

A Galois Connection for Weighted (Relational) Clones of Infinite Size

PETER FULLA and STANISLAV ŽIVNÝ, University of Oxford, United Kingdom

A Galois connection between clones and relational clones on a fixed finite domain is one of the cornerstones of the so-called algebraic approach to the computational complexity of non-uniform Constraint Satisfaction Problems (CSPs). Cohen et al. established a Galois connection between *finitely-generated* weighted clones and *finitely-generated* weighted relational clones [SICOMP'13], and asked whether this connection holds in general. We answer this question in the affirmative for weighted (relational) clones with *real* weights and show that the complexity of the corresponding valued CSPs is preserved.

CCS Concepts: • **Theory of computation** → **Problems, reductions and completeness**;

Additional Key Words and Phrases: discrete optimisation, Galois connection, universal algebra, valued constraint satisfaction problems, weighted polymorphisms

ACM Reference Format:

Peter Fulla and Stanislav Živný. 2016. A Galois Connection for Weighted (Relational) Clones of Infinite Size. *ACM Trans. Comput. Theory* V, N, Article A (January YYYY), 21 pages.
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

The constraint satisfaction problem (CSP) is a general framework capturing decision problems arising in many contexts of computer science [Hell and Nešetřil 2008]. The CSP is NP-hard in general but there has been much success in finding tractable fragments of the CSP by restricting the types of relations allowed in the constraints. A set of allowed relations has been called a *constraint language* [Feder and Vardi 1998]. For some constraint languages, the associated constraint satisfaction problems with constraints chosen from that language are solvable in polynomial-time, whilst for other constraint languages this class of problems is NP-hard [Feder and Vardi 1998]; these 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 constraint languages over two-element domains [Schaefer 1978], three-element domains [Bulatov 2006], for conservative (containing all unary relations) constraint languages [Bulatov 2011], for maximal constraint languages [Bulatov et al. 2001; Bulatov 2004], for graphs (corresponding to languages containing a single binary symmetric relation) [Hell and Nešetřil 1990], and for digraphs (corresponding to languages containing a single binary relation) without sources and sinks [Barto et al. 2009]. The most successful approach to classifying the complexity of constraint languages has been the algebraic approach [Jeavons et al. 1997; Bulatov et al. 2005; Barto and Kozik 2014]. The dichotomy conjecture of Feder and Vardi [Feder and Vardi 1998] asserts that every constraint language is either

The authors were supported by a Royal Society Research Grant. Stanislav Živný was supported by a Royal Society University Research Fellowship. A part of this work appeared in *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, 2015 [Fulla and Živný 2015].

Authors' address: P. Fulla and S. Živný, Department of Computer Science, University of Oxford, Oxford OX1 3QD, United Kingdom.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© YYYY ACM. 1942-3454/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

tractable or NP-hard, and the algebraic refinement of the conjecture specifies the precise boundary between tractable and NP-hard languages [Bulatov et al. 2005].

The *valued* constraint satisfaction problem (VCSP) is a general framework that captures not only feasibility problems but also optimisation problems [Cohen et al. 2006; Živný 2012; Jeavons et al. 2014]. A VCSP instance represents each constraint by a *weighted relation*, which is a \mathbb{Q} -valued function where $\mathbb{Q} = \mathbb{Q} \cup \{\infty\}$, and the goal is to find a labelling of variables minimising the sum of the values assigned by the constraints to that labelling. Tractable fragments of the VCSP have been identified by restricting the types of allowed weighted relations that can be used to define the valued constraints. A set of allowed weighted relations has been called a *valued constraint language* [Cohen et al. 2006]. Dichotomy theorems, which classify each possible valued constraint language as either tractable or NP-hard, have been established for valued constraint languages over two-element domains [Cohen et al. 2006], for conservative (containing all $\{0, 1\}$ -valued unary cost functions) valued constraint languages [Kolmogorov and Živný 2013], for finite-valued (all weighted relations are \mathbb{Q} -valued) constraint languages [Thapper and Živný 2015b]. Moreover, it has been shown that a dichotomy for constraint languages implies a dichotomy for valued constraint languages [Kolmogorov et al. 2015a]. Finally, the power of the basic linear programming relaxation [Thapper and Živný 2012; Kolmogorov et al. 2015b] and the power of the Sherali-Adams relaxation [Thapper and Živný 2015d; 2015c] for valued constraint languages have been completely characterised.

Cohen et al. have introduced an algebraic theory of weighted clones [Cohen et al. 2013], further extended in [Thapper and Živný 2015a; Kozik and Ochremiak 2015], for classifying the computational complexity of valued constraint languages. This theory establishes a one-to-one correspondence between valued constraint languages closed under expressibility (which does not change the complexity of the associated class of optimisation problems), called weighted relational clones, and weighted clones [Cohen et al. 2013]. This is an extension of (a part of) the algebraic approach to CSPs which relies on a one-to-one correspondence between constraint languages closed under pp-definability (which does not change the complexity of the associated class of decision problems), called relational clones, and clones [Bulatov et al. 2005], thus making it possible to use deep results from universal algebra. This theory has been developed primarily as an aid for studying the computational complexity of valued CSPs (and indeed, recent progress on valued CSPs [Thapper and Živný 2012; 2015b; Kolmogorov et al. 2015b; Thapper and Živný 2015d; Kolmogorov et al. 2015a] and on special cases of valued CSPs [Uppman 2013] heavily rely on the theory introduced in [Cohen et al. 2013]), but the theory is interesting in its own right [Thapper and Živný 2015a; Kozik and Ochremiak 2015; Vančura 2014; Vaiceniavičius 2014].

Contributions

The Galois connection between weighted clones and weighted relational clones established in [Cohen et al. 2013] was proved only for weighted (relational) clones generated by a *finite* set. The authors asked whether such a correspondence holds also for weighted (relational) clones in general. In this paper we answer this question in the affirmative.

Firstly, we show that the Galois connection from [Cohen et al. 2013] (using only rational weights) does *not* work for general weighted (relational) clones. Secondly, we alter the definition of weighted (relational) clones and establish a new Galois connection that holds even when the generating set has an infinite size. We allow weighted relations and weightings to assign real weights instead of rational, require weighted

relational clones to be closed under operator Opt and to be topologically closed, and prove that these changes preserve tractability of a constraint language (albeit only in an approximate sense).

Including the Opt operator (sometimes called argmin)¹ in the definition of weighted relational clones simplifies the structure of the space of all weighted clones, and guarantees that every non-projection polymorphism of a weighted relational clone Γ is assigned a positive weight by some weighted polymorphism of Γ . Indeed, including the Opt operator is very natural and can be used to simplify several results in [Kozik and Ochremiak 2015]. Real weights (as opposed to rational weights) were previously used, in the context of valued CSPs, in [Thapper and Živný 2015b] and our results confirm that real weights are necessary when studying infinite weighted (relational) clones.

The proof of the Galois connection in [Cohen et al. 2013] relies on results on linear programming duality; we use their generalisation from the theory of convex optimisation in order to establish the connection even for infinite sets.

2. BACKGROUND

2.1. Valued CSPs

Throughout the paper, let D be a fixed finite set of size at least two.

Definition 2.1. An m -ary relation² over D is any mapping $\phi : D^m \rightarrow \{c, \infty\}$ for some $c \in \mathbb{Q}$. We denote by $\mathbf{R}_D^{(m)}$ the set of all m -ary relations and let $\mathbf{R}_D = \bigcup_{m \geq 1} \mathbf{R}_D^{(m)}$.

Given an m -tuple $\mathbf{x} \in D^m$, we denote its i th entry by $\mathbf{x}[i]$ for $1 \leq i \leq m$.

Let $\overline{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$ denote the set of rational numbers with (positive) infinity.

Definition 2.2. An m -ary weighted relation over D is any mapping $\gamma : D^m \rightarrow \overline{\mathbb{Q}}$. We denote by $\Phi_D^{(m)}$ the set of all m -ary weighted relations and let $\Phi_D = \bigcup_{m \geq 1} \Phi_D^{(m)}$.

From Definition 2.2 we have that relations are a special type of weighted relations.

Example 2.3. An important example of a (weighted) relation is the binary equality $\phi_{=}$ on D defined by $\phi_{=}(x, y) = 0$ if $x = y$ and $\phi_{=}(x, y) = \infty$ if $x \neq y$.

Another example of a relation is the unary empty relation ϕ_{\emptyset} defined on D by $\phi_{\emptyset}(x) = \infty$ for all $x \in D$.

For any m -ary weighted relation $\gamma \in \Phi_D^{(m)}$, we denote by $\text{Feas}(\gamma) = \{\mathbf{x} \in D^m \mid \gamma(\mathbf{x}) < \infty\} \in \mathbf{R}_D^{(m)}$ the underlying *feasibility relation*, and by $\text{Opt}(\gamma) = \{\mathbf{x} \in \text{Feas}(\gamma) \mid \gamma(\mathbf{x}) \leq \gamma(\mathbf{y}) \text{ for every } \mathbf{y} \in D^m\} \in \mathbf{R}_D^{(m)}$ the relation of minimal-value tuples.

Definition 2.4. Let $V = \{x_1, \dots, x_n\}$ be a set of variables. A *valued constraint* over V is an expression of the form $\gamma(\mathbf{x})$ where $\gamma \in \Phi_D^{(m)}$ and $\mathbf{x} \in V^m$. The number m is called the *arity* of the constraint, the weighted relation γ is called the *constraint weighted relation*, and the tuple \mathbf{x} the *scope* of the constraint.

We call D the *domain*, the elements of D *labels* (for variables), and say that the weighted relations in Φ_D take *values* or *weights*.

¹Given a k -ary function $f : D^k \rightarrow \overline{\mathbb{Q}}$, $\text{Opt}(f)$ is the k -ary relation over D of minimal-value tuples of f .

²An m -ary relation R over D is commonly defined as a subset of D^m . For the corresponding mapping $\phi : D^m \rightarrow \{0, \infty\}$ it holds $\phi(\mathbf{x}) = 0$ when $\mathbf{x} \in R$ and $\phi(\mathbf{x}) = \infty$ otherwise. We shall use both definitions interchangeably. Because two mappings that differ only by a constant are usually equivalent for our purposes, we consider mappings $D^m \rightarrow \{c, \infty\}$ to be relations even if $c \neq 0$.

Definition 2.5. An instance of the *valued constraint satisfaction problem* (VCSP) is specified by a finite set $V = \{x_1, \dots, x_n\}$ of variables, a finite set D of labels, and an *objective function* I expressed as follows:

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

where each $\gamma_i(\mathbf{x}_i)$, $1 \leq i \leq q$, is a valued constraint over V . Each constraint can appear multiple times in I .

The goal is to find an *assignment* (or a *labelling*) of labels to the variables that minimises I .

CSPs are a special case of VCSPs using only (unweighted) relations with the goal to determine the existence of a feasible assignment.

Definition 2.6. Any set $\Gamma \subseteq \Phi_D$ is called a (*valued*) *constraint language* over D , or simply a *language*. We will denote by $\text{VCSP}(\Gamma)$ the class of all VCSP instances in which the constraint weighted relations are all contained in Γ .

Definition 2.7. A constraint language Γ is called *tractable* if $\text{VCSP}(\Gamma')$ can be solved (to optimality) in polynomial time for every finite subset $\Gamma' \subseteq \Gamma$, and Γ is called *intractable* if $\text{VCSP}(\Gamma')$ is NP-hard for some finite $\Gamma' \subseteq \Gamma$.

