

The complexity of finite-valued CSPs

[Extended Abstract]*

Johan Thapper[†]
Laboratoire de Recherche en Informatique (LRI)
Université Paris-Sud 11
France
thapper@lri.fr

Stanislav Živný
Department of Computer Science
University of Warwick
UK
S.Zivny@warwick.ac.uk

ABSTRACT

Let Γ be a set of rational-valued functions on a fixed finite domain; such a set is called a *finite-valued constraint language*. The valued constraint satisfaction problem, $\text{VCSP}(\Gamma)$, is the problem of minimising a function given as a sum of functions from Γ . We establish a dichotomy theorem with respect to exact solvability for *all* finite-valued languages defined on domains of *arbitrary* finite size.

We show that every *core* language Γ either admits a binary idempotent and symmetric fractional polymorphism in which case the basic linear programming relaxation solves any instance of $\text{VCSP}(\Gamma)$ exactly, or Γ satisfies a simple hardness condition that allows for a polynomial-time reduction from Max-Cut to $\text{VCSP}(\Gamma)$. In other words, there is a single algorithm for all tractable cases and a single reason for intractability. Our results show that for exact solvability of VCSPs the basic linear programming relaxation suffices and semidefinite relaxations do not add any power. The proof uses a variation of Motzkin's Transposition Theorem, hyperplane arrangements, and a technique recently introduced by Kolmogorov [arXiv:1207.7213] reformulated using Markov chains.

Our results generalise *all previous partial classifications* of finite-valued languages: the classifications of $\{0, 1\}$ -valued languages on two-element, three-element, and four-element domains obtained by Creignou [JCSS'95], Jonsson et al. [SICOMP'06], and Jonsson et al. [CP'11], respectively; the classification of $\{0, 1\}$ -valued languages containing all unary functions obtained by Deineko et al. [JACM'06]; the classifications of finite-valued languages on two-element and three-element domains obtained by Cohen et al. [AIJ'06] and Huber et al. [SODA'13], respectively; the classification of finite-

valued languages containing all $\{0, 1\}$ -valued unary functions obtained by Kolmogorov and Živný [SODA'12]; and the classification of Min-0-Ext problems obtained recently by Hirai [SODA'13].

Categories and Subject Descriptors

F.2.0 [Analysis of Algorithms and Problem Complexity]: General

General Terms

Algorithms, Theory

Keywords

discrete optimisation, complexity, valued constraint satisfaction problems, dichotomy

1. INTRODUCTION

In this paper we study the following problem: what classes of discrete explicitly-represented functions can be minimised exactly in polynomial time? Such problems can be readily described as (finite-)valued constraint satisfaction problems. We provide a complete answer to this question for rational-valued functions defined on arbitrary finite domains.

The constraint satisfaction problem, or CSP for short, provides a common framework for many theoretical and practical problems in computer science. An instance can be vaguely described as a set of variables to be assigned values from the domains of the variables so that all constraints are satisfied [42]. The CSP is NP-complete in general and thus we are interested in restrictions which give rise to tractable classes of problems. Following Feder and Vardi [22], we restrict the constraint language, that is, all constraint relations in a given instance must belong to a fixed, finite set of relations on the domain. The most successful approach to classifying language-restricted CSPs is the so-called algebraic approach [30, 29, 7], which has led to several complexity classifications [6, 8, 5, 2] and algorithmic characterisations [3, 28] going beyond the seminal work of Schaefer on Boolean CSPs [45].

There are several natural optimisation variants of CSPs that have been studied in the literature such as Max-CSP, where the goal is to maximise the number of satisfied constraints (or, equivalently, minimise the number of unsatisfied constraints) [9, 19, 31, 34, 21], and Max-Ones [19, 33] and Min-Cost-Hom [48], where all constraints have to be satisfied and some additional function of the assignment is

*A full version of this paper is available at [arXiv:1210.2987](https://arxiv.org/abs/1210.2987).

[†]The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 257039).

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC '13, June 1-4, 2013, Palo Alto, CA, USA.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

optimised. The most general variant is the valued constraint satisfaction problem, or VCSP for short [12]. A valued constraint language Γ is a set of functions on a fixed domain and a VCSP instance over Γ is given by a sum of functions from Γ with the goal to minimise the sum. The VCSP framework is very robust and has also been studied under different names such as Min-Sum problems, Gibbs energy minimisation, Markov Random Fields, Conditional Random Fields and others in different contexts in computer science [40, 50, 16]. The VCSP in its full generality considers functions with the range being the rationals with positive infinity [12]; this includes both CSPs and Max-CSPs as special cases with functions with the range being $\{0, \infty\}$ and $\{0, 1\}$, respectively. In this work we will focus on finite-valued CSPs, that is, the range of the functions is the set of rationals. (Finite-valued CSPs are called generalised CSPs in [43].)

Given the generality of the VCSP, it is not surprising that only few complexity classifications are known. In the general-valued case, only languages on a two-element domain [12, 17] and conservative (containing all $\{0, 1\}$ -valued unary functions) languages [38] have been completely classified with respect to exact solvability. In the finite-valued case, languages on two-element domains [12], three-element domains [27], and conservative languages [38] have been completely classified with respect to exact solvability. In the special case of $\{0, 1\}$ -valued languages, which correspond to Max-CSPs, languages on two-element domains [18], three-element domains [31], and four-element domains [34], and conservative (containing all $\{0, 1\}$ -valued unary functions) languages [21] have been classified with respect to exact solvability. Generalising the algebraic approach to CSPs [7], algebraic properties called multimorphisms [12], fractional polymorphism [11], and weighted polymorphisms [13, 10] have been invented for the study of the computational complexity of classes of VCSPs. The authors have previously shown that any VCSP instance over a valued constraint language Γ can be solved by the basic linear programming (BLP) relaxation if, and only if, Γ admits symmetric fractional polymorphisms of all arities [49]. This result established the tractability of several (previously open) classes of VCSPs including submodular VCSPs on non-distributive lattices and trees. Moreover, results from [49] have been used in the complexity classification of minimum 0-extension problems, which can be cast as special binary VCSPs [26].

1.1 Contribution

We study the computational complexity of finite-valued constraint languages on arbitrary finite domains. We characterise all tractable finite-valued languages as those admitting a binary idempotent and symmetric fractional polymorphism. (This is a polynomial-time checkable condition.) Tractability follows from the results in [37, 49] that show that all instances over such languages are solvable by the BLP. In the other direction, we show that instances over languages not admitting such a fractional polymorphism are NP-hard by a reduction from Max-Cut [23].

THEOREM 1.1. *Let D be an arbitrary finite set and let Γ be a core finite-valued language defined on D . VCSP(Γ) is tractable if, and only if, BLP solves VCSP(Γ). Otherwise, VCSP(Γ) is NP-hard.*

The core condition on Γ is defined in Section 2.4 and the precise statement of the main result is in Section 3.

The proof of our main result is a combination of a technique recently introduced by Kolmogorov [37], a variation of Motzkin’s Transposition Theorem, and hyperplane arrangements. Furthermore, we find a natural way to derive the result from [37], which says that having a binary symmetric fractional polymorphism implies having symmetric fractional polymorphisms of all arities, using a Markov chain argument.

1.2 Related work

Apart from language-based restrictions on (V)CSPs, also structure-based restrictions [25, 41, 24] and hybrid restrictions [14, 15] have been studied. Not only exact solvability, but also approximability of Max-CSPs and VCSPs has attracted a lot of attention [19, 35, 32]. Moreover, the robust approximability of Max-CSPs has also been studied [20, 39, 4]. Under the assumption of the *unique games conjecture* [36], Raghavendra showed that the basic SDP relaxation solves all tractable finite-valued VCSPs (without a characterisation of the tractable cases) [43]. Moreover, Chapters 6 and 7 of [44] imply that if a finite-valued language Γ admits a cyclic fractional polymorphism of some arity $k \geq 2$ then the basic SDP relaxation solves any VCSP instance over Γ . Our results show, assuming $P \neq NP$, that for exact solvability the BLP relaxation suffices.

Our results demonstrate that (i) only a binary fractional polymorphism of a certain type is sufficient for tractability, and (ii) only cores and constants are required for the hardness condition (details are explained in Section 2). This is in contrast with the $\{0, \infty\}$ -valued CSPs (that is, decision problems), where the hardness condition also requires an equivalence relation and the conjectured tractable cases are characterised by polymorphisms of arity higher than two [7].

2. PRELIMINARIES

We use the following notation: any name with a bar denotes a tuple. We denote by \bar{x}_i the i th component of a tuple \bar{x} . Superscripts are used for collections of tuples. When both a superscript and a subscript are used we drop the bar; e.g., we write x_i^j for \bar{x}_i^j .

2.1 Valued CSPs

Let D be a finite *domain*. The set of non-negative rational numbers is denoted by $\mathbb{Q}_{\geq 0}$. A *valued (constraint) language*¹ is a set of (*cost*) *functions* $f : D^m \rightarrow \mathbb{Q}_{\geq 0}$ where $m = ar(f)$ is called the *arity* of f and it may be different for different functions from Γ .²

DEFINITION 1. *An instance I of the valued constraint satisfaction problem, or VCSP for short, is given by the set $V = \{x_1, \dots, x_n\}$ of variables and the objective function $f_I(x_1, \dots, x_n) = \sum_{i=1}^q w_i \cdot f_i(\bar{x}^i)$ where, for every $1 \leq i \leq q$, $f_i : D^{ar(f_i)} \rightarrow \mathbb{Q}_{\geq 0}$, $\bar{x}^i \in V^{ar(f_i)}$, and $w_i \in \mathbb{Q}_{\geq 0}$ is a weight. A solution to I is a function $h : V \rightarrow D$, its measure given by $\sum_{i=1}^q w_i \cdot f_i(h(\bar{x}^i))$. The goal is to find a solution of minimum measure. This measure will be denoted by $\text{Opt}_\Gamma(I)$.*

¹These are known and were called in the abstract and the introduction as *finite-valued* languages. Since we deal with only finite-valued languages, we omit the prefix “finite-” in the rest of the paper.

²The range of the functions in Γ is $\mathbb{Q}_{\geq 0}$ for traditional reasons, but all results hold true with \mathbb{Q} as well.

We denote by $\text{VCSP}(\Gamma)$ the class of all instances in which all functions are from Γ .

A valued language Γ is called *tractable* if, for any finite $\Gamma' \subseteq \Gamma$, $\text{VCSP}(\Gamma')$ is tractable, that is, a solution of optimum measure can be found for any instance $I \in \text{VCSP}(\Gamma')$ in polynomial time; Γ is called *NP-hard* if $\text{VCSP}(\Gamma')$ is NP-hard for some finite $\Gamma' \subseteq \Gamma$.

