \text{Val}_{\text{VCSP}}(J, \alpha)$; (b) for every large enough k , feasible solution λ to the $\text{SA}(k, 2k)$ -relaxation of I , and assignment $\sigma: X_i \rightarrow D$ with positive support in λ , there exists a satisfying assignment α_i^σ of J_i such that $\phi_i(\sigma(\mathbf{x}_i)) \geq \text{Val}_{\text{VCSP}}(J_i, \alpha_i^\sigma)$; and (c) the assignments α_i^σ are “pairwise consistent”, i.e. $\alpha_i^{\sigma_i}$ and $\alpha_r^{\sigma_r}$ agree on the intersection of the variables of J_i and J_r whenever σ_i and σ_r are restrictions of some $\sigma: X \rightarrow D$ with positive support in λ .

We will also need the following technical lemmas.

Lemma 7. *Let Γ be a valued constraint language of finite size over domain D and let F be a finite set of operations over D . If $\text{supp}(\Gamma) \cap F = \emptyset$, then there exists a crisp constraint language Δ such that $\text{Pol}(\Delta) \cap F = \emptyset$ and $\Delta \leq_{\text{SA}} \Gamma$.*

Proof. By Lemma 2, for each $f \in F \cap \text{Pol}(\Gamma)$, there is an instance I_f of $\text{VCSP}(\Gamma)$ such that $f \notin \text{Pol}(\text{Opt}(\phi_{I_f}))$. Let $\Delta = \{\text{Opt}(\phi_{I_f}) \mid f \in F\} \cup \{\text{Feas}(\phi) \mid \phi \in \Gamma\}$. For $f \in F \cap \text{Pol}(\Gamma)$, we have $f \notin \text{Pol}(\text{Opt}(\phi_{I_f})) \supseteq \text{Pol}(\Delta)$. For $f \in F \setminus \text{Pol}(\Gamma)$, we have $f \notin \text{Pol}(\phi)$, for some $\phi \in \Gamma$, so $f \notin \text{Pol}(\Delta)$. It follows that $\text{Pol}(\Delta) \cap F = \emptyset$. Finally, $\Delta \leq_{\text{SA}} \Gamma$ holds by Theorem 16(1) and (4). \square

Lemma 8. *Let Γ be a valued constraint language of finite size. If $\text{supp}(\Gamma)$ does not satisfy the BWC, then there is a crisp constraint language Δ of finite size such that $\text{Pol}(\Delta)$ does not satisfy the BWC, and $\Delta \leq_{\text{SA}} \Gamma$.*

Proof. Since $\text{supp}(\Gamma)$ does not satisfy the BWC, there exists an $m \geq 3$ such that $\text{supp}(\Gamma)$ does not contain any m -ary WNU. Let F be the (finite) set of all m -ary WNUs. The result follows by applying Lemma 7 to Γ and F . \square

We now have everything that is needed to prove Theorem 5.

Proof of Theorem 5. Suppose that $\text{supp}(\Gamma)$ does not satisfy the BWC. By Lemma 8, there exists a crisp constraint language Δ such that $\text{Pol}(\Delta)$ does not satisfy the BWC and $\Delta \leq_{\text{SA}} \Gamma$. Since $\mathcal{C}_D \subseteq \Gamma$, we may assume, without loss of generality, that $\mathcal{C}_D \subseteq \Delta$.

By Theorem 15, there exists a finite non-trivial Abelian group \mathcal{G} and an interpretation of $E_{\mathcal{G},3}$ in Δ . By Theorem 14, $E_{\mathcal{G},3}$ does not have bounded valued relational width. By Theorem 16(3), we have $E_{\mathcal{G},3} \leq_{\text{SA}} \Delta \leq_{\text{SA}} \Gamma$, so Γ does not have bounded valued relational width. \square

6 Reductions: Proof of Theorem 16

Theorem 16 follows from Lemmas 10–15 proved in this section.

For a valued constraint language Γ , let $\text{ar}(\Gamma)$ denote $\max\{\text{ar}(\phi) \mid \phi \in \Gamma\}$.

It will sometimes be convenient to add null constraints to a VCSP instance as placeholders, to ensure that they have a scope, even if these relations may not necessarily be members of the corresponding constraint language Γ . In order to obtain an equivalent instance that is formally in $\text{VCSP } \Gamma$, the null constraints can simply be dropped, as they are always satisfied and do not influence the value of the objective function.

We extend the convention of denoting the set of variables in \mathbf{x}_i by X_i to tuples $\mathbf{y}_i, \mathbf{y}'_i$, and \mathbf{v} , whose sets are denoted by Y_i, Y'_i , and V_i , respectively.

The following technical lemma is the basis for most of the reductions.

Lemma 9. *Let Δ and Δ' be valued constraint languages over domains D and D' , respectively.*

Let $(I, i) \mapsto J_i$ be a map that to each instance I of $\text{VCSP}(\Delta)$ with variables V and objective function $\sum_{i=1}^q \phi_i(\mathbf{x}_i)$, and index $i \in [q]$, associates an instance J_i of $\text{VCSP}(\Delta')$ with variables Y_i and objective function ϕ_{J_i} . Let J be the $\text{VCSP}(\Delta')$ instance with variables $V' = \bigcup_{i=1}^q Y_i$ and objective function $\sum_{i=1}^q \phi_{J_i}$.

Suppose that the following holds:

- (a) *For every satisfying and optimal assignment α of J , there exists a satisfying assignment σ^α of I such that*

$$\text{Val}_{\text{VCSP}}(I, \sigma^\alpha) \leq \text{Val}_{\text{VCSP}}(J, \alpha).$$

Furthermore, suppose that for $k \geq \text{ar}(\Delta)$, and any feasible solution λ of the $\text{SA}(k, 2k)$ -relaxation of I , the following properties hold:

- (b) *For $i \in [q]$, and $\sigma: X_i \rightarrow D$ with positive support in λ , there exists a satisfying assignment α_i^σ of J_i such that*

$$\phi_i(\sigma(\mathbf{x}_i)) \geq \text{Val}_{\text{VCSP}}(J_i, \alpha_i^\sigma);$$

- (c) *for $i, r \in [q]$, any $X \subseteq V$ with $X_i \cup X_r \subseteq X$, and $\sigma: X \rightarrow D$ with positive support in λ ,*

$$\alpha_i^{\sigma_i}|_{Y_i \cap Y_r} = \alpha_r^{\sigma_r}|_{Y_i \cap Y_r},$$

where $\sigma_i = \sigma|_{X_i}$ and $\sigma_r = \sigma|_{X_r}$.

Then, $I \mapsto J$ is a many-one reduction from $\text{VCSP}(\Delta)$ to $\text{VCSP}(\Delta')$, and for any $1 \leq k' \leq \ell'$, there exist $1 \leq k \leq \ell$ such that if I is a gap instance for $\text{SA}(k, \ell)$, then J is a gap instance for $\text{SA}(k', \ell')$. In particular, the reduction preserves bounded valued relational width.

Proof. From condition (a), if J is satisfiable, then so is I and $\text{Opt}_{\text{VCSP}}(I) \leq \text{Opt}_{\text{VCSP}}(J)$. Conversely, if I is satisfiable, and σ is an optimal assignment, then the LP solution λ that assigns probability 1 to $\sigma|_X$ for every X is feasible, so by (b), $\text{Opt}_{\text{VCSP}}(I) \geq \text{Opt}_{\text{VCSP}}(J)$.

Let $1 \leq k' \leq \ell'$ be arbitrary, and let $k = \max\{\ell', \text{ar}(\Delta')\} \cdot \text{ar}(\Delta)$, $\ell = 2k$. Assume that I is a gap instance for the $\text{SA}(k, \ell)$ -relaxation of $\text{VCSP}(\Delta)$, and let λ be a feasible solution such that $\text{Val}_{\text{LP}}(I, \lambda) < \text{Opt}_{\text{VCSP}}(I)$ (where $\text{Opt}_{\text{VCSP}}(I)$ may be ∞ , i.e. I may be unsatisfiable). We show that there is a feasible solution κ to the $\text{SA}(k', \ell')$ -relaxation of J such that $\text{Val}_{\text{LP}}(J, \kappa) \leq \text{Val}_{\text{LP}}(I, \lambda)$. Then, by condition (a), we have $\text{Val}_{\text{LP}}(J, \kappa) \leq \text{Val}_{\text{LP}}(I, \lambda) < \text{Opt}_{\text{VCSP}}(I) \leq \text{Opt}_{\text{VCSP}}(J)$, so J is a gap instance for the $\text{SA}(k', \ell')$ -relaxation of $\text{VCSP}(\Delta')$. Since k' and ℓ' were chosen arbitrarily, the result then follows.

To this end, augment I with null constraints on $X_{q+1}, \dots, X_{q'}$ so that for every at most ℓ -subset $X \subseteq V$, there exists an $i \in [q']$ such that $X_i = X$. Rewrite the objective function of J as $\sum_{j=1}^p \phi'_j(\mathbf{y}'_j)$, $\phi' \in \Delta'$, where, by possibly first adding extra null constraints to J , we will

assume that for every at most ℓ' -subset $Y \subseteq V'$, there exists a $j \in [p]$ such that $Y'_j = Y$. For each $i \in [q]$, let C_i be the set of indices $j \in [p]$ corresponding to the weighted constraints in the instance J_i .

Extend α_i^σ to $i \in [q'] \setminus [q]$ as follows. For $X \subseteq V$, let $Y_X = \bigcup_{i \in [q]: X_i \subseteq X} Y_i$. For $\sigma \in \text{supp}(\lambda_i)$, and any $r, s \in [q]$ such that $X_r \cup X_s \subseteq X_i$ and $y \in Y_r \cap Y_s$, by (c), it holds that $\alpha_r^{\sigma_r}(y) = \alpha_s^{\sigma_s}(y)$. Therefore, we can uniquely define $\alpha_i^\sigma: Y_{X_i} \rightarrow D'$ by letting $\alpha_i^\sigma(y) = \alpha_r^{\sigma_r}(y)$ for any choice of $r \in [q]$ with $X_r \subseteq X_i$ and $y \in Y_r$. Furthermore, this definition is consistent with α_i^σ for $i \in [q]$ in the sense that (c) now holds for all $i, r \in [q']$.

For $m \geq 1$, let $X_{(\leq m)} = \{X = \bigcup_{i \in S} X_i \mid S \subseteq [q], |X| \leq m\}$, and for $Y \subseteq V'$ with $|Y| \leq \ell'$, let $X_{(\leq m)}(Y) = \{X \in X_{(\leq m)} \mid Y \subseteq Y_X\}$.