We refer the reader to a recent survey [Jeavons et al. 2014] for more information on the computational complexity of constraint languages.

2.2. Weighted relational clones

Definition 2.8. A weighted relation γ of arity r can be obtained by *addition* from the weighted relation γ_1 of arity s and the weighted relation γ_2 of arity t if γ satisfies the identity

$$\gamma(x_1, \dots, x_r) = \gamma_1(y_1, \dots, y_s) + \gamma_2(z_1, \dots, z_t) \quad (2)$$

for some (fixed) choice of y_1, \dots, y_s and z_1, \dots, z_t from amongst x_1, \dots, x_r .

Definition 2.9. A weighted relation γ of arity r can be obtained by *minimisation* from the weighted relation γ' of arity $r + s$ if γ satisfies the identity

$$\gamma(x_1, \dots, x_r) = \min_{(y_1, \dots, y_s) \in D^s} \gamma'(x_1, \dots, x_r, y_1, \dots, y_s). \quad (3)$$

Definition 2.10. A constraint language $\Gamma \subseteq \Phi_D$ is called a *weighted relational clone* if it contains the binary equality relation $\phi_=_$ and the unary empty relation ϕ_\emptyset ,³ and is closed under addition, minimisation, scaling by non-negative rational constants, and addition of rational constants.

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

Note that for any weighted relational clone Γ , if $\gamma \in \Gamma$ then $\text{Feas}(\gamma) \in \Gamma$ as $\text{Feas}(\gamma) = 0\gamma$ (we define $0 \cdot \infty = \infty$).

Definition 2.11. Let $\Gamma \subseteq \Phi_D$ be a constraint language, $I \in \text{VCSP}(\Gamma)$ an instance with variables V , and $L = (v_1, \dots, v_r)$ a list of variables from V . The *projection* of I onto L , denoted $\pi_L(I)$, is the r -ary weighted relation on D defined as

$$\pi_L(I)(x_1, \dots, x_r) = \min_{\{s: V \rightarrow D \mid (s(v_1), \dots, s(v_r)) = (x_1, \dots, x_r)\}} I(s). \quad (4)$$

³Although the definition in [Cohen et al. 2013] does not require the inclusion of ϕ_\emptyset , the proofs there implicitly assume its presence in any weighted relational clone.

We say that a weighted relation γ is *expressible* over a constraint language Γ if $\gamma = \pi_L(I)$ for some $I \in \text{VCSP}(\Gamma)$ and list of variables L . We call the pair (I, L) a *gadget* for expressing γ over Γ .

The list of variables L in a gadget may contain repeated entries. The minimum over an empty set is ∞ .

Example 2.12. For any $\Gamma \subseteq \Phi_D$, we can express the binary equality relation $\phi_=$ on D over language Γ using the following gadget. Let $I \in \text{VCSP}(\Gamma)$ be the instance with a single variable v and no constraints, and let $L = (v, v)$. Then, by Theorem 2.11, $\pi_L(I) = \phi_=$.

We may equivalently define a weighted relational clone as a set $\Gamma \subseteq \Phi_D$ that contains the unary empty relation ϕ_\emptyset and is closed under expressibility, scaling by non-negative rational constants, and addition of rational constants [Cohen et al. 2013, Proposition 4.5].

The following result has been shown in [Cohen et al. 2013].

THEOREM 2.13. *A constraint language Γ is tractable if and only if $\text{wRelClone}(\Gamma)$ is tractable, and Γ is intractable if and only if $\text{wRelClone}(\Gamma)$ is intractable.*

Consequently, when trying to identify tractable constraint languages, it is sufficient to consider only weighted relational clones.

2.3. Weighted clones

Any mapping $f : D^k \rightarrow D$ is called a *k-ary operation*. We will apply a *k-ary operation* f to *k m-tuples* $\mathbf{x}_1, \dots, \mathbf{x}_k \in D^m$ coordinatewise, that is,

$$f(\mathbf{x}_1, \dots, \mathbf{x}_k) = (f(\mathbf{x}_1[1], \dots, \mathbf{x}_k[1]), \dots, f(\mathbf{x}_1[m], \dots, \mathbf{x}_k[m])) \in D^m. \quad (5)$$

Definition 2.14. Let γ be an *m-ary weighted relation* on D and let f be a *k-ary operation* on D . Then f is a *polymorphism* of γ if, for any $(\mathbf{x}_1, \dots, \mathbf{x}_k) \in (\text{Feas}(\gamma))^k$, we have $f(\mathbf{x}_1, \dots, \mathbf{x}_k) \in \text{Feas}(\gamma)$.

For any 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 $\gamma \in \Gamma$. We write $\text{Pol}(\gamma)$ for $\text{Pol}(\{\gamma\})$.

A *k-ary projection* is an operation of the form $e_i^{(k)}(x_1, \dots, x_k) = x_i$ for some $1 \leq i \leq k$. Projections are (trivial) polymorphisms of all constraint languages.

Definition 2.15. The *superposition* of a *k-ary operation* $f : D^k \rightarrow D$ with *k ℓ -ary operations* $g_i : D^\ell \rightarrow D$ for $1 \leq i \leq k$ is the *ℓ -ary function* $f[g_1, \dots, g_k] : D^\ell \rightarrow D$ defined by

$$f[g_1, \dots, g_k](x_1, \dots, x_\ell) = f(g_1(x_1, \dots, x_\ell), \dots, g_k(x_1, \dots, x_\ell)). \quad (6)$$

Definition 2.16. A *clone* of operations, C , is a set of operations on D that contains all projections and is closed under superposition. The *k-ary operations* in a clone C will be denoted by $C^{(k)}$.

Example 2.17. For any D , let \mathbf{J}_D be the set of all projections on D . By Definition 2.16, \mathbf{J}_D is a clone.

It is well known that $\text{Pol}(\Gamma)$ is a clone for all constraint languages Γ .

Definition 2.18. A *k-ary weighting* of a clone C is a function $\omega : C^{(k)} \rightarrow \mathbb{Q}$ such that $\omega(f) < 0$ only if f is a projection and

$$\sum_{f \in C^{(k)}} \omega(f) = 0. \quad (7)$$

We will call a function $\omega : C^{(k)} \rightarrow \mathbb{Q}$ that satisfies Equation (7) but assigns a negative weight to some operation $f \notin \mathbf{J}_D^{(k)}$ an *improper weighting*. In order to emphasise the distinction we may also call a weighting a *proper weighting*.

When specifying a weighting, we often write it as a weighted sum of operations (i.e. $\sum \omega(f) \cdot f$) without any zero terms.

Definition 2.19. For any clone C , any k -ary weighting ω of C , and any $g_1, \dots, g_k \in C^{(\ell)}$, the *superposition* of ω and g_1, \dots, g_k is the function $\omega[g_1, \dots, g_k] : C^{(\ell)} \rightarrow \mathbb{Q}$ defined by

$$\omega[g_1, \dots, g_k](f') = \sum_{f \in C^{(k)} \wedge f[g_1, \dots, g_k] = f'} \omega(f). \quad (8)$$

By convention, the value of an empty sum is 0.

If the result of a superposition is a proper weighting (that is, negative weights are only assigned to projections), then that superposition will be called a *proper superposition*.

Definition 2.20. A *weighted clone*, Ω , is a non-empty set of weightings of some fixed clone C , called the *support clone* of Ω , which is closed under scaling by non-negative rational constants, addition of weightings of equal arity, and proper superposition with operations from C .

We now link weightings and weighted relations by the concept of weighted polymorphism, which will allow us to establish a correspondence between weighted clones and weighted relational clones.

Definition 2.21. Let γ be an m -ary weighted relation on D and let ω be a k -ary weighting of a clone C of operations on D . We call ω a *weighted polymorphism* of γ if $C \subseteq \text{Pol}(\gamma)$ and for any $(\mathbf{x}_1, \dots, \mathbf{x}_k) \in (\text{Feas}(\gamma))^k$, we have

$$\sum_{f \in C^{(k)}} \omega(f) \cdot \gamma(f(\mathbf{x}_1, \dots, \mathbf{x}_k)) \leq 0. \quad (9)$$

If ω is a weighted polymorphism of γ , we say that γ is *improved* by ω .

Example 2.22. Let $D = \{0, 1\}$ with ordering $0 < 1$. Binary operations \min and \max return the smaller and larger of their two arguments respectively. A function $\gamma : D^k \rightarrow \mathbb{Q}$ is *submodular* if it satisfies $\gamma(\mathbf{x}_1) + \gamma(\mathbf{x}_2) \geq \gamma(\min(\mathbf{x}_1, \mathbf{x}_2)) + \gamma(\max(\mathbf{x}_1, \mathbf{x}_2))$ for all $\mathbf{x}_1, \mathbf{x}_2$. Clearly, submodular functions are improved by the binary weighting $\omega = -e_1^{(2)} - e_2^{(2)} + \min + \max$.

Definition 2.23. For any $\Gamma \subseteq \Phi_D$, we define $\text{wPol}(\Gamma)$ to be the set of all weightings of $\text{Pol}(\Gamma)$ which are weighted polymorphisms of all weighted relations $\gamma \in \Gamma$. We write $\text{wPol}(\gamma)$ for $\text{wPol}(\{\gamma\})$.

Definition 2.24. We denote by \mathbf{W}_C the set of all possible (proper) weightings of clone C , and define \mathbf{W}_D to be the union of the sets \mathbf{W}_C over all clones C on D .

Any $\Omega \subseteq \mathbf{W}_D$ may contain weightings of *different* clones over D . We can then extend each of these weightings with zeros, as necessary, so that they are weightings of the same clone C , where C is the smallest clone containing all the clones associated with weightings in Ω .

Definition 2.25. We define $\text{wClone}(\Omega)$ to be the smallest weighted clone containing this set of extended weightings obtained from Ω .

For any $\Omega \subseteq \mathbf{W}_D$, we denote by $\text{Imp}(\Omega)$ the set of all weighted relations in Φ_D which are improved by all weightings $\omega \in \Omega$.

The main result in [Cohen et al. 2013] establishes a 1-to-1 correspondence between weighted relational clones and weighted clones.

THEOREM 2.26 ([COHEN ET AL. 2013]).

- (1) For any finite D and any finite $\Gamma \subseteq \Phi_D$, $\text{Imp}(\text{wPol}(\Gamma)) = \text{wRelClone}(\Gamma)$.
- (2) For any finite D and any finite $\Omega \subseteq \mathbf{W}_D$, $\text{wPol}(\text{Imp}(\Omega)) = \text{wClone}(\Omega)$.

Thus, when trying to identify tractable constraint languages, it is sufficient to consider only languages of the form $\text{Imp}(\Omega)$ for some weighted clone Ω .

3. RESULTS

First we show that Theorem 2.26 can be slightly extended to certain constraint languages and sets of weightings of infinite size.

THEOREM 3.1.

- (1) Let $\Gamma \subseteq \Phi_D$. Then $\text{Imp}(\text{wPol}(\Gamma)) = \text{wRelClone}(\Gamma)$ if and only if $\text{wRelClone}(\Gamma) = \text{Imp}(\Omega)$ for some $\Omega \subseteq \mathbf{W}_D$.
- (2) Let $\Omega \subseteq \mathbf{W}_D$. Then $\text{wPol}(\text{Imp}(\Omega)) = \text{wClone}(\Omega)$ if and only if $\text{wClone}(\Omega) = \text{wPol}(\Gamma)$ for some $\Gamma \subseteq \Phi_D$.

