



HAL
open science

An Algebraic Theory of Complexity for Discrete Optimization

Davis Cohen, Martin Cooper, Paidi Creed, Peter Jeavons, Stanislas Zivny

► **To cite this version:**

Davis Cohen, Martin Cooper, Paidi Creed, Peter Jeavons, Stanislas Zivny. An Algebraic Theory of Complexity for Discrete Optimization. *SIAM Journal on Computing*, 2013, vol. 42 (n° 5), pp. 1915-1939. <10.1137/130906398>. <hal-01122748>

HAL Id: hal-01122748

<https://hal.science/hal-01122748v1>

Submitted on 4 Mar 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization



Open Archive TOULOUSE Archive Ouverte (OATAO)

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible.

This is an author-deposited version published in : <http://oatao.univ-toulouse.fr/>
Eprints ID : 12741

To link to this article : DOI :10.1137/130906398
URL : <http://dx.doi.org/10.1137/130906398>

To cite this version : Cohen, Davis and Cooper, Martin C. and Creed, Paidi and Jeavons, Peter and Zivny, Stanislas *[An Algebraic Theory of Complexity for Discrete Optimization](#)*. (2013) SIAM Journal on Computing, vol. 42 (n° 5). pp. 1915-1939. ISSN 0097-5397

Any correspondence concerning this service should be sent to the repository administrator: staff-oatao@listes-diff.inp-toulouse.fr

AN ALGEBRAIC THEORY OF COMPLEXITY FOR DISCRETE OPTIMIZATION*

DAVID A. COHEN[†], MARTIN C. COOPER[‡], PÁIDÍ CREED[§], PETER G. JEAUVONS[¶],
AND STANISLAV ŽIVNÝ^{||}

Abstract. Discrete optimization problems arise in many different areas and are studied under many different names. In many such problems the quantity to be optimized can be expressed as a sum of functions of a restricted form. Here we present a unifying theory of complexity for problems of this kind. We show that the complexity of a finite-domain discrete optimization problem is determined by certain algebraic properties of the objective function, which we call *weighted polymorphisms*. We define a Galois connection between sets of rational-valued functions and sets of weighted polymorphisms and show how the closed sets of this Galois connection can be characterized. These results provide a new approach to studying the complexity of discrete optimization. We use this approach to identify certain maximal tractable subproblems of the general problem and hence derive a complete classification of complexity for the Boolean case.

Key words. Galois connection, constraint optimization, discrete optimization, valued constraint satisfaction problems, weighted clones, weighted polymorphisms

1. Introduction. Discrete optimization problems arise in many different areas and are studied under many different names, including min-sum problems, Gibbs energy minimization, Markov random fields, conditional random fields, 0/1 integer programming, pseudo-Boolean function minimization, constraint optimization, and valued constraint satisfaction [39, 12, 50, 49, 23, 19, 7, 6, 21, 22, 51].

Here we adopt a very general framework where each problem instance is specified by a set of variables, a set of possible values for those variables, and a set of constraints. Each combination of values allowed by each constraint has an associated *cost*, and the goal is to find an assignment with minimal total cost. This simple abstract mathematical framework can be used to express discrete optimization problems arising in a wide variety of fields, including operational research (scheduling, resource utilization, transportation), computer vision (region segmentation, object recognition, image enhancement), automated reasoning (MAX-SAT, MIN-ONES), graph theory (MIN-CUT, maximum independent set), and many others.

[†]Department of Computer Science, Royal Holloway, University of London, UK (d.cohen@rhul.ac.uk).

[‡]IRIT, University of Toulouse III, France (cooper@irit.fr).

[§]School of Mathematical Sciences, Queen Mary, University of London, UK (p.creed@qmul.ac.uk).

[¶]Department of Computer Science, University of Oxford, UK (peter.jeaavons@cs.ox.ac.uk).

^{||}Department of Computer Science, University of Oxford, UK (standa@cs.ox.ac.uk). This author was supported by a Junior Research Fellowship at University College, Oxford.

In the special case when all defined costs are zero, the problem we are studying collapses to the standard constraint satisfaction problem (CSP). The general CSP is NP-hard and so is unlikely to have a polynomial-time algorithm. However, there has been much success in finding tractable fragments of the CSP by restricting the types of relation allowed in the constraints. A set of allowed relations has been called a *constraint language* [35]. For some constraint languages the associated CSPs with constraints chosen from that language are solvable in polynomial time, while for other constraint languages this class of problems is NP-hard [36, 35, 29]; these two cases are referred to as *tractable languages* and *NP-hard languages*, respectively. Dichotomy theorems, which classify each possible constraint language as either tractable or NP-hard, have been established for languages over 2-element domains [42] and 3-element domains [14], for conservative languages [16, 1], and for maximal languages [17, 13].

The more general framework we consider here, which allows nonzero costs, is also NP-hard, but again we can try to identify tractable fragments by restricting the types of allowed constraints. Each type of constraint is specified by a rational-valued function defined on a set of tuples, which specifies the cost associated with each allowed tuple of values. Such a function is called a *weighted relation*, and a set of weighted relations will be called a *valued constraint language* [19]. Much less is known about the complexity of the optimization problems associated with different valued constraint languages, although some results have been obtained for certain special cases. In particular, a complete characterization of complexity has been obtained for valued constraint languages over a 2-element domain [19]. This result generalizes a number of earlier results for particular optimization problems such as MAX-SAT [25] and MIN-ONES [26]. A complete classification has also been obtained for valued constraint languages containing all unary $\{0, 1\}$ -valued weighted relations (such languages are called *conservative*) [38]. This result generalizes a number of earlier results such as GRAPH MIN-COST-HOM [32] and DIGRAPH MIN-COST-HOM [46].

One class of weighted relations that has been extensively studied and shown to be tractable is the class of *submodular* functions [26, 19, 37, 28, 38, 52].

In the classical CSP framework it has been shown that the complexity of any constraint language over any finite domain is determined by certain algebraic properties known as *polymorphisms* [36, 35]. This result has reduced the problem of identification of tractable constraint languages to that of identification of suitable sets of polymorphisms. In other words, it has been shown to be enough to study just those constraint languages which are characterized by having a given set of polymorphisms. Using this algebraic approach, considerable progress has now been made toward a complete characterization of the complexity of constraint languages over finite domains of arbitrary size [29, 15, 4, 2, 3, 5].

In this paper, we introduce a new algebraic construct which we call a *weighted polymorphism*. We are able to show that the weighted polymorphisms of a valued constraint language are sufficient to determine the complexity of that language. In addition, we are able to define a Galois connection between valued constraint languages and sets of weighted polymorphisms and to characterize the closed sets on both sides.

This paper builds on and extends our earlier attempts to develop an algebraic theory of complexity for valued constraints [19, 18]. These earlier papers introduced the notions of *multimorphism* [19] and *fractional polymorphism* [18] and provided a number of new tools for analyzing the complexity of this class of problems. These earlier approaches have led to considerable progress, including a characterization of languages solvable by a standard LP relaxation in terms of fractional polymorphisms [47], a complete complexity classification for valued constraint languages including all unary

$\{0, 1\}$ -valued weighted relations [38], and a complete complexity classification for valued constraint languages whose weighted relations allow all tuples [48]. However, the algebraic objects introduced in these earlier approaches could not be successfully combined to obtain a suitable closure operation and hence could not be used to define a Galois connection. The notion of weighted polymorphism introduced here subsumes all the earlier notions and finally allows the development of a fully algebraic approach.

The Galois connection we establish here can be used in the search for tractable valued constraint languages in a very similar way to the use of polymorphisms in the search for tractable constraint languages in the classical CSP. First, we need only consider valued constraint languages characterized by weighted polymorphisms. This greatly simplifies the search for a characterization of all tractable valued constraint languages. Second, we will show below that any tractable valued constraint language must have an associated nontrivial weighted polymorphism. Hence the results of this paper provide a powerful new set of tools for analyzing the complexity of finite-domain discrete optimization problems.

The structure of the paper is as follows. In section 2 we introduce the general framework of the valued constraint satisfaction problem (VCSP) and define the notion of expressibility. In section 3 we focus on the classical CSP and show how it fits in the VCSP framework as a special case. We briefly recall the notion of a polymorphism and the Galois connection that has been so fruitful in the study of the complexity of the classical CSP. In sections 4 and 5 we introduce *weighted relational clones* (valued constraint languages closed under expressibility and certain other operations) and the corresponding closed sets of weighted polymorphisms, which we call *weighted clones*. We then state our main result: weighted relational clones are in one-to-one correspondence with weighted clones. In section 6 we give proofs of the theorems establishing this Galois connection. In section 7 we use the algebraic theory to establish necessary conditions that must be satisfied by any tractable valued constraint language. Using these results, we obtain a complete classification for the Boolean case in section 8. Finally, in section 9, we state some conclusions and outline directions for future work.

2. The VCSP. For any set D , the set of tuples of length r over D is denoted by D^r , and a subset of D^r is called a *relation* over D of arity r . A (partial) function ϱ from D^r to \mathbb{Q} associates a rational¹ weight with each of the tuples in some subset of D^r and so will be called a *weighted relation* on D of arity r .

The idea of a weighted relation is very general and can be used to define a wide variety of discrete optimization problems. The general framework we shall use for such problems is defined as follows.

DEFINITION 2.1. *An instance of the VCSP is a triple $\mathcal{P} = \langle V, D, C \rangle$, where V is a finite set of variables, D is a finite set of possible values, and C is a finite multi-set of constraints. Each element of C is a pair $c = \langle \sigma, \varrho \rangle$, where σ is a tuple of variables called the scope of c , and ϱ is a weighted relation on D of arity $|\sigma|$.*

An assignment for \mathcal{P} is a mapping $s : V \rightarrow D$. The cost of an assignment s , denoted $\text{Cost}_{\mathcal{P}}(s)$, is given by the sum of the weights assigned to the restrictions of s onto each constraint scope, that is,

$$\text{Cost}_{\mathcal{P}}(s) \stackrel{\text{def}}{=} \sum_{\langle (v_1, v_2, \dots, v_m), \varrho \rangle \in C} \varrho(s(v_1), s(v_2), \dots, s(v_m)).$$

¹To avoid representational issues, we restrict ourselves to rational rather than real-valued weights. The resulting framework is sufficiently general to encode very many standard optimization problems; for examples, see [19].

If $\varrho(s(v_1), s(v_2), \dots, s(v_m))$ is undefined for some $\langle\langle v_1, v_2, \dots, v_m \rangle, \varrho \rangle \in C$, then the assignment s is said to be infeasible and $\text{Cost}_{\mathcal{P}}(s)$ is undefined.

A solution to \mathcal{P} is a feasible assignment with minimal cost.

In many earlier treatments of the VCSP (e.g., [19]) infeasible assignments are given an infinite cost, whereas here the cost of an infeasible assignment is simply undefined.

Example 2.1 (MAX-CUT). In graph theory a *cut* of a graph is a partition of the vertices into two disjoint sets. The size of a cut is the number of edges of the graph that intersect both sides of this partition. The MAX-CUT problem for a graph is to find a cut with the largest possible size. This problem is NP-hard [30].

The MAX-CUT problem for the graph (V, E) can be expressed as the VCSP instance $\langle V, \{0, 1\}, C \rangle$, where $C = \{ \langle e, \varrho_{\neq} \rangle \mid e \in E \}$, and ϱ_{\neq} is the binary weighted relation on the set $\{0, 1\}$ defined by

$$\varrho_{\neq}(x, y) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } x \neq y, \\ 1 & \text{otherwise.} \end{cases}$$

Any assignment for this VCSP instance partitions the elements of V into those assigned the value 0 and those assigned the value 1. The cost of the assignment is equal to the number of edges minus the size of the corresponding cut.

Example 2.2 (DIGRAPH MIN-COST-HOM). Given two directed graphs (digraphs) $G = (V_G, E_G)$ and $H = (V_H, E_H)$, a homomorphism from G to H is a mapping $f : V_G \rightarrow V_H$ that preserves edges, that is, $(u, v) \in E_G$ implies $(f(u), f(v)) \in E_H$. Assume that for any $u \in V_G$ and $v \in V_H$ a rational cost $c_u(v)$ is given. The cost of a homomorphism f from G to H is then defined to be $\sum_{u \in V_G} c_u(f(u))$. The DIGRAPH MIN-COST-HOM problem is to find a homomorphism from G to H of minimum cost [32, 46].