Let $j \in [p]$ be arbitrary and let $X = \bigcup_{i \in S} X_i \in X_{(\leq n)}(Y'_j)$, for some $S \subseteq [q]$. For each $y \in Y'_j$, let $i(y) \in S$ be an index such that $y \in Y_{i(y)}$ and let $X' = \bigcup_{y \in Y'_j} X_{i(y)}$. Then, $Y'_j \subseteq X' \subseteq X$, and $|X'| \leq \ell' \cdot \text{ar}(\Delta) \leq k$, so $X' \in X_{(\leq k)}(Y'_j)$. In other words,

$$\text{for all } X \in X_{(\leq n)}(Y'_j), \text{ there exists } X' \subseteq X \text{ such that } X' \in X_{(\leq k)}(Y'_j). \quad (19)$$

In particular (19) shows that $X_{(\leq k)}(Y'_j) \neq \emptyset$, since $\bigcup_{i \in [q]} X_i \in X_{(\leq n)}(Y'_j)$ for all j .

For $j \in [p]$, $\alpha: Y'_j \rightarrow D'$, and an $i \in [q']$ such that $X_i \in X_{(\leq \ell)}(Y'_j)$, define

$$\mu_j^i(\alpha) = \sum_{\substack{\sigma \in \text{supp}(\lambda_i) \\ \alpha_i^\sigma|_{Y'_j} = \alpha}} \lambda_i(\sigma). \quad (20)$$

Claim: Definition (20) is independent of the choice of $X_i \in X_{(\leq \ell)}(Y'_j)$.

It suffices to show $\mu_j^r = \mu_j^i$ for every $r, i \in [q']$ with $X_r \in X_{(\leq k)}(Y'_j)$ and $X_i \in X_{(\leq \ell)}(Y'_j)$.

First, we prove this equality in the case when $X_r \subseteq X_i$. Then,

$$\mu_j^r(\alpha) = \sum_{\substack{\tau \in \text{supp}(\lambda_r) \\ \alpha_r^\tau|_{Y'_j} = \alpha}} \sum_{\substack{\sigma \in \text{supp}(\lambda_i) \\ \sigma|_{X_r} = \tau}} \lambda_i(\sigma) = \sum_{\substack{\sigma \in \text{supp}(\lambda_i) \\ \alpha_r^{\sigma_r}|_{Y'_j} = \alpha}} \lambda_i(\sigma) = \sum_{\substack{\sigma \in \text{supp}(\lambda_i) \\ \alpha_i^\sigma|_{Y'_j} = \alpha}} \lambda_i(\sigma) = \mu_j^i(\alpha),$$

where the first equality follows by (20) and (10) for λ since $|X_r| \leq k$, and the second equality follows by interchanging the order of summation and noting that $\sigma \in \text{supp}(\lambda_i)$ implies that $\sigma_r \in \text{supp}(\lambda_r)$, again by (10) for λ . The third equality follows by (c) extended to $i, r \in [q']$.

Next, let $X_r \in X_{(\leq k)}(Y'_j)$ and $X_i \in X_{(\leq \ell)}(Y'_j)$ be arbitrary. From (19), it follows that X_i contains a subset $X_s \in X_{(\leq k)}(Y'_j)$. Since $|X_r \cup X_s| \leq 2k = \ell$, there exists an index u such that $X_u = X_r \cup X_s$. The claim now follows by a repeated application of the first case: $\mu_j^r = \mu_j^u = \mu_j^s = \mu_j^i$.

By the claim, we can pick an arbitrary $X_i \in X_{(\leq \ell)}(Y'_j)$ and uniquely define $\kappa_j = \mu_j^i$. We now show that this definition of κ satisfies the equations (10)–(13).

- μ_j^i can be seen as the probability distribution obtained by drawing σ according to λ_i and mapping it to the assignment $\alpha_i^\sigma|_{Y'_j}$. Hence, κ_j satisfies the equations (11) and (13).
- The equations (12) hold trivially if ϕ'_j is a null constraint. Otherwise, $j \in C_i$ for some $i \in [q]$. This implies that $X_i \in X_{(\leq k)}(Y'_j)$, and by the claim that $\kappa_j = \mu_j^i$. Then, $\alpha \in \text{supp}(\kappa_j)$ implies that there is a $\sigma \in \text{supp}(\lambda_i)$ such that $\alpha_i^\sigma|_{Y'_j} = \alpha$. By condition (b) and equation (12) for λ_i , the tuple $\alpha_i^\sigma(\mathbf{y}'_j) \in \text{Feas}(\phi'_j)$, so κ_j satisfies (12).
- To verify that the equations (10) hold, let $s, j \in [p]$ be such that $Y'_s \subseteq Y'_j$, and $\beta: Y'_s \rightarrow D'$. Let $X_i \in X_{(\leq k)}(Y'_j)$. We now have:

$$\sum_{\substack{\alpha: Y'_j \rightarrow D' \\ \alpha|_{Y'_s} = \beta}} \kappa_j(\alpha) = \sum_{\substack{\alpha: Y'_j \rightarrow D' \\ \alpha|_{Y'_s} = \beta}} \mu_j^i(\alpha) = \mu_s^i(\beta) = \kappa_s(\beta),$$

where the second equality follows from $Y'_s \subseteq Y'_j$ and a rearrangement of terms, and the last equality follows from the claim since $Y'_s \subseteq Y'_j \subseteq Y_{X_i}$, so $X_i \in X_{(\leq k)}(Y'_s)$.

We conclude that κ is a feasible solution to the $\text{SA}(k', l')$ -relaxation of J .

Let $i \in [q]$ and note that by the claim, for every $j \in C_i$, we have $\kappa_j = \mu_j^i$. Therefore,

$$\begin{aligned}
\sum_{j \in C_i} \sum_{\alpha \in \text{Feas}(\phi'_j)} \kappa_j(\alpha) \phi'_j(\alpha(\mathbf{y}'_j)) &= \sum_{j \in C_i} \sum_{\alpha \in \text{Feas}(\phi'_j)} \sum_{\substack{\sigma \in \text{supp}(\lambda_i) \\ \alpha_i^\sigma|_{Y'_j} = \alpha}} \lambda_i(\sigma) \phi'_j(\alpha(\mathbf{y}'_j)) \\
&= \sum_{\sigma \in \text{supp}(\lambda_i)} \lambda_i(\sigma) \sum_{j \in C_i} \sum_{\substack{\alpha \in \text{Feas}(\phi'_j) \\ \alpha_i^\sigma|_{Y'_j} = \alpha}} \phi'_j(\alpha(\mathbf{y}'_j)) \\
&= \sum_{\sigma \in \text{supp}(\lambda_i)} \lambda_i(\sigma) \sum_{j \in C_i} \phi'_j(\alpha_i^\sigma(\mathbf{y}'_j)) \leq \sum_{\sigma \in \text{supp}(\lambda_i)} \lambda_i(\sigma) \phi_i(\sigma),
\end{aligned} \tag{21}$$

where the inequality follows from assumption (b). Summing inequality (21) over $i \in [q]$ shows that $\text{Val}_{\text{LP}}(J, \kappa) \leq \text{Val}_{\text{LP}}(I, \lambda)$ and the lemma follows. \square

Lemma 10. *Let Γ be a valued constraint language of finite size and let ϕ be a weighted relation expressible over Γ . Then, $\Gamma \cup \{\phi\} \leq_{\text{SA}} \Gamma$.*

Proof. Let I be an instance of $\text{VCSP}(\Gamma \cup \{\phi\})$ with variables $V = \{x_1, \dots, x_n\}$ and objective function $\phi_I(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i)$, where $\phi_i \in \Gamma \cup \{\phi\}$ and \mathbf{x}_i is such that $X_i \subseteq V$. Let I' be an instance of $\text{VCSP}(\Gamma)$ such that $\phi(x_1, \dots, x_m) = \min_{v_i \in D} \phi_{I'}(x_1, \dots, x_m, v_1, \dots, v_p)$.

For $i \in [q]$ such that $\phi_i \in \Gamma$, let J_i be the instance on variables $Y_i = X_i$ with $\phi_{J_i}(Y_i) = \phi_i(\mathbf{x}_i)$. For $i \in [q]$ such that $\phi_i = \phi$, let \mathbf{v}_i be a copy of the variables v_1, \dots, v_p , and let J_i be the instance on variables $Y_i = X_i \cup V_i$ with objective function $\phi_{J_i}(Y_i) = \phi_{I'}(\mathbf{x}_i, \mathbf{v}_i)$. Let J be the $\text{VCSP}(\Gamma)$ instance with variables $\bigcup_i Y_i$ and objective function $\sum_i \phi_{J_i}$.

We verify properties (a)–(c) of Lemma 9.

(a) Let α be any satisfying assignment of J and define $\sigma^\alpha = \alpha|_V$. For $i \in [q]$ such that $\phi_i \in \Gamma$, we have $\phi_i(\sigma^\alpha(\mathbf{x}_i)) = \phi_{J_i}(\alpha(Y_i)) < \infty$. For $i \in [q]$ such that $\phi_i = \phi$, we have $\phi_i(\sigma^\alpha(\mathbf{x}_i)) \leq \phi_{J_i}(\alpha(Y_i)) = \phi_{I'}(\alpha(\mathbf{x}_i, \mathbf{v}_i)) < \infty$. Summing over all $i \in [q]$ gives $\text{Val}_{\text{VCSP}}(I, \sigma^\alpha) \leq \text{Val}_{\text{VCSP}}(J, \alpha) < \infty$.

Let $k \geq \text{ar}(\Gamma \cup \{\phi\})$ and suppose that λ is a feasible solution to the $\text{SA}(k, 2k)$ -relaxation of I . For all $i \in [q]$ and $\sigma: X_i \rightarrow D$ with positive support in λ , define α_i^σ as follows. If $\phi_i \in \Gamma$, then define $\alpha_i^\sigma = \sigma$. Otherwise, $\phi_i = \phi$. Let $\gamma_i^\sigma: V_i \rightarrow D$ be any assignment such that $\phi_i(\sigma(\mathbf{x}_i)) = \phi_J(\sigma(\mathbf{x}_i), \gamma_i^\sigma(\mathbf{v}_j))$, and define $\alpha_i^\sigma: X_i \cup V_i \rightarrow D$ by $\alpha_i^\sigma = \sigma \cup \gamma_i^\sigma$.

(b) For all $i \in [q]$, $\text{Val}_{\text{VCSP}}(J_i, \alpha_i^\sigma) = \phi_{J_i}(\alpha_i^\sigma(Y_i)) = \phi_i(\sigma(\mathbf{x}_i)) < \infty$, where the equalities hold by construction, and the inequality follows from the feasibility of λ .

(c) Let $i, r \in [q]$ and $X \subseteq V$ be as in the lemma and suppose that $\sigma: X \rightarrow D$ has positive support in λ . If $i = r$, then there is nothing to show. Otherwise, $Y_i \cap Y_r = X_i \cap X_r$, so $\alpha_i^{\sigma_i}|_{Y_i \cap Y_r} = \sigma_i|_{X_i \cap X_r} = \sigma|_{X_i \cap X_r} = \sigma_r|_{X_i \cap X_r} = \alpha_r^{\sigma_r}|_{Y_i \cap Y_r}$.