2.2 Expressive power

DEFINITION 2. For a valued language Γ , we let $\langle \Gamma \rangle$ be the set of all functions $f(x_1, \dots, x_m)$ such that for some instance $I \in \text{VCSP}(\Gamma)$ with objective function f_I in variables $x_1, \dots, x_m, x_{m+1}, \dots, x_n$, we have

$$f(x_1, \dots, x_m) = \min_{x_{m+1}, \dots, x_n} f_I(x_1, \dots, x_m, x_{m+1}, \dots, x_n).$$

We then say that Γ expresses f and call $\langle \Gamma \rangle$ the expressive power of Γ .

In other words, $\langle \Gamma \rangle$ is the closure of Γ under addition, multiplication by constants, and minimisation over extra variables. For two functions f and f' , we write $f \equiv f'$ if $f = a \cdot f' + b$ for some $a \in \mathbb{Q}_{>0}$ and $b \in \mathbb{Q}$. (In other words, $f \equiv f'$ if f can be obtained from f' by *scaling* and *translating*.) For a valued language Γ , let $\Gamma_{\equiv} = \{f \mid f \equiv f' \text{ for some } f' \in \Gamma\}$. It has been shown that with respect to exact solvability, we only need to consider valued languages closed under expressibility:

THEOREM 2.1 ([12, 11]). *Let Γ and Γ' be valued languages defined on domain D such that $\Gamma' \subseteq \langle \Gamma \rangle_{\equiv}$. Then $\text{VCSP}(\Gamma')$ polynomial-time reduces to $\text{VCSP}(\Gamma)$.*

We define the following condition:

There exist distinct $a, b \in D$ such that $\langle \Gamma \rangle$ contains a binary function h with $\text{argmin } h = \{(a, b), (b, a)\}$. (MC)

A slightly different condition³ was formulated in [27]:

There exist distinct $a, b \in D$ such that $\langle \Gamma \rangle$ contains a unary function u with $\text{argmin } u = \{a, b\}$ and a binary function h with $h(a, b) = h(b, a) < h(a, a) = h(b, b)$. (MC')

Observe that (MC') implies (MC). In fact, we will prove in Section 4 that the two conditions are equivalent.

LEMMA 2.2. *For any valued language Γ , (MC) holds if, and only, if (MC') holds.*

It is known that condition (MC') and thus, by Lemma 2.2, condition (MC) implies intractability (via a reduction from Max-Cut [23]):

LEMMA 2.3 ([12]). *If a valued language Γ satisfies condition (MC) then Γ is NP-hard.*

2.3 Fractional polymorphisms

Let Γ be a valued language defined on D . For an n -ary function $f \in \Gamma$ and $\bar{a}^1, \dots, \bar{a}^m \in D^n$, we define f^m by $f^m(\bar{a}^1, \dots, \bar{a}^m) = \frac{1}{m}(f(\bar{a}^1) + \dots + f(\bar{a}^m))$. An m -ary operation on D is a function $g : D^m \rightarrow D$. Let $\mathcal{O}_D^{(m)}$ denote the set of all m -ary operations on D . An m -ary fractional operation is a function $\omega : \mathcal{O}_D^{(m)} \rightarrow \mathbb{Q}_{\geq 0}$. Define $\|\omega\|_1 := \sum_g \omega(g)$. An

³Condition (MC') was called (MC) in [27].

m -ary fractional operation ω is called an m -ary fractional polymorphism [11] of Γ if $\|\omega\|_1 = 1$ and for every function $f \in \Gamma$ and tuples $\bar{a}^1, \dots, \bar{a}^m \in D^{\text{ar}(f)}$, it holds that

$$\sum_{g \in \mathcal{O}_D^{(m)}} \omega(g) f(g(\bar{a}^1, \dots, \bar{a}^m)) \leq f^m(\bar{a}^1, \dots, \bar{a}^m), \quad (1)$$

where the operations g are applied componentwise. The set $\{g \mid \omega(g) > 0\}$ of operations is called the *support* of ω and is denoted by $\text{supp}(\omega)$. It is known and easy to show that expressibility preserves fractional polymorphisms: if ω is a fractional polymorphism of Γ then ω is also a fractional polymorphism of $\langle \Gamma \rangle$ [11].

An operation g is *idempotent* if $g(x, \dots, x) = x$. Let S_m be the symmetric group on $\{1, \dots, m\}$. An m -ary operation g is *symmetric* if for every permutation $\pi \in S_m$, we have $g(x_1, \dots, x_m) = g(x_{\pi(1)}, \dots, x_{\pi(m)})$. An m -ary operation g is *cyclic* if $g(a_1, a_2, \dots, a_m) = g(a_2, \dots, a_m, a_1)$ for every $a_1, \dots, a_m \in D$. Note that in the case of $m = 2$ both definitions coincide. A fractional operation is called idempotent, symmetric, or cyclic if all operations in its support are idempotent, symmetric, or cyclic, respectively.

Let $\mathcal{O}_D^{(m \rightarrow k)}$ denote the set of all mappings $\mathbf{g} : D^m \rightarrow D^k$. A (generalised) fractional operation of arity $m \rightarrow k$ is a function $\omega : \mathcal{O}_D^{(m \rightarrow k)} \rightarrow \mathbb{Q}_{\geq 0}$. As before, we define $\|\omega\|_1 := \sum_{\mathbf{g}} \omega(\mathbf{g})$. A (generalised) fractional operation ω of arity $m \rightarrow k$ is called a (generalised) fractional polymorphism [37] (of arity $m \rightarrow k$) of Γ if $\|\omega\|_1 = 1$ and for every function $f \in \Gamma$ and tuples $\bar{a}^1, \dots, \bar{a}^m \in D^{\text{ar}(f)}$, it holds that

$$\sum_{\mathbf{g} \in \mathcal{O}_D^{(m \rightarrow k)}} \omega(\mathbf{g}) f^k(\mathbf{g}(\bar{a}^1, \dots, \bar{a}^m)) \leq f^m(\bar{a}^1, \dots, \bar{a}^m). \quad (2)$$

Note that a fractional polymorphism of arity m is the same as a generalised fractional polymorphism of arity $m \rightarrow 1$. In fact a generalised fractional operation of arity $m \rightarrow k$ is just a sequence of k fractional operations of arity $m \rightarrow 1$; however, this viewpoint, introduced in [37], turns out to be very useful. For brevity, we will often omit the word ‘‘generalised’’.

For an operation g , we denote by χ_g the fractional operation that takes the value 1 on the operation g and 0 on all other operations. For a generalised operation \mathbf{g} , $\chi_{\mathbf{g}}$ is defined analogously.

2.4 Cores

Let $S \subseteq D$. The sub-language $\Gamma[S]$ of Γ induced by S is the valued language defined on domain S and containing the restriction of every function $f \in \Gamma$ onto S .

DEFINITION 3. A valued language Γ is a core if for every unary fractional polymorphism ω of Γ , $\text{supp}(\omega)$ contains only injective operations. A valued language Γ' is a core of Γ if Γ' is a core and $\Gamma' = \Gamma[h(D)]$ for some $h \in \text{supp}(\omega)$ with ω a unary fractional polymorphism of Γ .

LEMMA 2.4. *If Γ' is a core of Γ then $\text{Opt}_{\Gamma}(I) = \text{Opt}_{\Gamma'}(I')$ for all instances $I \in \text{VCSP}(\Gamma)$, where I' is obtained from I by substituting each function in Γ for its restriction in Γ' .*

By Lemma 2.4, proved in Section 4, we may assume that Γ is a core valued language.

For a valued language Γ , let Γ_c denote the language obtained from Γ by adding all functions obtained from func-

tions in Γ by fixing a subset of the variables to domain values.

We will use the following result, which says that we can restrict our attention to core valued languages whose expressive power contain certain unary functions.

PROPOSITION 2.5 ([27]). *Let Γ be a core valued language defined on a finite domain D .*

1. *For each $a \in D$, $\langle \Gamma_c \rangle$ contains a unary function u_a such that $\text{argmin } u_a = a$.*
2. *Γ is NP-hard if, and only if, Γ_c is NP-hard.*

Proposition 2.5 is proved for a different definition of a core in [27] but we will show in Section 4 that Definition 3 coincides with the definition of the core in [27].

It follows readily from Proposition 2.5 that every (generalised) fractional polymorphism of Γ_c for a core valued language Γ is idempotent.

3. RESULTS

3.1 Complexity classification

The computational complexity of valued constraint languages has attracted a lot of attention in the literature. Partial classifications obtained so far are the following:

- $\{0, 1\}$ -valued languages on $|D| = 2$ [18, 19].
- $\{0, 1\}$ -valued languages on $|D| = 3$ [31].
- $\{0, 1\}$ -valued languages on $|D| = 4$ [34].
- $\{0, 1\}$ -valued languages containing (special types of) unary functions [21].
- valued languages on $|D| = 2$ [12].
- valued languages on $|D| = 3$ [27].
- valued languages containing $\{0, 1\}$ -valued unary functions [38].
- valued languages containing unary functions and certain special binary functions [26].

In all the classifications obtained so far, the tractable cases were explained by certain specific binary symmetric fractional polymorphisms, and the hardness result essentially came from the condition (MC).

A recent result of the authors characterised the power of the basic linear programming (BLP) relaxation (cf. [49]), which is defined in Appendix A. An equivalent simplified condition was subsequently given by Kolmogorov [37].

THEOREM 3.1 ([49, 37]). *Let Γ be a valued language. Then BLP solves $VCSP(\Gamma)$ if, and only if, Γ has a binary symmetric fractional polymorphism.*

An alternative proof, based on Markov chains, of Kolmogorov's result (showing that a binary symmetric fractional polymorphism implies symmetric fractional polymorphisms of all arities) is given in Appendix B.

The main technical contribution of this paper is the following result.

THEOREM 3.2. *Let D be an arbitrary finite set and let Γ be a core valued language defined on D . If Γ_c does not satisfy (MC), then Γ admits a binary idempotent and symmetric fractional polymorphism.*

Theorem 3.2 implies our main result, Theorem 3.3, which shows that having a binary symmetric fractional polymorphism is the *only* reason for tractability. This provides a complexity classification of *all* valued languages defined on arbitrary finite domains, thus generalising all classifications mentioned above.

THEOREM 3.3 (MAIN). *Let D be an arbitrary finite set and let Γ be a core valued language defined on D .*

- *Either Γ has a binary idempotent and symmetric fractional polymorphism and BLP solves $VCSP(\Gamma)$;*
- *or (MC) holds for Γ_c and $VCSP(\Gamma)$ is NP-hard.*

PROOF. Assuming that Γ is a core, if Γ_c satisfies (MC), then $VCSP(\Gamma_c)$ is NP-hard by Lemma 2.3. In this case $VCSP(\Gamma)$ is NP-hard by Proposition 2.5(2). Otherwise, by Theorem 3.2, Γ admits a binary idempotent and symmetric fractional polymorphism and it follows from Theorem 3.1 that BLP solves $VCSP(\Gamma)$. \square

Deciding whether a valued language Γ is a core and deciding the tractability of a $VCSP(\Gamma)$ instance is discussed in Section 3.3.