Given a fixed digraph $H = (V_H, E_H)$, we denote by ϱ_H the binary weighted relation on the set V_H defined by

$$\varrho_H(x, y) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } (x, y) \in E_H, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

The DIGRAPH MIN-COST-HOM problem for input graph $G = (V_G, E_G)$ and fixed target graph $H = (V_H, E_H)$ can be expressed as an instance $\langle V_G, V_H, C \rangle$ of VCSP by setting $C = \{ \langle e, \varrho_H \rangle \mid e \in E_G \} \cup \{ \langle u, c_u \rangle \mid u \in V_G \}$.

A *valued constraint language* is any set Γ of weighted relations on some fixed set D . We define $\text{VCSP}(\Gamma)$ to be the class of all VCSP instances in which all weighted relations in all constraints belong to Γ .

Example 2.1 shows that $\text{VCSP}(\{\varrho_{\neq}\})$ includes the MAX-CUT problem. In fact the class of instances $\text{VCSP}(\{\varrho_{\neq}\})$ corresponds very closely to the MAX-CUT problem, in the sense that any instance from $\text{VCSP}(\{\varrho_{\neq}\})$ can be interpreted as an instance of the MAX-CUT problem on the graph defined by the constraint scopes.

Valued constraint languages may be infinite, but it will be convenient to follow [19] and define the complexity of a valued constraint language in terms of its finite subsets.

DEFINITION 2.2. A *valued constraint language* Γ is called *tractable* if for every finite subset $\Gamma' \subseteq \Gamma$, there exists an algorithm solving any instance $\mathcal{P} \in \text{VCSP}(\Gamma')$ in polynomial time. Conversely, Γ is called *NP-hard* if there is some finite subset $\Gamma' \subseteq \Gamma$ for which $\text{VCSP}(\Gamma')$ is NP-hard.

One advantage of defining tractability in terms of finite subsets is that the tractability of a valued constraint language is independent of whether the cost functions are

represented explicitly (via tables of values) or implicitly (via oracles) since in a finite subset the weighted relations necessarily have bounded arity.

Example 2.1 shows that the valued constraint language $\{\varrho_{\neq}\}$ is NP-hard.

We now define a closure operator on weighted relations, which adds to a given set of weighted relations all other weighted relations which can be *expressed* using that set, in the sense defined below.

DEFINITION 2.3. *For any VCSP instance $\mathcal{P} = \langle V, D, C \rangle$ and any list $L = \langle v_1, \dots, v_r \rangle$ of variables of \mathcal{P} , the projection of \mathcal{P} onto L , denoted $\pi_L(\mathcal{P})$, is the weighted relation on D of arity r defined as follows:*

$$\pi_L(\mathcal{P})(x_1, \dots, x_r) \stackrel{\text{def}}{=} \min_{\{s: V \rightarrow D \mid \langle s(v_1), \dots, s(v_r) \rangle = \langle x_1, \dots, x_r \rangle\}} \text{Cost}_{\mathcal{P}}(s).$$

We say that a weighted relation ϱ is expressible over a valued constraint language Γ if there exists a VCSP instance $\mathcal{P} \in \text{VCSP}(\Gamma)$ and a list L of variables of \mathcal{P} such that $\pi_L(\mathcal{P}) = \varrho$. We call the pair $\langle \mathcal{P}, L \rangle$ a gadget for expressing ϱ over Γ .

We define $\text{Express}(\Gamma)$ to be the expressive power of Γ , that is, the set of all weighted relations expressible over Γ .

Note that the list of variables L in a gadget may contain repeated entries, the sum over an empty set is zero, and the minimum over an empty set is undefined.

Example 2.3. Let \mathcal{P} be the VCSP instance with a single variable v and no constraints, and let $L = \langle v, v \rangle$. Then, by Definition 2.3,

$$\pi_L(\mathcal{P})(x, y) = \begin{cases} 0 & \text{if } x = y, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Hence for any valued constraint language Γ , over any set D , $\text{Express}(\Gamma)$ contains this binary weighted relation, which will be called the *weighted equality relation*.

Example 2.4. Let \mathcal{P} be the VCSP instance with domain $\{0, 1\}$, variables v_1, v_2, v_3 and constraints $\langle \langle v_1, v_2 \rangle, \varrho_{\neq} \rangle$ and $\langle \langle v_2, v_3 \rangle, \varrho_{\neq} \rangle$, and let $L = \langle v_1, v_3 \rangle$. Then, by Definition 2.3,

$$\pi_L(\mathcal{P})(x, y) = \begin{cases} 0 & \text{if } x = y, \\ 1 & \text{otherwise.} \end{cases}$$

Hence $\text{Express}(\{\varrho_{\neq}\})$ contains this binary weighted relation, which will be denoted $\varrho_{=}$.

However, using the results of this paper, we will be able to show, for example, that $\varrho_{\neq} \notin \text{Express}(\{\varrho_{=}\})$ (see Example 8.2).

The next result shows that expressibility preserves tractability.

THEOREM 2.4. *A valued constraint language Γ is tractable if and only if $\text{Express}(\Gamma)$ is tractable; similarly, Γ is NP-hard if and only if $\text{Express}(\Gamma)$ is NP-hard.*

Proof. By the definition of a tractable valued constraint language, it is sufficient to show that for any finite subset Γ' of $\text{Express}(\Gamma)$ there exists a polynomial-time reduction from $\text{VCSP}(\Gamma')$ to $\text{VCSP}(\Gamma'')$, where Γ'' is a finite subset of Γ .

Let Γ' be a finite subset of $\text{Express}(\Gamma)$ and let \mathcal{P}' be any instance of $\text{VCSP}(\Gamma')$. By Definition 2.3, any weighted relation $\varrho' \in \text{Express}(\Gamma)$ can be constructed by using some gadget $\langle \mathcal{P}_{\varrho'}, L \rangle$, where $\mathcal{P}_{\varrho'}$ is an instance of $\text{VCSP}(\Gamma)$. Hence we can simply replace each constraint in \mathcal{P}' which has a weighted relation ϱ' not already in Γ with the corresponding gadget to obtain an instance \mathcal{P} of $\text{VCSP}(\Gamma)$ which has exactly the same solutions as \mathcal{P}' . The maximum size of any of the gadgets used is a constant

determined by the finite set Γ' , so this construction can be carried out in polynomial time in the size of \mathcal{P}' . \square

This result shows that, when trying to identify tractable valued constraint languages, it is sufficient to consider only languages of the form $\text{Express}(\Gamma)$. In the following sections, we will show that such languages can be characterized using certain algebraic properties.

3. Classical constraint satisfaction. In this section we consider the special case when the weights are all zero.

DEFINITION 3.1. *We denote by \mathbf{R}_D the set of all zero-valued weighted relations on a set D .*

There is a one-to-one correspondence between the set of zero-valued weighted relations \mathbf{R}_D and the set of all relations over D . In this correspondence each weighted relation ϱ in \mathbf{R}_D is associated with the relation $R(\varrho)$ containing precisely those tuples on which ϱ is defined. Similarly, each zero-valued weighted relation in \mathbf{R}_D is associated with the *predicate* which is true for precisely those tuples where the weighted relation is defined. Subsets of \mathbf{R}_D are sometimes referred to as *crisp* constraint languages [19] and $\text{VCSP}(\mathbf{R}_D)$ is equivalent to the classical CSP, where each assignment is either allowed (cost 0) or disallowed (infeasible, or cost undefined).

DEFINITION 3.2. *A weighted relation ϱ of arity r can be obtained by addition from the weighted relation ϱ_1 of arity s and the weighted relation ϱ_2 of arity t if ϱ satisfies the identity $\varrho(x_1, \dots, x_r) = \varrho_1(y_1, \dots, y_s) + \varrho_2(z_1, \dots, z_t)$ for some (fixed) choice of y_1, \dots, y_s and z_1, \dots, z_t from among the x_1, \dots, x_r .*

For zero-valued weighted relations this notion of addition corresponds to performing a relational join operation on the associated relations $R(\varrho_1)$ and $R(\varrho_2)$ [33]. It also corresponds to taking a conjunction of the associated predicates [15]. Moreover, minimizing a weighted relation $\varrho \in \mathbf{R}_D$ over one of its arguments corresponds to taking a relational projection of $R(\varrho)$ onto its remaining coordinates. It also corresponds to existential quantification of the associated predicate over that argument.

DEFINITION 3.3. *A set $\Gamma \subseteq \mathbf{R}_D$ is called a relational clone if it contains the weighted equality relation and is closed under addition and minimization over arbitrary arguments.*

For each $\Gamma \subseteq \mathbf{R}_D$ we define $\text{RelClone}(\Gamma)$ to be the smallest relational clone containing Γ .

It is a straightforward consequence of Definitions 2.3 and 3.3 that the expressive power of a crisp constraint language is given by the smallest relational clone containing it, as the next result indicates.

PROPOSITION 3.4. *For any $\Gamma \subseteq \mathbf{R}_D$, $\text{Express}(\Gamma) = \text{RelClone}(\Gamma)$.*

This alternative characterization for the expressive power of a crisp constraint language was first observed in [35] and used to study the complexity of such languages using tools from universal algebra. We now give a brief summary of this algebraic approach.

For any set D , a function $f : D^k \rightarrow D$ is called a *k-ary operation* on D .

DEFINITION 3.5. *We denote by \mathbf{O}_D the set of all finitary operations on D and by $\mathbf{O}_D^{(k)}$ the *k-ary operations* in \mathbf{O}_D .*

DEFINITION 3.6. *The *k-ary projections* on D are the operations $e_i^{(k)} : D^k \rightarrow D$ such that $(a_1, \dots, a_k) \mapsto a_i$.*

DEFINITION 3.7. *Let $f \in \mathbf{O}_D^{(k)}$ and $g_1, \dots, g_k \in \mathbf{O}_D^{(\ell)}$. The *superposition* of f and g_1, \dots, g_k is the ℓ -ary operation $f[g_1, \dots, g_k] : D^\ell \rightarrow D$ such that $(x_1, \dots, x_\ell) \mapsto f(g_1(x_1, \dots, x_\ell), \dots, g_k(x_1, \dots, x_\ell))$.*

DEFINITION 3.8. A set $F \subseteq \mathbf{O}_D$ is called a clone of operations if it contains all the projections on D and is closed under superposition. For each $F \subseteq \mathbf{O}_D$ we define $\text{Clone}(F)$ to be the smallest clone containing F .

We can extend k -ary operations to operate on tuples in a natural way, as follows. Let $\mathbf{x}_1, \dots, \mathbf{x}_k$ be tuples of length r over a set D , where each $\mathbf{x}_i = \langle x_{i,1}, x_{i,2}, \dots, x_{i,r} \rangle$. We can obtain another element of D^r by applying f to the tuples \mathbf{x}_i coordinatewise, as follows:

$$f(\mathbf{x}_1, \dots, \mathbf{x}_k) \stackrel{\text{def}}{=} \langle f(x_{1,1}, \dots, x_{k,1}), f(x_{1,2}, \dots, x_{k,2}), \dots, f(x_{1,r}, \dots, x_{k,r}) \rangle.$$

DEFINITION 3.9. Let ϱ be a weighted relation of arity r on a set D and let $f \in \mathbf{O}_D^{(k)}$. We say that f is a polymorphism of ϱ if for any $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k \in D^r$ such that $\varrho(\mathbf{x}_i)$ is defined for $i = 1, \dots, k$, we have that $\varrho(f(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k))$ is also defined.

If f is a polymorphism of ϱ we say ϱ is invariant under f .

DEFINITION 3.10. For any valued constraint language Γ over a set D , we denote by $\text{Pol}(\Gamma)$ the set of all operations on D which are polymorphisms of all weighted relations $\varrho \in \Gamma$ and by $\text{Pol}^{(k)}(\Gamma)$ the k -ary operations in $\text{Pol}(\Gamma)$.

DEFINITION 3.11. For any $F \subseteq \mathbf{O}_D$, we denote by $\text{Inv}(F)$ the set of all weighted relations in \mathbf{R}_D that are invariant under all operations $f \in F$.

For any set D , the mappings Pol and Inv form a Galois connection between \mathbf{O}_D and \mathbf{R}_D [11]. A characterization of this Galois connection for finite sets D is given by the following two theorems, originally obtained for sets of relations [31, 10].