PROOF. We will only prove the first case as the second one is analogous.

Suppose that $\text{wRelClone}(\Gamma) = \text{Imp}(\Omega)$ for some $\Omega \subseteq \mathbf{W}_D$. As $\Gamma \subseteq \text{wRelClone}(\Gamma)$, every weighting in Ω improves Γ , hence $\Omega \subseteq \text{wPol}(\Gamma)$ and $\text{Imp}(\text{wPol}(\Gamma)) \subseteq \text{Imp}(\Omega) = \text{wRelClone}(\Gamma)$. The inclusion $\text{wRelClone}(\Gamma) \subseteq \text{Imp}(\text{wPol}(\Gamma))$ follows from the fact that $\text{Imp}(\text{wPol}(\Gamma))$ is a weighted relational clone [Cohen et al. 2013, Proposition 6.2] that contains Γ .

The converse implication holds trivially for $\Omega = \text{wPol}(\Gamma)$. \square

We remark that any *finitely generated* weighted relational clone on a finite domain satisfies, by Theorem 2.26 (1), the condition of Theorem 3.1 (1). Similarly, any finitely generated weighted clone on a finite domain, by Theorem 2.26 (2), satisfies the condition of Theorem 3.1 (2).

However, our next result shows that Theorem 2.26 does *not* hold for all infinite constraint languages and infinite sets of weightings.

THEOREM 3.2.

- (1) There is a finite D and an infinite $\Gamma \subseteq \Phi_D$ with $\text{Imp}(\text{wPol}(\Gamma)) \neq \text{wRelClone}(\Gamma)$.
- (2) There is a finite D and an infinite $\Omega \subseteq \mathbf{W}_D$ with $\text{wPol}(\text{Imp}(\Omega)) \neq \text{wClone}(\Omega)$.

Our aim is to establish a Galois connection even for infinite sets of weighted relations and weightings. As we demonstrate in the proof of Theorem 3.2, this cannot be done when restricted to rational weights; hence we allow weighted relations and weightings to assign *real-valued* weights. To distinguish them from their formerly defined rational-valued counterparts, we will use a subscript/superscript \mathbb{R} .

We will show in Theorem 5.9 that $\text{wPol}_{\mathbb{R}}(\Gamma)$ is topologically closed (in a natural topology defined later) for any set of weighted relations Γ ; analogously, in Theorem 5.10 we will show that $\text{Imp}_{\mathbb{R}}(\Omega)$ is topologically closed for any set of weightings Ω . Therefore, our new definitions of weighted (relational) clones require them to be topologically closed.

Inspired by weighted pp-definitions [Thapper 2010], we extend the notion of weighted relational clones: we require them to be closed also under operator Opt . This change is justified by Theorem 5.13 in which we prove that the inclusion of Opt

preserves tractability. In order to retain the one-to-one correspondence with weighted clones, we need to alter their definition too: weightings now assign weights to all operations and hence are independent of the support clone (which becomes meaningless and we discard it).

Including the Opt operator brings two advantages to the study of weighted clones. Firstly, it slightly simplifies the structure of the space of all weighted clones. According to the original definition, a weighted clone is determined by its support clone and the set of weightings it consists of; by our definition a weighted clone equals the set of its weightings. Secondly, any non-projection polymorphism of a weighted relational clone Γ is assigned a positive weight by some weighted polymorphism of Γ (see Theorem 5.8).

Our main result is the following theorem, which holds for our new definition of real-valued weightings and weighted relations.

THEOREM 3.3 (MAIN).

- (1) For any finite D and any $\Gamma \subseteq \Phi_D^{\mathbb{R}}$, $\text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma)) = \text{wRelClone}_{\mathbb{R}}(\Gamma)$.
- (2) For any finite D and any $\Omega \subseteq \mathbf{W}_D^{\mathbb{R}}$, $\text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega)) = \text{wClone}_{\mathbb{R}}(\Omega)$.

Finally, we show that taking the weighted relational clone of a constraint language preserves solvability with an absolute error bounded by ϵ (for any $\epsilon > 0$), and demonstrate certain difficulties with proving that it preserves exact solvability.

4. PROOF OF THEOREM 3.2

In this section we will prove Theorem 3.2, which we state here as two lemmas.

LEMMA 4.1. *There is a finite D and an infinite $\Gamma \subseteq \Phi_D$ with $\text{Imp}(\text{wPol}(\Gamma)) \neq \text{wRelClone}(\Gamma)$.*

PROOF. We set the domain to be $D = \{0, 1, 2\}$ and choose a positive *irrational* number t . Let $U \subseteq \Phi_D^{(1)}$ be the set of unary weighted relations ρ such that

$$\rho(2) - \rho(0) \geq (1 + t) \cdot (\rho(1) - \rho(0)) \quad (10)$$

holds whenever $\rho(0), \rho(1), \rho(2)$ are all finite. It is easy to show that U is closed under addition, scaling by non-negative rational constants, and addition of rational constants.

For any rational $u < t$, we define a unary weighted relation $\mu_u^- \in U$ such that $\mu_u^-(0) = 0$, $\mu_u^-(1) = -1$, and $\mu_u^-(2) = -1 - u$. For any rational $v > t$, we define a unary weighted relation $\mu_v^+ \in U$ such that $\mu_v^+(0) = 0$, $\mu_v^+(1) = 1$, and $\mu_v^+(2) = 1 + v$. It is easy to verify that these weighted relations belong to U . Set U also contains any unary (unweighted) relation.

Let us define $\Gamma \subseteq \Phi_D$ as the set of weighted relations γ that can be written as

$$\gamma(x_1, \dots, x_r) = \sum_{i=1}^r \rho_i(x_i) + \sum_{(i,j) \in S} \phi_{=} (x_i, x_j), \quad (11)$$

where r equals the arity of γ , $\rho_i \in U$ for all i , $\phi_{=}$ is the binary equality relation, and S is an equivalence relation on $\{1, \dots, r\}$. We claim that Γ is a weighted relational clone. It certainly contains $\phi_{=}$ and ϕ_{\emptyset} , and is closed under addition, scaling by non-negative rational constants, and addition of rational constants (as set U is closed under these operations). It is also closed under minimisation. Without loss of generality, let us assume we minimise an r -ary weighted relation γ ($r \geq 2$) over the last variable (x_r). If the equivalence class of r in S is a singleton, we simply add the value of $\min_{x_r \in D} \rho_r(x_r)$ to (say) ρ_1 . Otherwise, we can pick any $i \neq r$ such that $(i, r) \in S$ and replace weighted relation ρ_i with $\rho_i + \rho_r$.

We want to determine which weightings improve Γ . Let $\omega \in \text{wPol}^{(k)}(\Gamma)$ be a k -ary weighting and $\mathbf{x} \in D^k$. For any $a \in D$, we will denote by s_a the sum of weights $\omega(f)$ of all operations f such that $f(\mathbf{x}) = a$. Note that $s_0 + s_1 + s_2 = 0$. For any rational $v > t$, weighting ω improves $\mu_v^+ \in \Gamma$, so we get $s_1 + (1+v) \cdot s_2 \leq 0$ and therefore $v \cdot s_2 \leq s_0$. Similarly, for any rational $u < t$, weighting ω improves $\mu_u^- \in \Gamma$, and therefore $s_0 \leq u \cdot s_2$. We can choose both u and v arbitrarily close to t , so it must hold $s_0 = t \cdot s_2$. However, s_0 and s_2 are rational while t is not. Therefore, we must have $s_0 = s_1 = s_2 = 0$ for any weighting $\omega \in \text{wPol}^{(k)}(\Gamma)$ and any $\mathbf{x} \in D^k$.

Now, let us consider the unary weighted relation ρ defined as $\rho(0) = 0$ and $\rho(1) = \rho(2) = 1$. It follows from the previous paragraph that any weighting $\omega \in \text{wPol}(\Gamma)$ improves ρ , i.e. $\rho \in \text{Imp}(\text{wPol}(\Gamma))$. However, $\rho \notin \Gamma = \text{wRelClone}(\Gamma)$, so we get $\text{Imp}(\text{wPol}(\Gamma)) \neq \text{wRelClone}(\Gamma)$. \square

LEMMA 4.2. *There is a finite D and an infinite $\Omega \subseteq \mathbf{W}_D$ with $\text{wPol}(\text{Imp}(\Omega)) \neq \text{wClone}(\Omega)$.*

PROOF. We set the domain to be $D = \{0, 1, 2\}$ and choose a positive *irrational* number t . Let C be the set of all operations f such that $f(0, \dots, 0) = 0$. Clearly, C contains all projections and is closed under superposition; hence it is a clone. Let us define a set of weightings $\Omega \subseteq \mathbf{W}_D$ of the support clone C . For any arity $k \geq 1$, $\Omega^{(k)}$ consists of weightings ω such that for all $\mathbf{x} \in D^k$,

$$t \cdot \sum_{f(\mathbf{x})=2} \omega(f) \leq \sum_{f(\mathbf{x})=0} \omega(f). \quad (12)$$

It is easy to check that Ω is closed under addition of weightings and non-negative scaling. To show that it is also closed under superposition, let us consider any sequence of ℓ -ary operations g_1, \dots, g_k and $\mathbf{x} \in D^\ell$. For any $a \in D$ we have

$$\sum_{f(\mathbf{x})=a} \omega[g_1, \dots, g_k](f) = \sum_{f[g_1, \dots, g_k](\mathbf{x})=a} \omega(f) = \sum_{f(\mathbf{y})=a} \omega(f), \quad (13)$$

where $\mathbf{y} = (g_1(\mathbf{x}), \dots, g_k(\mathbf{x}))$. As ω satisfies Inequality (12) for vector \mathbf{y} , the superposition $\omega[g_1, \dots, g_k]$ satisfies it for vector \mathbf{x} . Therefore, Ω is a weighted clone.

Let us denote by c_0 the unary constant zero operation, by f, g the unary operations and by h the binary operation such that

$$f(x) = \begin{cases} 0 & \text{for } x = 0 \\ 0 & \text{for } x = 1 \\ 2 & \text{for } x = 2 \end{cases}, \quad g(x) = \begin{cases} 0 & \text{for } x = 0 \\ 2 & \text{for } x = 1 \\ 2 & \text{for } x = 2 \end{cases}, \quad h(x, y) = \begin{cases} 1 & \text{for } x = 0 \wedge y = 2 \\ 2 & \text{for } x = 2 \wedge y = 2 \\ 0 & \text{otherwise} \end{cases}. \quad (14)$$

We denote by ω_0 the unary weighting $-e_1^{(1)} + c_0$. For any rational $v > t$, we define a unary weighting $\mu_v^{(1)} = -(1+v) \cdot e_1^{(1)} + v \cdot f + g$. For any positive rational $u < t$, we define a binary weighting $\mu_u^{(2)} = -u \cdot e_1^{(2)} - e_2^{(2)} + (1+u) \cdot h$. It is easy to show that all these weightings belong to Ω .