COROLLARY 3.4 (OF THEOREM 3.2). *Let D be an arbitrary finite set and let Γ be a core valued language defined on D . The following are equivalent:*

1. *Γ_c does not satisfy (MC);*
2. *Γ admits an idempotent and cyclic fractional polymorphism of some arity $k > 1$;*
3. *Γ admits an idempotent and symmetric fractional polymorphism of some arity $k > 1$;*
4. *Γ admits a binary idempotent and symmetric fractional polymorphism;*
5. *BLP solves $VCSP(\Gamma)$.*

PROOF. Theorem 3.1 gives (4) \Leftrightarrow (5). The implications (4) \Rightarrow (3) \Rightarrow (2) are trivial, and it is not hard to show that (2) \Rightarrow (1). Finally, Theorem 3.2 gives the implication (1) \Rightarrow (4). \square

Corollary 3.4 answers Problem 1 from [27] that asked about the relationship between the complexity of a valued language Γ and the existence of various types of fractional polymorphisms of Γ . Note that Corollary 3.4 holds unconditionally. Problem 1 from [27] also involved the solvability by the basic SDP relaxation [43], which at the time was known to be implied by (2) and imply (1), provided that $P \neq NP$. Under the same assumption, we conclude that solvability by the basic SDP relaxation is also characterised by any of the equivalent statements of Corollary 3.4.

3.2 Proof overview

In order to prove Theorem 3.2, we proceed as follows. Using a variant of Motzkin's Transposition Theorem, we prove, in Section 5, the following:

LEMMA 3.5. *If Γ does not satisfy (MC) then Γ has a binary fractional polymorphism ω such that for each $\{a, b\} \subseteq D$, there exists $g \in \text{supp}(\omega)$ with $\{g(a, b), g(b, a)\} \neq \{a, b\}$.*

Therefore, if Γ_c does not satisfy (MC), we know that Γ_c has a fractional polymorphism ω with the properties given in the lemma. Furthermore, by Proposition 2.5(1), we may always assume that $\langle \Gamma_c \rangle$ contains a unary function u_a for each $a \in D$ such that $\text{argmin } u_a = \{a\}$. This means that ω is idempotent. To finish the proof, we will massage ω into a binary symmetric fractional polymorphism. We will use a technique introduced by Kolmogorov [37], based on a graph of generalised operations.

For a binary operation $g \in \mathcal{O}_D^{(2)}$, define \bar{g} by $\bar{g}(x, y) = g(y, x)$. For a mapping $\mathbf{g} \in \mathcal{O}_D^{(2 \rightarrow 2)}$, where $\mathbf{g} = (g, \bar{g})$, and a binary operation $h \in \mathcal{O}_D^{(2)}$, let $\mathbf{g}^h(x, y) = (h \circ (g, \bar{g}), h \circ (\bar{g}, g))$. Let $\mathbf{1}$ be the identity mapping in $\mathcal{O}_D^{(2 \rightarrow 2)}$. Let $\mathcal{V} = \{\mathbf{1}^{h_1 \dots h_k} \mid h_i \in \text{supp}(\omega), k \geq 0\}$. Note that all $\mathbf{g} \in \mathcal{V}$ are of the form $\mathbf{g} = (g, \bar{g})$.

Let $G = (V, E)$ be the directed graph with

- $V = \mathcal{V}$;
- $E = \{(\mathbf{g}, \mathbf{g}^h) \mid \mathbf{g} \in \mathcal{V}, h \in \text{supp}(\omega)\}$.

Let $\mathcal{R} \subseteq \mathcal{V}$ denote the set of all vertices \mathbf{g} with the property that for any other vertex $\mathbf{g}' \in \mathcal{V}$, if there is a path from \mathbf{g} to \mathbf{g}' , then there is a path from \mathbf{g}' to \mathbf{g} . The following result is proved in Section 7. A very similar result was proved in [37] using a different proof technique.

THEOREM 3.6. *There exists a $2 \rightarrow 2$ fractional polymorphism ρ of Γ_c with support contained in \mathcal{R} .*

Let f be an arbitrary function from Γ and let $n = \text{ar}(f)$. Let $\text{Range}_n(\mathbf{g}) = \{\mathbf{g}(\bar{x}^1, \bar{x}^2) \mid \bar{x}^1, \bar{x}^2 \in D^n\}$. The following lemma establishes the main result:

LEMMA 3.7 (KEY LEMMA). *For any $\mathbf{g} \in \text{supp}(\rho)$ and $(\bar{x}^1, \bar{x}^2) \in \text{Range}_n(\mathbf{g})$ and any $1 \leq i \leq n$, $f^2(\bar{x}^1, \bar{x}^2)$ is invariant under exchanging \bar{x}_i^1 and \bar{x}_i^2 .*

COROLLARY 3.8. *For any $\mathbf{g} \in \text{supp}(\rho)$ and $(\bar{x}^1, \bar{x}^2) \in \text{Range}_n(\mathbf{g})$ and any \mathbf{p} such that $\mathbf{p} \in \mathcal{O}_D^{(2 \rightarrow 2)}$ is acts as a permutation on all inputs, we have $f^2(\bar{x}^1, \bar{x}^2) = f^2(\mathbf{p}(\bar{x}^1, \bar{x}^2))$.*

COROLLARY 3.9. *Let $\mathbf{p} \in \mathcal{O}_D^{(2 \rightarrow 2)}$ be a mapping that orders its inputs according to some fixed total order on D . Then $\rho' = \sum_{\mathbf{g}} \rho(\mathbf{g}) \chi_{\mathbf{p} \circ \mathbf{g}}$ is a binary symmetric fractional polymorphism of Γ_c of arity $2 \rightarrow 2$.*

Since ρ' is a fractional polymorphism of Γ_c of arity $2 \rightarrow 2$, it follows that

$$\sum_{\mathbf{g}=(g_1, g_2) \in \mathcal{O}_D^{(2 \rightarrow 2)}} \rho'(\mathbf{g}) \frac{1}{2} (\chi_{g_1} + \chi_{g_2})$$

is a binary fractional polymorphism of Γ_c , and hence of Γ , this finishes the proof of Theorem 3.2.

We now sketch how to derive a proof of Lemma 3.7; the details are in Sections 6 and 7. Let ρ be a $2 \rightarrow 2$ fractional polymorphism of Γ_c as given by Theorem 3.6.

DEFINITION 4. *Let $w_a = \sum_{\mathbf{g}: \mathbf{g}(a, b) = (a, a)} \rho(\mathbf{g})$ and $w_b = \sum_{\mathbf{g}: \mathbf{g}(a, b) = (b, b)} \rho(\mathbf{g})$. We say that ρ is submodular on the pair $\{a, b\} \subseteq D$ if $w_a = w_b = \frac{1}{2}$.*

Let $S = (V(S), E(S))$ be the undirected graph with:

- $V(S) = D$;
- $E(S) = \{\{a, b\} \mid \rho \text{ is submodular on } \{a, b\}\}$.

To establish Lemma 3.7, we need the following two results, which are proved in Section 6 and 7, respectively.

LEMMA 3.10. *S is connected.*

LEMMA 3.11. *Assume that ρ is submodular on $\{a_1, a_2\}$ and that $(\bar{x}^1, \bar{x}^2) \in \text{Range}_{n-1}(\mathbf{g})$ for some $\mathbf{g} \in \text{supp}(\rho)$. Then $f^2((a_1, \bar{x}^1), (a_2, \bar{x}^2)) = f^2((a_2, \bar{x}^1), (a_1, \bar{x}^2))$.*

We are now ready to prove Lemma 3.7

PROOF (OF LEMMA 3.7). Without loss of generality, assume that $i = 1$ and let $((a, \bar{x}^1), (b, \bar{x}^2)) \in \text{Range}_n(\mathbf{g})$. We need to show that $f^2((a, \bar{x}^1), (b, \bar{x}^2)) = f^2((b, \bar{x}^1), (a, \bar{x}^2))$. Let $a = a_0, a_1, \dots, a_l = b$ be a path from a to b in the graph S . By Lemma 3.11, we have

$$f^2((a_i, \bar{x}^1), (a_{i+1}, \bar{x}^2)) = f^2((a_{i+1}, \bar{x}^1), (a_i, \bar{x}^2)), \quad (3)$$

for all $0 \leq i < l$.

Summing (3) over $0 \leq i < l$, we obtain

$$\sum_{0 \leq i < l} f^2((a_i, \bar{x}^1), (a_{i+1}, \bar{x}^2)) = \sum_{0 \leq i < l} f^2((a_{i+1}, \bar{x}^1), (a_i, \bar{x}^2)). \quad (4)$$

Finally, by cancelling terms in (4),

$$\frac{1}{2} f^2((a_0, \bar{x}^1)) + \frac{1}{2} f^2((a_l, \bar{x}^2)) = \frac{1}{2} f^2((a_l, \bar{x}^1)) + \frac{1}{2} f^2((a_0, \bar{x}^2)),$$

which establishes the result. \square

3.3 Meta problems

Let Γ be a valued language defined on D . In this section, we study two *meta problems* relevant to our classification. The first is to determine whether Γ is a core, and if not, to find a core of Γ . The second is to decide whether $\text{VCSP}(\Gamma)$ is tractable or NP-hard. We show that both of these problems are decidable.

To check whether Γ is a core, it suffices to verify that the following system of linear inequalities is unsatisfiable:

$$\begin{aligned} \sum_{g \in \mathcal{O}_D^{(1)}} \omega(g) f(g(\bar{x})) &\leq \|\omega\|_1 f(\bar{x}) \quad \forall f \in \Gamma, \bar{x} \in D^{\text{ar}(f)} \\ \sum_{g \in \Omega} \omega(g) &> 0 \\ \omega(g) &\geq 0 \quad \forall g \in \mathcal{O}_D^{(1)}, \end{aligned} \quad (5)$$

where Ω is the set of non-injective operations in $\mathcal{O}_D^{(1)}$. Note that for any such solution ω , $\omega/\|\omega\|_1$ is a unary fractional polymorphism with non-empty support on Ω . To find a core of Γ , we determine an inclusion-minimal subset $S \subseteq D$ such that $\Gamma[S]$ is a core.