It follows that Lemma 9 is applicable, so $\Gamma \cup \{\phi\} \leq_{\text{SA}} \Gamma$. \square

Lemma 11. *Let Γ be a valued constraint language of finite size over domain D . Then, $\Gamma \cup \{\phi_{=}^D\} \leq_{\text{SA}} \Gamma$.*

Proof. Let I be an instance of $\text{VCSP}(\Gamma \cup \{\phi_{=}^D\})$ with variables $V = \{x_1, \dots, x_n\}$ and objective function $\phi_I(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i)$, where $\phi_i \in \Gamma \cup \{\phi_{=}^D\}$ and \mathbf{x}_i is such that $X_i \subseteq V$. Define the undirected graph $G = (V, E)$, where E contains an edge between u and v if, and only if, there is a constraint $\phi_{=}^D(u, v)$ in I . Let \sim be the equivalence relation on V defined by $u \sim v$ if u

and v are in the same connected component of G . For $v \in V$, let \tilde{v} denote the equivalence class of \sim containing v . For a tuple of variables $\mathbf{x} = (v_1, \dots, v_m) \in V^m$, define $\tilde{\mathbf{x}} = (\tilde{v}_1, \dots, \tilde{v}_m)$.

For $i \in [q]$, let $\mathbf{y}_i = \tilde{\mathbf{x}}_i$ and let J_i be an instance on variables Y_i . If $\phi_i \in \Gamma$, then let the objective function be $\phi_{J_i}(Y_i) = \phi_i(\mathbf{y}_i)$. Otherwise, let ϕ_{J_i} be a null-constraint on Y_i . Let J be the VCSP(Γ) instance with variables $\bigcup_i Y_i$ and objective function $\sum_i \phi_{J_i}$.

We verify properties (a)–(c) of Lemma 9.

(a) Let α be satisfying assignment of J and define $\sigma^\alpha(v) = \alpha(\tilde{v})$ for all $v \in V$. It is clear that $\phi_i(\sigma^\alpha(\mathbf{x}_i)) = \phi_{J_i}(\alpha(Y_i))$ for all $i \in [q]$. Summing over all i gives $\text{Val}_{\text{VCSP}}(I, \sigma^\alpha) = \text{Val}_{\text{VCSP}}(J, \alpha)$.

Let $k \geq \text{ar}(\Gamma \cup \{\phi_\perp^D\}) \geq 2$ and suppose that λ is a feasible solution to the SA($k, 2k$)-relaxation of I . We claim that, for all $i \in [q]$ and $\sigma: X_i \rightarrow D$ with positive support in λ ,

$$u \sim v \implies \sigma(u) = \sigma(v). \quad (22)$$

For $\tilde{v} \in Y_i$, let $\alpha_i^\sigma(\tilde{v}) = \sigma(u)$ for some $u \in \tilde{v} \cap X_i$. By the claim, the definition of α_i^σ is actually independent of the choice of $u \in \tilde{v} \cap X_i$. The justification of the claim follows at the end of the proof.

(b) For all $i \in [q]$, $\text{Val}_{\text{VCSP}}(J_i, \alpha_i^\sigma) = \phi_{J_i}(\alpha_i^\sigma(Y_i)) = \phi_i(\sigma(\mathbf{x}_i)) < \infty$, where the second equality holds by (22), and the inequality follows from the feasibility of λ .

(c) Let $i, r \in [q]$ and $X \subseteq V$ be as in the lemma and suppose that $\sigma: X \rightarrow D$ has positive support in λ . Let $\tilde{v} \in Y_i \cap Y_r$, let $v_1 \in \tilde{v} \cap X_i$ and $v_2 \in \tilde{v} \cap X_r$. By (22), $\sigma(v_1) = \sigma(v_2)$, so $\alpha_i^{\sigma_i}(\tilde{v}) = \sigma_i(v_1) = \sigma(v_1) = \sigma(v_2) = \sigma_r(v_2) = \alpha_r^{\sigma_r}(\tilde{v})$.

It remains to prove that (22) holds for all $\sigma: X_i \rightarrow D$ with positive support in λ . The proof is by induction over the length of a shortest path, $u = u_0, \dots, u_d = v$, between u and v in the graph G . If $d = 0$, then $u = v$, so there is nothing to prove. Assume therefore that $d > 0$ and that the claim holds for all assignments with positive support and all $u' \sim v'$ with a shortest path of length strictly smaller than d . Let $X' = \{u_0, u_d\}$ and note that since $|X'| = 2 \leq k$, there exists an assignment $\tau': X' \rightarrow D$ with positive support in λ such that $\tau' = \sigma|_{X'}$. Now, let $X = X' \cup \{u_{d-1}\}$. Since $|X| \leq 3 \leq 2k$, it follows that λ has a distribution over assignments to X , so there exists an assignment $\tau: X \rightarrow D$ with positive support in λ such that $\tau|_{X'} = \tau' = \sigma|_{X'}$. In particular, $\tau(u_0) = \sigma(u_0)$ and $\tau(u_d) = \sigma(u_d)$.

By assumption, there is an equality constraint on $X'' = \{u_{d-1}, u_d\}$ in J , so any assignment $\tau'': X'' \rightarrow D$ with positive support in λ must have $\tau''(u_{d-1}) = \tau''(u_d)$. Since equation (10) holds for $X'' \subseteq X$, it follows that $\tau|_{X''}$ has positive support in λ , and hence $\tau(u_{d-1}) = \tau(u_d)$.

By the induction hypothesis applied to τ and the path u_0, \dots, u_{d-1} , we now have $\sigma(u_0) = \tau(u_0) = \tau(u_{d-1}) = \tau(u_d) = \sigma(u_d)$. It follows that Lemma 9 is applicable, so $\Gamma \cup \{\phi\} \leq_{\text{SA}} \Gamma$. \square

Lemma 12. *Let Γ and Δ be constraint languages of finite size and assume that Γ interprets Δ . Then, $\Delta \leq_{\text{SA}} \Gamma$.*

Proof. Let (d, S, h) be an interpretation of Δ in Γ . By Lemma 10 and Lemma 11, we may assume that Γ contains the d -ary weighted relation ϕ_S , and for each $\phi_i \in \Delta$, that Γ contains $h^{-1}(\phi_i)$.

Let I be an instance of VCSP(Δ) with variables $V = \{x_1, \dots, x_n\}$ and objective function $\phi_I(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i)$. Assume that ϕ_I contains a distinguished unary null-constraint for each singleton subset $\{v_j\} \subseteq V$, i.e. that for each $1 \leq j \leq n$, there exists an $i \in [q]$ such that ϕ_i is a null constraint, and $\mathbf{x}_i = (x_j)$. Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be d -tuples of distinct variables and let V_j be the set of variables in \mathbf{v}_j , for each $1 \leq j \leq n$, and let $V' = \bigcup_{j=1}^n V_j$. Define $\hat{x}_j = \mathbf{v}_j$ and extend this to tuples: For $\mathbf{x} = (x_{i_1}, \dots, x_{i_m}) \in V^m$, let $\hat{\mathbf{x}} = (\mathbf{v}_{i_1}, \dots, \mathbf{v}_{i_m})$.

For $i \in [q]$, let J_i be an instance on variables $Y_i = \bigcup_{j \in X_i} V_j$. If $\phi_i(x_j)$ is one of the distinguished null-constraints, then let $\phi_{J_i}(Y_i) = \phi_S(\mathbf{v}_j)$. Otherwise, let $\phi_{J_i}(Y_i) = h^{-1}(\phi_i)(\hat{\mathbf{x}}_i)$. Let J be the VCSP(Γ) instance with variables $\bigcup_i Y_i$ and objective function $\sum_i \phi_{J_i}$.

We verify properties (a)–(c) of Lemma 9.

(a) Let α be satisfying assignment of J and define $\sigma^\alpha(x_j) = a$ if $h(\alpha(\mathbf{v}_j)) = a$. This is well-defined since $\alpha(\mathbf{v}_j)$ must be in S . By construction, all constraints in I are satisfied.

For each $x_j \in V$ and $a \in D$, let $\tau_{j,a}$ be an assignment such that $\tau_{j,a}(\mathbf{v}_j) \in S$. Let $k \geq \text{ar}(\Delta)$ and suppose that λ is a feasible solution to the $\text{SA}(k, 2k)$ -relaxation of I . For all $i \in [q]$ and $\sigma: X_i \rightarrow D$ with positive support in λ , define $\alpha_i^\sigma = \bigcup_{j: x_j \in X_i} \tau_{j, \sigma(x_j)}$.

(b) For all $i \in [q]$, $\text{Val}_{\text{VCSP}}(J_i, \alpha_i^\sigma) = \phi_{J_i}(\alpha_i^\sigma(Y_i)) = \phi_i(\sigma(\mathbf{x}_i)) < \infty$, where the equalities hold by construction, and the inequality follows from the feasibility of λ .

(c) Let $i, r \in [q]$ and $X \subseteq V$ be as in Lemma 9 and suppose that $\sigma: X \rightarrow D$ has positive support in λ . Let $v \in Y_i \cap Y_r$ and let \mathbf{v}_j be the tuple of variables that contains v . Then, $x_j \in X_i \cap X_r$, so $\alpha_i^{\sigma_i}(v) = \tau_{j, \sigma(x_j)}(v) = \alpha_r^{\sigma_r}(v)$.

It follows that Lemma 9 is applicable, so $\Delta \leq_{\text{SA}} \Gamma$. \square

Lemma 13. *Let Γ be a valued constraint language of finite size and $\phi \in \Gamma$. Then, $\Gamma \cup \{\text{Opt}(\phi)\} \leq_{\text{SA}} \Gamma$.*

Proof. To avoid trivial cases, we will assume that all weighted relations in Γ take at least one finite value. Let I be an arbitrary instance of $\text{VCSP}(\Gamma \cup \{\text{Opt}(\phi)\})$ with variables V and objective function $\sum_{i=1}^q \phi_i(\mathbf{x}_i)$. We create an instance J of $\text{VCSP}(\Gamma)$ as follows. The variables of J are the same as the variables in I . Every weighted constraint $\phi_i(\mathbf{x}_i)$ in I , where $\phi_i \neq \text{Opt}(\phi)$ appears also in J . Every weighted constraint $\text{Opt}(\phi)(\mathbf{x}_i)$ is replaced by C copies of $\phi(\mathbf{x}_i)$ in J , where the value of the constant C is chosen as follows: If ϕ only takes a single distinct finite value, then let $C = 1$. Otherwise, let $L = \sum_{i=1}^q \min(\phi_i)$ and $U = \sum_{i=1}^q \max(\phi_i)$, where $\min(\phi_i)$ and $\max(\phi_i)$ select, respectively, the smallest and largest *finite* value of the weighted relation ϕ_i . Let δ be the difference between the next to least and the least value of ϕ . Now, let $C = \lceil (U - L + 1)/\delta \rceil$. Both U and L can be computed in polynomial time and the value of C depends linearly on the number of constraints in I , so the size of J is polynomial in the size of I .