We will show that all weighted relations improved by Ω are relations, i.e. $\text{Imp}(\Omega) \subseteq \mathbf{R}_D$. Suppose, to the contrary, that there is an r -ary weighted relation $\gamma \in \text{Imp}(\Omega)$ that is not a relation. First, we obtain from it a ternary weighted relation with the same property. Weighting ω_0 improves γ , so we have $\gamma(\mathbf{0}) \leq \gamma(\mathbf{x})$ for all $\mathbf{x} \in \text{Feas}(\gamma)$, where $\mathbf{0} = c_0(\mathbf{x})$ is the zero r -tuple. As γ is not a relation, there must be an r -tuple $\mathbf{z} = (z_1, \dots, z_r) \in \text{Feas}(\gamma)$ for which $\gamma(\mathbf{0}) < \gamma(\mathbf{z})$. Let us define a ternary weighted relation ρ so that $\rho(x_0, x_1, x_2) = \gamma(x_{z_1}, \dots, x_{z_r})$. It holds that $\rho(0, 0, 0) = \gamma(\mathbf{0})$ and $\rho(0, 1, 2) = \gamma(\mathbf{z})$, so $\rho(0, 0, 0) < \rho(0, 1, 2) < \infty$. Moreover, $\rho \in \text{Imp}(\Omega)$.

For any rational $v > t$, weighting $\mu_v^{(1)}$ improves ρ . As $(0, 1, 2) \in \text{Feas}(\rho)$, we also have $(f(0), f(1), f(2)) = (0, 0, 2) \in \text{Feas}(\rho)$, $(g(0), g(1), g(2)) = (0, 2, 2) \in \text{Feas}(\rho)$, and the inequality

$$\rho(0, 2, 2) - \rho(0, 1, 2) \leq v \cdot (\rho(0, 1, 2) - \rho(0, 0, 2)). \quad (15)$$

For any positive rational $u < t$, weighting $\mu_u^{(2)}$ improves ρ . As $(0, 0, 2), (0, 2, 2) \in \text{Feas}(\rho)$, we get

$$\rho(0, 2, 2) - \rho(0, 1, 2) \geq u \cdot (\rho(0, 1, 2) - \rho(0, 0, 2)). \quad (16)$$

We can choose both u and v arbitrarily close to t , so it must hold

$$\rho(0, 2, 2) - \rho(0, 1, 2) = t \cdot (\rho(0, 1, 2) - \rho(0, 0, 2)). \quad (17)$$

However, weights assigned by ρ are rational while t is not. Therefore, $\rho(0, 0, 2) = \rho(0, 1, 2) = \rho(0, 2, 2)$. Similarly, by applying weightings $\mu_u^{(2)}$ to $(0, 0, 0), (0, 0, 2)$ and weightings $\mu_v^{(1)}$ to $(0, 0, 1)$ we obtain $\rho(0, 0, 0) = \rho(0, 0, 1) = \rho(0, 0, 2)$, which contradicts $\rho(0, 0, 0) < \rho(0, 1, 2)$. Therefore, $\text{Imp}(\Omega)$ contains only (unweighted) relations.

Now, let us consider the unary weighting $\omega = -e_1^{(1)} + g$. Although it does not belong to $\Omega = \text{wClone}(\Omega)$ (it violates Inequality (12) for $\mathbf{x} = (1)$), ω certainly improves any $\gamma \in \text{Imp}(\Omega)$. Therefore, $\text{wPol}(\text{Imp}(\Omega)) \neq \text{wClone}(\Omega)$. \square

5. NEW GALOIS CONNECTION

In this section we will prove our main results. In Section 5.1, we will describe the differences between the previous definitions of weighted (relational) clones (as they were defined in [Cohen et al. 2013] and presented in Section 2 and the first part of Section 3) and our new definitions. Section 5.2 proves the main result, which establishes a 1-to-1 correspondence between weighted relational clones and weighted clones. Finally, Section 5.3 is devoted to computational-complexity consequences of our results.

5.1. Preliminaries

Let $\overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$ denote the set of real numbers with (positive) infinity. We will allow weights in relations and weighted relations, as defined in Definition 2.1 and 2.2 respectively, to be real numbers. In other words, an m -ary weighted relation γ on D is a mapping $\gamma : D^m \rightarrow \overline{\mathbb{R}}$. We will add a subscript/superscript \mathbb{R} to the notation introduced in Section 2 in order to emphasise the use of real weights.

For any fixed arity m and any $F \subseteq D^m$, consider the set of all m -ary weighted relations $\gamma \in \Phi_D^{\mathbb{R}}$ with $\text{Feas}(\gamma) = F$. Let us denote this set by H and equip it with the inner product defined as

$$\langle \alpha, \beta \rangle = \sum_{\mathbf{x} \in F} \alpha(\mathbf{x}) \cdot \beta(\mathbf{x}) \quad (18)$$

for any $\alpha, \beta \in H$; H is then a real Hilbert space. Set $\Phi_D^{\mathbb{R}}$ is a disjoint union of such Hilbert spaces for all m and F , and therefore a topological space with the disjoint union topology induced by inner products on the underlying Hilbert spaces. When we say a set of weighted relations is open/closed, we will be referring to this topology.

Definition 5.1. A constraint language $\Gamma \subseteq \Phi_D^{\mathbb{R}}$ is called a *weighted relational clone* if it contains the binary equality relation $\phi_{=}$ and the unary empty relation ϕ_{\emptyset} , is closed under addition, minimisation, scaling by non-negative real constants, addition of real constants, and operator Opt , and is topologically closed.

For any Γ , we define $\text{wRelClone}_{\mathbb{R}}(\Gamma)$ to be the smallest weighted relational clone containing Γ .

As opposed to Theorem 2.10, our new definition requires weighted relational clones to be closed under operator Opt . In order to establish a Galois connection now, we need to make an adjustment to the definition of weighted clone too. We will discard the explicit underlying support clone; instead, $(k\text{-ary})$ weightings will assign weights to all $(k\text{-ary})$ operations. The role of the support clone of a weighted clone Ω is then taken over by $\text{supp}(\Omega)$ (see Theorem 5.7).

We denote by $\mathcal{O}_D^{(k)}$ the set of all $k\text{-ary}$ operations on D and let $\mathcal{O}_D = \bigcup_{k \geq 0} \mathcal{O}_D^{(k)}$.

Definition 5.2. A $k\text{-ary}$ weighting is a function $\omega : \mathcal{O}_D^{(k)} \rightarrow \mathbb{R}$ such that $\omega(f) < 0$ only if f is a projection and

$$\sum_{f \in \mathcal{O}_D^{(k)}} \omega(f) = 0. \quad (19)$$

We define $\text{supp}(\omega)$ as

$$\text{supp}(\omega) = \mathbf{J}_D^{(k)} \cup \left\{ f \in \mathcal{O}_D^{(k)} \mid \omega(f) > 0 \right\}. \quad (20)$$

We will call a function $\omega : \mathcal{O}_D^{(k)} \rightarrow \mathbb{R}$ that satisfies Equation (19) but assigns a negative weight to some operation $f \notin \mathbf{J}_D^{(k)}$ an *improper weighting*. In order to emphasise the distinction we may also call a weighting a *proper weighting*.

We denote by $\mathbf{W}_D^{\mathbb{R}}$ the set of all weightings on domain D . For any fixed arity k , consider the set H of all functions $\mathcal{O}_D^{(k)} \rightarrow \mathbb{R}$ equipped with the inner product defined as

$$\langle \alpha, \beta \rangle = \sum_{f \in \mathcal{O}_D^{(k)}} \alpha(f) \cdot \beta(f) \quad (21)$$

for any $\alpha, \beta \in H$; H is then a real Hilbert space. Set $\mathbf{W}_D^{\mathbb{R}}$ lies in the disjoint union of such Hilbert spaces for all k , which is a topological space with the disjoint union topology induced by inner products on the underlying Hilbert spaces. When we say a set of weightings is open/closed, we will be referring to this topology. Clearly, any closure point of a set of weightings is itself a weighting.

Definition 5.3. Let Ω be a non-empty set of weightings on a fixed domain D . We define $\text{supp}(\Omega) = \mathbf{J}_D \cup \bigcup_{\omega \in \Omega} \text{supp}(\omega)$.

We call Ω a *weighted clone* if it is closed under scaling by non-negative real constants, addition of weightings of equal arity, and proper superposition with operations from $\text{supp}(\Omega)$, and is topologically closed.

It is often convenient to build a desired proper weighting by taking a sum of (possibly) improper superpositions. The following lemma, which is an analogue of [Cohen et al. 2013, Lemma 6.4], shows that weighted clones are closed under such constructions.

LEMMA 5.4. *Let Ω be a weighted clone, $\omega_1, \dots, \omega_n \in \Omega$, and $c_1, \dots, c_n \geq 0$. We will denote the arity of weighting ω_i by ℓ_i . For any $1 \leq i \leq n$ and $1 \leq j \leq \ell_i$, let $g_{i,j} \in \text{supp}(\Omega)$ be a $k\text{-ary}$ operation (for some fixed arity k). If the $k\text{-ary}$ weighting μ defined as*

$$\mu = \sum_{i=1}^n c_i \cdot \omega_i[g_{i,1}, \dots, g_{i,\ell_i}] \quad (22)$$

is proper, then $\mu \in \Omega$.

PROOF. We show that weighting μ can be constructed using proper superpositions only.

Let us denote $\sum_{1 \leq i \leq n} \ell_i$ by t . For any $1 \leq m \leq n$, let $s_m = \sum_{1 \leq i < m} \ell_i$. A superposition with projections is always proper (as all negative weights are transferred to projections), and therefore the t -ary weighting μ' defined as

$$\mu' = \sum_{i=1}^n c_i \cdot \omega_i \left[e_{s_i+1}^{(t)}, \dots, e_{s_i+\ell_i}^{(t)} \right] \quad (23)$$

belongs to Ω . Since $\mu = \mu'[g_{1,1}, \dots, g_{1,\ell_1}, g_{2,1}, \dots, g_{n,\ell_n}]$, we get $\mu \in \Omega$. \square

The following lemma has also been observed in [Kozik and Ochremiak 2015; Thapper and Živný 2015d].

LEMMA 5.5. *Let Ω be a weighted clone. Then $\text{supp}(\Omega)$ is a clone.*

PROOF. We will denote $\text{supp}(\Omega)$ by C . Since it contains all projections, we only need to show that it is closed under superposition.

Let $f \in C^{(k)}$ and $g_1, \dots, g_k \in C^{(\ell)}$. If $f[g_1, \dots, g_k]$ is a projection or is equal to g_i for some i , then it clearly belongs to C . Otherwise, f is not a projection and therefore there is a k -ary weighting $\omega \in \Omega$ for which $\omega(f) > 0$. Weighting $\omega[g_1, \dots, g_k]$ certainly assigns a positive weight to $f[g_1, \dots, g_k]$ (we are using the fact that only operations g_1, \dots, g_k may receive negative weight from projections in ω). However, it is possibly improper, as it may assign a negative weight to some g_i that is not a projection.

We denote by G the set of such operations $g \in \{g_1, \dots, g_k\}$ that are not projections and $\omega[g_1, \dots, g_k](g) < 0$. For any $g \in G$, there is an ℓ -ary weighting $\mu_g \in \Omega$ for which $\mu_g(g) > 0$. Then the ℓ -ary weighting defined as

$$\omega[g_1, \dots, g_k] + \sum_{g \in G} \frac{-\omega[g_1, \dots, g_k](g)}{\mu_g(g)} \cdot \mu_g \quad (24)$$

is proper, belongs to Ω (by Theorem 5.4), and assigns a positive weight to $f[g_1, \dots, g_k]$. \square

Again, we link weightings and weighted relations by the concept of weighted polymorphism.