To check whether $\text{VCSP}(\Gamma)$ is tractable, it suffices, by Theorem 3.3, to check whether it has a binary idempotent and symmetric fractional polymorphism. This is the case

if, and only if, the following system of linear inequalities is satisfiable

$$\sum_{g \in \Omega} \omega(g) f(g(\bar{x}, \bar{y})) \leq \|\omega(g)\|_1 f^2(\bar{x}, \bar{y}), \quad \forall f \in \Gamma, \bar{x}, \bar{y} \in D^{ar(f)}$$

$$\omega(g) \geq 0, \quad \forall g \in \Omega,$$

where Ω is the set of binary operations $g \in \mathcal{O}_D^{(2)}$ on D that are idempotent and symmetric.

4. EQUIVALENCE OF HARDNESS CONDITIONS AND EQUIVALENCE OF CORES

We begin by proving Lemma 2.4.

PROOF (OF LEMMA 2.4). Clearly $\text{Opt}_\Gamma(I) \leq \text{Opt}_{\Gamma'}(I')$. For the other direction, let ω be a unary fractional polymorphism of Γ with $h' \in \text{supp}(\omega)$ such that $\Gamma' = \Gamma[h'(D)]$. Let $V = V(I)$ be the set of variables of I and let $s : V \rightarrow D$ be an optimal solution to I . Then,

$$\begin{aligned} \sum_i w_i \cdot f_i(s(\bar{x}^i)) &\geq \sum_i w_i \sum_{h \in \mathcal{O}_D^{(1)}} \omega(h) f_i(h(s(\bar{x}^i))) \\ &= \sum_{h \in \mathcal{O}_D^{(1)}} \omega(h) \sum_i w_i \cdot f_i((h \circ s)(\bar{x}^i)). \end{aligned}$$

It follows that $h \circ s$ must also be an optimal solution to I , for each $h \in \text{supp}(\omega)$. So $h' \circ s$ is a solution to I of measure $\text{Opt}_\Gamma(I)$. But $h' \circ s$ is also a solution to I' , so $\text{Opt}_\Gamma(I) \geq \text{Opt}_{\Gamma'}(I')$. \square

In [27], a valued language Γ is defined to be a core if, for each $a \in D$, there is an instance I_a of $\text{VCSP}(\Gamma)$ such that a appears in every optimal solution to I_a . We now show that this condition is equivalent to Definition 3. We will use the following variation of Motzkin's Transposition Theorem [46].

LEMMA 4.1. *For any $A \in \mathbb{Q}^{m \times n}$, $B \in \mathbb{Q}^{p \times n}$, exactly one of the following holds:*

- $Ay > 0, By \geq 0$, for some $y \in \mathbb{Q}_{\geq 0}^n$; or
- $A^\top z_1 + B^\top z_2 \leq 0$, for some $0 \neq z_1 \in \mathbb{Q}_{\geq 0}^m, z_2 \in \mathbb{Q}_{\geq 0}^p$.

LEMMA 4.2. *For a valued language Γ , the following are equivalent:*

1. All unary fractional polymorphisms of Γ are injective.
2. There is an instance I of $\text{VCSP}(\Gamma)$ such that every $a \in D$ appears in every optimal solution to I .
3. For each $a \in D$, there is an instance I_a of $\text{VCSP}(\Gamma)$ such that a appears in every optimal solution to I_a .

PROOF. The system (5) is unsatisfiable if, and only if, $h \in \text{supp}(\omega)$ is injective for every unary fractional polymorphism ω of Γ . According to Lemma 4.1, this is true if, and only if, the following system is satisfiable:

$$\begin{aligned} \sum_{f \in \Gamma, \bar{x} \in D^{ar(f)}} z_2(f, \bar{x})(f(\bar{x}) - f(g(\bar{x}))) &\leq 0, \quad \forall g \in \mathcal{O}_D^{(1)}, \\ z_1 + \sum_{f \in \Gamma, \bar{x} \in D^{ar(f)}} z_2(f, \bar{x})(f(\bar{x}) - f(g(\bar{x}))) &\leq 0, \quad \forall g \in \Omega, \\ z_1 &> 0, \\ z_2(f, \bar{x}) &\geq 0, \quad \forall f \in \Gamma, \bar{x} \in D^{ar(f)}. \end{aligned}$$

That is, if, and only if, there exists an instance I of $\text{VCSP}(\Gamma)$, with variables D and objective function $\sum_{f, \bar{x}} z_2(f, \bar{x}) f(\bar{x})$, for which every non-injective g is non-optimal. This can be the case if, and only if, every $a \in D$ appears in every optimal solution to I . Hence, conditions (1) and (2) are equivalent. Clearly, condition (2) implies condition (3). Finally, given instances I_a , $a \in D$, such that a appears in every optimal solution, we can take the disjoint union of the variable sets of I_a as the variable set of I , and the sum of the objective function of I_a as the objective function of I . Then I satisfies condition (2). This establishes the lemma. \square

Finally we prove the equivalence of condition (MC) and condition (MC').

PROOF (OF LEMMA 2.2). We need to prove that (MC) implies (MC'). Let Γ be a valued language with a function $h \in \langle \Gamma \rangle$ such that $\text{argmin } h = \{(a, b), (b, a)\}$. Note that $u(x) = \min_y h(x, y)$ is a unary function with $\text{argmin } u = \{a, b\}$. If $h(a, a) = h(b, b)$, then u and h satisfy (MC'). Otherwise, assume without loss of generality that $h(a, b) = h(b, a) = 0, h(x, y) \geq 1$ for $\{x, y\} \neq \{a, b\}$, and that $h(a, a) < h(b, b)$. Let $C = \max_x h(a, a) - h(x, x)$, and define $u'(x) = \min_y C \cdot u(y) + h(y, y) + h(x, y)$. Note that $u(x) = 0$ for $x = a, b$ and $u(x) \geq 1$ otherwise. Also note that $\min_y h(y, y) = h(a, a) - C$. The three arguments in the following min-expressions correspond to the cases $y \neq a, b, y = a$, and $y = b$, respectively.

$$u'(x) \geq \min\{C + (h(a, a) - C) + 1, 0 + h(a, a) + 1, 0 + h(b, b) + 1\} > h(a, a) \quad (x \neq a, b)$$

$$u'(a) \geq \min\{C + (h(a, a) - C) + 1, 0 + h(a, a) + h(a, a), 0 + h(b, b) + 0\} > h(a, a)$$

$$u'(b) \leq C \cdot h(a, b) + h(a, a) + h(b, a) = h(a, a)$$

Thus $\text{argmin } u' = \{b\}$.

Now, let $\delta = h(b, b) - h(a, a) > 0$ and define

$$h'(x, y) = h(x, y) + \frac{\delta u'(x) + u'(y)}{2 u'(a) - u'(b)}.$$

It is straightforward to verify that $h'(a, b) = h'(b, a) < h'(a, a) = h'(b, b)$, so u and h' satisfy (MC'). \square

5. PROOF OF LEMMA 3.5

PROOF (OF LEMMA 3.5). Let $\pi_1(x, y) = x$ and $\pi_2(x, y) = y$ be the two binary projections on D . Let $\Omega(a, b)$ be the set of operations $g \in \mathcal{O}_D^{(2)}$ for which $\{g(a, b), g(b, a)\} \neq \{a, b\}$. Assume that there exist rational values $y(f, \bar{x}) \geq 0$, for $f \in \Gamma, \bar{x} \in (D \times D)^{ar(f)}$, such that for $i = 1, 2$,

$$\sum_{f, \bar{x}} y(f, \bar{x}) f(g(\bar{x})) \geq \sum_{f, \bar{x}} y(f, \bar{x}) f(\pi_i(\bar{x})), \quad \forall g \in \mathcal{O}_D^{(2)}, \quad (6)$$

$$\sum_{f, \bar{x}} y(f, \bar{x}) f(g(\bar{x})) > \sum_{f, \bar{x}} y(f, \bar{x}) f(\pi_i(\bar{x})), \quad \forall g \in \Omega(a, b). \quad (7)$$

Let x_1, \dots, x_n be an enumeration of $D \times D$ with $x_1 = (a, b)$ and $x_2 = (b, a)$. Let I be the instance of $\text{VCSP}(\Gamma)$ with variables x_1, \dots, x_n and objective function $f_I(x_1, \dots, x_n) = \sum_{f, \bar{x}} y(f, \bar{x}) f(\bar{x})$.

Define $h(x, y) = \min_{x_3, \dots, x_n} f_I(x, y, x_3, \dots, x_n)$. The equations (6) imply that π_1 and π_2 are among the optimal solutions to I , and equations (7) imply that π_1 and π_2 have

strictly smaller measure than any solution $g \in \Omega(a, b)$, so $h(a, b) = h(b, a) < h(x, y)$ for all $\{x, y\} \neq \{a, b\}$.

We conclude that if (MC) cannot be satisfied, then there is no solution to the system (6)+(7). By Lemma 4.1, there is a solution $z_1(g), z_2(g) \geq 0$ to the following system of equations:

$$\begin{aligned} \sum_{g \in \Omega(a, b)} z_1(g)(2f(g(\bar{x})) - f(\pi_1(\bar{x})) - f(\pi_2(\bar{x}))) + \\ \sum_{g \in \mathcal{O}_D^{(2)}} z_2(g)(2f(g(\bar{x})) - f(\pi_1(\bar{x})) - f(\pi_2(\bar{x}))) \leq 0, \\ \forall f \in \Gamma, \bar{x} \in (D \times D)^{ar(f)}, \end{aligned} \quad (8)$$

with $z_1(g) \neq 0$ for some $g \in \Omega(a, b)$. Now let $z = z_1 + z_2$ (let $z_1(g) = 0$ for $g \notin \Omega(a, b)$) and normalise so that $\|z\|_1 = 1$. A solution to (8) then implies a solution to the following system of inequalities:

$$\sum_{g \in \mathcal{O}_D^{(2)}} z(g)f(g(\bar{x})) \leq f^2(\pi_1(\bar{x}), \pi_2(\bar{x})), \quad \forall f, \bar{x},$$

with $z(g) \geq 0$ and $z(g) > 0$ for some $g \in \Omega(a, b)$. Denote this solution by $z_{a,b}(g)$. Now, if (MC) cannot be satisfied for any distinct $a, b \in D$, then we have solutions $z_{a,b}(g)$ for all $a \neq b \in D$. The lemma follows with $\omega(g) = (|D|^2 - |D|)^{-1} \sum_{a \neq b} z_{a,b}(g)$. \square

6. PROOF OF LEMMA 3.10

The aim of this section is to prove that the graph S of submodular pairs is connected. In order to do so, we introduce yet another graph T that records the ‘‘definable 2-subsets of D ’’. We then show that T is a subgraph of S and that T is connected.

Let $T = (V(T), E(T))$ be the undirected graph with:

- $V(T) = D$;
- $E(T) = \{\{a, b\} \mid \text{there exists a unary function } u \in \langle \Gamma_c \rangle \text{ such that } \text{argmin } u = \{a, b\}\}$.

