

Piecewise Linear Valued CSPs Solvable by Linear Programming Relaxation*

Manuel Bodirsky[†]
Institut für Algebra
Technische Universität Dresden

Marcello Mamino[†]
Dipartimento di Matematica
Università di Pisa

Caterina Viola^{†‡}
Institut für Algebra
Technische Universität Dresden

Abstract

Valued constraint satisfaction problems (VCSPs) are a large class of combinatorial optimisation problems. The computational complexity of VCSPs depends on the set of allowed cost functions in the input. Recently, the computational complexity of all VCSPs for finite sets of cost functions over finite domains has been classified. Many natural optimisation problems, however, cannot be formulated as VCSPs over a finite domain. We initiate the systematic investigation of infinite-domain VCSPs by studying the complexity of VCSPs for piecewise linear homogeneous cost functions. Such VCSPs can be solved in polynomial time if the cost functions are improved by fully symmetric fractional operations of all arities. We show this by reducing the problem to a finite-domain VCSP which can be solved using the basic linear program relaxation. It follows that VCSPs for submodular PLH cost functions can be solved in polynomial time; in fact, we show that submodular PLH functions form a maximally tractable class of PLH cost functions.

1 Introduction

In a *valued constraint satisfaction problem (VCSP)* we are given a finite set of variables, a finite set of cost functions that depend on these variables, and

*This article is the full extended version of the conference paper “Submodular Functions and Valued Constraint Satisfaction Problems over Infinite Domains” [5] of the same authors. Although the conclusions of the present work are more general than those contained in the conference version, some of the methods used, the algorithms, and the intermediate results obtained are different and incomparable.

[†]The authors have received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 681988, CSP-Infinity).

[‡]This author has been supported by DFG Graduiertenkolleg 1763 (QuantLA).

a cost u ; the task is to find values for the variables such that the sum of the cost functions is at most u . By restricting the set of possible cost functions in the input, a great variety of computational optimisation problems can be modelled as a valued constraint satisfaction problem. By allowing the cost functions to evaluate to $+\infty$, we can even model ‘crisp’ (i.e., hard) constraints on the variable assignments, and hence the class of (classical) constraint satisfaction problems (CSPs) is a subclass of the class of all VCSPs.

If the domain is finite, the computational complexity of VCSPs has recently been classified for all sets of cost functions, assuming the Feder-Vardi conjecture for classical CSPs [16, 15, 17]. Even more recently, two solutions to the Feder-Vardi conjecture have been announced [23, 7]. These fascinating achievements settle the complexity of VCSPs over finite domains.

Several outstanding combinatorial optimisation problems cannot be formulated as VCSPs over a finite domain, but they can be formulated as VCSPs over the domain \mathbb{Q} , the set of rational numbers. One example is the famous linear programming problem, where the task is to optimise a linear function subject to linear inequalities. This can be modelled as a VCSP by allowing unary linear cost functions and cost functions of higher arity to express the crisp linear inequalities. Another example is the minimisation problem for sums of piecewise linear convex cost functions (see, e.g., [6]). Both of these problems can be solved in polynomial time, e.g. by the ellipsoid method (see, e.g., [11]).

Despite the great interest in such concrete VCSPs over the rational numbers in the literature, VCSPs over infinite domains have not yet been studied systematically. In order to obtain general results we need to restrict the class of cost functions that we investigate, because without any restriction it is already hopeless to classify the complexity of infinite-domain CSPs (any language over a finite alphabet is polynomial-time Turing equivalent to an infinite domain CSP [2]). One restriction that captures a variety of optimisation problems of theoretical and practical interest is the class of all *piecewise linear homogeneous* cost functions over \mathbb{Q} , defined below. We first illustrate by an example the type of cost functions that we want to capture in our framework.

Example 1.1. (LEAST CORRELATION CLUSTERING WITH PARTIAL INFORMATION) We consider the following problem: we are given a graph on n vertices such that the set of edges E is partitioned in two classes, E_- and E_+ . An edge $(x, y) \in E$ is either in E_+ or in E_- depending on whether x and y have been deemed to be similar or different. The goal is to (decide whether it is possible to) produce a partition of the vertices, namely a clustering, that agrees with the edge partition on at least l edges, where l is a given (rational) number between 0 and $|E|$. That is, we want a clustering that bounds the number of disagreements, i.e., the number of edges from E_+ between clusters plus the number of edges from E_- inside clusters. This

problem can be seen as an instance of a VCSP with variables x_1, \dots, x_n , with objective function

$$\phi(x_1, \dots, x_n) = \sum_{(x_i, x_j) \in E_+} f_1(x_i, x_j) + \sum_{(x_i, x_j) \in E_-} f_2(x_i, x_j),$$

and with threshold $u := |E| - l$, where the cost functions $f_1, f_2: \mathbb{Q}^2 \rightarrow \mathbb{Q}$ are defined by

$$f_1(x_i, x_j) = \begin{cases} 0 & \text{if } x_i = x_j \\ 1 & \text{otherwise} \end{cases}, \text{ and } f_2(x_i, x_j) = \begin{cases} 1 & \text{if } x_i = x_j \\ 0 & \text{otherwise} \end{cases}.$$