THEOREM 3.12. For any finite set D and any finite $\Gamma \subseteq \mathbf{R}_D$, $\text{Inv}(\text{Pol}(\Gamma)) = \text{RelClone}(\Gamma)$.

THEOREM 3.13. For any finite set D and any finite $F \subseteq \mathbf{O}_D$, $\text{Pol}(\text{Inv}(F)) = \text{Clone}(F)$. As with any Galois connection [11], this means that there is a one-to-one correspondence between clones and relational clones. Together with Proposition 3.4, this result shows that the expressive power of any crisp constraint language Γ on a finite set D corresponds to a particular clone of operations on D . Hence, by Theorem 2.4, the search for tractable crisp constraint languages corresponds to a search for suitable clones of operations [35, 15]. This key observation paved the way for applying deep results from universal algebra in the search for tractable constraint languages [17, 16, 13, 14, 4, 2, 3, 5].

4. Weighted relational clones. In this section we return to the general case of weighted relations taking arbitrary values in \mathbb{Q} in order to define the notion of a weighted relational clone.

DEFINITION 4.1. We denote by Φ_D the set of all weighted relations on D taking values in \mathbb{Q} and by $\Phi_D^{(r)}$ the weighted relations in Φ_D of arity r .

We now define a closure operator on weighted relations, which adds to a set of weighted relations all other weighted relations which can be obtained from the original set by nonnegative scaling and addition of a constant.

DEFINITION 4.2. A weighted relation $\varrho' \in \Phi_D$ can be obtained from a weighted relation $\varrho \in \Phi_D$ by nonnegative scaling and addition of constants if there exist $\alpha, \beta \in \mathbb{Q}$ with $\alpha \geq 0$ such that $\varrho' \equiv \alpha\varrho + \beta$. We denote by Γ_{\sim} the smallest set of weighted relations containing Γ which is closed under nonnegative scaling and addition of constants.

The next result shows that adding weighted relations that can be obtained by nonnegative scaling and addition of constants preserves tractability.

THEOREM 4.3. A valued constraint language Γ is tractable if and only if Γ_{\sim} is tractable; similarly, Γ is NP-hard if and only if Γ_{\sim} is NP-hard.

Proof. By Definition 2.2, it is sufficient to show that for each finite subset Γ' of Γ_{\sim} there exists a polynomial-time reduction from $\text{VCSP}(\Gamma')$ to $\text{VCSP}(\Gamma)$.

Let Γ' be a fixed finite subset of Γ_{\sim} . We will show that each weighted relation $\varrho' \in \Gamma'$ that is not in Γ can be replaced by a weighted relation from Γ to obtain a polynomial-time reduction from $\text{VCSP}(\Gamma')$ to $\text{VCSP}(\Gamma' \setminus \{\varrho'\} \cup \Gamma)$. By performing each of these reductions in sequence we can obtain the desired polynomial-time reduction from $\text{VCSP}(\Gamma')$ to $\text{VCSP}(\Gamma)$.

Let \mathcal{P}' be any instance of $\text{VCSP}(\Gamma')$, and choose any weighted relation $\varrho' \in \Gamma'$. By Definition 4.2, ϱ' can be obtained by nonnegative scaling and addition of constants from some weighted relation $\varrho \in \Gamma$. Hence, we can replace each constraint of the form $\langle \sigma, \varrho' \rangle$ in \mathcal{P}' with a new constraint $\langle \sigma, \varrho \rangle$, where $\varrho \in \Gamma$ and $\varrho' = \alpha\varrho + \beta$ for some nonnegative rational value α and some arbitrary rational constant β , to obtain an instance of $\text{VCSP}(\Gamma' \setminus \{\varrho'\} \cup \Gamma)$. It only remains to ensure that this new instance has the same solutions as \mathcal{P}' .

The constant β is added to the cost of all assignments and so does not affect the choice of solution.

Since α is a nonnegative rational value, it can be expressed as p/q for some nonnegative integer p and positive integer q .

If p is nonzero, then the effect of the scale factor p/q can be simulated by taking p copies of the new constraints and q copies of all other constraints. The values of p, q are constants determined by the choice of ϱ' , so this construction can be carried out in polynomial time in the size of \mathcal{P}' .

It only remains to deal with the case where p is zero. In this case we still need to replace each constraint $\langle \sigma, \varrho' \rangle$ with one copy of the corresponding constraint $\langle \sigma, \varrho \rangle$, to ensure that infeasible assignments to ϱ' are still excluded. Assume that \mathcal{P}' contains k such constraints. Let M be the maximum weight assigned by ϱ' , and let m be the minimum difference between any two distinct weights assigned by any other weighted relations in Γ' . The cost of any feasible assignment after replacing the k constraints $\langle \sigma, \varrho' \rangle$ is greater by at most kM than the cost of the same assignment for \mathcal{P}' . Hence if we also take $\lceil \frac{Mk}{m} + 1 \rceil$ copies of all the remaining constraints of \mathcal{P}' , then we obtain an instance of $\text{VCSP}(\Gamma' \setminus \{\varrho'\} \cup \Gamma)$ with the same solutions as \mathcal{P}' . Since M and m are constants determined by the finite set Γ' , this construction can again be carried out in polynomial time in the size of \mathcal{P}' . \square

DEFINITION 4.4. *A set $\Gamma \subseteq \Phi_D$ is a weighted relational clone if it contains the weighted equality relation and is closed under nonnegative scaling and addition of constants, addition, and minimization over arbitrary arguments.*

For each $\Gamma \subseteq \Phi_D$ we define $\text{wRelClone}(\Gamma)$ to be the smallest weighted relational clone containing Γ .

It is a straightforward consequence of Definitions 2.3, 4.2, and 4.4 that for any valued constraint language $\Gamma \subseteq \Phi_D$, the set of weighted relations that can be expressed using weighted relations obtained from Γ by nonnegative scaling and addition of constants is given by the smallest weighted relational clone containing Γ , as the next result indicates.

PROPOSITION 4.5. *For any $\Gamma \subseteq \Phi_D$, $\text{Express}(\Gamma_{\sim}) = \text{wRelClone}(\Gamma)$.*

Hence, by Theorems 2.4 and 4.3, the search for tractable valued constraint languages corresponds to a search for suitable weighted relational clones.

In the next section we establish an alternative characterization for weighted relational clones which facilitates this search.

5. Weighted clones. To obtain a suitable alternative characterization for weighted relational clones we now generalize the notion of a clone of operations, introduced in Definition 3.8, by introducing the notion of a *weighted* clone.

Recall from Definition 3.8 that a clone of operations, C , is a set of operations on some fixed set D that contains all projections and is closed under superposition. The k -ary operations in a clone C will be denoted $C^{(k)}$.

DEFINITION 5.1. *We define a k -ary weighting of a clone C to be a function $\omega : C^{(k)} \rightarrow \mathbb{Q}$ such that $\omega(f) < 0$ only if f is a projection and*

$$\sum_{f \in C^{(k)}} \omega(f) = 0.$$

We denote by \mathbf{W}_C the set of all possible weightings of C and by $\mathbf{W}_C^{(k)}$ the set of k -ary weightings of C .

For any weighting ω , we denote by $\mathbf{dom}(\omega)$ the set of operations $C^{(k)}$ on which ω is defined. We denote by $\text{ar}(\omega)$ the arity of ω .

Since a weighting is simply a rational-valued function satisfying certain linear inequalities it can be scaled by any nonnegative rational to obtain a new weighting. Similarly, any two weightings of the same clone of the same arity can be added to obtain a new weighting of that clone.

The notion of superposition from Definition 3.7 can also be extended to weightings in a natural way by forming a superposition with each argument of the weighting, as follows.

DEFINITION 5.2. *For any clone C , any $\omega \in \mathbf{W}_C^{(k)}$ and any $g_1, g_2, \dots, g_k \in C^{(\ell)}$, we define the superposition of ω and g_1, \dots, g_k to be the function $\omega[g_1, \dots, g_k] : C^{(\ell)} \rightarrow \mathbb{Q}$ defined by*

$$(5.1) \quad \omega[g_1, \dots, g_k](f') \stackrel{\text{def}}{=} \sum_{\substack{f \in C^{(k)} \\ f[g_1, \dots, g_k] = f'}} \omega(f).$$

Example 5.1. Let D be a totally ordered set, and let $C = \text{Clone}(\{\max\})$, where \max is the binary maximum operation on D . Note that $C^{(2)}$ contains just three binary operations: $e_1^{(2)}, e_2^{(2)}$, and \max . Let ω be the 2-ary weighting of C given by

$$\omega(f) \stackrel{\text{def}}{=} \begin{cases} -1 & \text{if } f = e_1^{(2)}, \\ +1 & \text{if } f = e_2^{(2)}, \\ 0 & \text{if } f = \max, \end{cases}$$

and let

$$\langle g_1, g_2 \rangle = \langle e_2^{(2)}, \max \rangle.$$

Note that $e_1^{(2)}[g_1, g_2] = g_1 = e_2^{(2)}$ and $e_2^{(2)}[g_1, g_2] = g_2 = \max$, so, applying Definition 5.2, we have

$$\omega[g_1, g_2](f) = \begin{cases} 0 & \text{if } f = e_1^{(2)}, \\ -1 & \text{if } f = e_2^{(2)}, \\ +1 & \text{if } f = \max. \end{cases}$$

Note that $\omega[g_1, g_2]$ satisfies the conditions of Definition 5.1 and hence is a 2-ary weighting of C .

Example 5.2. Let C be a clone on some totally ordered set D and assume that C contains the binary maximum and minimum operations on D , which are denoted by \max and \min . Note that $C^{(4)}$ contains operations such as $\max[e_i^{(4)}, e_j^{(4)}]$ which returns the maximum of the i th and j th argument values. Operations of this form will be denoted $\max(x_i, x_j)$.

Let ω be the 4-ary weighting of C given by

$$\omega(f) \stackrel{\text{def}}{=} \begin{cases} -1 & \text{if } f \text{ is a projection, that is, } f \in \{e_1^{(4)}, e_2^{(4)}, e_3^{(4)}, e_4^{(4)}\}, \\ +1 & \text{if } f \in \{\max(x_1, x_2), \min(x_1, x_2), \max(x_3, x_4), \min(x_3, x_4)\}, \\ 0 & \text{otherwise,} \end{cases}$$

and let

$$\langle g_1, g_2, g_3, g_4 \rangle = \langle e_1^{(3)}, e_2^{(3)}, e_3^{(3)}, \max(x_1, x_2) \rangle.$$

Then, by Definition 5.2 we have

$$\omega[g_1, g_2, g_3, g_4](f) = \begin{cases} -1 & \text{if } f \text{ is a projection, that is, } f \in \{e_1^{(3)}, e_2^{(3)}, e_3^{(3)}\}, \\ +1 & \text{if } f \in \{\max(x_1, x_2, x_3), \min(x_1, x_2), \min(x_3, \max(x_1, x_2))\}, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\omega[g_1, g_2, g_3, g_4]$ satisfies the conditions of Definition 5.1 and hence is a 3-ary weighting of C .

Example 5.3. Let C and ω be the same as in Example 5.2 but now consider

$$\langle g'_1, g'_2, g'_3, g'_4 \rangle = \langle e_1^{(4)}, \max(x_2, x_3), \min(x_2, x_3), e_4^{(4)} \rangle.$$

By Definition 5.2 we have

$$\omega[g'_1, g'_2, g'_3, g'_4](f) = \begin{cases} -1 & \text{if } f \in \{e_1^{(4)}, \max(x_2, x_3), \min(x_2, x_3), e_4^{(4)}\}, \\ +1 & \text{if } f \in \left\{ \begin{array}{l} \max(x_1, x_2, x_3), \min(x_1, \max(x_2, x_3)), \\ \max(\min(x_2, x_3), x_4), \min(x_2, x_3, x_4) \end{array} \right\}, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\omega[g'_1, g'_2, g'_3, g'_4]$ does not satisfy the conditions of Definition 5.1 because, for example, we have that $\omega[g'_1, g'_2, g'_3, g'_4](f) < 0$ when $f = \max(x_2, x_3)$, which is not a projection. Hence $\omega[g'_1, g'_2, g'_3, g'_4]$ is not a valid weighting of C .

It follows immediately from Definition 3.7 that the sum of the weights in any superposition $\omega[g_1, \dots, g_k]$ is equal to the sum of the weights in ω , which is zero, by Definition 5.1. However, as we have seen in Example 5.3, it is not always the case that an arbitrary superposition satisfies the other condition in Definition 5.1 that negative weights are only assigned to projections. Hence we make the following definition.