First, we prove that $\text{Opt}_{\text{VCSP}}(J)$ determines $\text{Opt}_{\text{VCSP}}(I)$. Any satisfying assignment to I is also a satisfying assignment to J , so

$$\text{Opt}_{\text{VCSP}}(J) \leq CN \min(\phi) + \text{Opt}_{\text{VCSP}}(I), \quad (23)$$

where N is the number of occurrences of $\text{Opt}(\phi)$ in I .

If J has a satisfying assignment σ , then we distinguish two cases. First, assume that σ assigns an optimal value to every copy of ϕ . Then, σ is also a satisfying assignment of I , so

$$\text{Opt}_{\text{VCSP}}(I) \leq \text{Val}_{\text{VCSP}}(I, \sigma) = \text{Val}_{\text{VCSP}}(J, \sigma) - CN \min(\phi). \quad (24)$$

From (23) and (24), we see that if σ is an optimal assignment to J , so that $\text{Val}_{\text{VCSP}}(J, \sigma) = \text{Opt}_{\text{VCSP}}(J)$, then $\text{Val}_{\text{VCSP}}(I, \sigma) = \text{Opt}_{\text{VCSP}}(J) - CN \min(\phi) \leq \text{Opt}_{\text{VCSP}}(I)$, so σ is also an optimal assignment to I .

Otherwise, σ assigns a sub-optimal value to at least C copies of ϕ , so

$$\text{Val}_{\text{VCSP}}(J, \sigma) \geq CN \min(\phi) + C\delta + \text{Opt}_{\text{VCSP}}(I) \geq CN \min(\phi) + (U - L + 1) + L \geq CN \min(\phi) + U + 1.$$

In this case, $\text{Opt}_{\text{VCSP}}(J) > CN \min(\phi) + U \geq CN \min(\phi) + \text{Opt}_{\text{VCSP}}(I)$, so (23) does not hold, and hence I is unsatisfiable.

In summary, if J is unsatisfiable, or if $\text{Opt}_{\text{VCSP}}(J) > CN \min(I) + U$, then I is unsatisfiable, and otherwise $\text{Opt}_{\text{VCSP}}(I) = \text{Opt}_{\text{VCSP}}(J) - CN \min(\phi)$.

Next, we prove that, for any given parameters $1 \leq k \leq \ell$, if $\Gamma \cup \{\text{Opt}(\phi)\}$ does not have valued relational width (k, ℓ) , then Γ does not have valued relational width (k, ℓ) . Let I be an

instance of $\text{VCSP}(\Gamma \cup \{\text{Opt } \phi\})$ and λ a feasible solution to the $\text{SA}(k, \ell)$ -relaxation of I , with $\text{Val}_{\text{LP}}(I, \lambda) < \text{Opt}_{\text{VCSP}}(I)$, where $\text{Opt}_{\text{VCSP}}(I)$ could be ∞ . We will assume that I has been augmented with null constraints so that, for every subset $V' \subseteq V$ with $|V'| \leq \ell$, there is some $i \in [q]$ with $X_i = V'$. Let J be the instance of $\text{VCSP}(\Gamma)$ constructed above.

Let λ' be the feasible solution to the $\text{SA}(k, \ell)$ -relaxation of J obtained from λ by letting $\lambda'_j = \lambda_i$ for all ϕ -constraints of J with index j that were introduced as copies of the $\text{Opt}(\phi)$ -constraint of I with index i . Then, λ' assigns an optimal value to each ϕ -constraint, so $\text{Val}_{\text{LP}}(J, \lambda') = \text{Val}_{\text{LP}}(I, \lambda) + CN \min(\phi)$.

If J is unsatisfiable, then $\text{Opt}_{\text{LP}}(J) < \text{Opt}_{\text{VCSP}}(J)$, so Γ does not have valued relational width (k, ℓ) . If J is satisfiable, then it was shown above that I is satisfiable and $\text{Opt}_{\text{VCSP}}(J) = \text{Opt}_{\text{VCSP}}(I) + CN \min(\phi)$, so $\text{Opt}_{\text{LP}}(J) \leq \text{Opt}_{\text{LP}}(I) + CN \min(\phi) < \text{Opt}_{\text{VCSP}}(I) + CN \min(\phi) = \text{Opt}_{\text{VCSP}}(J)$, and again Γ does not have valued relational width (k, ℓ) . Since k and ℓ were chosen arbitrarily, the result follows. \square

Lemma 14. *Let Γ be a valued constraint language of finite size and $\phi \in \Gamma$. Then, $\Gamma \cup \{\text{Feas}(\phi)\} \leq_{\text{SA}} \Gamma$.*

Proof. To avoid trivial cases, we will assume that all weighted relations in Γ take at least one finite value. Let I be an arbitrary instance of $\text{VCSP}(\Gamma \cup \{\text{Feas}(\phi)\})$ with variables V and objective function $\sum_{i=1}^q \phi_i(\mathbf{x}_i)$. We create an instance J of $\text{VCSP}(\Gamma)$ as follows. The variables of J are the same as the variables in I . For every weighted constraint $\phi_i(\mathbf{x}_i)$ in I with $\phi_i \in \Gamma$, we add C copies of $\phi_i(\mathbf{x}_i)$ in J . Every weighted constraint $\text{Feas}(\phi)(\mathbf{x}_i)$ is replaced by $\phi(\mathbf{x}_i)$ in J . The value of the constant C is chosen as follows: If ϕ only takes a single distinct finite value, then let $C = 1$. Otherwise, let L be the smallest value of ϕ and U the largest *finite* value of ϕ . Let δ be less than or equal to the least possible difference between any two satisfying assignments of I . Now, let $C = \lceil N(U - L + 1)/\delta \rceil$, where N is the number of occurrences of $\text{Feas}(\phi)$ in I . Both U and L can be computed in polynomial time. As a bound on the least possible difference between any two satisfying assignments of I , we take $\delta = 1/M$ where $M > 0$ is any constant such that $M \cdot \phi_i$ is an integer for each i . The value of C depends linearly on the number of constraints in I , so the size of J is polynomial in the size of I .

An assignment $\sigma: V \rightarrow D$ satisfies I if, and only if, it satisfies J , and

$$C \cdot \text{Val}_{\text{VCSP}}(I, \sigma) + NL \leq \text{Val}_{\text{VCSP}}(J, \sigma) \leq C \cdot \text{Val}_{\text{VCSP}}(I, \sigma) + NU, \quad (25)$$

Let σ be an optimal assignment to J and suppose that there exists an assignment σ' to I such that $\text{Val}_{\text{VCSP}}(I, \sigma') < \text{Val}_{\text{VCSP}}(I, \sigma)$. Then,

$$\begin{aligned} \text{Val}_{\text{VCSP}}(J, \sigma') &\leq C \cdot \text{Val}_{\text{VCSP}}(I, \sigma') + NU \\ &\leq C \cdot (\text{Val}_{\text{VCSP}}(I, \sigma) - \delta) + NU \\ &\leq C \cdot \text{Val}_{\text{VCSP}}(I, \sigma) + NL + N(U - L) - C \cdot \delta \\ &< \text{Val}_{\text{VCSP}}(J, \sigma), \end{aligned}$$

which contradicts σ being optimal. Hence, σ is also an optimal assignment to I .

Next, we prove that, for any given parameters $1 \leq k \leq \ell$, if $\Gamma \cup \{\text{Feas}(\phi)\}$ does not have valued relational width (k, ℓ) , then Γ does not have valued relational width (k, ℓ) . Let I be an instance of $\text{VCSP}(\Gamma \cup \{\text{Feas}(\phi)\})$ and λ a feasible solution to the $\text{SA}(k, \ell)$ -relaxation of I with $\text{Val}_{\text{LP}}(I, \lambda) < \text{Opt}_{\text{VCSP}}(I)$. We will assume that I has been augmented with null constraints so that, for every subset $V' \subseteq V$ with $|V'| \leq \ell$, there is some $i \in [q]$ with $X_i = V'$. If I is unsatisfiable, let J be the instance of $\text{VCSP}(\Gamma)$ constructed as above. Otherwise, let $\epsilon = \text{Opt}_{\text{VCSP}}(I) - \text{Val}_{\text{LP}}(I, \lambda) > 0$, and let J be the instance constructed as above, but with $C = \max\{\lceil N(U - L + 1)/\delta \rceil, \lceil N(U - L + 1)/\epsilon \rceil\}$.

Let λ' be the feasible solution to the $\text{SA}(k, \ell)$ -relaxation of J obtained from λ by letting $\lambda'_j = \lambda_i$ for every constraint of J with index j that was introduced as a (possibly single) copy of the constraint $\phi_i(\mathbf{x}_i)$ of I . Then, $\text{Val}_{\text{LP}}(J, \lambda') \leq C \cdot \text{Val}_{\text{LP}}(I, \lambda) + NU$.

The instance J is unsatisfiable if, and only if, I is unsatisfiable, and in this case, $\text{Opt}_{\text{LP}}(J) < \text{Opt}_{\text{VCSP}}(J)$, so Γ does not have valued relational width (k, ℓ) . Otherwise, J is satisfiable, and $\text{Opt}_{\text{VCSP}}(J) \geq C \cdot \text{Opt}_{\text{VCSP}}(I) + NL$, so $\text{Opt}_{\text{LP}}(J) \leq C \cdot \text{Opt}_{\text{LP}}(I) + NU < C \cdot (\text{Opt}_{\text{VCSP}}(I) - \epsilon) + NU \leq C \cdot (\text{Opt}_{\text{VCSP}}(I) - \epsilon) + NL + N(U - L) \leq \text{Opt}_{\text{VCSP}}(J) + N(U - L) - C \cdot \epsilon < \text{Opt}_{\text{VCSP}}(J)$, so Γ does not have valued relational width (k, ℓ) . Since k and ℓ were chosen arbitrarily, the result follows. \square

Lemma 15. *Let Γ be a valued constraint language of finite size over domain D and let Γ' be a core of Γ with $D' \subseteq D$. Then, $\Gamma' \cup \mathcal{C}_{D'} \leq_{\text{SA}} \Gamma$.*

Proof. Let I' be an instance of $\text{VCSP}(\Gamma')$, and let I be the instance of $\text{VCSP}(\Gamma)$ obtained from I' by substituting every restricted weighted relation in Γ' by its corresponding weighted relation in Γ . Then, by Lemma 3, $\text{Opt}_{\text{VCSP}}(I') = \text{Opt}_{\text{VCSP}}(I)$. Fix $1 \leq k \leq \ell$ and assume that I' is a gap instance for the $\text{SA}(k, \ell)$ -relaxation of $\text{VCSP}(\Gamma')$. Then, $\text{Opt}_{\text{LP}}(I) \leq \text{Opt}_{\text{LP}}(I') < \text{Opt}_{\text{VCSP}}(I') = \text{Opt}_{\text{VCSP}}(I)$, where the first inequality follows from the fact that I' is a restriction of I . This establishes $\Gamma' \leq_{\text{SA}} \Gamma$.