Observe that this problem cannot be modelled as a VCSP over a finite domain, as we do not want to bound the possible number of clusters and we want to allow graphs with any finite number of vertices as input. The least correlation clustering with partial information that we defined above is a generalisation of the well-known min-correlation clustering problem (see [1]). As the min-correlation clustering problem is known to be NP-hard, so is our problem. \triangle

A partial function $f: \mathbb{Q}^n \rightarrow \mathbb{Q}$ is called *piecewise linear homogeneous* (PLH) if it is first-order definable over the structure $\mathfrak{L} := (\mathbb{Q}; <, 1, (c \cdot)_{c \in \mathbb{Q}})$; being undefined at $(x_1, \dots, x_n) \in \mathbb{Q}^n$ is interpreted as $f(x_1, \dots, x_n) = +\infty$. The structure \mathfrak{L} has quantifier elimination (see Section 3.4) and hence there are finitely many regions such that f is a homogeneous linear polynomial in each region; this is the motivation for the name *piecewise linear homogeneous*. The cost function from Example 1.1 is PLH.

In this article we present a sampling technique to solve VCSPs for PLH cost functions in polynomial time. The technique consists of a polynomial-time many-one reduction from the VCSP for a finite set of PLH cost functions to the VCSP for a *sample*, i.e., the same set of cost functions interpreted over a suitable finite domain. We present this technique in two steps: first, in Section 4 we present the reduction for the *feasibility problem*, i.e., the problem of deciding whether there exists an assignment of values for the variables with finite cost; second, in Section 5 we extend this method to solve the *optimisation problem*, i.e., the task of deciding whether there exists an assignment of values for the variables with cost at most equal to a given threshold. In Section 6 we present a sufficient condition under which infinite-domain VCSPs that admit an efficient sampling algorithm can be solved in polynomial time using linear programming relaxation. The condition is given in terms of algebraic properties of the cost functions, namely the existence of *fully symmetric fractional polymorphisms of all arities*. We combine the results mentioned above to solve the VCSP for sets of PLH cost functions having such fractional polymorphisms. In particular, we apply the combination to show containment in P of the VCSP for *submodular* PLH cost functions,

for *convex* PLH cost functions, and for *componentwise increasing* PLH cost functions.

Submodular cost functions naturally appear in several scientific fields such as, for example, economics, game theory, machine learning, and computer vision, and play a key role in operational research and combinatorial optimisation (see, e.g., [10]). Submodularity also plays an important role for the computational complexity of VCSPs over finite domains, and guided the research on VCSPs for some time (see, e.g., [8, 14]), even though this might no longer be visible in the final classification obtained in [16, 15, 17]. The polynomial-time tractability of VCSPs for submodular PLH cost functions is the main result of an extended abstract [5] which announced some of the results presented here: we first provided a polynomial-time many one reduction to a finite-domain VCSP with costs in the ring $\mathbb{Q}^* := \mathbb{Q}((\epsilon))$ of formal Laurent power series in one indeterminate. Then we used a fully combinatorial polynomial-time algorithm to solve the VCSP for the \mathbb{Q}^* -valued submodular sample.

In the present article, starting from the same polynomial-time many-one reduction, we provide a polynomial-time many-one reduction to a finite-domain VCSP with values in \mathbb{Q} and then use linear programming relaxation to solve the finite-domain problem. The approach we present here is more general: it can be applied to all VCSPs for PLH cost functions that have fully symmetric fractional polymorphisms of all arities and it does not require the existence of fully combinatorial algorithms solving the finite-domain VCSP computed by our reduction. The concept of fractional polymorphisms comes from universal algebra and it has been fruitfully used to characterise the computational complexity of finite-domain VCSPs (see, e.g., [16], [21], [17], [15]). At the end of Section 7.1, we also show that submodularity defines a *maximal tractable class* of PLH valued structures: informally, adding any cost function that is not submodular leads to an NP-hard VCSP. Section 8 closes with some problems and challenges for future research.

2 Valued Constraint Satisfaction Problems

A *valued structure* Γ (over D) consists of

- a signature τ consisting of function symbols f , each equipped with an arity $\text{ar}(f)$,
- a set $D = \text{dom}(\Gamma)$ (the *domain*),
- for each $f \in \tau$ a *cost function*, i.e., a function $f^\Gamma : D^{\text{ar}(f)} \rightarrow \mathbb{Q} \cup \{+\infty\}$.