DEFINITION 5.3. *If the result of a superposition is a valid weighting, then that superposition will be called a proper superposition.*

Remark 5.1. The superposition of a projection operation and other projection operations is always a projection operation. So, by Definition 5.2, for any clone C and any $\omega \in \mathbf{W}_C^{(k)}$, if $g_1, \dots, g_k \in C^{(\ell)}$ are projections, then the function $\omega[g_1, \dots, g_k]$ can take negative values only on projections and hence is a valid weighting. This means that a superposition with any list of projections is always a proper superposition.

We are now ready to define weighted clones.

DEFINITION 5.4. *A weighted clone, W , is a nonempty set of weightings of some fixed clone C which is closed under nonnegative scaling, addition of weightings of equal arity, and proper superposition with operations from C . The clone C is called the support of W .*

Example 5.4. For any clone, C , the set \mathbf{W}_C containing all possible weightings of C is a weighted clone with support C .

Example 5.5. For any clone, C , the set \mathbf{W}_C^0 containing all *zero-valued* weightings of C is a weighted clone with support C . Note that \mathbf{W}_C^0 contains exactly one weighting of each possible arity, which assigns the value 0 to all operations in C of that arity.

We now establish a link between weightings and weighted relations, which will allow us to link weighted clones and weighted relational clones.

DEFINITION 5.5. *Let ϱ be a weighted relation of arity r on some set D and let ω be a k -ary weighting of some clone of operations C on the set D .*

We say that ω is a weighted polymorphism of ϱ if for any $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k \in D^r$ such that $\varrho(\mathbf{x}_i)$ is defined for $i = 1, \dots, k$, we have that $\varrho(f(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k))$ is defined for all $f \in C^{(k)}$, and

$$(5.2) \quad \sum_{f \in C^{(k)}} \omega(f) \varrho(f(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k)) \leq 0.$$

If ω is a weighted polymorphism of ϱ we say ϱ is improved by ω .

Note that by Definition 3.9, if ϱ is improved by the weighting $\omega \in \mathbf{W}_C^{(k)}$, then every element of $C^{(k)}$ must be a polymorphism of ϱ .

Example 5.6. Consider the class of submodular functions [40]. These are precisely the functions ϱ on an ordered domain which satisfy the following identity:

$$\varrho(\min(\mathbf{x}_1, \mathbf{x}_2)) + \varrho(\max(\mathbf{x}_1, \mathbf{x}_2)) - \varrho(\mathbf{x}_1) - \varrho(\mathbf{x}_2) \leq 0.$$

In other words, the set of submodular functions is the set of weighted relations with a 2-ary weighted polymorphism ω_{sub} , defined by

$$\omega_{sub}(f) \stackrel{\text{def}}{=} \begin{cases} -1 & \text{if } f \in \{e_1^{(2)}, e_2^{(2)}\}, \\ +1 & \text{if } f \in \{\min(x_1, x_2), \max(x_1, x_2)\}, \\ 0 & \text{otherwise.} \end{cases}$$

Submodular function minimization is known to be tractable [44, 34].

DEFINITION 5.6. *For any $\Gamma \subseteq \Phi_D$, we denote by $\text{wPol}(\Gamma)$ the set of all weightings of $\text{Pol}(\Gamma)$ which are weighted polymorphisms of all weighted relations $\varrho \in \Gamma$. The set of k -ary weightings in $\text{wPol}(\Gamma)$ will be denoted $\text{wPol}^{(k)}(\Gamma)$.*

To define a mapping in the other direction, we need to consider the union of the sets \mathbf{W}_C over all clones C on some fixed set D , which will be denoted \mathbf{W}_D . If we have a set $W \subseteq \mathbf{W}_D$ which may contain weightings of *different* clones over D , then we can extend each of these weightings with zeros, as necessary, so that they are weightings of the same clone C , given by

$$C = \text{Clone} \left(\bigcup_{\omega \in W} \text{dom}(\omega) \right).$$

This set of extended weightings obtained from W will be denoted \overline{W} . For any set $W \subseteq \mathbf{W}_D$, we define $\text{wClone}(W)$ to be the smallest weighted clone containing \overline{W} .

DEFINITION 5.7. *For any $W \subseteq \mathbf{W}_D$, we denote by $\text{Imp}(W)$ the set of all weighted relations in Φ_D which are improved by all weightings $\omega \in W$. The set of r -ary weighted relations in $\text{Imp}(W)$ will be denoted $\text{Imp}^{(r)}(W)$.*

It follows immediately from the definition of a Galois connection [11] that for any set D , the mappings wPol and Imp form a Galois connection between \mathbf{W}_D and Φ_D , as illustrated in Figure 5.1. A characterization of this Galois connection for finite sets D is given by the following two theorems, which are proved in section 6.

THEOREM 5.8. For any finite D and any finite $\Gamma \subseteq \Phi_D$, $\text{Imp}(\text{wPol}(\Gamma)) = \text{wRelClone}(\Gamma)$.

THEOREM 5.9. For any finite D and any finite $W \subseteq \mathbf{W}_D$, $\text{wPol}(\text{Imp}(W)) = \text{wClone}(W)$.

As with any Galois connection [11], this means that there is a one-to-one correspondence between weighted clones and weighted relational clones. Hence, by Proposition 4.5, Theorem 2.4, and Theorem 4.3, the search for tractable valued constraint languages over a finite set corresponds to a search for suitable weighted clones of operations.

6. Proofs of Theorems 5.8 and 5.9. Our proofs of Theorems 5.8 and 5.9 will both use the following result, which is a variant of the well-known Farkas lemma used in linear programming [40, 43].

LEMMA 6.1 (Farkas, 1894). *Let S and T be finite sets of indices, where T is the disjoint union of two subsets, T_{\geq} and $T_{=}$. For all $i \in S$ and all $j \in T$, let $a_{i,j}$ and b_j be rational numbers. Exactly one of the following holds:*

- *Either there exists a set of nonnegative rational numbers $\{x_i \mid i \in S\}$ and a rational number c such that*




FIG. 5.1. Galois connection between Φ_D and \mathbf{W}_D .

$$\begin{aligned} \text{for each } j \in T_{\geq}, \quad & \sum_{i \in S} a_{i,j} x_i \geq b_j + c, \quad \text{and} \\ \text{for each } j \in T_{=}, \quad & \sum_{i \in S} a_{i,j} x_i = b_j + c. \end{aligned}$$

- Or else there exists a set of integers $\{y_j \mid j \in T\}$ such that $\sum_{j \in T} y_j = 0$ and

$$\begin{aligned} \text{for each } j \in T_{\geq}, \quad & y_j \geq 0, \\ \text{for each } i \in S, \quad & \sum_{j \in T} y_j a_{i,j} \leq 0, \quad \text{and} \\ & \sum_{j \in T} y_j b_j > 0. \end{aligned}$$

Such a set is called a certificate of unsolvability.

We note that there is an effective procedure for deciding which of the cases mentioned in Lemma 6.1 holds for any instance and for calculating the values of the corresponding coefficients x_i or y_j [43].

We will prove Theorem 5.8 in two parts. First, we show in Proposition 6.2 that the set of all weighted relations improved by any given set of weightings is always a weighted relational clone. Then we show that for any finite set Γ the set of weighted relations improved by all weightings in $\text{wPol}(\Gamma)$ is precisely the weighted relational clone $\text{wRelClone}(\Gamma)$.

PROPOSITION 6.2. *For any finite set D and any $W \subseteq \mathbf{W}_D$, $\text{Imp}(W)$ is a weighted relational clone.*

Proof. Certainly $\text{Imp}(W)$ contains the weighted equality relation (defined in Example 2.3), since this weighted relation satisfies inequality (5.2) in Definition 5.5 for all tuples \mathbf{x}_i on which it is defined. Similarly, $\text{Imp}(W)$ is closed under nonnegative scaling, addition of constants, and addition and rearrangement of arguments, since all of these operations preserve inequality (5.2). Hence, to show $\text{Imp}(W)$ is a weighted relational clone we only need to show $\text{Imp}(W)$ is closed under minimization.

Let $\varrho \in \text{Imp}^{(r)}(W)$ and assume that ϱ' is obtained from ϱ by minimizing over the last argument. In other words, $\varrho'(x_1, x_2, \dots, x_{r-1}) = \min_{x_r}(\varrho(x_1, x_2, \dots, x_r))$. We will now show that $\varrho' \in \text{Imp}(W)$.

Let $\omega \in W$ be a k -ary weighting of a clone C . Since $\varrho \in \text{Imp}(W)$, we know that ϱ and ω satisfy inequality (5.2) for all $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k$ such that $\varrho(\mathbf{x}_i)$ is defined. Now consider any $\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_k$ for which each $\varrho'(\mathbf{x}'_i)$ is defined. Extend each \mathbf{x}'_i to a tuple \mathbf{x}''_i of arity r in such a way that $\varrho(\mathbf{x}''_i)$ is minimized. Since all negative values of ω are associated with projections, we have

$$\sum_{f \in C^{(k)}} \omega(f) \varrho'(f(\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_k)) \leq \sum_{f \in C^{(k)}} \omega(f) \varrho(f(\mathbf{x}''_1, \mathbf{x}''_2, \dots, \mathbf{x}''_k)) \leq 0. \quad \square$$

We now prove Theorem 5.8, which states that for any finite set D and any finite $\Gamma \subset \Phi_D$, $\text{Imp}(\text{wPol}(\Gamma)) = \text{wRelClone}(\Gamma)$.

Proof. We first establish that for any $\Gamma \subset \Phi_D$ we have the inclusion $\text{wRelClone}(\Gamma) \subseteq \text{Imp}(\text{wPol}(\Gamma))$. To see this, observe that $\Gamma \subseteq \text{Imp}(\text{wPol}(\Gamma))$ and, hence, $\text{wRelClone}(\Gamma) \subseteq \text{wRelClone}(\text{Imp}(\text{wPol}(\Gamma)))$, which is equal to $\text{Imp}(\text{wPol}(\Gamma))$ by Proposition 6.2.

We will prove the reverse inclusion, $\text{Imp}(\text{wPol}(\Gamma)) \subseteq \text{wRelClone}(\Gamma)$, as follows. Given a weighted relation ϱ of arity r , we will show that either there exists a weighted

operation $\omega \in \text{wPol}(\Gamma)$ such that $\varrho \notin \text{Imp}(\{\omega\})$ or else ϱ plus some constant can be obtained by minimization from a nonnegative weighted sum of weighted relations in Γ , and hence $\varrho \in \text{wRelClone}(\Gamma)$.

We now give the details of this argument. Let k be the number of r -tuples for which ϱ is defined and fix an arbitrary order, $\mathbf{x}_1, \dots, \mathbf{x}_k$, for these tuples. This list of tuples can be viewed as (the rows of) a matrix with k rows and r columns, which we will call S_ϱ .

By Proposition 4.5, $\varrho \in \text{wRelClone}(\Gamma)$ if and only if it can be expressed using weighted relations from Γ_\sim . By Definition 2.3, a weighted relation ϱ' is expressible over Γ_\sim if and only if there exists an instance $\mathcal{P} \in \text{VCSP}(\Gamma_\sim)$ and a list L of variables of \mathcal{P} such that $\pi_L(\mathcal{P}) = \varrho'$.

We consider instances \mathcal{P} with $|D^k|$ variables, where each variable is associated with a distinct tuple from D^k . Each constraint of \mathcal{P} is a pair $\langle S, \gamma \rangle$ for some $\gamma \in \Gamma$ and some list of variables S . Each such S can be viewed as a list of k -tuples over D , and hence as a matrix over D , whose columns are these k -tuples. Since we are using \mathcal{P} to express the defined values of ϱ , it is sufficient to consider only matrices S with rows $\mathbf{t}_1, \dots, \mathbf{t}_k$ such that $\gamma(\mathbf{t}_i)$ is defined for $i = 1, \dots, k$. For any $\gamma \in \Gamma$, a pair $\langle S, \gamma \rangle$ with this property will be called a k -match to Γ .