Let F be the set of unary operations on D' that are not in $\text{supp}(\Gamma')$ and apply Lemma 7 to Γ' and F . This provides a crisp core constraint language Δ on D' such that $\Delta \leq_{\text{SA}} \Gamma'$ and such that every unary operation in $\text{Pol}(\Delta)$ is also in $\text{supp}(\Delta)$. We finish the proof by showing that $\Gamma' \cup \mathcal{C}_{D'} \leq_{\text{SA}} \Gamma' \cup \Delta$, using Lemma 9.

Let I_Δ be the instance on variables $V_\Delta = \{x_a \mid a \in D\}$ and containing, for every $\phi \in \Delta$, and $\mathbf{a} \in D^{\text{ar}(\phi)}$, a constraint $\phi(\mathbf{x}_\mathbf{a})$, where $\mathbf{x}_\mathbf{a}[i] = x_{a[i]}$ for $1 \leq i \leq \text{ar}(\phi)$. Every satisfying assignment α to I_Δ defines an operation $f_\alpha: D \rightarrow D$ by the map $a \mapsto \alpha(x_a)$. The instance I_Δ is sometimes called the *indicator instance* [35] and has the following property:

α is a satisfying assignment of I_Δ if, and only if, f_α is a unary polymorphism of Δ . (26)

Let I be an arbitrary instance of $\text{VCSP}(\Gamma' \cup \mathcal{C}_{D'})$ with variables $V = \{x_1, \dots, x_n\}$ and objective function $\phi_I(x_1, \dots, x_n) = \sum_{i=1}^q \phi_i(\mathbf{x}_i)$, where $\phi_i \in \Gamma' \cup \mathcal{C}_{D'}$ and \mathbf{x}_i such that $X_i \subseteq V$. Assume without loss of generality that $V \cap V_\Delta = \emptyset$. For $v \in V$, define $\hat{v} := x_a$ if there is a unary constraint $v = a$ in I , and define $\hat{v} := v$, otherwise. For a tuple of variables $\mathbf{x} = (v_1, \dots, v_m) \in V^m$, define $\hat{\mathbf{x}} = (\hat{v}_1, \dots, \hat{v}_m)$.

For $i \in [q]$ such that $\phi_i \in \Gamma'$, let $\mathbf{y}_i = \hat{\mathbf{x}}_i$ and let J_i be the instance on variables Y_i with objective function $\phi_{J_i}(Y_i) = \phi_i(\mathbf{y}_i)$. For $i \in [q]$ such that $\phi_i(x_i)$ is a unary constraint $x_i = a$, let J_i be the instance I_Δ on variables $Y_i = V_\Delta$. Note that each J_i corresponding to a unary constraint $x_i = a$ is the same instance I_Δ on the same variables V_Δ . Let J be the $\text{VCSP}(\Gamma' \cup \Delta)$ instance with variables $\bigcup_i Y_i$ and objective function $\sum_i \phi_{J_i}$.

We verify properties (a)–(c) of Lemma 9.

(a) Let α be an optimal assignment to J and consider the operation f_α in (26) obtained from the unique copy of I_Δ in J . Since the unary operations in $\text{Pol}(\Delta)$ are bijections and closed under composition, it follows that f_α^{-1} is also in $\text{Pol}(\Delta)$ and therefore in $\text{supp}(\Gamma)$. Hence, by Lemma 2, $\beta := f_\alpha^{-1} \circ \alpha$ is also an optimal assignment to J and f_β is the identity operation. We define $\sigma^\alpha(x) = a$ if $\hat{x} = x_a$ for some $a \in D'$, and $\sigma^\alpha(x) = \beta(x)$ otherwise. All unary constraints $x = a$ in I are satisfied by σ^α and all other constraints take the same value as in J , hence $\text{Val}_{\text{VCSP}}(I, \sigma^\alpha) = \text{Val}_{\text{VCSP}}(J, \alpha)$.

Let $k \geq \text{ar}(\Gamma' \cup \mathcal{C}_{D'})$ and suppose that λ is a feasible solution to the $\text{SA}(k, 2k)$ -relaxation of I . Let γ be the satisfying assignment of I_Δ that assigns a to x_a for all $a \in D'$. For all $i \in [q]$ and $\sigma: X_i \rightarrow D$ with positive support in λ , define $\alpha_i^\sigma = (\sigma \cup \gamma)|_{Y_i}$.

(b) For all $i \in [q]$, $\text{Val}_{\text{VCSP}}(J_i, \alpha_i^\sigma) = \phi_{J_i}(\alpha_i^\sigma(Y_i)) = \phi_i(\sigma(\mathbf{x}_i)) < \infty$, where the equalities hold by construction, and the inequality follows from the feasibility of λ .

(c) Let $i, r \in [q]$ and $X \subseteq V$ be as in the lemma and suppose that $\sigma: X \rightarrow D$ has positive support in λ . Let $y \in Y_i \cap Y_r$. If $y = x_a$ for some $a \in D'$, then $\alpha_i^{\sigma_i}(y) = \alpha_r^{\sigma_r}(y) = \gamma(x_a) = a$. Otherwise, $y \in X_i \cap X_r$, so $\alpha_i^{\sigma_i}(y) = \sigma_i(y) = \sigma(y) = \sigma_r(y) = \alpha_r^{\sigma_r}(y)$.

It follows that Lemma 9 is applicable, so $\Gamma' \cup \mathcal{C}_{D'} \leq_{\text{SA}} \Gamma' \cup \Delta$. \square

7 Gap Instances for SA-relaxations of $\text{VCSP}(E_{\mathcal{G},3})$

In this section, we give a construction of gap instances for SA-relaxations of $\text{VCSP}(E_{\mathcal{G},3})$, which shows that $E_{\mathcal{G},3}$ does not have bounded valued relational width. This result can also be derived from results in [56] using additional non-trivial results. We provide here a direct, elementary proof for constant level LP relaxations, whereas [56] deals with linear level SDP relaxations.

Let \mathcal{G} be an Abelian group over a finite set G and let g be a non-zero element in G . Let $R_0 = \{(x, y, z) \in G^3 \mid x = y + z + 0\}$ and $R_g = \{(x, y, z) \in G^3 \mid x = y + z + g\}$ and $\Delta = \{R_0, R_g\}$. Both R_0 and R_g are expressible in $E_{\mathcal{G},3}$, so by Theorem 16(1), it suffices to prove that Δ does not have bounded valued relational width.

Let $k \geq 3$. We construct an unsatisfiable instance I of $\text{VCSP}(\Delta)$ and a feasible solution to its $\text{SA}(k, k)$ -relaxation. The construction is similar to the one in [28, Theorem 31] where it is used to show that constraint languages without “the ability to count” do not have bounded width. Our theorem is a strengthening of this result.

Let $n \geq 1$ be a positive integer. Let $T_{n \times n}$ be the torus grid graph on $n \times n$ vertices resulting from taking the square grid graph on $(n+1) \times (n+1)$ vertices and identifying the topmost with the bottommost vertices as well as the leftmost with the rightmost vertices.




Figure 1: Variables in the torus $T_{3,3}$ obtained from the 4×4 grid graph.

The instance I_n contains one variable for each vertex and one variable for each edge in $T_{n \times n}$. For $0 \leq a, b < n$, let $x_{a,b}$, $y_{a,b}$, and $z_{a,b}$, be the variables corresponding to vertices, horizontal edges, and vertical edges, respectively; cf., Figure 1. Let I_n contain the following constraints:

$$y_{a,b+1} = y_{a,b} + x_{a,b} + c_{a,b} \quad (27)$$

$$z_{a+1,b} = z_{a,b} + x_{a,b} + d_{a,b}, \quad (28)$$

where indices are taken modulo n , and the elements $c_{a,b}, d_{a,b} \in \{0, g\}$ are chosen so that

$$\sum_{a,b} c_{a,b} - \sum_{a,b} d_{a,b} = g. \quad (29)$$

The following result establishes Theorem 14. We note that it actually shows that I_n is a gap instance for $\text{SA}(k(n), k(n))$, with $k = \Theta(\sqrt{n})$.

Theorem 17. *For every $k \geq 3$ and $n > 2k$, the instance I_n is a gap instance for $\text{SA}(k, k)$.*

Proof. The instance I_n is unsatisfiable by construction: Summing the equations (27) over a and b and simplifying implies the equation $0 = \sum_{a,b} (x_{a,b} + c_{a,b})$. Similarly, the equations (28) imply $0 = \sum_{a,b} (x_{a,b} + d_{a,b})$. By taking the difference of these two equations, it follows that $0 = \sum_{a,b} (c_{a,b} - d_{a,b}) = g$ by (29), a contradiction. Hence, the constraints of I_n cannot be simultaneously satisfied.

Next, we define a feasible solution to the $\text{SA}(k, k)$ -relaxation of I_n . Denote by V the set all variables of I_n . Let X_1, \dots, X_m be an enumeration of all subsets $X \subseteq V$ with $|X| \leq k$. For each j , we define a set $\bar{X}_j \subseteq V$ of variables as follows: Let $V_x = \{x_{a,b} \mid 0 \leq a, b < n\}$. Let S be the smallest superset of the vertices in $V_x \cap X_j$ such that if $y_{a,b} \in X_j$, then S contains $x_{a,b}$ and $x_{a,b+1}$, and if $z_{a,b} \in X_j$, then S contains $x_{a,b}$ and $x_{a+1,b}$. Since $|S| \leq 2k < n$, it follows that there is an a' such that $x_{a',b} \notin S$ for all $0 \leq b < n$, and a b' such that $x_{a,b'} \notin S$ for all $0 \leq a < n$. This implies that a cycle in S can only “enclose” relatively small connected components in $T_{n \times n}[V_x \setminus S]$ with a total number of vertices less than k^2 . These small components will be referred to as *holes*. In addition $T_{n \times n}[V_x \setminus S]$ always contains one large component with more than $n^2 - 2k - k^2 > \frac{3}{4}n^2 - n$ vertices. Let S' be the union of S and the vertices of the holes and define \bar{X}_j to the vertices and edges of the subgraph induced by S' in $T_{n \times n}$.

For $i \in [m]$, define N_i to be the set of assignments $\bar{\sigma} : \bar{X}_i \rightarrow G$ that satisfy every constraint in I_n whose scope is contained in \bar{X}_i , and let U_i be the uniform distribution on N_i . Let λ_i be the following probability distribution on assignments $\sigma : X_i \rightarrow G$:

$$\lambda_i(\sigma) = \Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_{X_i} = \sigma]. \quad (30)$$

The result follows by showing that λ is a feasible solution to the $\text{SA}(k, k)$ -relaxation of I_n , which is shown below in Lemma 16. \square

Lemma 16. *For $i \in [m]$, let λ_i be the following probability distribution on X_i :*

$$\lambda_i(\sigma) = \Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_{X_i} = \sigma], \quad (31)$$

where U_i is the uniform distribution on N_i . Then, λ is a feasible solution to the $\text{SA}(k, k)$ -relaxation of I_n .