Here, $+\infty$ is an extra element with the expected properties that for all $c \in \mathbb{Q} \cup \{+\infty\}$

$$\begin{aligned} (+\infty) + c &= c + (+\infty) = +\infty \\ \text{and } c < +\infty &\text{ iff } c \in \mathbb{Q}. \end{aligned}$$

Let Γ be a valued structure with a finite signature τ . The *valued constraint satisfaction problem* for Γ , denoted by $\text{VCSP}(\Gamma)$, is the following computational problem.

Definition 2.1. An *instance* I of $\text{VCSP}(\Gamma)$ consists of

- a finite set of variables V_I ,
- an expression ϕ_I of the form

$$\sum_{i=1}^m f_i(x_1^i, \dots, x_{\text{ar}(f_i)}^i)$$

where $f_1, \dots, f_m \in \tau$ and all the x_j^i are variables from V_I , and

- a value $u_I \in \mathbb{Q}$.

The task is to decide whether there exists an assignment $\alpha: V_I \rightarrow \text{dom}(\Gamma)$, whose *cost*, defined as

$$\sum_{i=1}^m f_i^\Gamma(\alpha(x_1^i), \dots, \alpha(x_{\text{ar}(f_i)}^i))$$

is finite, and if so, whether there is one whose cost is at most u_I .

Note that, given a valued structure Γ , if the signature τ of Γ is finite, it is inessential for the computational complexity of $\text{VCSP}(\Gamma)$ how the function symbols in ϕ_I are represented.

The function from $\phi_I^\Gamma: \text{dom}(\Gamma)^{|V_I|} \rightarrow \mathbb{Q} \cup \{+\infty\}$ described by the expression ϕ_I is also called the *objective function*. The problem of deciding whether there exists an assignment $\alpha: V \rightarrow \text{dom}(\Gamma)$ with finite cost is called the *feasibility problem*, which can also be modelled as a (classical) constraint satisfaction problem (cf. Section 3.2). The choice of defining the VCSP as a decision problem and not as an optimisation problem as it is common for VCSPs over finite domains is motivated by two major issues that do not occur in the finite-domain case: first, in the infinite-domain setting one needs to capture the difference between a proper minimum and an infimum value that the cost of the assignment can be arbitrarily close to but never reach; second, our definition allows to model the case in which the infimum is $-\infty$, i.e., when there are assignments for the variables of arbitrarily small cost.

By a *finite-valued structure* we refer to a valued structure whose cost functions are finite-valued, i.e., they assume a finite value on every point with rational coordinates. The VCSP for a finite-valued structure is merely an optimisation problem and it does not involve any associated feasibility problem, i.e., the problem asking wheather there is an assignment with finite cost.

VCSPs have been studied intensively when $D = \text{dom}(\Gamma)$ is finite, and as mentioned in the introduction, in this case a complete classification of the computational complexity of $\text{VCSP}(\Gamma)$ has been obtained recently. However, many well-known optimisation problems can only be formulated when we allow infinite domains D .

Example 2.2. Let Γ be the valued structure with domain $D := \mathbb{Q}$ and the signature $\tau = \{R_+, R_1, \leq, \text{id}\}$ where

- R_+ is ternary, and

$$R_+^\Gamma(x, y, z) = \begin{cases} 0 & \text{if } x + y = z \\ +\infty & \text{otherwise;} \end{cases}$$

- R_1 is unary and

$$R_1^\Gamma(x) := \begin{cases} 0 & \text{if } x = 1 \\ +\infty & \text{otherwise;} \end{cases}$$

- \leq is binary and

$$\leq^\Gamma(x, y) := \begin{cases} 0 & \text{if } x \leq y \\ +\infty & \text{otherwise;} \end{cases}$$

- id is unary and

$$\text{id}^\Gamma(x) := x.$$

Then instances of $\text{VCSP}(\Gamma)$ can be viewed as instances of the linear program feasibility problem, i.e., the problem of deciding whether there exists a solution satisfying finitely many linear inequality constraints. \triangle

We give another example to illustrate the flexibility of the VCSP framework for formulating optimisation problems; the valued structure in this example contains non-convex cost functions, but, as we will see later, can be solved in polynomial time.

Example 2.3. Let Γ be the valued structure with signature $\tau = \{g_1, g_2, g_3\}$ and the cost functions

- $g_1^\Gamma : \mathbb{Q} \rightarrow \mathbb{Q}$ defined by $g_1(x) = -x$,

- $g_2^\Gamma: \mathbb{Q}^2 \rightarrow \mathbb{Q}$ defined by $g_2(x, y) := \min(x, -y)$, and
- $g_3^\Gamma: \mathbb{Q}^3 \rightarrow \mathbb{Q}$ defined by $g_3(x, y, z) := \max(x, y, z)$.