LEMMA 6.1. $E(T) \subseteq E(S)$.

PROOF. Take an arbitrary edge $\{a, b\} \in E(T)$ and let u_a, u_b , and u_{ab} be unary functions in $\langle \Gamma_c \rangle$ such that $\text{argmin } u_a = \{a\}$, $\text{argmin } u_b = \{b\}$, and $\text{argmin } u_{ab} = \{a, b\}$, respectively. Since u_{ab} minimizes on $\{a, b\}$ and is invariant under both ω and ρ , we have $g(a, b), g(b, a) \in \{a, b\}$ for every $g \in \text{supp}(\omega)$ and every $\mathbf{g} = (g, \bar{g}) \in \text{supp}(\rho)$. By construction of ω , there is an operation $h \in \text{supp}(\omega)$ for which $(h(a, b), h(b, a)) \notin \{(a, b), (b, a)\}$, so by our previous observation, we must have either $h(a, b) = h(b, a) = a$ or $h(a, b) = h(b, a) = b$. Suppose that $\mathbf{g}(a, b) \in \{(a, b), (b, a)\}$ for some $\mathbf{g} \in \text{supp}(\rho)$. Then $\mathbf{g}^h(a, b) = (h(a, b), h(b, a))$ or $(h(b, a), h(a, b))$. In either case, \mathbf{g}^h is symmetric on $\{a, b\}$. So \mathbf{g}^h is reachable from \mathbf{g} in the graph G and every \mathbf{g}' reachable from \mathbf{g}^h is symmetric on $\{a, b\}$. Therefore \mathbf{g} cannot be in $\mathcal{R} \supseteq \text{supp}(\rho)$, a contradiction. We conclude that every $\mathbf{g} \in \text{supp}(\rho)$ is symmetric on $\{a, b\}$ and maps (a, b) to either (a, a) or (b, b) .

Let $w_a = \sum_{\mathbf{g}: \mathbf{g}(a, b) = (a, a)} \rho(\mathbf{g})$ and $w_b = \sum_{\mathbf{g}: \mathbf{g}(a, b) = (b, b)} \rho(\mathbf{g})$. By the previous argument, we have $w_a + w_b = 1$. By the fractional polymorphism inequality applied to u_a , we have

$$\frac{1}{2}(u_a(a) + u_a(b)) \geq w_a u_a(a) + w_b u_a(b). \quad (9)$$

Since $u_a(a) < u_a(b)$, we have $w_a \geq w_b$. But inequality (9) holds for u_b as well, hence $w_a \leq w_b$, and therefore $w_a = w_b = \frac{1}{2}$. \square

LEMMA 6.2. T is connected.

To prove this lemma, we will introduce some terminology from the study of hyperplane arrangements which will facilitate our reasoning about the edges of T . For a more thorough treatment of this subject, see Abramenko and Brown [1] and Stanley [47].

DEFINITION 5. Let $\{\bar{v}^i\}_{i \in I}$ be a finite set of vectors in \mathbb{R}^n . The set of hyperplanes $\mathcal{A} = \{H_i\}_{i \in I}$, where $H_i = \{\bar{x} \in \mathbb{R}^n \mid \bar{v}^i \cdot \bar{x} = 0\}$, is called a (linear) hyperplane arrangement.

To each vector $\bar{x} \in \mathbb{R}^n$, we associate a *sign vector*, $\sigma(\bar{x}) \in \{-1, 0, +1\}^I$, where the i th component is given by the sign of $\bar{v}^i \cdot \bar{x}$ for each $i \in I$. For a sign vector $\bar{v} \in \{-1, 0, +1\}^I$, a non-empty set $A = \sigma^{-1}(\bar{v}) = \{\bar{x} \in \mathbb{R}^n \mid \sigma(\bar{x}) = \bar{v}\}$ is called a *cell* of \mathcal{A} . We denote the defining sign vector, \bar{v} of A , by $\sigma(A)$.

A cell C with $\sigma(C)_i \neq 0$ for all $i \in I$ is called a *chamber*. The chambers are the connected full-dimensional regions of $\mathbb{R}^n \setminus \bigcup_{i \in I} H_i$. A cell P with $\sigma(P)_i = 0$ for exactly one $i \in I$ is called a *panel*. We say that P is a *panel of a chamber* C if the panel P is contained in the topological closure $\text{cl}(C)$ of C . Each panel is a panel of precisely two chambers.

The *chamber graph* of \mathcal{A} is the undirected graph with the chambers of \mathcal{A} as vertices and an edge between two chambers C_1 and C_2 if $\sigma(C_1)$ and $\sigma(C_2)$ differ by a single sign change, or equivalently, if C_1 and C_2 share a common panel. We will use the following properties of the chamber graph that can be found in [1, Proposition 1.54].

PROPOSITION 6.3. *The chamber graph of \mathcal{A} is connected and the minimal length of a path between C_1 and C_2 in the chamber graph is equal to the number of positions at which $\sigma(C_1)$ and $\sigma(C_2)$ differ.*

We are now ready to prove Lemma 6.2.

PROOF (OF LEMMA 6.2). For each $a \in D$, we have a unary function $u_a \in \langle \Gamma_c \rangle$ with $\text{argmin } u_a = \{a\}$. For $\bar{x} \in \mathbb{R}^D$, with components \bar{x}_c , consider the linear combination $f_{\bar{x}} = \sum_{c \in D} \bar{x}_c u_c$. Note that if \bar{x} is rational and non-negative, then $f_{\bar{x}} \in \langle \Gamma_c \rangle$. The inequality $f_{\bar{x}}(a) < f_{\bar{x}}(b)$ is equivalent to $\sum_{c \in D} \bar{x}_c (u_c(a) - u_c(b)) < 0$, i.e., $f_{\bar{x}}$ takes a strictly smaller value on a than on b precisely when the vector \bar{x} is on the negative side of the hyperplane H^{ab} defined by the normal \bar{v}^{ab} with components $v_c^{ab} = u_c(a) - u_c(b)$. Hence, by determining the sign of $\bar{x} \cdot \bar{v}^{ab}$, we can decide whether $f_{\bar{x}}(a) < f_{\bar{x}}(b)$ or $f_{\bar{x}}(a) > f_{\bar{x}}(b)$. If \bar{x} lies on the hyperplane, then $f_{\bar{x}}(a) = f_{\bar{x}}(b)$.

For each $a \in D$, let H^a be the hyperplane defined by the *unit vector* \bar{e}^a , i.e., $e_a^a = 1$ and $e_c^a = 0$ for $a \neq c$. Fix a strict total order $<_D$ on D . Let $\mathcal{A} = \{H^{ab} \mid a <_D b\} \cup \{H^a \mid a \in D\}$ be a hyperplane arrangement in \mathbb{R}^D . Let \mathcal{C} be the set of chambers C that has a positive sign for each H^a , i.e., each $C \in \mathcal{C}$ is contained in the positive (open) orthant of \mathbb{R}^D . Since all remaining components of $C \in \mathcal{C}$ are also non-zero, they determine a strict order on the values of the functions $f_{\bar{x}}$, $\bar{x} \in C$. For each $a \in D$, let $U_a = \{C \in \mathcal{C} \mid \forall \bar{x} \in C : \text{argmin } f_{\bar{x}} = \{a\}\}$. Each U_a is non-empty since the vector \bar{x}

given by $x_c = \epsilon$ for $c \neq a$ and $x_a = 1$ determines a function minimizing on a when $\epsilon > 0$ is chosen small enough.

Fix $a, b \in D$ and pick any $C_a \in U_a, C_b \in U_b$. Let $C_a = C_0, C_1, \dots, C_l = C_b$ be a minimal-length path from C_a to C_b in the chamber graph. Consider the sign vectors along this path: $\sigma(C_0), \sigma(C_1), \dots, \sigma(C_l)$. By Proposition 6.3 the sign of a fixed component changes at most once along this sequence. In particular, since C_a and C_b both have positive signs for the hyperplanes H^a , it follows that C_i is contained in the positive orthant for every i . Hence, for each i , there is a $a_i \in D$ such that $C_i \in U_{a_i}$. For each i with $a_i \neq a_{i+1}$, the path moves from a chamber where $f_{\bar{x}}$ minimizes on a_i to a chamber where it minimizes on a_{i+1} . This means that C_i and C_{i+1} share a panel P_i with a sign vector $\sigma(P_i)$ obtained from either $\sigma(C_i)$ or $\sigma(C_{i+1})$ by setting the component corresponding to $H^{a_i a_{i+1}}$ to 0 (assuming $a_i <_D a_{i+1}$). Since all other components of $\sigma(P_i)$ have the same sign as in $\sigma(C_i)$ and $\sigma(C_{i+1})$, we have $f_{\bar{x}}(a_i) = f_{\bar{x}}(a_{i+1}) < f_{\bar{x}}(c)$, for every $\bar{x} \in P_i$ and $c \neq a_i, a_{i+1}$. For a hyperplane arrangement, such as \mathcal{A} , that is defined in terms of rational normal vectors, each cell is defined as the solutions to a set of linear equalities and inequalities with rational coefficients. Every cell therefore contains at least one rational vector. In particular, there exists a non-negative rational vector $\bar{x} \in P_i$ with $\operatorname{argmin} f_{\bar{x}} = \{a_i, a_{i+1}\}$, so $\{a_i, a_{i+1}\} \in E(T)$. This holds for all $0 \leq i < l$ with $a_i \neq a_{i+1}$, so we conclude that a subsequence of $a = a_0, a_1, \dots, a_l = b$ is a path in T from a to b . \square

7. PROOFS OF THEOREM 3.6 AND LEMMA 3.11

A (time-homogeneous) finite-state Markov chain M is given by a set of states and conditional probabilities $p(i, j)$ for M to be in state j at time $t + 1$ given that it was in state i at time t . Let $p^{(k)}(i, j)$ denote the probability that M proceeds from state i to state j in exactly k transitions. M is called *irreducible* if, for every pair of states (i, j) , there exists $r \geq 1$ with $p^{(r)}(i, j) > 0$. A state i is called *transient* if, for some state j , there is a path from i to j but not from j to i . A state that is not transient is called *recurrent*. A state has *periodicity* r if $r = \gcd\{k \mid p^{(k)}(i, i) > 0\}$. M is called *aperiodic* if all states have periodicity 1. A *stationary distribution* of M is a probability distribution λ on the set of states of M such that $\lambda(i) = \sum_j \lambda(j)p(j, i)$ for all states i . The following is well known.