Note that $X_j \subseteq X_i \implies \bar{X}_j \subseteq \bar{X}_i$. However, we will need a stronger property, namely that it is possible to move within the set $\{\bar{X}_i\}_{i=1}^m$ (in particular without creating holes) from \bar{X}_j to \bar{X}_i by adding vertices one at a time. More formally, define the binary relation \rightarrow on $\{\bar{X}_i\}_{i=1}^m$ by letting $\bar{X}_j \rightarrow \bar{X}_i$ if, and only if, $V_x \cap (\bar{X}_i \setminus \bar{X}_j) = \{x_{a,b}\}$ for some $0 \leq a, b < n$. Let \preceq be the reflexive transitive closure of \rightarrow .

Lemma 17. $X_j \subseteq X_i \implies \bar{X}_j \preceq \bar{X}_i$.

Proof. Assume to the contrary that there are i and j such that $X_j \subseteq X_i$ but $\bar{X}_j \not\preceq \bar{X}_i$. Let X be the set of vertices in $\bar{X}_i \setminus \bar{X}_j$. It is easy to see that i and j can be chosen so that for every $x_{a,b} \in X$, removing $x_{a,b}$ from \bar{X}_i creates a hole. Therefore, every $x_{a,b} \in X$ must have four neighbours in \bar{X}_i . Consider the set Y of vertices in $T_{n \times n} \setminus X$ that are neighbours to some $x_{a,b} \in X$. If $Y \subseteq \bar{X}_j$, then \bar{X}_j also contains a hole; a contradiction. Otherwise, $Y \setminus \bar{X}_i \neq \emptyset$. But then X contains a vertex with fewer than four neighbours in \bar{X}_i ; again a contradiction. \square

For $\bar{\tau} \in N_j$ and i such that $\bar{X}_j \subseteq \bar{X}_i$, let $N_{j,i}(\bar{\tau})$ denote the set of assignments $\bar{\sigma} \in N_i$ such that $\bar{\tau} = \bar{\sigma}|_{\bar{X}_j}$, i.e. the set of extensions of $\bar{\tau}$ to an assignment in N_i . Next, we give an expression for the size of the sets $N_{j,i}(\bar{\tau})$ that is independent of the choice of $\bar{\tau}$.

A *horizontal component* of \bar{X}_j is a set of edges $\{y_{a,b}, y_{a,b+1}, \dots, y_{a,b+r}\} \subseteq \bar{X}_j$ such that $y_{a,b-1}, y_{a,b+r+1} \notin \bar{X}_j$. A *vertical component* of \bar{X}_j is defined analogously. Let C_j , H_j , and V_j be the number of vertices, horizontal components, and vertical components, respectively, in \bar{X}_j . By the construction of \bar{X}_j , the horizontal components and the vertical components partition the y and z -vertices of \bar{X}_j . Therefore, an assignment is precisely determined by freely choosing the value of every vertex, and of one edge in each horizontal component and one edge in each vertical component:

$$|N_j| = |G|^{C_j+H_j+V_j}. \quad (32)$$

Lemma 18. For $X_j \subseteq X_i$ and all $\bar{\tau} \in N_j$,

$$|N_{j,i}(\bar{\tau})| = \frac{|N_i|}{|N_j|}. \quad (33)$$

Proof. First assume that $\bar{X}_j \rightarrow \bar{X}_i$ and let $x_{a,b}$ be the unique vertex in $\bar{X}_i \setminus \bar{X}_j$. Since \bar{X}_j does not contain a hole, it follows that $x_{a,b}$ must have less than four neighbours in \bar{X}_j . We consider the following three possible cases:

1. $x_{a,b}$ has a single neighbour in \bar{X}_j . Without loss of generality, assume that this neighbour is $x_{a,b+1}$ so that $\bar{X}_i \setminus \bar{X}_j = \{x_{a,b}, y_{a,b}\}$. Choose the value of $x_{a,b}$ arbitrarily. If $y_{a,b+1} \in \bar{X}_j$, then the equation $y_{a,b+1} = y_{a,b} + x_{a,b} + c_{a,b}$ forces the value of $y_{a,b}$. In this case, we have $|G|$ possible extensions and $C_i = C_j + 1$, $H_i = H_j$, and $V_i = V_j$, so (33) holds. Otherwise, $y_{a,b+1} \notin \bar{X}_j$, so the value of $y_{a,b}$ can be chosen arbitrarily. In this case, we have $|G|^2$ possible extensions and $C_i = C_j + 1$, $H_i = H_j + 1$, and $V_i = V_j$, so (33) holds.
2. $x_{a,b}$ has two neighbours in \bar{X}_j . If $x_{a,b}$ has one horizontal and one vertical neighbour, then we can argue as in case (1). Otherwise, without loss of generality, assume that $\bar{X}_i \setminus \bar{X}_j = \{y_{a,b-1}, x_{a,b}, y_{a,b}\}$. We have three possible cases, depending on the size of the intersection $\{y_{a,b-2}, y_{a,b+1}\} \cap \bar{X}_j$. If this intersection contains both y -variables, then the values of $y_{a,b-1}$, $x_{a,b}$, and $y_{a,b}$ are all forced by the equations. In this case we have 1 possible extension, $C_i = C_j + 1$, $H_i = H_j - 1$, and $V_i = V_j$, so (33) holds. If the intersection contains one or zero y -variables, then we can choose the value of $x_{a,b}$ arbitrarily and proceed similarly to case (1).
3. $x_{a,b}$ has three neighbours in \bar{X}_j . This case follows by extending the argument in (2) for two vertical neighbours.

We now prove by induction that the general expression in (33) holds. By Lemma 17, there exists an i' such that $\bar{X}_j \preceq \bar{X}_{i'} \rightarrow \bar{X}_i$. We have just shown that (33) holds for i' , i , and all $\bar{\sigma}' \in N_{i'}$. Assume by induction that (33) holds for j , i' , and all $\bar{\tau} \in N_j$. Then,

$$|N_{j,i}(\bar{\tau})| = \sum_{\bar{\sigma}' \in N_{j,i'}(\bar{\tau})} |N_{i',i}(\bar{\sigma}')| = \sum_{\bar{\sigma}' \in N_{j,i'}(\bar{\tau})} \frac{|N_i|}{|N_{i'}|} = \frac{|N_{i'}|}{|N_j|} \frac{|N_i|}{|N_{i'}|} = \frac{|N_i|}{|N_j|},$$

which proves the lemma. \square

We are now ready to prove Lemma 16.

Proof of Lemma 16. By construction, λ satisfies (11) and (12) for the SA(k, k)-relaxation of I_n . It remains to show that it also satisfies (10).

Let $X_j \subseteq X_i$ and $\tau: X_j \rightarrow G$. Let X be a subset of variables such that $X_j \subseteq X \subseteq \bar{X}_i$. Then,

$$\begin{aligned} \Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_{X_j} = \tau] &= \sum_{\sigma: X \rightarrow G} \Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_{X_j} = \tau \text{ and } \bar{\sigma}|_X = \sigma] \\ &= \sum_{\substack{\sigma: X \rightarrow G \\ \sigma|_{X_j} = \tau}} \Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_X = \sigma]. \end{aligned} \quad (34)$$

For $X = \bar{X}_j$, equation (34) implies the following.

$$\Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_{X_j} = \tau] = \sum_{\substack{\bar{\tau}: \bar{X}_j \rightarrow G \\ \bar{\tau}|_{X_j} = \tau}} \Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_{\bar{X}_j} = \bar{\tau}] = \sum_{\substack{\bar{\tau} \in N_j \\ \bar{\tau}|_{X_j} = \tau}} \frac{|N_{j,i}(\bar{\tau})|}{|N_i|} = \sum_{\bar{\tau}|_{X_j} = \tau} \frac{1}{|N_j|} = \lambda_j(\tau), \quad (35)$$

where the next-to-last inequality follows from Lemma 18. Hence,

$$\lambda_j(\tau) = \Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_{X_j} = \tau] \quad (36)$$

$$= \sum_{\substack{\sigma: X_i \rightarrow G \\ \sigma|_{X_j} = \tau}} \Pr_{\bar{\sigma} \sim U_i} [\bar{\sigma}|_{X_i} = \sigma] \quad (37)$$

$$= \sum_{\substack{\sigma: X_i \rightarrow G \\ \sigma|_{X_j} = \tau}} \lambda_i(\sigma), \quad (38)$$

where (37) follows by (34) with $X = X_i$. It follows that λ satisfies (10), hence it is a feasible solution to the $\text{SA}(k, k)$ -relaxation of I_n . \square

8 Proof of Lemma 5

Lemma (Lemma 5 restated). *Let Γ be a valued constraint language of finite size on domain D and Γ' a core of Γ on domain $D' \subseteq D$. Then, $\text{supp}(\Gamma)$ satisfies the BWC if, and only if, $\text{supp}(\Gamma' \cup \mathcal{C}_{D'})$ satisfies the BWC.*

Proof. Let μ be a unary fractional polymorphism of Γ with an operation g in its support such that $g(D) = D'$. We begin by constructing a unary fractional polymorphism μ' of Γ such that every operation in $\text{supp}(\mu')$ has an image in D' . We will use a technique for generating fractional polymorphisms described in [42, Lemma 10]. It takes a fractional polymorphism, such as μ , a set of *collections* \mathbb{G} , which in our case will be the set of operations in the clone of $\text{supp}(\mu)$, a set of *good* collections \mathbb{G}^* , which will be operations from \mathbb{G} with an image in D' , and an *expansion operator* Exp which assigns to every collection a probability distribution on \mathbb{G} .

The procedure starts by generating each collection $f \in \text{supp}(\mu)$ with probability $\mu(f)$, and subsequently the expansion operation Exp maps $f \in \mathbb{G}$ to the probability distribution that assigns probability $\Pr_{h \sim \mu}[h \circ f = f']$ to each operation $f' \in \mathbb{G}$. The expansion operator is required to be *non-vanishing*, which means that starting from any collection $f \in \mathbb{G}$, repeated expansion must assign non-zero probability to a good collection in \mathbb{G}^* . In our case, this is immediate, since starting from a collection f , the good collection $g \circ f$ gets probability at least $\mu(g)$ which is non-zero by assumption. By [42, Lemma 10], it now follows that Γ has a fractional polymorphism μ' with $\text{supp}(\mu') \subseteq \mathbb{G}^*$. So every operation in $\text{supp}(\mu')$ has an image in D' .