Each assignment to the variables of \mathcal{P} can be seen as a mapping from k -tuples over D to D and hence is associated with an operation $f \in \mathbf{O}_D^{(k)}$. For any list of variables S of \mathcal{P} , we will write $f(S)$ to denote the assignment to those variables obtained by applying f to the columns of S , viewed as a matrix. With this notation, we have that $\pi_L(\mathcal{P}) = \varrho$ with $L = S_\varrho$, if we can find nonnegative rationals $x_{S,\gamma}$ for all k -matches to Γ , and a constant $c \in \mathbb{Q}$, such that the following system of inequalities and equations is satisfied:

$$\text{for each } f \in \mathbf{O}_D^{(k)}, \quad \sum_{\gamma \in \Gamma} \sum_{\{\text{all } k\text{-matches } \langle S, \gamma \rangle\}} x_{S,\gamma} \gamma(f(S)) \geq \varrho(f(S_\varrho)) + c,$$

$$\text{for each projection } e \in \mathbf{O}_D^{(k)}, \quad \sum_{\gamma \in \Gamma} \sum_{\{\text{all } k\text{-matches } \langle S, \gamma \rangle\}} x_{S,\gamma} \gamma(e(S)) = \varrho(e(S_\varrho)) + c.$$

Moreover, if $f \notin \text{Pol}^{(k)}(\Gamma)$, then the left-hand side of the corresponding inequality will be undefined, by Definition 3.9, so it is sufficient to consider only $f \in \text{Pol}^{(k)}(\Gamma)$. This gives us a system of inequalities and equations with rational coefficients.

If this system has a solution, then $\varrho \in \text{wRelClone}(\Gamma)$. On the other hand, if this system of equations and inequalities has no solution, then we appeal to Lemma 6.1 to get a certificate of unsolvability. That is, in this case we know that there exists a set of integers $\{y_f \mid f \in \text{Pol}^{(k)}(\Gamma)\}$ such that $\sum_{f \in \text{Pol}^{(k)}(\Gamma)} y_f = 0$, $y_f < 0$ only if f is a projection, and

$$(6.1) \quad \text{for each } k\text{-matching } \langle S, \gamma \rangle \text{ of } \Gamma, \quad \sum_{f \in \text{Pol}^{(k)}(\Gamma)} y_f \gamma(f(S)) \leq 0 \quad \text{and}$$

$$(6.2) \quad \sum_{f \in \text{Pol}^{(k)}(\Gamma)} y_f \varrho(f(S_\varrho)) > 0.$$

Now, consider the k -ary weighting ω of the clone $\text{Pol}(\Gamma)$ defined by $\omega(f) = y_f$ for each $f \in \text{Pol}^{(k)}(\Gamma)$. From (6.1), we can see that ω is a weighted polymorphism of every $\gamma \in \Gamma$. On the other hand, (6.2) shows that ω is not a weighted polymorphism of ϱ . \square

Remark 6.1. The proof of Theorem 5.8 demonstrates the decidability of the following question: for any finite $\Gamma \subseteq \Phi_D$ and any weighted relation ϱ defined on D , does ϱ belong to $\text{wRelClone}(\Gamma)$?

We will prove Theorem 5.9 in two parts. First, we show in Proposition 6.3 that the set of weighted polymorphisms of any given set of weighted relations is always a weighted clone. Then we show that for any finite set W the set of weightings that improve all weighted relations in $\text{Imp}(W)$ is precisely the weighted clone $\text{wClone}(W)$.

PROPOSITION 6.3. *For any finite set D , and any $\Gamma \subseteq \Phi_D$, $\text{wPol}(\Gamma)$ is a weighted clone.*

Proof. By Definition 5.6, $\text{wPol}(\Gamma)$ is a set of weightings of $\text{Pol}(\Gamma)$. Similarly, $\text{wPol}(\Gamma)$ is closed under addition and nonnegative scaling, since both these operations preserve inequality (5.2). Hence, to show $\text{wPol}(\Gamma)$ is a weighted clone we only need to show $\text{wPol}(\Gamma)$ is closed under proper superposition by members of $\text{Pol}(\Gamma)$.

Let $\omega \in \text{wPol}^{(k)}(\Gamma)$ and suppose $\omega' = \omega[g_1, \dots, g_k]$ is a proper superposition of ω , where $g_1, g_2, \dots, g_k \in \text{Pol}^{(\ell)}(\Gamma)$. We will now show that $\omega' \in \text{wPol}^{(\ell)}(\Gamma)$. Suppose ϱ is a weighted relation of arity r satisfying $\omega \in \text{wPol}(\{\varrho\})$, i.e., ϱ and ω satisfy inequality (5.2) for all $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k$ such that each $\varrho(\mathbf{x}_i)$ is defined. Given any $\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_\ell$ for which each $\varrho(\mathbf{x}'_i)$ is defined, set $\mathbf{x}_i = g_i(\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_\ell)$ for $i = 1, 2, \dots, k$. Then, if we set $f' = f[g_1, \dots, g_k]$, we have $f'(\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_\ell) = f(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k)$ for any $f \in \text{Pol}^{(k)}(\Gamma)$. Hence, by Definition 5.2, we have

$$\sum_{f' \in \text{Pol}^{(\ell)}(\Gamma)} \omega'(f') \varrho(f'(\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_\ell)) = \sum_{f \in \text{Pol}^{(k)}(\Gamma)} \omega(f) \varrho(f(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k)) \leq 0. \quad \square$$

We will make use of the following technical lemma, which shows that any weighted sum of arbitrary superpositions of a pair of weightings ω_1 and ω_2 can be obtained by taking a weighted sum of superpositions of ω_1 and ω_2 with projection operations and then taking a superposition of the result. This result implies that any weighting which can be expressed as a weighted sum of arbitrary superpositions can also be expressed as a superposition of a weighted sum of *proper* superpositions.

LEMMA 6.4. *Let C be a clone, and let ω_1 and ω_2 be weightings of C , of arity k and ℓ respectively. For any $g_1, \dots, g_k \in C^{(m)}$ and any $g'_1, \dots, g'_\ell \in C^{(m)}$,*

$$(6.3) \quad c_1 \omega_1[g_1, \dots, g_k] + c_2 \omega_2[g'_1, \dots, g'_\ell] = \omega[g_1, \dots, g_k, g'_1, \dots, g'_\ell],$$

where $\omega = c_1 \omega_1[e_1^{(k+\ell)}, \dots, e_k^{(k+\ell)}] + c_2 \omega_2[e_{k+1}^{(k+\ell)}, \dots, e_{k+\ell}^{(k+\ell)}]$

Proof. For any $f \in C^{(m)}$, the result of applying the right-hand-side expression in (6.3) to f is

$$\sum_{f' \in C^{(k+\ell)}} \left(\sum_{\substack{h' \in C^{(k)} \\ h'[e_1^{(k+\ell)}, \dots, e_k^{(k+\ell)}] = f'}} c_1 \omega_1(h') + \sum_{\substack{h' \in C^{(\ell)} \\ h'[e_{k+1}^{(k+\ell)}, \dots, e_{k+\ell}^{(k+\ell)}] = f'}} c_2 \omega_2(h') \right).$$

Replacing each f' by the equivalent superposition of h' with projections, we obtain

$$\sum_{\substack{h' \in C^{(k)} \\ h'[g_1, \dots, g_k] = f}} c_1 \omega_1(h') + \sum_{\substack{h' \in C^{(\ell)} \\ h'[g'_1, \dots, g'_\ell] = f}} c_2 \omega_2(h'),$$

which is the result of applying the left-hand side of (6.3) to f . \square

We now prove Theorem 5.9, which states that for any finite set D and any finite $W \subset \mathbf{W}_D$, $\text{wPol}(\text{Imp}(W)) = \text{wClone}(W)$.

Proof. We first establish that for any $W \subset \mathbf{W}_D$ we have the inclusion $\text{wClone}(W) \subseteq \text{wPol}(\text{Imp}(W))$. To see this, observe that every operation in $C = \text{Clone}(\bigcup_{\omega \in W} \text{dom}(\omega))$ is a polymorphism of $\text{Imp}(W)$, by Definition 3.9, so $\overline{W} \subseteq \text{wPol}(\text{Imp}(W))$. Hence, we have that $\text{wClone}(W) \subseteq \text{wClone}(\text{wPol}(\text{Imp}(W)))$, which is equal to $\text{wPol}(\text{Imp}(W))$ by Proposition 6.3.

We will prove the reverse inclusion, $\text{wPol}(\text{Imp}(W)) \subseteq \text{wClone}(W)$, as follows. Given any weighting $\omega_0 \in \mathbf{W}_D$, we will show that either there exists a weighted relation $\varrho \in \text{Imp}(W)$ such that $\omega_0 \notin \text{wPol}(\{\varrho\})$ or else ω_0 is equal to a nonnegative weighted sum of superpositions of weightings in W , and hence $\omega_0 \in \text{wClone}(W)$.

We now give the details of this argument. Let k be the arity of ω_0 , and let $M = |D|^k$. We first observe that it is sufficient to consider weighted relations of arity M in $\text{Imp}(W)$. To see this, suppose there exists a weighted relation $\varrho \in \text{Imp}(W)$ with arity $N > M$ such that $\omega_0 \notin \text{wPol}(\{\varrho\})$ and let $\mathbf{x}_1, \dots, \mathbf{x}_k \in D^N$ be any set of tuples for which inequality (5.2) fails to hold for ω_0 and ϱ . Let \mathbf{X} be the $k \times N$ matrix whose rows are the tuples $\mathbf{x}_1, \dots, \mathbf{x}_k$. Since $N > M$ it follows that some of the columns in this matrix must be equal. Moreover, if the i th and j th columns of \mathbf{X} are equal, then so will be the i th and j th entries of the tuple $f(\mathbf{x}_1, \dots, \mathbf{x}_k)$ obtained by applying any $f \in \mathbf{O}_D^{(k)}$ to these k tuples.

Now let ϱ' be the weighted relation of arity $\leq M$ that depends only on the first of each of these repeated columns and takes the same values as ϱ takes on arguments with the appropriate entries repeated. Let \mathbf{X}' be the reduced form of \mathbf{X} (with repeated columns deleted). By this approach, we can construct ϱ' so that $\varrho' \in \text{Imp}(W)$, but \mathbf{X}' gives a certificate to show that $\omega_0 \notin \text{wPol}(\{\varrho'\})$, i.e., the rows of \mathbf{X}' form a list of tuples for which (5.2) fails to hold for ω_0 and ϱ' .

Moreover, if we have a weighted relation $\varrho \in \text{Imp}(W)$ with arity $N < M$ such that $\omega_0 \notin \text{wPol}(\{\varrho\})$, then ϱ can be extended to a weighted relation ϱ' of arity M that does not depend on the $M - N$ added inputs and, hence, is also contained in $\text{Imp}(W)$ and is such that $\omega_0 \notin \text{wPol}(\{\varrho'\})$.

By the argument given above, there exists a weighted relation $\varrho \in \text{Imp}(W)$ such that $\omega_0 \notin \text{wPol}(\{\varrho\})$ if and only if there exists a weighted relation ϱ_M of arity M in $\text{Imp}(W)$ such that $\omega_0 \notin \text{wPol}(\{\varrho_M\})$. Furthermore, by reordering the arguments of ϱ_M if necessary, we can assume that ϱ_M and ω_0 violate (5.2) on the particular list of tuples $\mathbf{x}_1, \dots, \mathbf{x}_k$ given by taking the rows of a matrix, \mathbf{X}_M , whose columns are precisely the tuples in D^k , ordered lexicographically.

By Definition 5.5, such a weighted relation ϱ_M exists if and only if the following system of inequalities can be satisfied for all $\omega \in W$ and all $\mathbf{t}_1, \dots, \mathbf{t}_{\text{ar}(\omega)} \in D^M$ such that $\varrho_M(\mathbf{t}_i)$ is defined for $i = 1, \dots, \text{ar}(\omega)$:

$$(6.4) \quad \sum_{g \in \text{dom}(\omega)} \omega(g) \varrho_M(g(\mathbf{t}_1, \dots, \mathbf{t}_{\text{ar}(\omega)})) \leq 0;$$

and, for the tuples $\mathbf{x}_1, \dots, \mathbf{x}_k$ forming the rows of \mathbf{X}_M , $\varrho_M(\mathbf{x}_i)$ is defined for $i = 1, \dots, k$ and

$$(6.5) \quad \sum_{f \in \text{dom}(\omega_0)} \omega_0(f) \varrho_M(f(\mathbf{x}_1, \dots, \mathbf{x}_k)) > 0.$$

There is a one-to-one correspondence between operations $g : D^k \rightarrow D$ and tuples $\mathbf{t}_g \in D^M$, where the tuple \mathbf{t}_g contains the list of values returned by the operation g applied to the columns of \mathbf{X}_M .