Definition 5.6. Let γ be an m -ary weighted relation on D and let ω be a k -ary weighting on D . We call ω a *weighted polymorphism* of γ if $\text{supp}(\omega) \subseteq \text{Pol}(\gamma)$ and for any $(\mathbf{x}_1, \dots, \mathbf{x}_k) \in (\text{Feas}(\gamma))^k$, we have

$$\sum_{f \in \text{supp}(\omega)} \omega(f) \cdot \gamma(f(\mathbf{x}_1, \dots, \mathbf{x}_k)) \leq 0. \quad (25)$$

If ω is a weighted polymorphism of γ we say that γ is *improved* by ω . We will denote the set of weighted polymorphisms of Γ by $\text{wPol}_{\mathbb{R}}(\Gamma)$ and the set of weighted relations improved by Ω by $\text{Imp}_{\mathbb{R}}(\Omega)$.

The next lemma shows that $\text{supp}(\Omega)$ consists of all polymorphisms of $\text{Imp}_{\mathbb{R}}(\Omega)$ and hence fulfills the same role as the support clone in Theorem 2.20.

LEMMA 5.7. *Let $\Omega \subseteq \mathbf{W}_D^{\mathbb{R}}$ be a weighted clone. Then $\text{supp}(\Omega) = \text{Pol}(\text{Imp}_{\mathbb{R}}(\Omega))$.*

PROOF. We will denote $\text{supp}(\Omega)$ by C . Projections are polymorphisms of every weighted relation, and any operation f with $\omega(f) > 0$ for some $\omega \in \Omega$ is a polymorphism of $\text{Imp}_{\mathbb{R}}(\Omega)$ by the definition of weighted polymorphism. Therefore, $C \subseteq \text{Pol}(\text{Imp}_{\mathbb{R}}(\Omega))$.

Let $\text{Inv}(C)$ be the set of (unweighted) relations over D that are invariant under all operations from C (i.e. operations from C are their polymorphisms). As any relation invariant under $\text{supp}(\omega)$ is also improved by ω , we have $\text{Inv}(C) \subseteq \text{Imp}_{\mathbb{R}}(\Omega)$ and thus $\text{Pol}(\text{Imp}_{\mathbb{R}}(\Omega)) \subseteq \text{Pol}(\text{Inv}(C)) = C$ (the last equality follows from the Galois connection between relational clones and clones of operations [Bodnarčuk et al. 1969; Geiger 1968] and Theorem 5.5). \square

The following corollary has been observed in the context of Min-Sol-Hom and Min-Cost-Hom [Uppman 2013] by Hannes Uppman.⁴

COROLLARY 5.8. *Let $\Gamma \subseteq \Phi_D^{\mathbb{R}}$ be a weighted relational clone. Then $\text{supp}(\text{wPol}_{\mathbb{R}}(\Gamma)) = \text{Pol}(\Gamma)$.*

PROOF. We are going to use the Galois connection established later in Section 5.2.

As $\text{wPol}_{\mathbb{R}}(\Gamma)$ is a weighted clone (Theorem 5.9), by Theorem 5.7 we have $\text{supp}(\text{wPol}_{\mathbb{R}}(\Gamma)) = \text{Pol}(\text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma))) = \text{Pol}(\Gamma)$ (the last equality follows from Theorem 5.11). \square

Finally, we introduce some notation that will be used throughout Section 5.2. A sequence of k m -tuples over D can be written as $X = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in (D^m)^k$. By X^T we will denote the transpose of X , i.e. the sequence of m k -tuples $(\mathbf{y}_1, \dots, \mathbf{y}_m) \in (D^k)^m$ such that $\mathbf{y}_i = (\mathbf{x}_1[i], \dots, \mathbf{x}_k[i])$. Let f be a k -ary operation; we denote by $f(X)$ the m -tuple obtained by applying f coordinatewise to $\mathbf{x}_1, \dots, \mathbf{x}_k$, i.e. $f(X) = f(\mathbf{x}_1, \dots, \mathbf{x}_k) = (f(\mathbf{y}_1), \dots, f(\mathbf{y}_m))$.

Let $\gamma \in \Phi_D^{\mathbb{R}}$ be a weighted relation and $\omega \in \mathbf{W}_D^{\mathbb{R}}$ a k -ary weighting with $\text{supp}(\omega) \subseteq \text{Pol}(\gamma)$. Let us denote by H the Hilbert space of functions $\text{Pol}^{(k)}(\gamma) \rightarrow \mathbb{R}$ with the inner product analogous to (21). As weighting ω assigns non-zero weights only to operations from $\text{supp}(\omega) \subseteq \text{Pol}^{(k)}(\gamma)$, we can identify ω with its restriction to $\text{Pol}^{(k)}(\gamma)$. For any $X \in (\text{Feas}(\gamma))^k$, we denote by $\gamma[X]$ the vector in H such that $\gamma[X](f) = \gamma(f(X))$ for all $f \in \text{Pol}^{(k)}(\gamma)$. Inequality (25) can be then written as $\langle \omega, \gamma[X] \rangle \leq 0$.

The (internal) polar cone K° of a set $K \subseteq H$ is defined as

$$K^\circ = \left\{ \alpha \in H \mid \langle \alpha, \beta \rangle \leq 0 \text{ for all } \beta \in K \right\}. \quad (26)$$

It is well known ([Boyd and Vandenberghe 2004; Hiriart-Urruty and Lemaréchal 2001]) that K° is a convex cone, i.e. K° is closed under addition of vectors and scaling by non-negative constants. Moreover, K° is topologically closed, and $K^{\circ\circ} = (K^\circ)^\circ$ is the closure of the smallest convex cone containing K .⁵

Let

$$K = \{ \gamma[X] \mid X \in (\text{Feas}(\gamma))^k \}. \quad (27)$$

Weighting ω is then a weighted polymorphism of γ if and only if $\omega \in K^\circ$.

5.2. Main Proofs

LEMMA 5.9. *For any finite D and any $\Gamma \subseteq \Phi_D^{\mathbb{R}}$, $\text{wPol}_{\mathbb{R}}(\Gamma)$ is a weighted clone.*

PROOF. Let $k \geq 1$ be a fixed arity. We denote by H the Hilbert space of functions $\text{Pol}^{(k)}(\Gamma) \rightarrow \mathbb{R}$ and define a set $K \subseteq H$ as

$$K = \{ \gamma[X] \mid \gamma \in \Gamma \wedge X \in (\text{Feas}(\gamma))^k \}. \quad (28)$$

⁴Private communication, 2014.

⁵If K is a finite set, the smallest convex cone containing K is topologically closed. This is why the former definitions of weighted (relational) clones did not have to require topological closedness explicitly.

A k -ary weighting ω with $\text{supp}(\omega) \subseteq \text{Pol}^{(k)}(\Gamma)$ is a weighted polymorphism of Γ if and only if its restriction to $\text{Pol}^{(k)}(\Gamma)$ belongs to K° . Set $\text{wPol}_{\mathbb{R}}^{(k)}(\Gamma)$ is therefore closed under addition as K° is convex, it is closed under non-negative scaling as K° is a cone, and it is topologically closed as K° is.

It remains to show that $\text{wPol}_{\mathbb{R}}(\Gamma)$ is closed under superposition. Let $\omega \in \text{wPol}_{\mathbb{R}}(\Gamma)$ be a k -ary weighting and $g_1, \dots, g_k \in \text{Pol}(\Gamma)$ be ℓ -ary operations. For any $\gamma \in \Gamma$ and $X \in (\text{Feas}(\gamma))^\ell$, we have $Y = (g_1(X), \dots, g_k(X)) \in (\text{Feas}(\gamma))^k$ and

$$\sum_{f \in \text{supp}(\omega[g_1, \dots, g_k])} \omega[g_1, \dots, g_k](f) \cdot \gamma(f(X)) = \sum_{f \in \text{supp}(\omega)} \omega(f) \cdot \gamma(f[g_1, \dots, g_k](X)) \quad (29)$$

$$= \sum_{f \in \text{supp}(\omega)} \omega(f) \cdot \gamma(f(Y)) \quad (30)$$

$$\leq 0. \quad (31)$$

Therefore, weighting $\omega[g_1, \dots, g_k]$ is a weighted polymorphism of Γ . \square

LEMMA 5.10. *For any finite D and any $\Omega \subseteq \mathbf{W}_D^{\mathbb{R}}$, $\text{Imp}_{\mathbb{R}}(\Omega)$ is a weighted relational clone.*

PROOF. Both $\phi_{=}$ and ϕ_{\emptyset} are improved by any weighting and hence belong to $\text{Imp}_{\mathbb{R}}(\Omega)$. Addition, non-negative scaling, and addition of a constant preserve Inequality (25), and therefore $\text{Imp}_{\mathbb{R}}(\Omega)$ is closed under these operations.

We need to prove that $\text{Imp}_{\mathbb{R}}(\Omega)$ is closed under minimisation. Let $\gamma \in \text{Imp}_{\mathbb{R}}(\Omega)$ be an r -ary weighted relation and consider γ' obtained from γ by minimising over the last argument, i.e.

$$\gamma'(x_1, \dots, x_{r-1}) = \min_{x_r \in D} \gamma(x_1, \dots, x_{r-1}, x_r). \quad (32)$$

Let $\omega \in \Omega$ be a k -ary weighting and $X' = (\mathbf{x}'_1, \dots, \mathbf{x}'_k) \in (\text{Feas}(\gamma'))^k$. For any $i \in \{1, \dots, k\}$, we can extend $(r-1)$ -tuple \mathbf{x}'_i to an r -tuple $\mathbf{x}_i \in \text{Feas}(\gamma)$ so that $\gamma'(\mathbf{x}'_i) = \gamma(\mathbf{x}_i)$; we will denote the list of these extended r -tuples by $X = (\mathbf{x}_1, \dots, \mathbf{x}_k)$. Note that $f(X)$ is an extension of $f(X')$ for any k -ary operation $f \in \text{supp}(\omega)$, and therefore $\gamma'(f(X')) \leq \gamma(f(X))$. Moreover, $\gamma'(f(X')) = \gamma(f(X))$ whenever f is a projection. As γ satisfies Inequality (25) and only projections may be assigned a negative weight, we have

$$\sum_{f \in \text{supp}(\omega)} \omega(f) \cdot \gamma'(f(X')) \leq \sum_{f \in \text{supp}(\omega)} \omega(f) \cdot \gamma(f(X)) \leq 0, \quad (33)$$

and thus $\gamma' \in \text{Imp}_{\mathbb{R}}(\Omega)$.

Now we prove that $\text{Imp}_{\mathbb{R}}(\Omega)$ is closed under operator Opt . Let $\gamma \in \text{Imp}_{\mathbb{R}}(\Omega)$ and $\rho = \text{Opt}(\gamma)$. We will assume that $\text{Feas}(\gamma)$ is non-empty (otherwise $\gamma = \rho$) and denote by c the minimum weight assigned by γ . Let $\omega \in \Omega$ be a k -ary weighting. As ρ is a relation, we only need to show that all operations in the support of ω are polymorphisms of ρ . Let $X \in (\text{Feas}(\rho))^k \subseteq (\text{Feas}(\gamma))^k$. For any operation f in the support of ω , it holds $\gamma(f(X)) \geq c$. Moreover, $\gamma(f(X)) = c$ whenever f is a projection. We have

$$0 \geq \sum_{f \in \text{supp}(\omega)} \omega(f) \cdot \gamma(f(X)) \geq \sum_{f \in \text{supp}(\omega)} \omega(f) \cdot c = 0, \quad (34)$$

which for all $f \in \text{supp}(\omega)$ implies $\gamma(f(X)) = c$ and hence $f(X) \in \text{Feas}(\rho)$. Therefore, $\rho \in \text{Imp}_{\mathbb{R}}(\Omega)$.