Now, we show that if $\text{supp}(\Gamma)$ contains an m -ary WNU t , then $\text{supp}(\Gamma' \cup \mathcal{C}_{D'})$ also contains an m -ary WNU. Let ω be a fractional polymorphism of Γ with t in its support. Define ω' by $\omega'(f') = \Pr_{h \sim \mu', f \sim \omega}[h \circ f = f']$. Then, ω' is a fractional polymorphism of Γ in which every operation has an image in D' , so ω' is a fractional polymorphism of Γ' . Furthermore, for any

unary operation $h \in \text{supp}(\mu')$, $h \circ t$ is again a WNU, so $\text{supp}(\Gamma')$ contains an m -ary WNU t' . Next, let $h(x) = t'(x, \dots, x)$. Since Γ' is a core, the set of unary operations in $\text{supp}(\Gamma')$ contains only bijections and is closed under composition (Lemma 1). It follows that h has an inverse $h^{-1} \in \text{supp}(\Gamma')$, and since $\text{supp}(\Gamma')$ is a clone, $h^{-1} \circ t'$ is an idempotent WNU in $\text{supp}(\Gamma')$. We conclude that $h^{-1} \circ t' \in \text{supp}(\Gamma' \cup \{\mathcal{C}_{D'}\})$.

For the opposite direction, let t' be an m -ary WNU in $\text{supp}(\Gamma' \cup \{\mathcal{C}_{D'}\})$, and let ω' be a fractional polymorphism of $\Gamma' \cup \{\mathcal{C}_{D'}\}$ with t' in its support. Then, ω' is also a fractional polymorphism of Γ' . Define ω by $\omega(f) = \Pr_{h \sim \mu', f' \sim \omega'}[f'[h, \dots, h] = f]$. Then, ω is a fractional polymorphism of Γ , and, for every $h \in \text{supp}(\mu')$, the operation $t[h, \dots, h]$ is an m -ary WNU in $\text{supp}(\omega)$. We conclude that $t \in \text{supp}(\Gamma)$, which finishes the proof. \square

9 Proofs of Theorems 7 and 10

Theorem (Theorem 7 restated). *Let D be an arbitrary finite domain and let Γ be an arbitrary valued constraint language of finite size on D with $\mathcal{C}_D \subseteq \Gamma$. Assume that Γ expresses a unary finite-valued weighted relation ν that is injective on D . Then, either $\text{supp}(\Gamma)$ satisfies the BWC, in which case Γ has valued relational width $(2, 3)$, or $\text{VCSP}(\Gamma)$ is NP-hard.*

Proof. If Γ satisfies the BWC then the result follows from Theorem 4. If Γ does not satisfy the BWC then, by Lemma 8, there exists a crisp constraint language Δ such that $\text{Pol}(\Delta)$ does not satisfy the BWC and $\Delta \leq_{\text{SA}} \Gamma$. By Theorem 15, there exists a non-trivial Abelian group \mathcal{G} over a finite set G and an interpretation of $E_{\mathcal{G},3}$ in Δ with parameters (d, S, h) . By Theorem 16(3), we have $E_{\mathcal{G},3} \leq_{\text{SA}} \Delta \leq_{\text{SA}} \Gamma$.

Let C be larger than $\max_{a \in D} \nu(a)$. The d -ary weighted relation $\phi(x_1, \dots, x_d)$ defined by

$$\phi(x_1, \dots, x_d) = \nu(x_1) + C\nu(x_2) + C^2\nu(x_3) + \dots + C^{d-1}\nu(x_d)$$

is injective on the set of d -tuples over D . By Theorem 16(1), $\{\phi\} \leq_{\text{SA}} \Gamma$. We define

$$\phi'(x_1, \dots, x_d) = \min_{y_1, \dots, y_d} h^{-1}(\phi_{\text{SA}}^G)(x_1, \dots, x_d, y_1, \dots, y_d) + \phi(y_1, \dots, y_d).$$

By Theorem 16(1) and (3), $\{\phi'\} \leq_{\text{SA}} \Gamma$. Thus, we have an injective unary weighted relation ϕ' on the interpreted $E_{\mathcal{G},3}$. For every $x \in G$, let $h_x \in D^d$ be an arbitrarily chosen element of $h^{-1}(x)$. Finally, define the unary finite-valued weighted relation $\phi'' : G \rightarrow \mathbb{Q}$ by $\phi''(x) = \phi'(h_x)$. (Note that the choice of h_x does not affect the value of $\phi''(x)$.)

We denote by $E'_{\mathcal{G}}$ the crisp constraint language on domain G with, for every $r \geq 1$, $a \in G$, and $\mathbf{c} = (c_1, \dots, c_r) \in \mathbb{Z}^r$ with $\sum_{i=1}^r c_i = 0$, a relation $S'_{a,\mathbf{c}} = \{(x_1, \dots, x_r) \in G^r \mid \sum_{i=1}^r c_i x_i = a\}$. By [60, Theorem 3.18], $\text{VCSP}(E'_{\mathcal{G}} \cup \{\phi''\})$ is APX-hard, and thus NP-hard since ϕ'' is injective and thus non-constant on G . We will finish the proof by showing how to reduce, in polynomial time, any instance I' of $\text{VCSP}(E'_{\mathcal{G}} \cup \{\phi''\})$ to an instance I of $\text{VCSP}(E_{\mathcal{G},3} \cup \{\phi''\})$.

Let V denote the set of variables of I' . The variables of I will include V and a set of new auxiliary variables for each constraint of I' not involving ϕ'' . Let $\phi''(x)$ be a constraint of I' for some $x \in V$. Then we include the constraint $\phi''(x)$ in I . Let $S'_{a,\mathbf{c}}(\mathbf{x})$ be a constraint of I' for some $r \geq 1$, $a \in G$, $\mathbf{c} = (c_1, \dots, c_r) \in \mathbb{Z}^r$ with $\sum_{i=1}^r c_i = 0$, and $\mathbf{x} = (x_1, \dots, x_r) \in V^r$. Since $|G|x = 0$ in \mathcal{G} , for all $x \in G$ we can, without loss of generality, assume that $0 \leq c_i < |G|$. Thus $S'_{a,\mathbf{c}}$ is equivalent to an m -ary relation S' over G where $m = \sum_{i=1}^r c_i \leq r|G|$. The relation S' can be expressed with $O(m)$ relations from $E_{\mathcal{G},3}$ using $O(m)$ auxiliary variables. \square

Theorem (Theorem 10 restated). *Let Γ be a conservative valued constraint language. Either $\text{VCSP}(\Gamma)$ is NP-hard, or $\text{supp}(\Gamma)$ contains a majority operation and hence Γ has valued relational width $(2, 3)$.*

Proof. If $\text{Pol}(\Gamma)$ does not contain a majority operation then Γ is NP-hard by Theorem 9. If $\text{supp}(\Gamma)$ contains a majority operation then, by Corollary 1, Γ has valued relational width $(2, 3)$.

Let F be the set of majority operations in $\text{Pol}(\Gamma) \setminus \text{supp}(\Gamma)$. By Lemma 2, for each $f \in F$, there is an instance I_f of $\text{VCSP}(\Gamma)$ such that $f \notin \text{Pol}(\text{Opt}(I_f))$. Let $\Gamma' = \Gamma \cup \{\text{Opt}(I_f) \mid f \in F\}$. If $\text{Pol}(\Gamma')$ does not contain a majority polymorphism, then, since Γ is conservative, so is Γ' , and hence Γ' is NP-hard by Theorem 9. Therefore, Γ is NP-hard by Theorem 16 (4). Assume that $\text{Pol}(\Gamma')$ contains a majority polymorphism f . Then, $f \notin F$, so $f \in \text{supp}(\Gamma)$. From Corollary 1, it follows that Γ has valued relational width $(2, 3)$. \square

10 Conclusions

Using techniques from the algebraic study of CSPs and the study of linear programming relaxations, we have given a precise characterisation of the power of constant level Sherali-Adams linear programming relaxations for exact solvability of valued constraint languages. Notably, we needed to prove that certain gadget constructions, such as going to the core and interpretations, common in the algebraic CSP literature but not commonly used in other areas of CSPs, such as approximation, preserve solvability by constant level Sherali-Adams relaxations.

It is quite plausible that our reductions, stated in Theorem 16, also work for SDPs, in which case, starting from [56] and extensions from [66] (and in particular replacing the grid construction by expanders, as already observed in [28]), VCSPs not solvable by third level Sherali-Adams LP relaxations would also not be solvable by linear levels of Lasserre SDP relaxations. We will address this in a follow-up paper.

The complexity of Min-Ones problems with respect to exact solvability and approximability was established in [22, 40]. Minimum-Solutions problems are a generalisation of Min-Ones problems to larger domains, including integer programs over bounded domains [37]. Following our characterisation of the power of Sherali-Adams, we have given a complete complexity classification of exact solvability of Minimum-Solution problems over arbitrary finite domains.

Acknowledgements

The authors wish to thank Manuel Bodirsky for explaining the relation between varieties, pseudovarieties, and pp-interpretations in the context of finite domain CSPs. The authors are also grateful to Albert Atserias for useful discussions and for bringing [56] to our attention.

References

- [1] Mikhail Alekhovich, Sanjeev Arora, and Iannis Tzourakis. Towards Strong Nonapproximability Results in the Lovász-Schrijver Hierarchy. *Computational Complexity*, 20(4):615–648, 2011.
- [2] Albert Atserias, Andrei Bulatov, and Anuj Dawar. Affine systems of equations and counting infinitary logic. *Theoretical Computer Science*, 410:1666–1683, 2009.
- [3] Boaz Barak and David Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. In *Proceedings of International Congress of Mathematicians (ICM)*, 2014.
- [4] Libor Barto. The collapse of the bounded width hierarchy. *Journal of Logic and Computation*, 2014.
- [5] Libor Barto and Marcin Kozik. Robust Satisfiability of Constraint Satisfaction Problems. *SIAM Journal on Computing*. To appear.