Two examples of instances of $\text{VCSP}(\Gamma)$ are

$$\begin{aligned} & g_1(x) + g_1(y) + g_1(z) + g_2(x, y) \\ & + g_3(x, y, z) + g_3(x, x, x) + g_3(x, x, x) \end{aligned} \quad (1)$$

$$\begin{aligned} \text{and } & g_1(x) + g_1(y) + g_1(z) \\ & + g_3(x, y, z) + g_3(x, x, y) + g_3(y, z, z) \end{aligned} \quad (2)$$

We can make the cost function described by the expression in (1) arbitrarily small by fixing x to 0 and choosing y and z sufficiently large. On the other hand, the minimum for the cost function in (2) is 0, obtained by setting x, y, z to 0. Note that g_1 and g_3 are convex functions, but g_2 is not. \triangle

3 Cost Functions over the Rationals

The class of all valued structures over arbitrary infinite domains is too large to allow for complete complexity classifications, so we have to restrict our focus to subclasses. In this section we describe natural and large classes of cost functions over the domain $D = \mathbb{Q}$, the rational numbers. These classes are most naturally introduced using first-order definability, so we briefly fix the necessary logic concepts.

3.1 Logic Preliminaries

We fix some standard logic terminology; see, e.g., [12]. A *signature* is a set τ of function and relation symbols. Each function symbol f and each relation symbol R is equipped with an arity $\text{ar}(f), \text{ar}(R) \in \mathbb{N}$. A τ -*structure* \mathfrak{A} consists of

- a set $A = \text{dom}(\mathfrak{A})$, called the *domain* of \mathfrak{A} , whose elements are called the *elements* of the τ -structure;
- a relation $R^\mathfrak{A} \subseteq A^{\text{ar}(R)}$ for each relation symbol $R \in \tau$;
- a function $f^\mathfrak{A}: A^{\text{ar}(f)} \rightarrow A$ for each function symbol $f \in \tau$.

Function symbols of arity 0 are allowed and are called constant symbols. We give two examples of structures that play an important role in this article.

Example 3.1. Let \mathfrak{S} be the structure with domain \mathbb{Q} and signature $\sigma := \{+, 1, \leq\}$ where

- $+$ is a binary function symbol that denotes the usual addition over \mathbb{Q} ,

- 1 is a constant symbol that denotes $1 \in \mathbb{Q}$, and
- \leq is a binary relation symbol that denotes the usual linear order of the rationals. \triangle

Example 3.2. Let \mathcal{L} be the structure with domain \mathbb{Q} and (countably infinite) signature $\tau_0 := \{<, 1\} \cup \{c \cdot\}_{c \in \mathbb{Q}}$ where

- $<$ is a relation symbol of arity 2 and $<^{\mathcal{L}}$ is the strict linear order of \mathbb{Q} ,
- 1 is a constant symbol and $1^{\mathcal{L}} := 1 \in \mathbb{Q}$, and
- $c \cdot$ is a unary function symbol for every $c \in \mathbb{Q}$ such that $(c \cdot)^{\mathcal{L}}$ is the function $x \mapsto cx$ (multiplication by c). \triangle

A *relational structure* (or structure with *relational signature*) is a structure whose signature contains only relation symbols. A structure (or a valued structure) is called *finite* if its domain is finite.

3.2 Constraint Satisfaction Problems

Let Γ be a valued structure with signature τ . The question whether an instance of $\text{VCSP}(\Gamma)$ is *feasible*, that is, has an assignment with finite cost, can be viewed as a (classical) *constraint satisfaction problem*. Formally, the constraint satisfaction problem for a structure \mathfrak{A} with a finite relational signature τ is the following computational problem, denoted by $\text{CSP}(\mathfrak{A})$:

- the input is a finite conjunction ψ of atomic τ -formulas, and
- the question is whether ψ is satisfiable in \mathfrak{A} .

We can associate to Γ the following relational structure $\text{Feas}(\Gamma)$: for every function symbol f of arity n from τ the signature of $\text{Feas}(\Gamma)$ contains a relation symbol R_f of arity n such that $R_f^{\text{Feas}(\Gamma)} = \text{dom}(f^\Gamma)$. Every polynomial-time algorithm for $\text{VCSP}(\Gamma)$ in particular has to solve $\text{CSP}(\text{Feas}(\Gamma))$. In fact, an instance ϕ of $\text{VCSP}(\Gamma)$ can be translated into an instance ψ of $\text{CSP}(\text{Feas}(\Gamma))$ by replacing subexpressions of the form $f(x_1, \dots, x_n)$ in ϕ by $R_f(x_1, \dots, x_n)$ and by replacing $+$ by \wedge . It is easy to see that ϕ is a feasible instance of $\text{VCSP}(\Gamma)$ if, and only if, ψ is satisfiable in $\text{Feas}(\Gamma)$.