Finally, we show that $\text{Imp}_{\mathbb{R}}(\Omega)$ is topologically closed. Let r be a fixed arity and $F \subseteq D^r$; we claim that the set $\Gamma \subseteq \Phi_D^{\mathbb{R}}$ of r -ary weighted relations γ with $\text{Feas}(\gamma) = F$

which are *not improved* by Ω is an open set. Take any $\gamma \in \Gamma$. There must be a non-zero weighting $\omega \in \Omega$ (let us denote its arity by k) and $X \in F^k$ such that $\langle \omega, \gamma[X] \rangle = d$ for some positive d , i.e. ω violates Inequality (25) for γ and X . Then for every r -ary weighted relation γ' with $\text{Feas}(\gamma') = F$ and distance from γ less than

$$\frac{d}{\sum_{f \in \text{supp}(\omega)} |\omega(f)|} \quad (35)$$

it holds $\langle \omega, \gamma'[X] \rangle > 0$, so γ has a neighbourhood contained in Γ . \square

We are now ready to prove our main result, stated as Theorems 5.11 and 5.12.

THEOREM 5.11. *For any finite D and any $\Gamma \subseteq \Phi_D^{\mathbb{R}}$, $\text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma)) = \text{wRelClone}_{\mathbb{R}}(\Gamma)$.*

PROOF. First, we show that $\text{wRelClone}_{\mathbb{R}}(\Gamma) \subseteq \text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma))$. Surely $\Gamma \subseteq \text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma))$ and therefore $\text{wRelClone}_{\mathbb{R}}(\Gamma) \subseteq \text{wRelClone}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma)))$. By Theorem 5.10 we know that $\text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma))$ is a weighted relational clone. Hence, $\text{wRelClone}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma))) = \text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma))$.

Now we will prove the other inclusion, $\text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma)) \subseteq \text{wRelClone}_{\mathbb{R}}(\Gamma)$. Let $\rho \in \text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma))$ be a weighted relation and denote $|\text{Feas}(\rho)|$ by k . If $k = 0$, ρ is expressible from ϕ_{\emptyset} and hence $\rho \in \text{wRelClone}_{\mathbb{R}}(\Gamma)$. Otherwise, we will focus solely on the k -ary weighted polymorphisms of Γ . Let us denote the Hilbert space of functions $\text{Pol}^{(k)}(\Gamma) \rightarrow \mathbb{R}$ by H ; the k -ary weighted polymorphisms of Γ can be then seen as vectors from H . Denoting $m = |D|^k$, a k -ary operation on D is uniquely determined by the m -tuple of labels it assigns to its m possible inputs. Later we will define a correspondence between a subset of m -ary weighted relations and H .

The outline of the proof is as follows. We transform Γ into a set $M \subseteq \text{wRelClone}_{\mathbb{R}}(\Gamma)$ of m -ary weighted relations and consider their corresponding vectors in H . The polar cone of these vectors equals $\text{wPol}_{\mathbb{R}}^{(k)}(\Gamma)$, and its polar cone, in turn, consists of m -ary weighted relations improved by $\text{wPol}_{\mathbb{R}}^{(k)}(\Gamma)$. We know that the polar cone of the polar cone of a set is the closure of the smallest convex cone containing this set; therefore, any m -ary weighted relation improved by $\text{wPol}_{\mathbb{R}}^{(k)}(\Gamma)$ belongs to $\text{wRelClone}_{\mathbb{R}}(\Gamma)$. This also includes a particular m -ary weighted relation that we use to express ρ , so we get $\rho \in \text{wRelClone}_{\mathbb{R}}(\Gamma)$.

We begin by formally defining the correspondence between certain m -ary weighted relations and vectors from H . Let us denote by $(z_1, \dots, z_m) = Z^T$ the sequence of all k -tuples over D in an arbitrary fixed order; any k -ary operation f is then determined by the m -tuple $(f(z_1), \dots, f(z_m)) = f(Z)$. Let us define a set $F \subseteq D^m$ as

$$F = \left\{ f(Z) \mid f \in \text{Pol}^{(k)}(\Gamma) \right\}. \quad (36)$$

For any m -ary weighted relation γ with $\text{Feas}(\gamma) = F$, the corresponding vector in H is $\gamma[Z]$. Conversely, for any vector in H there is a corresponding m -ary weighted relation with finite weights precisely on F .

Now we transform Γ into a set of m -ary weighted relations $M \subseteq \text{wRelClone}_{\mathbb{R}}(\Gamma)$ that captures enough information to reconstruct the set of k -ary weighted polymorphisms of Γ . Let $\gamma \in \Gamma$ be an n -ary weighted relation and $X \in (\text{Feas}(\gamma))^k$; we will denote the k -tuples of X^T by (x_1, \dots, x_n) . We claim that there is an m -ary weighted relation $\mu_{\gamma, X} \in \text{wRelClone}_{\mathbb{R}}(\Gamma)$ with $\text{Feas}(\mu_{\gamma, X}) = F$ such that $\mu_{\gamma, X}(f(Z)) = \gamma(f(X))$ for all

$f \in \text{Pol}^{(k)}(\Gamma)$. First, we construct an m -ary weighted relation $\mu'_{\gamma, X}$ as

$$\mu'_{\gamma, X}(y_{z_1}, \dots, y_{z_m}) = \gamma(y_{x_1}, \dots, y_{x_n}), \quad (37)$$

where y_{z_i} are variables indexed by k -tuples over D . Clearly, $\mu'_{\gamma, X}(f(Z)) = \mu'_{\gamma, X}(f(z_1), \dots, f(z_m)) = \gamma(f(x_1), \dots, f(x_n)) = \gamma(f(X))$. However, we are not done yet, as $\mu'_{\gamma, X}$ assigns a finite weight to all m -tuples $f(Z)$ such that $f(X) \in \text{Feas}(\gamma)$, even if $f \notin \text{Pol}^{(k)}(\Gamma)$. We can easily fix this: Let f be an k -ary operation that is not a polymorphism of Γ ; then there is a weighted relation $\gamma_f \in \Gamma$ and $X_f \in (\text{Feas}(\gamma_f))^k$ such that $f(X_f) \notin \text{Feas}(\gamma_f)$. Adding $0 \cdot \mu'_{\gamma_f, X_f}$ to $\mu'_{\gamma, X}$ ensures that the weighted relation will assign infinity to m -tuple $f(Z)$ without changing other weights. This can be done for all (finitely many) such operations f , so we obtain a weighted relation $\mu_{\gamma, X}$ with $\text{Feas}(\mu_{\gamma, X}) = F$.

Similarly, there are m -ary weighted relations $\mu_\iota, \mu_{-\iota} \in \text{wRelClone}_{\mathbb{R}}(\Gamma)$ with $\text{Feas}(\mu_\iota) = \text{Feas}(\mu_{-\iota}) = F$ such that $\mu_\iota(f(Z)) = 1$ and $\mu_{-\iota}(f(Z)) = -1$ for all $f \in \text{Pol}^{(k)}(\Gamma)$. Again, we start with $\mu'_\iota(y_{z_1}, \dots, y_{z_m}) = 1$, $\mu'_{-\iota}(y_{z_1}, \dots, y_{z_m}) = -1$ and then add $0 \cdot \mu'_{\gamma_f, X_f}$ for all k -ary operations $f \notin \text{Pol}^{(k)}(\Gamma)$ to ensure that the resulting weighted relations $\mu_\iota, \mu_{-\iota}$ assign finite weights only to m -tuples from F .

Let $\iota \in H$ be the vector assigning every operation value 1. For any operation $f \in \text{Pol}^{(k)}(\Gamma)$ that is not a projection, let $\chi_f \in H$ be the vector such that $\chi_f(f) = 1$ and $\chi_f(g) = 0$ for all $g \neq f$. We define a set of m -ary weighted relations $M \subseteq \text{wRelClone}_{\mathbb{R}}(\Gamma)$, the set of corresponding vectors $V \subseteq H$, and an auxiliary set of vectors $W \subseteq H$ as follows:

$$M = \{\mu_{\gamma, X} \mid \gamma \in \Gamma \wedge X \in (\text{Feas}(\gamma))^k\} \cup \{\mu_\iota, \mu_{-\iota}\} \quad (38)$$

$$V = \{\gamma[X] \mid \gamma \in \Gamma \wedge X \in (\text{Feas}(\gamma))^k\} \cup \{\iota, -\iota\} \quad (39)$$

$$W = V \cup \{-\chi_f \mid f \in \text{Pol}^{(k)}(\Gamma) \setminus \mathbf{J}_D^{(k)}\}. \quad (40)$$

We claim that the polar cone W° consists of k -ary weighted polymorphisms of Γ . Let $\omega \in W^\circ$ be a vector. As $\langle \omega, \iota \rangle \leq 0$ and $\langle \omega, -\iota \rangle \leq 0$, we have $\langle \omega, \iota \rangle = 0$, i.e. the sum of weights of ω equals 0. For any non-projection f we have $\langle \omega, -\chi_f \rangle \leq 0$, i.e. $\omega(f)$ is non-negative. Finally, for any $\gamma \in \Gamma$ and $X \in (\text{Feas}(\gamma))^k$ it holds $\langle \omega, \gamma[X] \rangle \leq 0$; hence ω is a weighted polymorphism of Γ .

Let us now return to weighted relation ρ and denote the sequence of elements of $\text{Feas}(\rho)$ in an arbitrary fixed order by $R \in (\text{Feas}(\rho))^k$. As ρ is improved by $\text{wPol}_{\mathbb{R}}(\Gamma)$, any vector $\omega \in W^\circ$ satisfies Inequality (25) for ρ and any $X \in (\text{Feas}(\rho))^k$, in particular for $X = R$. Hence, we would like to claim that $\langle \omega, \rho[R] \rangle \leq 0$ for all $\omega \in W^\circ$ and thus $\rho[R] \in W^{\circ\circ}$. However, $\rho[R]$ might be ill-defined: Although $\rho \in \text{Imp}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\Gamma))$, not necessarily all operations $f \in \text{Pol}^{(k)}(\Gamma)$ are polymorphisms of ρ , and therefore possibly $\rho(f(R)) = \infty$. Let us denote the set of these problematic operations by

$$Q = \left\{ f \in \text{Pol}^{(k)}(\Gamma) \mid f(R) \notin \text{Feas}(\rho) \right\}. \quad (41)$$

On the other hand, every operation in the support of $\text{wPol}_{\mathbb{R}}^{(k)}(\Gamma)$ is a polymorphism of ρ . This implies that operations in Q must be assigned a zero weight by all $\omega \in W^\circ$. As $\rho[R]$ might not exist, let us define instead a substitute vector $\beta \in H$ such that $\beta(f) = \rho(f(R))$ for all $f \in \text{Pol}^{(k)}(\Gamma) \setminus Q$, with arbitrary values assigned to operations in Q . By the previous argument, $\beta \in W^{\circ\circ}$. Additionally, let $\beta_0 \in H$ be a vector such that $\beta_0(f) > 0$ if $f \in Q$, and $\beta_0(f) = 0$ otherwise. For any $\omega \in W^\circ$ it holds $\langle \omega, \beta_0 \rangle = 0$, so β_0 also belongs to $W^{\circ\circ}$.