Set $C = \text{Clone}(\bigcup_{\omega \in W} \text{dom}(\omega))$. We observe that to satisfy inequality (6.4), for any $\omega \in W$, if $\varrho_M(\mathbf{t}_i)$ is defined for $i = 1, \dots, \text{ar}(\omega)$, then $\varrho_M(g(\mathbf{t}_1, \dots, \mathbf{t}_{\text{ar}(\omega)}))$ must be defined for all $g \in \text{dom}(\omega)$. To achieve this, it is sufficient to ensure that for all $g \in C^{(k)}$, $\varrho_M(\mathbf{t}_g)$ is defined. All other values of ϱ_M can be left undefined, as this just reduces the number of inequalities in the system.

Using superposition (Definition 5.2), we can rewrite inequalities (6.4) to obtain the following equivalent system: for all $\omega \in \overline{W}$ and all $g_1, \dots, g_{\text{ar}(\omega)} \in C^{(k)}$,

$$(6.6) \quad \sum_{f \in \text{dom}(\omega[g_1, \dots, g_{\text{ar}(\omega)}])} \omega[g_1, \dots, g_{\text{ar}(\omega)}](f) \varrho_M(f(\mathbf{x}_1, \dots, \mathbf{x}_k)) \leq 0.$$

The values of $\varrho_M(f(\mathbf{x}_1, \dots, \mathbf{x}_k))$ may be scaled to obtain a corresponding set of integer values y_f for each $f \in C^{(k)}$. Moreover, if we add any constant value to each y_f , this will not affect inequalities (6.6) and (6.5), since the sum of the values of any weighting is zero. Hence we may assume that $\sum_{f \in C^{(k)}} y_f = 0$ and then apply Lemma 6.1 to the resulting system of inequalities with $T = C^{(k)}$ and $S = \{\omega[g_1, \dots, g_{\text{ar}(\omega)}] \mid \omega \in \overline{W}, g_1, \dots, g_{\text{ar}(\omega)} \in C^{(k)}\}$. We conclude that either a solution ϱ_M exists, in which case $\omega_0 \notin \text{wPol}(\text{Imp}(\overline{W}))$, or else there exists a set of nonnegative rational numbers $\{x_{\omega[g_1, \dots, g_{\text{ar}(\omega)}]} \mid \omega \in \overline{W}, g_1, \dots, g_{\text{ar}(\omega)} \in C^{(k)}\}$ such that for every $f \in C^{(k)}$,

$$(6.7) \quad \sum_{\omega \in \overline{W}} \sum_{\substack{g_1, \dots, g_{\text{ar}(\omega)} \\ g_i \in C^{(k)}}} x_{\omega[g_1, \dots, g_{\text{ar}(\omega)}]} \omega[g_1, \dots, g_{\text{ar}(\omega)}](f) \geq \omega_0(f) + c$$

for some constant value c .

By Definition 5.1, adding the left-hand side of these inequalities over all f gives 0, and adding the values of $\omega_0(f)$ over all f also gives zero, so the value of c must be zero. Moreover, each inequality in (6.7) must actually be an equality. In other words, ω_0 is equal to a nonnegative weighted sum of superpositions of weighting in \overline{W} .

Hence, by Lemma 6.4 and Remark 5.1, ω_0 is equal to a proper superposition of some element $\omega'_0 \in \text{wClone}(W)$, so $\omega_0 \in \text{wClone}(W)$. \square

Remark 6.2. The proof of Theorem 5.9 demonstrates the decidability of the following question: for any finite $W \subseteq \mathbf{W}_D$ and any weighting ω defined on D , does ω belong to $\text{wClone}(W)$?

7. Necessary conditions for tractability. In this section, we will start to investigate the structure of weighted clones and hence establish some necessary conditions for any valued constraint language to be tractable.

Note that by Definition 3.8, the smallest possible clone of operations over a fixed set D is the set of all projection operations on D , which is denoted \mathbf{J}_D .

PROPOSITION 7.1. *For any nonempty finite set D , there are precisely two weighted clones with support \mathbf{J}_D . These are $\mathbf{W}_{\mathbf{J}_D}$ and $\mathbf{W}_{\mathbf{J}_D}^0$.*

Proof. Let W be a weighted clone with support \mathbf{J}_D .

If the weights assigned by every weighting $\omega \in W$ are all zero, then W is the zero-valued weighted clone $\mathbf{W}_{\mathbf{J}_D}^0$ described in Example 5.5.

Otherwise, there is some $\omega \in W$ (of arity k) such that ω assigns positive weight to some k -ary projections and negative weights to some of the others. (The sum of the weights is zero, by Definition 5.1.) If we form the superposition of ω with the sequence of projections g_1, g_2, \dots, g_k , where $g_i = e_a^{(k)}$ if $\omega(e_i^{(k)})$ is positive and

$g_i = e_b^{(k)}$ otherwise, then we obtain a new weighting $\omega[g_1, g_2, \dots, g_k]$ of \mathbf{J}_D which assigns some positive weight w to $e_a^{(k)}$ and $-w$ to $e_b^{(k)}$.

By adding appropriate multiples of such functions for each successive pair of indices a and b , we can obtain any desired weighting of \mathbf{J}_D . Hence, in this case W contains all possible weightings of \mathbf{J}_D , so $W = \mathbf{W}_{\mathbf{J}_D}$. \square

Any weighting ω which is defined only for projection operations will be called a *trivial weighting*,

PROPOSITION 7.2. *For any finite set D with at least two elements and any set of trivial weightings $W \subseteq \mathbf{W}_D$, $\text{Imp}(W)$ is NP-hard.*

Proof. If W contains only trivial weightings, then $\text{wClone}(W)$ has support \mathbf{J}_D , so it is equal to $\mathbf{W}_{\mathbf{J}_D}$ or $\mathbf{W}_{\mathbf{J}_D}^0$ by Proposition 7.1.

Every weighting in $\mathbf{W}_{\mathbf{J}_D}^0$ is a weighted polymorphism of any possible weighted relation by Definition 5.5. Hence $\text{Imp}(\mathbf{W}_{\mathbf{J}_D}^0) = \Phi_D$.

The weighted relations that are improved by *all* weightings are precisely those which take at most one value. Hence $\text{Imp}(\mathbf{W}_{\mathbf{J}_D}) = (\mathbf{R}_D)_{\sim}$.

In both cases the resulting valued constraint language is NP-hard for any D with two or more elements. Hence $\text{Imp}(\text{wClone}(W))$ is NP-hard, and hence $\text{Imp}(W)$ is NP-hard. \square

Now consider weightings whose values are all 0.

PROPOSITION 7.3. *For any finite set D with at least two elements and any set of zero-valued weightings $W \subseteq \mathbf{W}_D$, $\text{Imp}(W)$ is NP-hard.*

Proof. By Definition 5.5, a zero-valued weighting will be a weighted polymorphism of *any* weighted relation which is a total function (i.e., any weighted relation where all assignments are feasible). Some valued constraint languages containing only total functions are NP-hard [19]. For example, consider the valued constraint language consisting of the following total function:

$$\varrho_{\neq}(x, y) = \begin{cases} 0 & \text{if } x \neq y, \\ 1 & \text{otherwise.} \end{cases}$$

We observed in Example 2.1 that on the domain $\{0, 1\}$ the problem $\text{VCSP}(\{\varrho_{\neq}\})$ corresponds to the MAX-CUT problem which is known to be NP-hard. Over domains of size $k > 2$ this problem corresponds to the problem $\text{MAX-}k\text{-CUT}$, which is also known to be NP-hard. \square

Using the Galois connection developed in the previous sections, these two results tell us that any valued constraint language that is *not* NP-hard must have a weighted polymorphism which is nontrivial and assigns at least some nonzero weights. A weighting which assigns positive weight to at least one operation that is not a projection will be called a *positive weighting*.

COROLLARY 7.4. *For any finite set D with at least two elements, and any $\Gamma \subseteq \Phi_D$, either Γ is NP-hard or else $\text{wPol}(\Gamma)$ is a weighted clone containing some positive weightings.*

Proof. By Proposition 6.3, in all cases $\text{wPol}(\Gamma)$ is a weighted clone.

By Theorem 5.8, for any finite $\Gamma' \subseteq \Gamma$, $\text{Imp}(\text{wPol}(\Gamma')) = \text{wRelClone}(\Gamma')$. By Proposition 4.5, Theorem 2.4, and Theorem 4.3, if Γ' is NP-hard, then $\text{wRelClone}(\Gamma')$ is also NP-hard, so $\text{Imp}(\text{wPol}(\Gamma'))$ must be NP-hard.

Conversely, if Γ' is not NP-hard, then the same argument shows that $\text{Imp}(\text{wPol}(\Gamma'))$ is not NP-hard, so by Propositions 7.2 and 7.3, $\text{wPol}(\Gamma')$ must contain some weightings that are nontrivial and some weightings that are not zero-valued.

Choose a weighting $\omega \in \text{wPol}(\Gamma)$ that is not zero-valued and a weighting ω' that is nontrivial (but may be zero-valued). If ω assigns positive weight to any nonprojection, then it is a positive weighting and we are done.

Otherwise, we have that ω assigns positive weight to some projections and negative weight to some other projections. Let f be an operation on which ω' is defined that is not a projection, and let k be the arity of f . If we form the superposition of ω with the sequence of functions $g_1, g_2, \dots, g_{\text{ar}(\omega)}$, where $g_i = f$ if $\omega(e_i^{\text{ar}(\omega)})$ is positive, and $g_i = e_1^{(k)}$ otherwise, then we obtain a new weighting $\omega[g_1, g_2, \dots, g_{\text{ar}(\omega)}]$ which assigns positive weight to f (see Example 5.1). \square

Assuming that $P \neq \text{NP}$, this result tells us that *tractable* valued constraint languages are associated with certain kinds of weighted clones.

To obtain further information about the weighted clones associated with tractable valued constraint languages, we now consider some special kinds of operations. For any $k \geq 2$, a k -ary operation f is called *sharp* if f is not a projection, but the operation obtained by equating any two inputs in f is a projection [27]. In other words, f is sharp if for all $i, j \in \{1, \dots, k\}$ with $i < j$, there exists an index $m \in \{1, \dots, k-1\}$ such that f satisfies the identity $f(x_1, x_2, \dots, x_{j-1}, x_i, x_j, x_{j+1}, \dots, x_{k-1}) = x_m$. For instance, when $k = 7$, choosing $i = 3$ and $j = 5$, we get that f satisfies the identity $f(x_1, x_2, x_3, x_4, x_3, x_5, x_6) = x_m$ for some $1 \leq m \leq 6$.

THEOREM 7.5. *Any weighted clone W containing positive weightings must contain a weighting that assigns positive weight to either*

1. *a set of unary operations that are not projections or*
2. *a set of sharp operations.*

Proof. Let ω be a positive weighting in W with the smallest possible arity, k . If $k = 1$, then we are done. Otherwise, we consider the weightings

$$\omega[e_1^{(k-1)}, e_2^{(k-1)}, \dots, e_{j-1}^{(k-1)}, e_i^{(k-1)}, e_j^{(k-1)}, \dots, e_{k-1}^{(k-1)}]$$

for all i, j with $1 \leq i < j \leq k$.

Each of these weightings has arity $k-1$ so, by the choice of ω , must not assign positive weight to any operation except (possibly) projections. Hence all nonprojection operations assigned positive weight by ω are sharp. \square

We can obtain further details about these weighted clones by considering the possible types of sharp operations.

First, we observe that all sharp operations must satisfy the identity $f(x, x, \dots, x) = x$; such operations are called *idempotent*.

Ternary sharp operations may be classified according to their values on tuples of the forms (x, x, y) , (x, y, x) , and (y, x, x) , which must be equal to either x or y . There are precisely eight possibilities, as listed in Table 7.1.

The first column in Table 7.1 corresponds to operations that satisfy the identities $f(x, x, y) = f(x, y, x) = f(y, x, x) = x$ for all $x, y \in D$; such operations are called *majority* operations. The last column in the table corresponds to operations that satisfy

TABLE 7.1
Sharp ternary operations.