3.3 Quantifier Elimination

We adopt the usual definition of first-order logic. A formula is *atomic* if it does not contain logical symbols (connectives or quantifiers). By convention, we have two special atomic formulas, \top and \perp , to denote truth and falsity.

Let τ be a signature. We say that a τ -structure \mathfrak{A} has *quantifier elimination* if every first-order τ -formula is equivalent to a quantifier-free τ -formula over \mathfrak{A} .

Theorem 3.3 ([9], Section 3, Theorem 1). *The structure \mathfrak{S} from Example 3.1 has quantifier elimination.*

Theorem 3.4. *The structure \mathfrak{L} from Example 3.2 has quantifier elimination.*

To prove Theorem 3.4 it suffices to prove the following lemma.

Lemma 3.5. *For every quantifier-free τ_0 -formula φ there exists a quantifier-free τ_0 -formula ψ such that $\exists x.\varphi$ is equivalent to ψ over \mathfrak{L} .*

Proof. Observe that every atomic τ_0 -formula ϕ has at most two variables:

- if ϕ has no variables, then it is equivalent to \top or \perp ,
- if ϕ has only one variable, say x , then it is equivalent to $c \cdot x \sigma d \cdot 1$ or to $d \cdot 1 \sigma c \cdot x$ for $\sigma \in \{<, =\}$ and $c, d \in \mathbb{Q}$. Moreover, if $c = 0$ then ϕ is equivalent to a formula without variables, and otherwise ϕ is equivalent to $x \sigma \frac{d}{c} \cdot 1$ or to $\frac{d}{c} \cdot 1 \sigma x$ for $\sigma \in \{<, =\}$, which we abbreviate by the more common $x < \frac{d}{c}$, $x = \frac{d}{c}$, and $\frac{d}{c} < x$, respectively.
- if ϕ has two variables, say x and y , then ϕ is equivalent to $c \cdot x \sigma d \cdot y$ or $c \cdot x \sigma d \cdot y$ for $\sigma \in \{<, =\}$. Moreover, if $c = 0$ or $d = 0$ then the formula ϕ is equivalent to a formula with at most one variable, and otherwise ϕ is equivalent to $x \sigma \frac{d}{c} \cdot y$ or to $\frac{d}{c} \cdot y \sigma x$.

We define ψ in five steps.

1. Rewrite φ , using De Morgan's laws, in such a way that all the negations are applied to atomic formulas.
2. Replace
 - $\neg(s = t)$ by $s < t \vee t < s$, and
 - $\neg(s < t)$ by $t < s \vee s = t$,
 where s and t are τ_0 -terms.
3. Write φ in disjunctive normal form in such a way that each of the clauses is a conjunction of non-negated atomic τ_0 -formulas (this can be done by distributivity).
4. Observe that $\exists x \bigvee_i \bigwedge_j \chi_{i,j}$, where the $\chi_{i,j}$ are atomic τ_0 -formulas, is equivalent to $\bigvee_i \exists x \bigwedge_j \chi_{i,j}$. Therefore, it is sufficient to prove the lemma for $\varphi = \bigwedge_j \chi_j$ where the χ_j are atomic τ_0 -formulas. As explained above, we can assume without loss of generality that the χ_j are of the form \top , \perp , $x \sigma c$, $c \sigma x$, or $x \sigma cy$, for $c \in \mathbb{Q}$ and $\sigma \in \{<, =\}$. If χ_j equals \perp , then φ is equivalent to \perp and there is nothing to be shown. If χ_j equals \top then it can simply be removed from φ . If χ_j equals $x = c$ or $x = cy$ then replace every occurrence of x by $c \cdot 1$ or by $c \cdot y$, respectively. Then φ does not contain the variable x anymore and thus $\exists x.\varphi$ is equivalent to φ .

5. We are left with the case that all atomic τ_0 -formulas involving x are (strict) inequalities, that is, $\varphi = \bigwedge_i \chi_i \wedge \bigwedge_i \chi'_i \wedge \bigwedge_i \chi''_i$, where

- the χ_i are atomic formulas not containing x ,
- the χ'_i are atomic formulas of the form $x > u_i$,
- the χ''_i are atomic formulas of the form $x < v_i$.

Then $\exists x.\varphi$ is equivalent to $\bigwedge_i \chi_i \wedge \bigwedge_{i,j} (u_i < v_j)$.