Any vector in $W^{\circ\circ}$ can be obtained from some vector in $V^{\circ\circ}$ by adding non-negative multiples of $-\chi_f$ for $f \in \text{Pol}^{(k)}(\Gamma) \setminus \mathbf{J}_D^{(k)}$. Therefore, there is a vector $\alpha \in V^{\circ\circ}$ such that $\alpha(f) \geq \beta(f) = \rho(f(R))$ for all $f \notin Q$, and $\alpha(f) = \beta(f) = \rho(f(R))$ when f is a projection. Also, there is a non-negative vector $\alpha_0 \in V^{\circ\circ}$ such that $\alpha_0(f) \geq \beta_0(f) > 0$ if $f \in Q$, and $\alpha_0(f) = \beta_0(f) = 0$ if f is a projection.

Vectors in V correspond to weighted relations in $M \subseteq \text{wRelClone}_{\mathbb{R}}(\Gamma)$. Set $V^{\circ\circ}$ is the closure of the smallest convex cone containing V , and therefore weighted relations corresponding to vectors in $V^{\circ\circ}$ also belong to $\text{wRelClone}_{\mathbb{R}}(\Gamma)$ (as it is closed under addition and non-negative scaling, and is topologically closed). Hence, there are m -ary weighted relations $\psi, \psi_0 \in \text{wRelClone}_{\mathbb{R}}(\Gamma)$ with $\text{Feas}(\psi) = \text{Feas}(\psi_0) = F$ such that $\psi[Z] = \alpha$ and $\psi_0[Z] = \alpha_0$. We are going to express ρ from them.

Let us denote the arity of ρ by n and the k -tuples of R^T by $(\mathbf{r}_1, \dots, \mathbf{r}_n)$. Consider the following gadget. Let I be a VCSP instance with variables $y_{\mathbf{z}_1}, \dots, y_{\mathbf{z}_m}$ and a single constraint $\psi(y_{\mathbf{z}_1}, \dots, y_{\mathbf{z}_m})$, and let $L = (y_{\mathbf{r}_1}, \dots, y_{\mathbf{r}_n})$. Then $\pi_L(I)$ is an n -ary weighted relation expressible over $\text{wRelClone}_{\mathbb{R}}(\Gamma)$; we will denote it by ρ' . For any n -tuple $\mathbf{x} \in D^n$, we have

$$\rho'(\mathbf{x}) = \min_{\{(y_{\mathbf{z}_1}, \dots, y_{\mathbf{z}_m}) \in D^m \mid (y_{\mathbf{r}_1}, \dots, y_{\mathbf{r}_n}) = \mathbf{x}\}} \psi(y_{\mathbf{z}_1}, \dots, y_{\mathbf{z}_m}) \quad (42)$$

$$= \min_{\{f: D^k \rightarrow D \mid (f(\mathbf{r}_1), \dots, f(\mathbf{r}_n)) = \mathbf{x}\}} \psi(f(\mathbf{z}_1), \dots, f(\mathbf{z}_m)) \quad (43)$$

$$= \min_{f(R)=\mathbf{x}} \psi(f(Z)) = \min_{f(R)=\mathbf{x}} \alpha(f). \quad (44)$$

Analogously, by replacing ψ with ψ_0 in the gadget we obtain an n -ary weighted relation ρ'_0 for which $\rho'_0(\mathbf{x}) = \min_{f(R)=\mathbf{x}} \alpha_0(f)$.

For any $\mathbf{x} \in \text{Feas}(\rho)$ and k -ary operation f such that $f(R) = \mathbf{x}$, it holds $\alpha(f) \geq \rho(f(R)) = \rho(\mathbf{x})$. As R is a list of all elements of $\text{Feas}(\rho)$, there is a projection f such that $f(R) = \mathbf{x}$; for it we have $\alpha(f) = \rho(f(R)) = \rho(\mathbf{x})$. Therefore, $\rho'(\mathbf{x}) = \rho(\mathbf{x})$. Similarly we get $\rho'_0(\mathbf{x}) = 0$ for any $\mathbf{x} \in \text{Feas}(\rho)$.

We are almost done; the last issue is that $\rho'(\mathbf{x})$ may be finite also for some $\mathbf{x} \notin \text{Feas}(\rho)$. But $f(R) \notin \text{Feas}(\rho)$ implies $f \in Q$, and in that case $\alpha_0(f)$ is positive. Therefore, $\text{Opt}(\rho'_0)$ is finite only on $\text{Feas}(\rho)$, and $\rho' + \text{Opt}(\rho'_0) = \rho$. \square

THEOREM 5.12. *For any finite D and any $\Omega \subseteq \mathbf{W}_D^{\mathbb{R}}$, $\text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega)) = \text{wClone}_{\mathbb{R}}(\Omega)$.*

PROOF. We begin with the inclusion $\text{wClone}_{\mathbb{R}}(\Omega) \subseteq \text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega))$.

Weightings in Ω are weighted polymorphisms of all weighted relations in $\text{Imp}_{\mathbb{R}}(\Omega)$, so $\Omega \subseteq \text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega))$, and hence $\text{wClone}_{\mathbb{R}}(\Omega) \subseteq \text{wClone}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega)))$. By Theorem 5.9, we have that $\text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega))$ is a weighted clone, so $\text{wClone}_{\mathbb{R}}(\text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega))) = \text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega))$.

Now we prove that for any $k \geq 1$ and any k -ary weighting $\mu \in \text{wPol}_{\mathbb{R}}(\text{Imp}_{\mathbb{R}}(\Omega))$, it holds $\mu \in \text{wClone}_{\mathbb{R}}(\Omega)$. First, let us establish the clone of operations we will be working with. Let C be the smallest clone containing $\text{supp}(\Omega)$. The support of $\text{wClone}_{\mathbb{R}}(\Omega)$ is itself a clone (by Theorem 5.5) so we also have $C = \text{supp}(\text{wClone}_{\mathbb{R}}(\Omega))$. As in the proof of Theorem 5.11, we will represent k -ary weightings by vectors of the Hilbert space $H = C^{(k)} \rightarrow \mathbb{R}$, and identify those vectors with certain m -ary weighted relations (where $m = |D|^k$).

The outline of the proof is as follows. We transform Ω into a set W of k -ary weightings. Although some of these weightings may be improper, any proper weighting obtained as their non-negative linear combination belongs to $\text{wClone}_{\mathbb{R}}(\Omega)$. The polar cone W° consists of m -ary weighted relations improved by Ω , and its polar cone $W^{\circ\circ}$ hence contains μ . As the polar cone of the polar cone of a set is the closure of the smallest convex cone containing this set, we get $\mu \in \text{wClone}_{\mathbb{R}}(\Omega)$.

Recall the correspondence between a subset of m -ary weighted relations and H from the proof of Theorem 5.11. This time, we are working with clone C , so we define F as

$$F = \left\{ f(Z) \mid f \in C^{(k)} \right\}. \quad (45)$$

Let Γ be the set of all m -ary weighted relations γ with $\text{Feas}(\gamma) = F$. Similarly as before, there is a bijection between Γ and H : the corresponding vector to a weighted relation $\gamma \in \Gamma$ is $\gamma[Z]$.

We show that k -ary polymorphisms of any $\gamma \in \Gamma$ are precisely the operations from $C^{(k)}$. Let $f \in C^{(k)}$. For any $X \in F^k$, there are k -ary operations $g_1, \dots, g_k \in C^{(k)}$ such that $X = (g_1(Z), \dots, g_k(Z))$. So we have $f(X) = f[g_1, \dots, g_k](Z) \in F$ because $f[g_1, \dots, g_k] \in C^{(k)}$. Conversely, let f be a k -ary operation not belonging to $C^{(k)}$. Certainly $Z = (e_1^{(k)}(Z), \dots, e_k^{(k)}(Z)) \in F^k$, but $f(Z) \notin F$. Therefore, f is not a polymorphism of γ .

Let us define a set $W \subseteq H$ as

$$W = \left\{ \omega[g_1, \dots, g_\ell] \in H \mid \ell \geq 1 \wedge \omega \in \Omega^{(\ell)} \wedge g_1, \dots, g_\ell \in C^{(k)} \right\}. \quad (46)$$

We claim that for any vector in the polar cone W° , the corresponding weighted relation is improved by Ω . Let $\gamma \in \Gamma$ be a weighted relation such that $\gamma[Z] \in W^\circ$, $\omega \in \Omega$ be an ℓ -ary weighting, and $X \in F^\ell$. Then there are k -ary operations $g_1, \dots, g_\ell \in C^{(k)}$ for which $X = (g_1(Z), \dots, g_\ell(Z))$, and we have

$$\sum_{f \in \text{supp}(\omega)} \omega(f) \cdot \gamma(f(X)) = \sum_{f \in C^{(\ell)}} \omega(f) \cdot \gamma(f[g_1, \dots, g_\ell](Z)) \quad (47)$$

$$= \sum_{f \in C^{(k)}} \omega[g_1, \dots, g_\ell](f) \cdot \gamma(f(Z)) \quad (48)$$

$$= \langle \omega[g_1, \dots, g_\ell], \gamma[Z] \rangle \leq 0. \quad (49)$$

Weighting μ is a weighted polymorphism of $\text{Imp}_{\mathbb{R}}(\Omega)$, so it improves any weighted relation γ corresponding to a vector in W° . Firstly, this implies $\text{supp}(\mu) \subseteq C^{(k)}$; we can therefore view μ as a vector of H . Secondly, μ satisfies Inequality (25) for γ and any $X \in F^k$. In particular, $Z \in F^k$, so we get $\langle \mu, \gamma[Z] \rangle \leq 0$ and thus $\mu \in W^{\circ\circ}$.

Set $W^{\circ\circ}$ is the closure of the smallest convex cone containing W . By Theorem 5.4, any proper weighting obtained as a non-negative linear combination of weightings from W belongs to $\text{wClone}_{\mathbb{R}}(\Omega)$. Therefore, $\mu \in \text{wClone}_{\mathbb{R}}(\Omega)$. \square

5.3. Complexity Consequences

When studying the computational complexity of constraint languages, the focus on weighted relational clones is justified by Theorem 2.13. The aim of this section is to discuss the consequences of our changes in the definition of weighted relational clones (allowing real weights and requiring weighted relational clones to be closed under operator Opt and to be topologically closed) on the validity of that theorem. We will assume the arithmetic model of computation, i.e. basic arithmetic operations with real numbers take a constant time.

First we show that adding operator Opt preserves the tractability of weighted relational clones.

THEOREM 5.13. *Let $\Gamma \subseteq \Phi_D^{\mathbb{R}}$ be a finite constraint language and $\gamma \in \Gamma$. Then $\text{VCSP}(\Gamma \cup \{\text{Opt}(\gamma)\})$ polynomial-time reduces to $\text{VCSP}(\Gamma)$.*

PROOF. Adding a constant to all weights of a weighted relation changes the value of every assignment by the same amount, and hence does not affect tractability. Without

loss of generality, we may therefore assume that all weighted relations in Γ assign non-negative weights and that the minimum weight assigned by γ equals 0. We will also assume that γ is not a relation, otherwise $\text{Opt}(\gamma) = \gamma$ so the claim would hold trivially. Let us denote by m the smallest positive weight assigned by γ , and by M the largest finite weight assigned by any $\gamma' \in \Gamma$.