THEOREM 7.1. *For any finite-state Markov chain M :*

1. *If M is irreducible, then there is a unique stationary distribution λ of M with $\lambda(i) > 0$ for all states i .*
2. *If M is aperiodic, then for any initial distribution π , there is a stationary distribution λ of M with $\sum_j \pi(j)p^{(k)}(j, i) \rightarrow \lambda(i)$ as $k \rightarrow \infty$, for all states i .*
3. *If i is transient, then $p^{(k)}(j, i) \rightarrow 0$ as $k \rightarrow \infty$, for all states j .*

We now define a Markov chain M on G . Let $w(\mathbf{g}, \mathbf{g}') = \sum_{h \in \operatorname{supp}(\omega)} \omega(h)$. The transition probabilities are given by

$$p(\mathbf{g}, \mathbf{g}') = \begin{cases} \frac{1}{2}w(\mathbf{g}, \mathbf{g}') + \frac{1}{2} & \text{if } \mathbf{g} = \mathbf{g}', \text{ and} \\ \frac{1}{2}w(\mathbf{g}, \mathbf{g}') & \text{otherwise.} \end{cases}$$

Note that the set $\mathcal{R} \subseteq \mathcal{V}$, defined in Section 3.2, is precisely the set of recurrent states of M . Let \mathcal{H} be a strongly connected component of $G[\mathcal{R}]$. Then, $M(\mathcal{H})$, the restriction of M to \mathcal{H} , is also a Markov chain.

LEMMA 7.2. *The Markov chains M and $M(\mathcal{H})$ are aperiodic and each $M(\mathcal{H})$ is irreducible.*

PROOF. Aperiodicity follows by construction as $p(\mathbf{g}, \mathbf{g}) \geq \frac{1}{2} > 0$ for all $\mathbf{g} \in \mathcal{V}$. Irreducibility follows since each \mathcal{H} is a strongly connected component of $G[\mathcal{R}]$. \square

LEMMA 7.3. *Let σ and λ be probability distributions on \mathcal{V} and assume that M converges to λ when starting in σ . Then, for every $\bar{x}^1, \bar{x}^2 \in D^n$,*

$$\sum_{\mathbf{g} \in \mathcal{V}} \sigma(\mathbf{g}) f^2(\mathbf{g}(\bar{x}^1, \bar{x}^2)) \geq \sum_{\mathbf{g} \in \mathcal{V}} \lambda(\mathbf{g}) f^2(\mathbf{g}(\bar{x}^1, \bar{x}^2)).$$

PROOF. By k times applying the $2 \rightarrow 2$ fractional polymorphism $\frac{1}{2}\chi_1 + \frac{1}{2}\sum_h \omega(h)\chi_{(h, \bar{h})}$, (where $\mathbf{1}$ is the identity in $\mathcal{O}_D^{(2 \rightarrow 2)}$) to the left-hand side, we have

$$\begin{aligned} \sum_{\mathbf{g} \in \mathcal{V}} \sigma(\mathbf{g}) f^2(\mathbf{g}(\bar{x}^1, \bar{x}^2)) &\geq \sum_{\mathbf{g} \in \mathcal{V}} \sigma(\mathbf{g}) \frac{1}{2} \left(f^2(\mathbf{g}(\bar{x}^1, \bar{x}^2)) + \right. \\ &\quad \left. \sum_{h \in \mathcal{O}_D^{(2)}} \omega(h) f^2(\mathbf{g}^h(\bar{x}^1, \bar{x}^2)) \right) \\ &= \sum_{\mathbf{g} \in \mathcal{V}} \sum_{\mathbf{g}' \in \mathcal{V}} \sigma(\mathbf{g}') p(\mathbf{g}', \mathbf{g}) f^2(\mathbf{g}(\bar{x}^1, \bar{x}^2)) \\ &\geq \dots \geq \sum_{\mathbf{g} \in \mathcal{V}} \sigma^{(k)}(\mathbf{g}) f^2(\mathbf{g}(\bar{x}^1, \bar{x}^2)), \end{aligned}$$

where $\sigma^{(k)}(\mathbf{g}) = \sum_{\mathbf{g}' \in \mathcal{V}} \sigma(\mathbf{g}') p^{(k)}(\mathbf{g}', \mathbf{g})$. By assumption, $\sigma^{(k)}(\mathbf{g}) \rightarrow \lambda(\mathbf{g})$ as $k \rightarrow \infty$. Since the right-hand side is a linear function in $\sigma^{(k)}(\mathbf{g})$, the lemma follows by continuity. \square

We are now ready to prove Theorem 3.6.

PROOF (OF THEOREM 3.6). Let σ be the $2 \rightarrow 2$ fractional polymorphism of Γ_c defined by $\sigma = \sum_h \omega(h)\chi_{(h, \bar{h})}$. Note that σ is a probability distribution on \mathcal{V} . By Theorem 7.1(2), there exists a stationary distribution ρ of M such that

$$\sum_{\mathbf{g}'} \sigma(\mathbf{g}') p^{(k)}(\mathbf{g}', \mathbf{g}) \rightarrow \rho(\mathbf{g})$$

as $k \rightarrow \infty$, for all \mathbf{g} . By Theorem 7.1(3), $\rho(\mathbf{g}) = 0$ for $\mathbf{g} \notin \mathcal{R}$, hence $\operatorname{supp}(\rho) \subseteq \mathcal{R}$. Finally, ρ is a $2 \rightarrow 2$ fractional polymorphism of Γ_c by Lemma 7.3. \square

In order to prove Lemma 3.11, we will need an additional observation.

LEMMA 7.4. *Let $x_1, \dots, x_m, c_1, \dots, c_m \in \mathbb{Q}_{\geq 0}$ be such that $c_i > 0$ for all i , $\sum_i c_i = 1$, and $x_j \geq \sum_i c_i x_i$ for all j . Then $x_1 = x_2 = \dots = x_m$.*

PROOF. Let $C = \sum_i c_i x_i$. We have $x_j \geq C$ for all j . If $x_j > C$ for some j , then $c_j x_j > c_j C$, so $C = \sum_i c_i x_i > \sum_i c_i C = C$, a contradiction. So, for all j , $x_j \leq C$, and hence $x_j = C$. \square

PROOF (OF LEMMA 3.11). Let \mathcal{H} be the strongly connected component of G containing \mathbf{g} . By Theorem 7.1 and Lemma 7.2, $M(\mathcal{H})$ has a unique stationary distribution λ ,

with $\lambda(\mathbf{h}) > 0$ for all $\mathbf{h} \in \mathcal{H}$. Let (\bar{y}^1, \bar{y}^2) be such that $(\bar{x}^1, \bar{x}^2) = \mathbf{g}(\bar{y}^1, \bar{y}^2)$, and for any $\hat{\mathbf{g}}$, let $(\bar{y}^{\hat{\mathbf{g}}1}, \bar{y}^{\hat{\mathbf{g}}2}) = \hat{\mathbf{g}}(\bar{y}^1, \bar{y}^2)$. For $i = 1, 2$, let $\Omega_i = \{\mathbf{g} \in \mathcal{V} \mid \mathbf{g}(a_1, a_2) = (a_i, a_i)\}$. For each $\hat{\mathbf{g}} \in \mathcal{H}$, we have

$$f^2((a_1, \bar{y}^{\hat{\mathbf{g}}1}), (a_2, \bar{y}^{\hat{\mathbf{g}}2})) \quad (10)$$

$$\geq \sum_{\mathbf{g}' \in \mathcal{V}} \rho(\mathbf{g}') f^2(\mathbf{g}'((a_1, \bar{y}^{\hat{\mathbf{g}}1}), (a_2, \bar{y}^{\hat{\mathbf{g}}2}))) \quad (11)$$

$$= \sum_{\mathbf{g}' \in \Omega_1} \rho(\mathbf{g}') f^2(\mathbf{g}' \circ \hat{\mathbf{g}}((a_1, \bar{y}^1), (a_1, \bar{y}^2))) \\ + \sum_{\mathbf{g}' \in \Omega_2} \rho(\mathbf{g}') f^2(\mathbf{g}' \circ \hat{\mathbf{g}}((a_2, \bar{y}^1), (a_2, \bar{y}^2))) \quad (12)$$

$$= \frac{1}{2} \sum_{\mathbf{h} \in \mathcal{H}} \sigma_1(\mathbf{h}) f^2(\mathbf{h}((a_1, \bar{y}^1), (a_1, \bar{y}^2))) \\ + \frac{1}{2} \sum_{\mathbf{h} \in \mathcal{H}} \sigma_2(\mathbf{h}) f^2(\mathbf{h}((a_2, \bar{y}^1), (a_2, \bar{y}^2))) \quad (13)$$

$$\geq \frac{1}{2} \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) f^2(\mathbf{h}((a_1, \bar{y}^1), (a_1, \bar{y}^2))) \\ + \frac{1}{2} \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) f^2(\mathbf{h}((a_2, \bar{y}^1), (a_2, \bar{y}^2))) \quad (14)$$

$$= \frac{1}{2} \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) (f^2((a_1, \bar{y}^{\mathbf{h}1}), (a_2, \bar{y}^{\mathbf{h}2})) \\ + f^2((a_1, \bar{y}^{\mathbf{h}2}), (a_2, \bar{y}^{\mathbf{h}1}))), \quad (15)$$

where (11) follows by applying ρ and (12) follows from ρ being idempotent and submodular on $\{a_1, a_2\}$. To get (13), define the probability distributions σ_i on \mathcal{V} by $\sigma_i(\mathbf{h}) = 2 \cdot \sum_{\mathbf{g}' \in \Omega_i: \mathbf{g}' \circ \hat{\mathbf{g}} = \mathbf{h}} \rho(\mathbf{g}')$ and remember that $\mathbf{g}' = \mathbf{1}^{h_1 \dots h_k}$, for some $h_i \in \text{supp}(\omega)$, so $\mathbf{g}' \circ \hat{\mathbf{g}} = \hat{\mathbf{g}}^{h_1 \dots h_k} \in \mathcal{H}$, i.e., $\text{supp}(\sigma_i) \subseteq \mathcal{H}$. From (13), one obtains (14) by two applications of Lemma 7.3 and (15) by rearranging the terms.

Also, $f^2((a_1, \bar{y}^{\hat{\mathbf{g}}2}), (a_2, \bar{y}^{\hat{\mathbf{g}}1})) = f^2((a_2, \bar{y}^{\hat{\mathbf{g}}1}), (a_1, \bar{y}^{\hat{\mathbf{g}}2})) \geq \sum_{\mathbf{g}'} \rho(\mathbf{g}') f^2(\mathbf{g}'((a_2, \bar{y}^{\hat{\mathbf{g}}1}), (a_1, \bar{y}^{\hat{\mathbf{g}}2})))$ which is equal to (11) since ρ is submodular on $\{a_1, a_2\}$. Hence,

$$f^2((a_1, \bar{y}^{\hat{\mathbf{g}}2}), (a_2, \bar{y}^{\hat{\mathbf{g}}1})) \geq \frac{1}{2} \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) (f^2((a_1, \bar{y}^{\mathbf{h}1}), (a_2, \bar{y}^{\mathbf{h}2})) \\ + f^2((a_1, \bar{y}^{\mathbf{h}2}), (a_2, \bar{y}^{\mathbf{h}1}))), \quad (16)$$

in the same way as before.