Each step of this procedure preserves the satisfying assignments for φ and the resulting formula is in the required form; this is obvious for all but the last step, and for the last step follows from the correctness of Fourier-Motzkin elimination for systems of linear inequalities (see, e.g., [19], Section 12.2). Therefore, the procedure is correct. \square

(of Theorem 3.4). Let φ be a τ_0 -formula. We prove that it is equivalent to a quantifier-free τ_0 -formula by induction on the number n of quantifiers of φ . For $n = 1$ we have two cases:

- If φ is of the form $\exists x.\varphi'$ (with φ' quantifier-free) then, by Lemma 3.5, it is equivalent to a quantifier-free τ_0 -formula ψ .
- If φ is of the form $\forall x.\varphi'$ (with φ' quantifier-free), then it is equivalent to $\neg\exists x.\neg\varphi'$. By Lemma 3.5, $\exists x.\neg\varphi'$ is equivalent to a quantifier-free τ_0 -formula ψ . Therefore, φ is equivalent to the quantifier-free τ_0 -formula $\neg\psi$.

Now suppose that φ is of the form $Q_1x_1Q_2x_2\cdots Q_nx_n.\varphi'$ for $n \geq 2$ and $Q_1, \dots, Q_n \in \{\forall, \exists\}$, and suppose that the statement is true for τ_0 -formulas with at most $n - 1$ quantifiers. In particular, $Q_2x_2\cdots Q_nx_n.\varphi'$ is equivalent to a quantifier-free τ_0 -formula ψ . Therefore, φ is equivalent to $Q_1x_1.\psi$, that is, a τ_0 -formula with one quantifier that is equivalent to a quantifier-free τ_0 -formula, again by the inductive hypothesis. \square

3.4 Piecewise Linear Homogeneous Functions

A *partial function* of arity $n \in \mathbb{N}$ over a set A is a function

$$f: \text{dom}(f) \rightarrow A \text{ for some } \text{dom}(f) \subseteq A^n.$$

Let \mathfrak{A} be a τ -structure. A partial function over A is called *first-order definable over \mathfrak{A}* if there exists a first-order τ -formula $\phi(x_0, x_1, \dots, x_n)$ such that for all $a_1, \dots, a_n \in A$

- if $(a_1, \dots, a_n) \in \text{dom}(f)$ then $\mathfrak{A} \models \phi(a_0, a_1, \dots, a_n)$ if, and only if, $a_0 = f(a_1, \dots, a_n)$, and

- if $f(a_1, \dots, a_n) \notin \text{dom}(f)$ then there is no $a_0 \in A$ such that $\mathfrak{A} \models \phi(a_0, a_1, \dots, a_n)$.

In the following, we consider *cost functions over \mathbb{Q}* , which will be functions from $\mathbb{Q}^n \rightarrow \mathbb{Q} \cup \{+\infty\}$. It is sometimes convenient to view a cost function as a partial function over \mathbb{Q} . If $t \in \mathbb{Q}^{\text{ar}(f)} \setminus \text{dom}(f)$ we interpret this as $f(t) = +\infty$.

Definition 3.6. A cost function $f: \mathbb{Q}^n \rightarrow \mathbb{Q} \cup \{+\infty\}$ (viewed as a partial function) is called

- *piecewise linear (PL)* if it is first-order definable over \mathfrak{S} , piecewise linear functions are sometimes called *semilinear* functions;
- *piecewise linear homogeneous (PLH)* if it is first-order definable over \mathfrak{L} (viewed as a partial function).

A valued structure Γ is called *piecewise linear (piecewise linear homogeneous)* if every cost function in Γ is PL (or PLH, respectively).

Definition 3.7. A relational structure \mathfrak{A} with domain \mathbb{Q} and relational signature τ is called *piecewise linear homogeneous (PLH)* if, for all $R \in \tau$, the interpretation $R^{\mathfrak{A}}$ is PLH.

Clearly, if Γ is a PLH valued structure then $\text{Feas}(\Gamma)$ is a PLH relational structure. Every PLH cost function is also PL, since all functions of the structure \mathfrak{L} are clearly first-order definable in \mathfrak{S} . The cost functions in the valued structure from Example 2.3 are PLH. The cost functions in the valued structure from Example 2.2 are PL, but not PLH.

We would like to point out that already the class of PLH cost functions is very large. In particular, it can be seen as a generalisation of the class of all cost functions over a finite domain. Indeed, every VCSP for a valued structure over a finite domain is also a VCSP for a valued structure that is PLH. To see this, suppose that $f: D^d \rightarrow \mathbb{Q} \cup \{+\infty\}$ is such a cost function, identifying D with a subset of \mathbb{Q} in an arbitrary way. Then the function $f': \mathbb{Q}^d \rightarrow \mathbb{Q} \cup \{+\infty\}$ defined by $f'(x_1, \dots, x_n) := f(x_1, \dots, x_n)$ if $x_1, \dots, x_n \in D$, and $f'(x_1, \dots, x_n) = +\infty$ otherwise, is PLH.