Input	1	2	3	4	5	6	7	8
(x, x, y)	x	x	x	x	y	y	y	y
(x, y, x)	x	x	y	y	x	x	y	y
(y, x, x)	x	y	x	y	x	y	x	y

the identities $f(x, x, y) = f(x, y, x) = f(y, x, x) = y$ for all $x, y \in D$; such operations are called *minority* operations. Columns 4, 6, and 7 in Table 7.1 correspond to operations that satisfy the identities $f(y, y, x) = f(x, y, x) = f(y, x, x) = y$ for all $x, y \in D$ (up to permutations of inputs); such operations are called *Pixley* operations [41].

For any $k \geq 3$, a k -ary operation f is called a *semiprojection* if it is not a projection, but there is an index $i \in \{1, \dots, k\}$ such that $f(x_1, \dots, x_k) = x_i$ for all $x_1, \dots, x_k \in D$ such that x_1, \dots, x_k are not pairwise distinct. In other words, a semiprojection is a particular form of sharp operation where the operation obtained by equating any two inputs is always the *same* projection. Columns 2, 3, and 5 in Table 7.1 correspond to semiprojections.

The following lemma shows that the only sharp operations of arity $k \geq 4$ are semiprojections.

LEMMA 7.6 (Świerczkowski's lemma [45]). *Given an operation of arity ≥ 4 , if every operation arising from the identification of two variables is a projection, then these projections coincide.*

Hence we may refine Theorem 7.5 to obtain the following corollary.

COROLLARY 7.7. *Any weighted clone W containing positive weightings must contain a weighting that assigns positive weight to either*

1. *a set of unary operations that are not projections; or*
2. *a set of binary idempotent operations that are not projections; or*
3. *a set of ternary operations that are majority operations, minority operations, Pixley operations, or semiprojections; or*
4. *a set of k -ary semiprojections (for some $k > 3$).*

Corollary 7.7 can be used to guide the search for tractable valued constraint languages, as we illustrate in the next section.

8. Classification of Boolean valued constraint languages. In this section, we consider the special case of valued constraint languages over a 2-valued domain, such as the Boolean domain $D = \{0, 1\}$.

There are only four unary operations on the Boolean domain, and one of these is the projection operation $e_1^{(1)}$, which is the identity operation. The remaining three unary operations are the operations given by $f(x) = 0$, $f(x) = 1$, and $f(x) = 1 - x$. These will be referred to as constant 0, constant 1, and inversion.

There are only two binary idempotent operations on the Boolean domain that are not projections: the operations min and max. The only sharp ternary operations are the unique majority operation (which we will call Mjrty), the unique minority operation (which we will call Mnrty), and three Pixley operations. There are no semiprojections.

Hence we can refine Corollary 7.7 even further in the special case of the Boolean domain, to limit the possibilities for weighted clones associated with tractable valued constraint languages to just nine cases.

THEOREM 8.1. *Any weighted clone W on the Boolean domain that contains positive weightings must contain a weighting ω that assigns positive weight to either*

1. *exactly one of the unary operations constant 0, constant 1, or inversion;*
2. *exactly one of the binary operations min and max or both of them equally; or*
3. *exactly one of the ternary operations Mjrty and Mnrty or both of them with $\omega(\text{Mjrty}) = 2\omega(\text{Mnrty})$.*

Proof. Let C be the support of W , and let ω be a positive weighting in W with minimal possible arity. Since there are no semiprojections on the Boolean domain, Corollary 7.7 tells us that ω is either unary, binary, or ternary.

Consider first the case when ω is unary. Since there are just three unary operations on the Boolean domain that are not projections, scale ω so it assigns weight -1 to the projection $e_1^{(1)}$, weight a to the constant 0 operation f_0 , weight b to the constant 1 operation f_1 , and weight c to the inversion operation f_- . If $c = 1$, then ω assigns positive weight only to f_- , and we are done. Otherwise, if $f_- \in C$, and hence c is defined, we consider the weighting $\omega' = \frac{1}{c+1}\omega + \frac{c}{c+1}\omega[f_-]$. It is straightforward to check that ω' assigns weight $c - 1$ to $e_1^{(1)}$, weight a to f_0 , weight b to f_1 , and weight 0 to f_- . By Lemma 6.4, ω' belongs to W .

If $a = 1$, then ω' assigns positive weight only to f_0 and we are done. Otherwise, if $f_0 \in C$, and hence a is defined, we consider the weighting $\omega'' = \omega' + \frac{a}{b}\omega'[f_0]$. It is straightforward to check that ω'' assigns positive weight only to f_1 . By Lemma 6.4, ω'' belongs to W .

Next consider the case when ω is binary. By Corollary 7.7 and our observations above about the possible binary idempotent operations on the Boolean domain, we know that ω assigns positive weight only to one or both of the operations min and max. If either of these weights is undefined (because the corresponding function does not belong to C) or zero, then we are done, so assume that ω assigns positive weight to both min and max. By taking the weighting $\omega + \omega[e_2^{(2)}, e_1^{(2)}]$, with a suitable scaling, we can obtain a weighting $\omega_a \in W$ that assigns weight -1 to both binary projections, weight a to min, and weight $2 - a$ to max for some $0 < a < 2$.

If $a < 1$, then the weighting $\omega_a + \frac{a}{1-a}\omega_a[\text{min}, \text{max}]$ assigns positive weight only to max. If $a > 1$, then the weighting $\omega_a + \frac{2-a}{a-1}\omega_a[\text{min}, \text{max}]$ assigns positive weight only to min. If $a = 1$, then ω_a assigns equal weight to min and max.

Finally, we consider the case when ω is ternary. By Corollary 7.7 and our observations above about the possible ternary sharp operations on the Boolean domain, we know that ω assigns positive weight to some subset of Mjrty, Mnrty, and the three Boolean Pixley operations f_1 , f_2 , and f_3 (corresponding to the fourth, sixth, and seventh columns of Table 7.1). We note that $f_1[e_2^{(3)}, e_3^{(3)}, e_1^{(3)}] = f_3$, $f_2[e_2^{(3)}, e_3^{(3)}, e_1^{(3)}] = f_1$, and $f_3[e_2^{(3)}, e_3^{(3)}, e_1^{(3)}] = f_2$. Hence, if ω assigns positive weight to any Pixley operation, then we have that W also contains the weighting $\omega' = \omega + \omega[e_2^{(3)}, e_3^{(3)}, e_1^{(3)}] + \omega[e_3^{(3)}, e_1^{(3)}, e_2^{(3)}]$, which assigns equal negative weight to each projection and equal positive weight to each Pixley operation. By a suitable scaling we shall assume that ω' assigns weight -1 to each projection.

Suppose first that ω' assigns positive weight to at least one of Mjrty and Mnrty and assigns weight $0 < w < 1$ to the three Pixley operations. We note that $f_i[f_1, f_2, f_3] = e_i^{(3)}$ for each $i = 1, 2, 3$. Moreover, $\text{Mjrty}[f_1, f_2, f_3] = \text{Mnrty}$ and $\text{Mnrty}[f_1, f_2, f_3] = \text{Mjrty}$. Thus, the weighting $\omega'' = \omega' + w\omega'[f_1, f_2, f_3]$ is nonzero and assigns weight 0 to each Pixley operation and equal negative weight to all projections. By Lemma 6.4, $\omega'' \in W$.

Assume that ω'' assigns positive weight to both Mnrty and Mjrty. By taking a suitable scaling, we can obtain a weighting $\omega_a \in W$ that assigns weight -1 to all three projections, weight a to Mnrty, and weight $3 - a$ to Mjrty for some $0 < a < 3$.

If $a < 1$, then the weighting $\omega_a + \frac{a}{1-a}\omega_a[\text{Mjrty}, \text{Mjrty}, \text{Mnrty}]$ assigns positive weight only to Mjrty. If $a > 1$, then the weighting $\omega_a + \frac{3-a}{a-1}\omega_a[\text{Mjrty}, \text{Mjrty}, \text{Mnrty}]$ assigns positive weight only to Mnrty. If $a = 1$, then ω_a assigns positive weight to both Mjrty and Mnrty in the ratio 2:1.

The only remaining case is when ω' assigns positive weight 1 just to the three Pixley operations. In this case we note that $f_1[e_1^{(3)}, e_2^{(3)}, f_1] = \text{Mjrty}$, $f_2[e_1^{(3)}, e_2^{(3)}, f_1] = e_1$,

and $f_3[e_1^{(3)}, e_2^{(3)}, f_1] = e_2$. Thus the function $\mu_1 = \omega'[e_1, e_2, f_1]$ assigns weight -1 to f_1 , $+1$ to Mjrty, and 0 otherwise. For $i = 2, 3$, we can obtain in a similar way a function μ_i , which assigns weight -1 to f_i and $+1$ to Mjrty. Then the weighting $\omega'' = \omega' + \mu_1 + \mu_2 + \mu_3$ will assign positive weight only to Mjrty. Again, by Lemma 6.4, $\omega'' \in W$. \square

Each of the nine types of weightings mentioned in Theorem 8.1 can be supported by a different clone, so these nine types of weightings can each generate different weighted clones.

Using the Galois connection developed above, this result tells us that any tractable valued constraint language over the Boolean domain must have as a weighted polymorphism one of nine specific kinds of weightings. Eight of these can be shown to be sufficient to ensure tractability using the results of [19]. The only remaining case is the unary weighting that assigns positive weight to the inversion operation only. Our next result shows that having a weighted polymorphism of this kind is *not* a sufficient condition for tractability on its own, but if a language has any additional positive weightings as weighted polymorphisms which are not implied by this one, then it will be tractable.

COROLLARY 8.2. *Any weighted clone W on the Boolean domain that contains positive weightings satisfies exactly one of the following:*

1. $W = \text{wClone}(\{\omega_{\neg}\} \cup \mathbf{W}_C^0)$ for some unary weighting ω_{\neg} that assigns positive weight to the inversion operation only, where C is the support of W ; in this case $\text{Imp}(W)$ is NP-hard.
2. W contains one of the eight other kinds of weighting listed in Theorem 8.1; in each of these eight cases $\text{Imp}(W)$ is tractable.

Proof. By Theorem 8.1, W must contain either a weighting ω_{\neg} that assigns positive weight to the inversion operation only or at least one of the eight other kinds of weighting listed in Theorem 8.1 (or both).

If $W = \text{wClone}(\{\omega_{\neg}\} \cup \mathbf{W}_C^0)$ for some unary weighting ω_{\neg} that assigns positive weight to the inversion operation only and C is the support of W , then we are in case 1. In this case, the weighted relation ϱ_{\neq} defined in Example 2.1 is an element of $\text{Imp}(W)$, so $\text{Imp}(W)$ is NP-hard (see the proof of Proposition 7.3).

If W contains a suitable weighting ω_{\neg} , but $W \neq \text{wClone}(\{\omega_{\neg}\} \cup \mathbf{W}_C^0)$, then W must also contain a nonzero weighting ω of minimal possible arity such that $\omega \notin \text{wClone}(\{\omega_{\neg}\} \cup \mathbf{W}_C^0)$.

If ω is unary, then we can use the same argument as in the proof of Theorem 8.1 to show that W must contain a unary operation that assigns positive weight to the constant 1 operation only or the constant 0 operation only.

If ω is not unary, then $\omega[e_1^{(1)}, \dots, e_1^{(1)}]$ must lie in $\text{wClone}(\omega_{\neg})$, so ω assigns positive weights only to operations f such that $f(x, \dots, x) = x$ or $f(x, \dots, x) = 1 - x$. If $\omega(f) = a > 0$ for some f such that $f(x, \dots, x) = 1 - x$, then we consider the weighting $\omega' = \omega + a\omega_{\neg}[f]$ and note that $\omega'(f) = 0$. By repeating this process we obtain a weighting ω'' which assigns positive weight only to operations f such that $f(x, \dots, x) = x$. Since ω has minimal arity, these must be sharp operations, so we can proceed as in the proof of Theorem 8.1 to show that case 2 holds.

Each of the eight types of weightings in case 2 is sufficient to ensure the tractability of $\text{Imp}(W)$ by the results of [19]. \square

The corresponding classification for valued constraint languages over the Boolean domain was obtained in [19] using a more intricate argument involving the explicit

construction of gadgets to express particular weighted relations. Here we have considered only the properties of weighted clones.

Example 8.1. The weighted relation $\varrho_{=}$ defined in Example 2.4 has as a weighted polymorphism the weighting ω_{sub} defined in Example 5.6 which assigns equal positive weight to max and min.