For $\mathbf{h} = (h_1, h_2) \in \mathcal{V}$, define $\bar{\mathbf{h}} = (h_2, h_1)$, and $\bar{\mathcal{H}} = \{\bar{\mathbf{h}} \mid \mathbf{h} \in \mathcal{H}\}$. For $\hat{\mathbf{g}} \in \mathcal{H} \cup \bar{\mathcal{H}}$, let $x_{\hat{\mathbf{g}}} = f^2((a_1, \bar{y}^{\hat{\mathbf{g}}1}), (a_2, \bar{y}^{\hat{\mathbf{g}}2}))$ and $c_{\hat{\mathbf{g}}} = \frac{1}{2}(\lambda(\hat{\mathbf{g}}) + \lambda(\bar{\hat{\mathbf{g}}}))$. It then follows from (10–15), (16), and Lemma 7.4 that $x_{\hat{\mathbf{g}}}$ is constant on $\mathcal{H} \cup \bar{\mathcal{H}}$. Since $\mathbf{g}, \bar{\mathbf{g}} \in \mathcal{H} \cup \bar{\mathcal{H}}$, we have $f^2((a_1, \bar{x}^1), (a_2, \bar{x}^2)) = x_{\mathbf{g}} = x_{\bar{\mathbf{g}}} = f^2((a_1, \bar{x}^2), (a_2, \bar{x}^1))$, as required. \square

8. CONCLUSIONS

In this work we have completely answered the question of which (finite-)valued constraint languages on finite domains are solvable exactly in polynomial time. In particular, we have characterised the tractable languages as those that admit a binary idempotent and symmetric fractional polymorphism, which we have shown to be a polynomial-time checkable condition. Moreover, all tractable languages are solvable by the basic linear programming relaxation. Thus, we have demonstrated that basic linear programming (BLP) relaxation suffices for exact solvability of (finite-)valued con-

straint languages and that, in this context, semidefinite programming relaxations do not add any power.

An intriguing open question is the precise boundary of the tractability of the minimisation problem in the value oracle model; that is, for objective functions that are not given explicitly as a sum of fixed-arity functions but only by an oracle. In particular, do the tractable cases (solvable by combinatorial algorithms) coincide or not?

9. ACKNOWLEDGMENTS

The authors wish to thank Hubie Chen for useful comments on an earlier draft of this paper.

10. REFERENCES

- [1] P. Abramenko and K. S. Brown. *Buildings – Theory and Applications*, volume 248 of *Graduate Texts in Mathematics*. Springer, 2008.
- [2] L. Barto. The dichotomy for conservative constraint satisfaction problems revisited. In *LICS'11*, pages 301–310. IEEE Computer Society, 2011.
- [3] L. Barto and M. Kozik. Constraint Satisfaction Problems of Bounded Width. In *FOCS'09*, pages 461–471. IEEE Computer Society, 2009.
- [4] L. Barto and M. Kozik. Robust Satisfiability of Constraint Satisfaction Problems. In *STOC'12*, pages 931–940. ACM, 2012.
- [5] 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 Journal on Computing*, 38(5):1782–1802, 2009.
- [6] A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *Journal of the ACM*, 53(1):66–120, 2006.
- [7] A. Bulatov, A. Krokhin, and P. Jeavons. Classifying the Complexity of Constraints using Finite Algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.
- [8] A. A. Bulatov. Complexity of conservative constraint satisfaction problems. *ACM Transactions on Computational Logic*, 12(4), 2011. Article 24.
- [9] D. Cohen, M. Cooper, P. Jeavons, and A. Krokhin. Supermodular Functions and the Complexity of MAX-CSP. *Discrete Applied Mathematics*, 149(1-3):53–72, 2005.
- [10] D. A. Cohen, M. C. Cooper, P. Creed, P. Jeavons, and S. Živný. An algebraic theory of complexity for discrete optimisation. Technical report, July 2012. arXiv:1207.6692.
- [11] D. A. Cohen, M. C. Cooper, and P. G. Jeavons. An Algebraic Characterisation of Complexity for Valued Constraints. In *CP'06*, volume 4204 of *LNCS*, pages 107–121. Springer, 2006.
- [12] D. A. Cohen, M. C. Cooper, P. G. Jeavons, and A. A. Krokhin. The Complexity of Soft Constraint Satisfaction. *Artificial Intelligence*, 170(11):983–1016, 2006.
- [13] D. A. Cohen, P. Creed, P. G. Jeavons, and S. Živný. An algebraic theory of complexity for valued constraints: Establishing a Galois connection. In *MFCS'11*, volume 6907 of *LNCS*, pages 231–242. Springer, 2011.

- [14] M. C. Cooper and S. Živný. Hybrid tractability of valued constraint problems. *Artificial Intelligence*, 175(9-10):1555–1569, 2011.
- [15] M. C. Cooper and S. Živný. Tractable triangles and cross-free convexity in discrete optimisation. *Journal of Artificial Intelligence Research*, 44:455–490, 2012.
- [16] Y. Crama and P. L. Hammer. *Boolean Functions - Theory, Algorithms, and Applications*. Cambridge University Press, 2011.
- [17] P. Creed and S. Živný. On minimal weighted clones. In *CP'11*, volume 6876 of *LNCS*, pages 210–224. Springer, 2011.
- [18] N. Creignou. A dichotomy theorem for maximum generalized satisfiability problems. *Journal of Computer and System Sciences*, 51(3):511–522, 1995.
- [19] N. Creignou, S. Khanna, and M. Sudan. *Complexity Classification of Boolean Constraint Satisfaction Problems*, volume 7 of *SIAM Monographs on Discrete Mathematics and Applications*. SIAM, 2001.
- [20] V. Dalmau and A. Krokhin. Robust satisfiability for CSPs: algorithmic and hardness results. 2011. In preparation.
- [21] V. Deineko, P. Jonsson, M. Klasson, and A. Krokhin. The approximability of Max CSP with fixed-value constraints. *Journal of the ACM*, 55(4), 2008. Article 16.
- [22] T. Feder and M. Y. Vardi. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM Journal on Computing*, 28(1):57–104, 1998.
- [23] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, 1979.
- [24] G. Gottlob, G. Greco, and F. Scarcello. Tractable Optimization Problems through Hypergraph-Based Structural Restrictions. In *ICALP'09*, volume 5556 of *LNCS*, pages 16–30. Springer, 2009.
- [25] M. Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1):1–24, 2007.
- [26] H. Hirai. Discrete Convexity and Polynomial Solvability in Minimum 0-Extension Problems. In *SODA'13*. SIAM, 2013. To appear.
- [27] A. Huber, A. Krokhin, and R. Powell. Skew Bisubmodularity and Valued CSPs. In *SODA'13*. SIAM, 2013. To appear.
- [28] P. M. Idziak, P. Markovic, R. McKenzie, M. Valeriote, and R. Willard. Tractability and learnability arising from algebras with few subpowers. *SIAM Journal on Computing*, 39(7):3023–3037, 2010.
- [29] P. G. Jeavons. On the Algebraic Structure of Combinatorial Problems. *Theoretical Computer Science*, 200(1-2):185–204, 1998.
- [30] P. G. Jeavons, D. A. Cohen, and M. Gyssens. Closure Properties of Constraints. *Journal of the ACM*, 44(4):527–548, 1997.
- [31] P. Jonsson, M. Klasson, and A. Krokhin. The Approximability of Three-valued MAX CSP. *SIAM Journal on Computing*, 35(6):1329–1349, 2006.
- [32] P. Jonsson, A. A. Krokhin, and F. Kuivinen. Hard constraint satisfaction problems have hard gaps at location 1. *Theoretical Computer Science*, 410(38-40):3856–3874, 2009.
- [33] P. Jonsson, F. Kuivinen, and G. Nordh. MAX ONES Generalized to Larger Domains. *SIAM Journal on Computing*, 38(1):329–365, 2008.
- [34] P. Jonsson, F. Kuivinen, and J. Thapper. Min CSP on Four Elements: Moving Beyond Submodularity. In *CP'11*, volume 6876 of *LNCS*, pages 438–453. Springer, 2011.
- [35] S. Khanna, M. Sudan, L. Trevisan, and D. Williamson. The approximability of constraint satisfaction problems. *SIAM Journal on Computing*, 30(6):1863–1920, 2001.
- [36] S. Khot. On the unique games conjecture (invited survey). In *CCC'10*, pages 99–121. IEEE Computer Society, 2010.
- [37] V. Kolmogorov. The power of linear programming for valued CSPs: a constructive characterization. Technical report, July 2012. arXiv:1207.7213.
- [38] V. Kolmogorov and S. Živný. The complexity of conservative valued CSPs. In *SODA'12*, pages 750–759. SIAM, 2012. Full version available on arXiv:1110.2809.
- [39] G. Kun, R. O'Donnell, S. Tamaki, Y. Yoshida, and Y. Zhou. Linear programming, width-1 CSPs, and robust satisfaction. In *ITCS'12*, pages 484–495. ACM, 2012.
- [40] S. L. Lauritzen. *Graphical Models*. Oxford University Press, 1996.
- [41] D. Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. In *STOC'12*, pages 735–744, 2010.
- [42] U. Montanari. Networks of Constraints: Fundamental properties and applications to picture processing. *Information Sciences*, 7:95–132, 1974.
- [43] P. Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *STOC'08*, pages 245–254. ACM, 2008.
- [44] P. Raghavendra. *Approximating NP-hard problems: Efficient algorithms and their limits*. PhD thesis, University of Washington, 2009.
- [45] T. J. Schaefer. The Complexity of Satisfiability Problems. In *STOC'78*, pages 216–226. ACM, 1978.
- [46] A. Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, Inc., 1986.
- [47] R. P. Stanley. An introduction to hyperplane arrangements. In E. Miller, V. Reiner, and B. Sturmfels, editors, *Geometric Combinatorics*, volume 13 of *IAS/Park City Mathematics Series*, pages 389–496. AMS, Providence, RI, 2007.
- [48] R. Takhanov. A Dichotomy Theorem for the General Minimum Cost Homomorphism Problem. In *STACS'10*, pages 657–668, 2010.
- [49] J. Thapper and S. Živný. The power of linear programming for valued CSPs. In *FOCS'12*. IEEE, 2012. Available on arXiv:1204.1079.
- [50] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.