Let $I \in \text{VCSP}(\Gamma \cup \{\text{Opt}(\gamma)\})$ be an instance with q constraints. We replace every constraint of the form $\text{Opt}(\gamma)(\mathbf{x})$ in I with $(q \cdot \lceil M/m \rceil + 1)$ copies of $\gamma(\mathbf{x})$, thus obtaining a polynomially larger instance $I' \in \text{VCSP}(\Gamma)$. Any feasible assignment for instance I is also feasible for I' with the same value, which does not exceed qM . On the other hand, any infeasible assignment for instance I is infeasible for I' , or it incurs an infinite value from a constraint of the form $\text{Opt}(\gamma)(\mathbf{x})$ in I and therefore a value of more than qM in I' . \square

For any constraint language $\Gamma \subseteq \Phi_D^{\mathbb{R}}$, we will denote by Γ_{\sim} the smallest set of weighted relations containing Γ that is closed under scaling by non-negative real constants and addition of real constants. Analogously to [Cohen et al. 2013, Theorem 4.3], we would like to show that Γ is tractable if and only if Γ_{\sim} is tractable. Their proof, however, does not apply to scaling by an *irrational* factor α , as it relies on the existence of integers p, q such that $\alpha = p/q$. In fact, we were not able to prove that real-valued scaling preserves tractability when insisting on exact solvability. If we consider solving VCSP with an absolute error bounded by ϵ (for any $\epsilon > 0$), then real-valued scaling *does preserve* tractability, as shown in the following theorem.

THEOREM 5.14. *Let $\Gamma, \Gamma' \subseteq \Phi_D^{\mathbb{R}}$ be finite constraint languages such that Γ contains only weighted relations of the form $c \cdot \gamma'$ for $c \geq 0, \gamma' \in \Gamma'$. For any $\epsilon > 0$, there is a polynomial-time reduction that for any instance $I \in \text{VCSP}(\Gamma)$ outputs an instance $I' \in \text{VCSP}(\Gamma')$ such that for any optimal assignment s' of I' it holds $I(s') \in [v, v + \epsilon]$, where v is the value of an optimal assignment of I .*

PROOF. Again, we may assume that all weighted relations in Γ and Γ' assign non-negative weights. Let us denote by M the largest finite weight assigned by any weighted relation in Γ' ; we may assume that M is well-defined and positive, otherwise we would have $\Gamma \subseteq \Gamma'$. We will denote by q the number of constraints in I and let $b = \lceil qM/\epsilon \rceil$.

Let instance $I' \in \text{VCSP}(\Gamma')$ have the same set of variables as I . For any constraint $c_i \cdot \gamma'_i(\mathbf{x}_i)$ of I , we add $(\lfloor bc_i \rfloor + 1)$ copies of constraint $\gamma'_i(\mathbf{x}_i)$ into I' . Note that any feasible assignment of I is a feasible assignment of I' , and vice versa. Let us assume that I admits a feasible solution. As $(\lfloor bc_i \rfloor + 1) - bc_i \in (0, 1]$, we have

$$b \cdot I(t) \leq I'(t) \leq b \cdot I(t) + qM \quad (50)$$

for any feasible assignment t . Let s be an optimal assignment of I ; we get

$$b \cdot I(s') \leq I'(s') \leq I'(s) \leq b \cdot I(s) + qM, \quad (51)$$

and therefore $I(s') \leq I(s) + \epsilon = v + \epsilon$. \square

Taking a topological closure of a language also preserves tractability with a bounded absolute error, as the following theorem shows.

THEOREM 5.15. *Let $\Gamma \subseteq \Phi_D^{\mathbb{R}}$ be a constraint language and denote by $\bar{\Gamma}$ its closure. For any $\epsilon > 0$, there is a polynomial-time reduction that for any instance $I \in \text{VCSP}(\bar{\Gamma})$ outputs an instance $I' \in \text{VCSP}(\Gamma)$ with the same variables such that any assignment t is either infeasible for both I and I' , or is feasible for both and $|I(t) - I'(t)| \leq \epsilon$.*

PROOF. Let us denote by q the number of constraints in I . For any $\gamma \in \bar{\Gamma}$, there is a weighted relation $\gamma' \in \Gamma$ of the same arity and with $\text{Feas}(\gamma) = \text{Feas}(\gamma')$ such that the distance between γ and γ' is at most ϵ/q . We obtain the sought instance I' by replacing all constraints γ from I with their counterparts γ' . \square

We finish this section with a discussion of the difficulty of improving Theorem 5.14 to exact solvability (to optimality). Let Γ be a finite constraint language and $\gamma \in \Gamma$. We would like to prove that $\text{VCSP}(\Gamma \cup \{c \cdot \gamma\})$ polynomial-time reduces to $\text{VCSP}(\Gamma)$, where $c \in \mathbb{R}_{\geq 0}$. Given $I \in \text{VCSP}(\Gamma \cup \{c \cdot \gamma\})$, let

$$\delta_I = \min \left\{ |I(s_1) - I(s_2)| \mid s_1, s_2 \text{ are solutions to } I \text{ with different values} \right\}. \quad (52)$$

If we choose an $\epsilon < \delta_I$, we obtain an optimal solution of I by Theorem 5.14. However, it is not clear how fast the value of δ_I approaches 0 as the size of I grows to infinity, and whether it is possible to compute it in polynomial time.

ACKNOWLEDGMENTS

The authors are grateful to the anonymous referees for their helpful comments and suggestions.

References

- Libor Barto and Marcin Kozik. 2014. Constraint Satisfaction Problems Solvable by Local Consistency Methods. *J. ACM* 61, 1 (2014). Article No. 3.
- Libor Barto, Marcin Kozik, and Todd Niven. 2009. 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, 5 (2009), 1782–1802.
- V.G. Bodnarčuk, L.A. Kalužnin, V.N. Kotov, and B.A. Romov. 1969. Galois theory for Post algebras. I. *Cybernetics and Systems Analysis* 5, 3 (1969), 243–252.
- S.P. Boyd and L. Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.
- Andrei Bulatov. 2004. A Graph of a Relational Structure and Constraint Satisfaction Problems. In *Proceedings 19th IEEE Symposium on Logic in Computer Science (LICS'04)*. IEEE Computer Society, 448–457.
- Andrei Bulatov. 2006. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM* 53, 1 (2006), 66–120.
- Andrei Bulatov, Andrei Krokhin, and Peter Jeavons. 2005. Classifying the Complexity of Constraints using Finite Algebras. *SIAM J. Comput.* 34, 3 (2005), 720–742.
- Andrei A. Bulatov. 2011. Complexity of conservative constraint satisfaction problems. *ACM Transactions on Computational Logic* 12, 4 (2011). Article 24.
- Andrei A. Bulatov, Andrei A. Krokhin, and Peter G. Jeavons. 2001. The complexity of maximal constraint languages. In *Proceedings 33rd ACM Symposium on Theory of Computing (STOC'01)*. 667–674.
- David A. Cohen, Martin C. Cooper, Páidí Creed, Peter Jeavons, and Stanislav Živný. 2013. An algebraic theory of complexity for discrete optimisation. *SIAM J. Comput.* 42, 5 (2013), 915–939.
- David A. Cohen, Martin C. Cooper, Peter G. Jeavons, and Andrei A. Krokhin. 2006. The Complexity of Soft Constraint Satisfaction. *Artificial Intelligence* 170, 11 (2006), 983–1016.
- Tomás Feder and Moshe Y. Vardi. 1998. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM J. Comput.* 28, 1 (1998), 57–104.
- Peter Fulla and Stanislav Živný. 2015. A Galois connection of valued constraint languages of infinite size. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP'15) (Lecture Notes in Computer Science)*, Vol. 9134. Springer, 517–528.
- David Geiger. 1968. Closed systems of functions and predicates. *Pacific Journal of Mathematics* 27, 1 (1968), 95–100.
- Pavol Hell and Jaroslav Nešetřil. 1990. On the Complexity of H -coloring. *Journal of Combinatorial Theory, Series B* 48, 1 (1990), 92–110.
- Pavol Hell and Jaroslav Nešetřil. 2008. Colouring, constraint satisfaction, and complexity. *Computer Science Review* 2, 3 (2008), 143–163.
- J.B. Hiriart-Urruty and C. Lemaréchal. 2001. *Fundamentals of Convex Analysis*. Springer Berlin Heidelberg.

- Peter Jeavons, Andrei Krokhin, and Stanislav Živný. 2014. The complexity of valued constraint satisfaction. *Bulletin of the European Association for Theoretical Computer Science (EATCS)* 113 (2014), 21–55.
- Peter G. Jeavons, David A. Cohen, and Marc Gyssens. 1997. Closure Properties of Constraints. *Journal of the ACM* 44, 4 (1997), 527–548.
- Vladimir Kolmogorov, Andrei A. Krokhin, and Michal Rolínek. 2015a. *The Complexity of General-Valued CSPs*. Technical Report. arXiv:1502.07327.
- Vladimir Kolmogorov, Johan Thapper, and Stanislav Živný. 2015b. The power of linear programming for general-valued CSPs. *SIAM J. Comput.* 44, 1 (2015), 1–36.
- Vladimir Kolmogorov and Stanislav Živný. 2013. The complexity of conservative valued CSPs. *J. ACM* 60, 2 (2013). Article No. 10.
- Marczin Kozik and Joanna Ochremiak. 2015. Algebraic Properties of Valued Constraint Satisfaction Problem. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP'15) (Lecture Notes in Computer Science)*, Vol. 9134. Springer, 846–858.
- Thomas J. Schaefer. 1978. The Complexity of Satisfiability Problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC'78)*. ACM, 216–226.
- Johan Thapper. 2010. *Aspects of a Constraint Optimisation Problem*. Ph.D. Dissertation. Department of Computer Science and Information Science, Linköping University.
- Johan Thapper and Stanislav Živný. 2012. The power of linear programming for valued CSPs. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'12)*. IEEE, 669–678.
- Johan Thapper and Stanislav Živný. 2015a. Necessary Conditions on Tractability of Valued Constraint Languages. *SIAM Journal on Discrete Mathematics* 29, 4 (2015), 2361–2384.
- Johan Thapper and Stanislav Živný. 2015b. The complexity of finite-valued CSPs. *J. ACM* (2015). To appear.
- Johan Thapper and Stanislav Živný. 2015c. The power of Sherali-Adams relaxations for valued CSPs. (2015). Submitted for publication.
- Johan Thapper and Stanislav Živný. 2015d. Sherali-Adams relaxations for valued CSPs. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP'15) (Lecture Notes in Computer Science)*. Springer.
- Hannes Uppman. 2013. The Complexity of Three-Element Min-Sol and Conservative Min-Cost-Hom. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP'13) (Lecture Notes in Computer Science)*, Vol. 7965. Springer, 804–815.
- Andrius Vaicenavičius. 2014. *A study of weighted clones*. Master's thesis. Mathematical Institute, University of Oxford.
- Jiří Vančura. 2014. *Weighted Clones*. Master's thesis. Department of Algebra, Charles University.
- Stanislav Živný. 2012. *The complexity of valued constraint satisfaction problems*. Springer.

Received April 2015; revised October 2015; accepted November 2015