In the following, we prove that the VCSP for a PLH valued structure with finite signature is polynomial-time many-one equivalent to the VCSP for a valued structure over a suitable finite domain. This will be done in two steps: firstly we will show the reduction for the feasibility problem, i.e., we will prove such a reduction for PLH relational structure; secondly we will extend this method to solve the *optimisation problem*, i.e., to find a solution of cost at most equal to the threshold given in the instance.

4 Efficiently Sampling a PLH Relational Structure

Throughout this section, we fix a PLH relational structure \mathfrak{A} with finite signature τ . We present an efficient sampling algorithm for \mathfrak{A} . Before defining the notion of sampling algorithm, we recall a well-known concept from universal algebra.

Definition 4.1. Let \mathfrak{A} and \mathfrak{B} be (relational) structures with the same signature τ with domain A and B respectively. A *homomorphism* from \mathfrak{A} to \mathfrak{B} is a function $h: A \rightarrow B$ such that for every relation symbol $R \in \tau$ and tuple $a \in D^{\text{ar}(R)}$

$$R^{\mathfrak{A}}(a) \text{ implies } R^{\mathfrak{B}}(g(a)),$$

where the function g is applied componentwise. We say that \mathfrak{A} is *homomorphic* to \mathfrak{B} and write $\mathfrak{A} \rightarrow \mathfrak{B}$ to indicate the existence of a homomorphism from \mathfrak{A} to \mathfrak{B} .

We now formally introduce the notion of a sampling algorithm for a relational structure.

Definition 4.2. Let \mathfrak{C} be a structure with finite relational signature τ . A *sampling algorithm* for \mathfrak{C} takes as input a positive integer d and computes a finite-domain structure \mathfrak{D} homomorphic to \mathfrak{C} such that every finite conjunction of atomic τ -formulas having at most d distinct free variables is satisfiable in \mathfrak{C} if, and only if, it is satisfiable in \mathfrak{D} . A sampling algorithm is called *efficient* if its running time is bounded by a polynomial in d . We refer to the output of a sampling algorithm by calling it the *sample*.

The definition above is a slight re-formulation of Definition 2.2 in [3], and it is easily seen to give the same results using the same proofs. We decided to bound the number of variables instead of the size of the conjunction of atomic τ -formulas because this is more natural in our context. These two quantities are polynomially related by the assumption that the signature τ is finite.

We give a formal definition of the *numerical data* in \mathfrak{A} , we will need it later on. By quantifier elimination (Theorem 3.4), each of the finitely many relations $R^{\mathfrak{A}}$ for $R \in \tau$ has a quantifier-free τ_0 -formula ϕ_R over \mathfrak{L} . As in the proof of Theorem 3.4, we can assume that all formulas ϕ_R are positive (namely contain no negations). From now on, we will fix one such representation. Let $\text{At}(\phi_R)$ denote the set of atomic subformulas of ϕ_R . Each atomic τ_0 -formula is of the form $t_1 \leq t_2$, where t_1 and t_2 are terms. We call an atomic formula *trivial* if it is equivalent to \perp or \top , and *non-trivial* otherwise. As in the proof of Theorem 3.4, we make the assumption that atomic formulas are of the form \perp or \top if they are trivial, and otherwise of the form either $c_1 \cdot 1 \leq x_i$, or $x_i \leq c_2 \cdot 1$, or $c_1 \cdot x_i \leq c_2 \cdot x_j$ with constants c_1 and c_2 not both negative and where function symbols $c_i \cdot$ are never composed. This

assumption can be made without loss of generality (again, see the proof of Theorem 3.4).

Given a set of non-trivial atomic formulas Φ , we define

$$H(\Phi) = \left\{ \frac{c_1}{c_2} \mid t_1 = c_1 \cdot x_i, t_2 = c_2 \cdot x_j, \text{ for some } t_1 \begin{smallmatrix} < \\ = \end{smallmatrix} t_2 \text{ in } \Phi \right\}$$

$$\begin{aligned} K(\Phi) = & \left\{ \frac{c_2}{c_1} \mid t_1 = c_1 \cdot x_i, t_2 = c_2 \cdot 1, \text{ for some } t_1 \begin{smallmatrix} < \\ = \end{smallmatrix} t_2 \text{ in } \Phi \right\} \\ & \cup \left\{ \frac{c_1}{c_2} \mid t_1 = c_1 \cdot 1, t_2 = c_2 \cdot x_j, \text{ for some } t_1 \begin{smallmatrix} < \\ = \end{smallmatrix} t_2 \text{ in } \Phi \right\} \end{aligned}$$

The efficient sampling algorithm for \mathfrak{A} works in two steps. First, the problem $\text{CSP}(\mathfrak{A})$ is transferred to the equivalent CSP for a suitable structure \mathfrak{A}^* that is an extension of an expansion (or an expansion of an extension) of \mathfrak{A} , and second we provide an efficient sampling algorithm for \mathfrak{A}^* , which is also an efficient sampling algorithm for \mathfrak{A} .