Hence the valued constraint language $\Gamma = \{\varrho_{=}\}$ is tractable and remains tractable if we add to Γ any other weighted relations that have this weighting as a weighted polymorphism. For example, we may add unary weighted relations with a single allowed value, which allow us to fix individual variables to a desired value, and still retain tractability.

Such an extended valued constraint language allows us to express, for example, the (s, t) -MIN-CUT problem [40].

Example 8.2. The weighted relation ϱ_{\neq} defined in Example 2.1 has a unary weighted polymorphism that assigns positive weight only to the inversion function. It has none of the other eight types of weightings listed in Theorem 8.1.

It follows that $\varrho_{\neq} \notin \text{Express}(\{\varrho_{=}\})$ (cf. Example 2.4).

9. Conclusions. We have presented an algebraic theory of valued constraint languages that generalizes and extends the algebraic theory developed over the past few years to study the complexity of the classical constraint satisfaction problem. We have shown that the complexity of any valued constraint language over a finite domain with rational-valued costs is determined by certain algebraic properties which we have called weighted polymorphisms.

When the weights are all zero, the optimization problem we are considering collapses to the classical CSP. In previous work [36, 35] it has been shown that every tractable constraint language for the CSP can be characterized by an associated clone of operations. That work initiated the use of algebraic properties in the search for tractable constraint languages, an area that has seen considerable activity in recent years; see, for instance, [17, 16, 15, 14, 37, 28, 4, 2, 3, 5].

The results in this paper show that a similar result holds for the valued constraint satisfaction problem: every tractable valued constraint language is characterized by an associated weighted clone. We therefore hope that our results here will provide a similar impetus for the investigation of a much broader class of discrete optimization problems. For example, a recent result on the power of linear programming for valued constraint languages [47] provides a characterization of languages solvable by a standard LP relaxation in terms of algebraic properties similar to weighted polymorphisms. Moreover, another recent result that classifies the complexity of all valued constraint languages consisting of total weighted relations [48] also relies heavily on algebraic properties of this kind.

Many questions about the complexity of discrete optimization problems over finite domains can now be translated into questions about the structure of weighted clones. This provides a new approach to tackling such questions by investigating the algebraic properties of weighted clones.

This work raises many open questions. In particular, we would like to know whether our key results, Theorems 5.8 and 5.9, can be extended to infinite valued constraint languages and infinite sets of weightings or to infinite domains of values [8, 9]. In the case of the CSP, considerable progress resulted from showing that clones of idempotent operations play a special role [15]. Moreover, it has been shown that rather than specific clones of operations one can consider large families of abstract algebras, known as varieties [15]. It remains to be seen whether such notions have interesting counterparts in the theory of weighted clones.

REFERENCES

- [1] L. BARTO, *The dichotomy for conservative constraint satisfaction problems revisited*, in Proceedings of the 26th IEEE Symposium on Logic in Computer Science (LICS'11), IEEE Computer Society, 2011, pp. 301–310.
- [2] L. BARTO AND M. KOZIK, *Constraint satisfaction problems of bounded width*, in Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS'09), IEEE Computer Society, 2009, pp. 461–471.
- [3] L. BARTO, M. KOZIK, M. MARÓTI, AND T. NIVEN, *CSP dichotomy for special triads*, Proc. Amer. Math. Soc., 137 (2009), pp. 2921–2934.
- [4] L. BARTO, M. KOZIK, AND T. NIVEN, *The CSP dichotomy holds for digraphs with no sources and no sinks (a positive answer to a conjecture of Bang-Jensen and Hell)*, SIAM J. Comput., 38 (2009), pp. 1782–1802.
- [5] J. BERMAN, P. IDZIAK, P. MARKOVIĆ, R. MCKENZIE, M. VALERIOTE, AND R. WILLARD, *Varieties with few subalgebras of powers*, Trans. Amer. Math. Soc., 362 (2010), pp. 1445–1473.
- [6] S. BISTARELLI, U. MONTANARI, AND F. ROSSI, *Semiring-based constraint satisfaction and optimisation*, J. ACM, 44 (1997), pp. 201–236.
- [7] S. BISTARELLI, U. MONTANARI, F. ROSSI, T. SCHIEX, G. VERFAILLIE, AND H. FARGIER, *Semiring-based CSPs and valued CSPs: Frameworks, properties, and comparison*, Constraints, 4 (1999), pp. 199–240.
- [8] M. BODIRSKY, *Constraint satisfaction problems with infinite templates*, in Complexity of Constraints, Lecture Notes in Comput. Sci. 5250, Springer, New York, 2008, pp. 196–228.
- [9] M. BODIRSKY AND J. KÁRA, *The complexity of temporal constraint satisfaction problems*, J. ACM, 57 (2010).
- [10] V. G. BODNARČUK, L. A. KALUŽNIN, V. N. KOTOV, AND B. A. ROMOV, *Galois theory for Post algebras I*, Cybernet. Systems Anal., 5 (1969), pp. 243–252.
- [11] F. BÖRNER, *Basics of Galois connections*, in Complexity of Constraints, Lecture Notes in Comput. Sci. 5250, Springer, New York, 2008, pp. 38–67.
- [12] E. BOROS AND P. L. HAMMER, *Pseudo-Boolean optimization*, Discrete Appl. Math., 123 (2002), pp. 155–225.
- [13] A. BULATOV, *A Graph of a Relational Structure and Constraint Satisfaction Problems*, in Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS'04), IEEE Computer Society, 2004, pp. 448–457.
- [14] A. BULATOV, *A dichotomy theorem for constraint satisfaction problems on a 3-element set*, J. ACM, 53 (2006), pp. 66–120.
- [15] A. BULATOV, A. KROKHIN, AND P. JEAVONS, *Classifying the complexity of constraints using finite algebras*, SIAM J. Comput., 34 (2005), pp. 720–742.
- [16] A. A. BULATOV, *Complexity of conservative constraint satisfaction problems*, ACM Trans. Comput. Log., 12 (2011).
- [17] A. A. BULATOV, A. A. KROKHIN, AND P. G. JEAVONS, *The complexity of maximal constraint languages*, in Proceedings of the 33rd ACM Symposium on Theory of Computing (STOC'01), 2001, pp. 667–674.
- [18] D. A. COHEN, M. C. COOPER, AND P. G. JEAVONS, *An algebraic characterisation of complexity for valued constraints*, in Proceedings of the 12th International Conference on Principles and Practice of Constraint Programming (CP'06), Lecture Notes in Comput. Sci. 4204, Springer, New York, 2006, pp. 107–121.
- [19] D. A. COHEN, M. C. COOPER, P. G. JEAVONS, AND A. A. KROKHIN, *The complexity of soft constraint satisfaction*, Artificial Intelligence, 170 (2006), pp. 983–1016.
- [20] D. A. COHEN, P. CREED, P. G. JEAVONS, AND S. ŽIVNÝ, *An algebraic theory of complexity for valued constraints: Establishing a Galois connection*, in Proceedings of the 36th International Symposium on Mathematical Foundations of Computer Science (MFCS'11), Lecture Notes in Comput. Sci. 6907, Springer, New York, 2011, pp. 231–242.
- [21] M. C. COOPER AND S. ŽIVNÝ, *Hybrid tractability of valued constraint problems*, Artificial Intelligence, 175 (2011), pp. 1555–1569.
- [22] M. C. COOPER AND S. ŽIVNÝ, *Tractable triangles and cross-free convexity in discrete optimisation*, J. Artificial Intelligence Res., 44 (2012), pp. 455–490.
- [23] Y. CRAMA AND P. L. HAMMER, *Boolean Functions: Theory, Algorithms, and Applications*, Cambridge University Press, Cambridge, UK, 2011.
- [24] P. CREED AND S. ŽIVNÝ, *On minimal weighted clones*, in Proceedings of the 17th International Conference on Principles and Practice of Constraint Programming (CP'11), Lecture Notes in Comput. Sci. 6876, Springer, New York, 2011, pp. 210–224.
- [25] N. CREIGNOU, *A dichotomy theorem for maximum generalized satisfiability problems*, J. Comput. System Sci., 51 (1995), pp. 511–522.

- [26] N. CREIGNOU, S. KHANNA, AND M. SUDAN, *Complexity Classification of Boolean Constraint Satisfaction Problems*, SIAM Monogr. Discrete Math. Appl. 7, SIAM, Philadelphia, 2001.
- [27] B. CSÁKÁNY, *Minimal clones—a minicourse*, Algebra Universalis, 54 (2005), pp. 73–89.
- [28] V. DEINEKO, P. JONSSON, M. KLASSON, AND A. KROKHIN, *The approximability of Max CSP with fixed-value constraints*, J. ACM, 55 (2008).
- [29] T. FEDER AND M. Y. VARDI, *The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory*, SIAM J. Comput., 28 (1998), pp. 57–104.
- [30] M. R. GAREY AND D. S. JOHNSON, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W.H. Freeman, San Francisco, 1979.
- [31] D. GEIGER, *Closed systems of functions and predicates*, Pacific J. Math., 27 (1968), pp. 95–100.
- [32] G. GUTIN, P. HELL, A. RAFIEY, AND A. YEO, *A dichotomy for minimum cost graph homomorphisms*, European J. Combin., 29 (2008), pp. 900–911.
- [33] M. GYSSENS, P. G. JEAUVONS, AND D. A. COHEN, *Decomposing constraint satisfaction problems using database techniques*, Artificial Intelligence, 66 (1994), pp. 57–89.
- [34] S. IWATA, L. FLEISCHER, AND S. FUJISHIGE, *A combinatorial strongly polynomial algorithm for minimizing submodular functions*, J. ACM, 48 (2001), pp. 761–777.
- [35] P. G. JEAUVONS, *On the algebraic structure of combinatorial problems*, Theoret. Comput. Sci., 200 (1998), pp. 185–204.
- [36] P. G. JEAUVONS, D. A. COHEN, AND M. GYSSENS, *Closure properties of constraints*, J. ACM, 44 (1997), pp. 527–548.
- [37] P. JONSSON, M. KLASSON, AND A. KROKHIN, *The approximability of three-valued MAX CSP*, SIAM J. Comput., 35 (2006), pp. 1329–1349.
- [38] V. KOLMOGOROV AND S. ŽIVNÝ, *The complexity of conservative valued CSPs*, J. ACM, 60 (2013).
- [39] S. L. LAURITZEN, *Graphical Models*, Oxford University Press, New York, 1996.
- [40] G. L. NEMHAUSER AND L. A. WOLSEY, *Integer and Combinatorial Optimization*, John Wiley & Sons, New York, 1988.
- [41] A. F. PIXLEY, *Distributivity and permutability of congruence relations in equational classes of algebras*, Proc. Amer. Math. Soc., 14 (1963), pp. 105–109.
- [42] T. J. SCHAEFER, *The complexity of satisfiability problems*, in Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC’78), ACM, 1978, pp. 216–226.
- [43] A. SCHRIJVER, *Theory of Linear and Integer Programming*, John Wiley, New York, 1986.
- [44] A. SCHRIJVER, *A combinatorial algorithm minimizing submodular functions in strongly polynomial time*, J. Combin. Theory Ser. B, 80 (2000), pp. 346–355.
- [45] S. ŚWIERCZKOWSKI, *Algebras which are independently generated by every n elements*, Fund. Math., 49 (1960), pp. 93–104.
- [46] R. 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), 2010, pp. 657–668.
- [47] J. THAPPER AND S. ŽIVNÝ, *The power of linear programming for valued CSPs*, in Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS’12), IEEE, 2012, pp. 669–678.
- [48] J. THAPPER AND S. ŽIVNÝ, *The complexity of finite-valued CSPs*, in Proceedings of the 45th ACM Symposium on the Theory of Computing (STOC’13), ACM, 2013, pp. 695–704.
- [49] M. J. WAINWRIGHT AND M. I. JORDAN, *Graphical models, exponential families, and variational inference*, Found. Trends Machine Learning, 1 (2008), pp. 1–305.
- [50] T. WERNER, *A Linear Programming Approach to Max-Sum Problem: A Review*, IEEE Trans. Pattern Analysis and Machine Intelligence, 29 (2007), pp. 1165–1179.
- [51] S. ŽIVNÝ, *The Complexity of Valued Constraint Satisfaction Problems*, Cognitive Technologies, Springer, New York, 2012.
- [52] S. ŽIVNÝ, D. A. COHEN, AND P. G. JEAUVONS, *The expressive power of binary submodular functions*, Discrete Appl. Math., 157 (2009), pp. 3347–3358.