APPENDIX

A. BASIC LP RELAXATION (BLP)

In this section we define the *basic LP relaxation* (BLP) (also known as the *standard*, or *canonical LP relaxation*) [49].

Let Γ be a valued language defined on D . Let I be a VCSP(Γ) instance given by the set $V = \{x_1, \dots, x_n\}$ of variables and the objective function

$$f_I(x_1, \dots, x_n) = \sum_{i=1}^q w_i \cdot f_i(\bar{x}^i)$$

where, for every $1 \leq i \leq q$, $f_i : D^{ar(f_i)} \rightarrow \mathbb{Q}_{\geq 0}$, $\bar{x}^i \in V^{ar(f_i)}$, and $w_i \in \mathbb{Q}_{\geq 0}$.

For a tuple \bar{x} , we denote by $\{\bar{x}\}$ the set of elements in \bar{x} .

The basic LP relaxation has variables λ_{i,σ_i} for every $1 \leq i \leq q$ and $\sigma_i : \{\bar{x}^i\} \rightarrow D$; and variables $\mu_{x,a}$ for every $x \in V$ and $a \in D$.

$$\begin{aligned} \min \quad & \sum_{i=1}^q \sum_{\sigma_i: \{\bar{x}^i\} \rightarrow D} w_i \cdot f_i(\sigma_i(\bar{x}^i)) \cdot \lambda_{i,\sigma_i} \\ \text{s.t.} \quad & \sum_{\sigma_i: \{\bar{x}^i\} \rightarrow D} \lambda_{i,\sigma_i} = \mu_{x,a} \quad \forall 1 \leq i \leq q, \forall x \in \bar{x}^i, \forall a \in D \\ & \sum_{a \in D} \mu_{x,a} = 1 \quad \forall x \in V \\ & 0 \leq \lambda, \mu \leq 1 \end{aligned}$$

Since Γ is fixed, this relaxation has polynomial size in I . Requiring λ_{i,σ_i} and $\mu_{x,a}$ to be in $\{0, 1\}$ provides an integer programming formulation of I with the meaning $\mu_{x,a} = 1$ if, and only if, variable x is assigned value a , and $\lambda_{i,\sigma_i} = 1$ if, and only if, \bar{x}^i is assigned $\sigma_i(\bar{x}^i)$.

For any VCSP instance, the basic LP relaxation gives a lower bound on the optimal solution. By Theorem 3.3, if Γ is a core valued language admitting a binary idempotent symmetric fractional polymorphism, then the optimal solution to the relaxation is integral and thus gives an optimal solution to the original VCSP instance.

B. FROM BINARY TO ALL ARITIES

THEOREM B.1 ([37]). *Suppose Γ is a finite-valued language with a symmetric fractional polymorphism of arity 2. Then Γ has symmetric fractional polymorphisms of all arities.*

In this section, we give an alternative proof of this theorem using a variation of the Markov chain argument from Section 7. The proof idea is identical to that of [37], but the technical content of Theorem 11 in [37] is naturally explained by our approach.

It suffices to prove that if Γ has a symmetric fractional polymorphism of arity $m-1 > 2$, then it has one of arity m . Let ω be an $(m-1)$ -ary symmetric fractional polymorphism of Γ . For $\mathbf{g} = (g_1, \dots, g_m) \in \mathcal{O}_D^{(m \rightarrow m)}$, let \mathbf{g}_{-i} denote the operation $(g_1, \dots, g_{i-1}, g_{i+1}, \dots, g_m) \in \mathcal{O}_D^{(m \rightarrow m-1)}$ obtained by omitting the i th operation from \mathbf{g} . For $h \in \text{supp}(\omega)$, let

$$\mathbf{g}^h = (h \circ \mathbf{g}_{-1}, \dots, h \circ \mathbf{g}_{-m}).$$

For $\bar{x}^1, \dots, \bar{x}^m$, let $(\bar{x}^{\mathbf{g}^1}, \dots, \bar{x}^{\mathbf{g}^m}) = \mathbf{g}(\bar{x}^1, \dots, \bar{x}^m)$. Let $\mathbf{1}$ be the identity mapping in $\mathcal{O}_D^{(m \rightarrow m)}$ and let $\mathcal{V}' = \{\mathbf{1}^{h_1 \dots h_k} \mid h_i \in \text{supp}(\omega), k \geq 0\}$.

Let $G' = (V, E)$ be the directed graph with

- $V = \mathcal{V}'$;
- $E = \{(\mathbf{g}, \mathbf{g}') \mid \mathbf{g} \in \mathcal{V}', h \in \text{supp}(\omega)\}$.

We now define a Markov chain M' on G' . Let $w(\mathbf{g}, \mathbf{g}') = \sum_{h \in \text{supp}(\omega): \mathbf{g}' = \mathbf{g}^h} \omega(h)$. The transition probabilities $p(\mathbf{g}, \mathbf{g}')$ are given by $\frac{1}{2}w(\mathbf{g}, \mathbf{g}')$ if $\mathbf{g} \neq \mathbf{g}'$ and $\frac{1}{2}w(\mathbf{g}, \mathbf{g}') + \frac{1}{2}$ otherwise. Let \mathcal{R}' denote the set of recurrent states of M' .

The properties of M' and $M'(\mathcal{H})$ for a strongly connected component \mathcal{H} in $G'[\mathcal{R}']$ are summarised in the following two lemmas. Their proofs are analogous to those of Lemma 7.2 and Lemma 7.3 and are therefore omitted.

LEMMA B.2. *The Markov chains M' and $M'(\mathcal{H})$ are aperiodic and each $M'(\mathcal{H})$ is irreducible.*

LEMMA B.3. *Let σ and λ be probability distributions on \mathcal{V}' and assume that M' converges to λ when starting in σ . Then, for every $\bar{x}^1, \dots, \bar{x}^m \in D^n$,*

$$\sum_{\mathbf{g} \in \mathcal{V}'} \sigma(\mathbf{g}) f^m(\mathbf{g}(\bar{x}^1, \dots, \bar{x}^m)) \geq \sum_{\mathbf{g} \in \mathcal{V}'} \lambda(\mathbf{g}) f^m(\mathbf{g}(\bar{x}^1, \dots, \bar{x}^m)).$$

PROOF (OF THEOREM B.1). As in Theorem 3.6, we can obtain a $(m \rightarrow m)$ -ary fractional polymorphism ρ of Γ with support on \mathcal{R}' . To finish the proof, we need an analog of Lemma 3.7. That is, we need to show that for any n -ary function $f \in \Gamma$, any $\mathbf{g} \in \mathcal{R}'$, and any $(\bar{x}^1, \dots, \bar{x}^m) \in \text{Range}_n(\mathbf{g})$, the expression $f^m(\bar{x}^1, \dots, \bar{x}^m)$ is invariant under exchanging $x_i^{j_1}$ and $x_i^{j_2}$ for all choices of $1 \leq j_1, j_2 \leq m$ and $1 \leq i \leq n$.

Let \mathcal{H} be the strongly connected component in $G'[\mathcal{R}']$ containing \mathbf{g} . By Lemma B.2 and Theorem 7.1, there is a unique stationary distribution, λ , of $M'(\mathcal{H})$, and $\text{supp}(\lambda) = \mathcal{H}$.

We make two claims.

- (a) $f^m(\hat{\mathbf{g}}(\bar{y}^1, \dots, \bar{y}^m))$ does not depend on the choice of $\hat{\mathbf{g}} \in \mathcal{H}$;
- (b) $\sum_{\mathbf{h} \in \mathcal{H}} f(\bar{x}^{\mathbf{h}^k})$ does not depend on the choice of k .

For every $\hat{\mathbf{g}} \in \mathcal{H}$, and all $\bar{y}^1, \dots, \bar{y}^m \in D^n$, by Lemma B.3, we have

$$f^m(\hat{\mathbf{g}}(\bar{y}^1, \dots, \bar{y}^m)) \geq \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) f^m(\mathbf{h}(\bar{y}^1, \dots, \bar{y}^m)).$$

Applying Lemma 7.4, we obtain claim (a).

Furthermore,

$$\begin{aligned} & \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) f^{m-1}(\mathbf{h}_{-k}(\bar{x}^1, \dots, \bar{x}^m)) \\ & \geq \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) \frac{1}{2} \left(f^{m-1}(\mathbf{h}_{-k}(\bar{x}^1, \dots, \bar{x}^m)) + \right. \\ & \quad \left. \sum_{h \in \text{supp}(\omega)} \omega(h) f(h \circ \mathbf{h}_{-k}(\bar{x}^1, \dots, \bar{x}^m)) \right) \\ & = \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) \sum_{\mathbf{h}' \in \mathcal{H}} p(\mathbf{h}, \mathbf{h}') f(\bar{x}^{\mathbf{h}'^k}) \\ & = \sum_{\mathbf{h}' \in \mathcal{H}} \lambda(\mathbf{h}') f(\bar{x}^{\mathbf{h}'^k}), \end{aligned} \tag{17}$$

where the inequality follows by applying inequality (1) for ω to half of $f^{m-1}(\mathbf{h}_{-k}(\bar{x}^1, \dots, \bar{x}^m))$ and the last equality follows from λ being a stationary distribution. As inequality (17) holds for all k , claim (b) follows.

Suppose that $(\bar{x}^1, \dots, \bar{x}^m) = \mathbf{g}(\bar{y}^1, \dots, \bar{y}^m)$. Then,

$$\begin{aligned}
\sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) f(\bar{x}^{\mathbf{h}_k}) &= \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) f^m(\mathbf{h}(\bar{x}^1, \dots, \bar{x}^m)) \\
&= \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) f^m(\mathbf{h} \circ \mathbf{g}(\bar{y}^1, \dots, \bar{y}^m)) \\
&= \sum_{\mathbf{h} \in \mathcal{H}} \lambda(\mathbf{h}) f^m(\mathbf{g}(\bar{y}^1, \dots, \bar{y}^m)) \\
&= f^m(\bar{x}^1, \dots, \bar{x}^m), \tag{18}
\end{aligned}$$

where the first equality follows from claim (b) and the third from claim (a), noting that $\mathbf{h} \circ \mathbf{g} \in \mathcal{H}$.

Since $m \geq 3$, we can pick $k \in \{1, \dots, m\} \setminus \{j_1, j_2\}$. Let $\mathbf{h} = \mathbf{1}^{h_1 \dots h_l}$ and let $\mathbf{p} \in \mathcal{O}_D^{(m \rightarrow m)}$ be the mapping that exchanges the elements in positions j_1 and j_2 . It is not hard to show, by induction over l , that $\mathbf{h} \circ \mathbf{p} = \mathbf{p} \circ \mathbf{h}$. Therefore, exchanging $x_i^{j_1}$ and $x_i^{j_2}$ does not affect the left-hand side of (18), so the exchange leaves the right-hand side unaffected as well. \square