- [6] Libor Barto and Marcin Kozik. Constraint Satisfaction Problems Solvable by Local Consistency Methods. *Journal of the ACM*, 61(1), 2014. Article No. 3.
- [7] Manuel Bodirsky. Constraint Satisfaction Problems with Infinite Templates. In *Complexity of Constraints*, volume 5250 of *Lecture Notes in Computer Science*, pages 196–228. Springer, 2008.
- [8] Endre Boros and Peter L. Hammer. Pseudo-Boolean optimization. *Discrete Applied Mathematics*, 123(1-3):155–225, 2002.
- [9] Andrei Bulatov. Combinatorial problems raised from 2-semilattices. *Journal of Algebra*, 298:321–339, 2006.
- [10] Andrei Bulatov. Bounded relational width. Unpublished manuscript, 2009.
- [11] Andrei Bulatov, Andrei Krokhin, and Peter Jeavons. Classifying the Complexity of Constraints using Finite Algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.
- [12] Siu On Chan, James R. Lee, Prasad Raghavendra, and David Steurer. Approximate constraint satisfaction requires large LP relaxations. *Journal of the ACM*. To appear.
- [13] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for maximum constraint satisfaction problems. *ACM Transactions on Algorithms*, 5(3), 2009.
- [14] Chandra Chekuri, Sanjeev Khanna, Joseph Naor, and Leonid Zosin. A linear programming formulation and approximation algorithms for the metric labeling problem. *SIAM Journal on Discrete Mathematics*, 18(3):608–625, 2004.
- [15] Eden Chlamtáč and Madhur Tulsiani. Convex relaxations and integrality gaps. In Miguel F. Anjos and Jean B. Lasserre, editors, *Handbook on Semidefinite, Conic and Polynomial Optimization*, volume 166 of *International Series in Operations Research & Management Science*, pages 139–169. Springer, 2012.
- [16] David A. Cohen, Martin C. Cooper, Páidí Creed, Peter Jeavons, and Stanislav Živný. An algebraic theory of complexity for discrete optimisation. *SIAM Journal on Computing*, 42(5):915–1939, 2013.
- [17] David A. Cohen, Martin C. Cooper, Peter Jeavons, and Stanislav Živný. Binarisation via dualisation for valued constraints. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI’15)*, pages 3731–3737. AAAI Press, 2015.
- [18] David A. Cohen, Martin C. Cooper, and Peter G. Jeavons. Generalising submodularity and Horn clauses: Tractable optimization problems defined by tournament pair multimorphisms. *Theoretical Computer Science*, 401(1-3):36–51, 2008.
- [19] David A. Cohen, Martin C. Cooper, Peter G. Jeavons, and Andrei A. Krokhin. The Complexity of Soft Constraint Satisfaction. *Artificial Intelligence*, 170(11):983–1016, 2006.
- [20] Yves Crama and Peter L. Hammer. *Boolean Functions - Theory, Algorithms, and Applications*. Cambridge University Press, 2011.
- [21] Nadia Creignou. A dichotomy theorem for maximum generalized satisfiability problems. *Journal of Computer and System Sciences*, 51(3):511–522, 1995.
- [22] Nadia Creignou, Sanjeev Khanna, and Madhu Sudan. *Complexity Classification of Boolean Constraint Satisfaction Problems*, volume 7 of *SIAM Monographs on Discrete Mathematics and Applications*. SIAM, 2001.

- [23] Víctor Dalmau. There are no pure relational width 2 constraint satisfaction problems. *Information Processing Letters*, 109(4):213–218, 2009.
- [24] Víctor Dalmau, Andrei Krokhin, and Rajsekar Manokaran. Towards a characterization of constant-factor approximable Min CSPs. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’15)*, pages 847–857. SIAM, 2015.
- [25] Víctor Dalmau and Andrei A. Krokhin. Robust Satisfiability for CSPs: Hardness and Algorithmic Results. *ACM Transactions on Computation Theory*, 5(4), 2013. Article No. 15.
- [26] Wenceslas Fernandez de la Vega and Claire Kenyon-Mathieu. Linear programming relaxations of maxcut. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’07)*, pages 53–61. SIAM, 2007.
- [27] Alina Ene, Jan Vondrák, and Yi Wu. Local distribution and the symmetry gap: Approximability of multiway partitioning problems. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’13)*, pages 306–325. SIAM, 2013.
- [28] Tomás Feder and Moshe Y. Vardi. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM Journal on Computing*, 28(1):57–104, 1998.
- [29] Peter Fulla and Stanislav Živný. A Galois connection of valued constraint languages of infinite size. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP’15)*, volume 9134 of *Lecture Notes in Computer Science*, pages 517–528. Springer, 2015.
- [30] Konstantinos Georgiou, Avner Magen, Toniann Pitassi, and Iannis Tourlakis. Integrality gaps of $2-o(1)$ for vertex cover SDPs in the Lovász–Schrijver hierarchy. *SIAM J. Comput.*, 39(8):3553–3570, 2010.
- [31] Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theor. Comput. Sci.*, 259(1-2):613–622, 2001.
- [32] Hiroshi Hirai. Discrete Convexity and Polynomial Solvability in Minimum 0-Extension Problems. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’13)*, pages 1770–1778. SIAM, 2013.
- [33] D. Hobby and R.N. McKenzie. *The Structure of Finite Algebras*, volume 76 of *Contemporary Mathematics*. American Mathematical Society, Providence, R.I., 1988.
- [34] Anna Huber, Andrei Krokhin, and Robert Powell. Skew bisubmodularity and valued CSPs. *SIAM Journal on Computing*, 43(3):1064–1084, 2014.
- [35] Peter Jeavons, David A. Cohen, and Marc Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, 1997.
- [36] Peter Jeavons, Andrei Krokhin, and Stanislav Živný. The complexity of valued constraint satisfaction. *Bulletin of the European Association for Theoretical Computer Science (EATCS)*, 113:21–55, 2014.
- [37] Peter Jonsson, Fredrik Kuivinen, and Gustav Nordh. MAX ONES Generalized to Larger Domains. *SIAM Journal on Computing*, 38(1):329–365, 2008.
- [38] Peter Jonsson and Gustav Nordh. Introduction to the MAXIMUM SOLUTION Problem. In *Complexity of Constraints*, volume 5250 of *Lecture Notes in Computer Science*, pages 255–282. Springer, 2008.

- [39] Peter Jonsson, Gustav Nordh, and Johan Thapper. The maximum solution problem on graphs. In *Proceedings of the 32nd International Symposium on Mathematical Foundations of Computer Science (MFCS'07)*, volume 4708 of *Lecture Notes in Computer Science*, pages 228–239. Springer, 2007.
- [40] Sanjeev Khanna, Madhu Sudan, Luca Trevisan, and David Williamson. The approximability of constraint satisfaction problems. *SIAM Journal on Computing*, 30(6):1863–1920, 2000.
- [41] Vladimir Kolmogorov, Andrei A. Krokhin, and Michal Rolínek. The complexity of general-valued CSPs. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS'15)*. IEEE Computer Society, 2015.
- [42] Vladimir Kolmogorov, Johan Thapper, and Stanislav Živný. The power of linear programming for general-valued CSPs. *SIAM Journal on Computing*, 44(1):1–36, 2015.
- [43] Vladimir Kolmogorov and Stanislav Živný. The complexity of conservative valued CSPs. *Journal of the ACM*, 60(2), 2013. Article No. 10.
- [44] Marcin Kozik, Andrei Krokhin, Matt Valeriote, and Ross Willard. Characterizations of several Maltsev Conditions. *Algebra Universalis*, 73(3-4):205–224, 2015.
- [45] Marcin Kozik and Joanna Ochremiak. Algebraic properties of valued constraint satisfaction problem. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP'15)*, volume 9134 of *Lecture Notes in Computer Science*, pages 846–858. Springer, 2015.
- [46] Amit Kumar, Rajsekar Manokaran, Madhur Tulsiani, and Nisheeth K. Vishnoi. On LP-based approximability for strict CSPs. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*, pages 1560–1573. SIAM, 2011.
- [47] Gábor Kun, Ryan O'Donnell, Suguru Tamaki, Yuichi Yoshida, and Yuan Zhou. Linear programming, width-1 CSPs, and robust satisfaction. In *Proceedings of the 3rd Innovations in Theoretical Computer Science (ITCS'12)*, pages 484–495. ACM, 2012.
- [48] Gábor Kun and Mario Szegedy. A New Line of Attack on the Dichotomy Conjecture. *European Journal of Combinatorics*, 52:338–367, 2016.
- [49] Benoit Larose and László Zádori. Bounded width problems and algebras. *Algebra Universalis*, 56:439–466, 2007.
- [50] Jean B. Lasserre. Global Optimization with Polynomials and the Problem of Moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.
- [51] Monique Laurent. A Comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre Relaxations for 0-1 programming. *Mathematics of Operations Research*, 28(3):470–496, 2003.
- [52] László Lovász and Alexander Schrijver. Cones of Matrices and Set-Functions and 0-1 Optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.
- [53] Miklós Maróti and Ralph McKenzie. Existence theorems for weakly symmetric operations. *Algebra Universalis*, 59(3-4):463–489, 2008.
- [54] Robert Powell and Andrei A. Krokhin. A reduction from valued CSP to min cost homomorphism problem for digraphs. Technical report, 2015. arXiv:1507.01776.

- [55] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC'08)*, pages 245–254. ACM, 2008.
- [56] Grant Schoenebeck. Linear level Lasserre lower bounds for certain k-CSPs. In *FOCS'08*, pages 593–602. IEEE Computer Society, 2008.
- [57] Grant Schoenebeck, Luca Trevisan, and Madhur Tulsiani. A linear round lower bound for Lovász-Schrijver SDP relaxations of vertex cover. In *Proc. of the 22nd Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 205–216, 2007.
- [58] H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal of Discrete Mathematics*, 3(3):411–430, 1990.
- [59] Rustem Takhanov. A Dichotomy Theorem for the General Minimum Cost Homomorphism Problem. In *Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science (STACS'10)*, pages 657–668, 2010.
- [60] Johan Thapper. *Aspects of a Constraint Optimisation Problem*. PhD thesis, Department of Computer Science and Information Science, Linköping University, 2010.
- [61] Johan Thapper and Stanislav Živný. The complexity of finite-valued CSPs. *Journal of the ACM*. To appear.
- [62] Johan Thapper and Stanislav Živný. The power of linear programming for valued CSPs. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'12)*, pages 669–678. IEEE, 2012.
- [63] Johan Thapper and Stanislav Živný. The complexity of finite-valued CSPs. In *Proceedings of the 45th ACM Symposium on the Theory of Computing (STOC'13)*, pages 695–704. ACM, 2013. Full version available at arXiv:1210.2977v3.
- [64] Johan Thapper and Stanislav Živný. Necessary Conditions on Tractability of Valued Constraint Languages. *SIAM Journal on Discrete Mathematics*, 29(4):2361–2384, 2015.
- [65] Johan Thapper and Stanislav Živný. Sherali-Adams relaxations for valued CSPs. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP'15)*, volume 9134 of *Lecture Notes in Computer Science*, pages 1058–1069. Springer, 2015.
- [66] Madhur Tulsiani. CSP gaps and reductions in the Lasserre hierarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC'09)*, pages 303–312. ACM, 2009.
- [67] Hannes Uppman. The Complexity of Three-Element Min-Sol and Conservative Min-Cost-Hom. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP'13)*, volume 7965 of *Lecture Notes in Computer Science*, pages 804–815. Springer, 2013.
- [68] Hannes Uppman. *On Some Combinatorial Optimization Problems*. PhD thesis, Department of Computer Science and Information Science, Linköping University, 2015.
- [69] Magnus Wahlström. Half-integrality, LP-branching and FPT Algorithms. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'14)*, pages 1762–1781. SIAM, 2014.

- [70] Yuichi Yoshida and Yuan Zhou. Approximation schemes via Sherali-Adams hierarchy for dense constraint satisfaction problems and assignment problems. In *Innovations in Theoretical Computer Science (ITCS'14)*, pages 423–438. ACM, 2014.