Definition 4.3. The ordered \mathbb{Q} -vector space \mathbb{Q}^* is defined as

$$\mathbb{Q}^* = \{x + y\epsilon \mid x, y \in \mathbb{Q}\}$$

where ϵ is merely a formal device, namely $x + y\epsilon$ represents the pair (x, y) . We define addition and multiplication by a scalar componentwise

$$\begin{aligned} (x_1 + y_1\epsilon) + (x_2 + y_2\epsilon) &= (x_1 + x_2) + (y_1 + y_2)\epsilon \\ c \cdot (x + y\epsilon) &= (cx) + (cy)\epsilon. \end{aligned}$$

Clearly, \mathbb{Q} is embedded in \mathbb{Q}^* , by mapping every rational number x into $x + 0\epsilon$. The order is induced by \mathbb{Q} extended with $0 < \epsilon \ll 1$, namely the lexicographical order of the components x and y

$$(x_1 + y_1\epsilon) < (x_2 + y_2\epsilon) \quad \text{iff} \quad \begin{cases} x_1 < x_2 & \text{or} \\ x_1 = x_2 \wedge y_1 < y_2. \end{cases}$$

Any τ_0 -formula has an obvious interpretation in any ordered \mathbb{Q} -vector space extending \mathbb{Q} , and, in particular, in \mathbb{Q}^* .

Theorem 4.4 ([22], Chapter 1, Remark 7.9). *The first-order theory of ordered \mathbb{Q} -vector spaces in the signature $\tau_0 \cup \{+, -\}$ is complete.*

Let \mathfrak{A}^* be the τ -structure obtained by interpreting each relation symbol $R \in \tau$ by the relation $R^{\mathfrak{A}^*}$ defined on \mathbb{Q}^* by the same (quantifier-free) τ_0 -formula ϕ_R that defines $R^{\mathfrak{A}}$ over \mathbb{Q} (the choice of equivalent τ_0 -formulas is immaterial). Theorem 4.4 has the following immediate consequence.

Corollary 4.5. *Let ϕ be a τ -formula. Then ϕ is satisfiable in \mathfrak{A} if, and only if, ϕ is satisfiable in \mathfrak{A}^* .*

Our goal is to prove the following theorem.

Theorem 4.6. *There is an efficient sampling algorithm for \mathfrak{A}^* .*

Before giving the proof of Theorem 4.6 we present some preliminary lemmas in which we explicitly define the domain of the wanted sample. Let ϕ be an atomic τ_0 -formula. We write $\bar{\phi}$ for the formula $t_1 \leq t_2$ if ϕ is of the form $t_1 < t_2$, and for the formula $t_1 = t_2$ if ϕ is of the form $t_1 = t_2$. We call $\bar{\phi}$ the *closure* of the formula ϕ . First, we investigate the positive solutions to the closures of finitely many atomic τ_0 -formulas. Then, in a second step that builds on the first one, we investigate the solutions to finitely many atomic τ_0 -formulas.

Lemma 4.7. *Let Φ be a finite set of atomic τ_0 -formulas having free variables in $\{v_1, \dots, v_d\}$. Assume that $\bar{\Phi} := \bigcup_{\phi \in \Phi} \bar{\phi}$ has a simultaneous solution $(x_1, \dots, x_d) \in \mathbb{Q}_{>0}$ in positive numbers. Then $\bar{\Phi}$ has a solution taking values in the set $C_{\Phi,d} \subset \mathbb{Q}$ defined as follows*

$$C_{\Phi,d} = \left\{ \left| k \prod_{i=1}^s |h_i|^{e_i} \right| \mid k \in K(\Phi), e_1, \dots, e_s \in \mathbb{Z}, \sum_{r=1}^s |e_r| < d \right\}$$

where h_1, \dots, h_s is an enumeration of the (finitely many) elements of $H(\Phi)$.

Proof. Let $\gamma \leq \beta$ be maximal such that there are Ψ_1, Ψ_2, Ψ_3 with

$$\begin{aligned} \bar{\Phi} &= \{s_1 = s'_1, \dots, s_\alpha = s'_\alpha\} \cup \{t_1 \leq t'_1, \dots, t_\beta \leq t'_\beta\} \\ \Psi_1 &= \{s_1 = s'_1, \dots, s_\alpha = s'_\alpha\} \\ \Psi_2 &= \{t_1 = t'_1, \dots, t_\gamma = t'_\gamma\} \\ \Psi_3 &= \{t_{\gamma+1} \leq t'_{\gamma+1}, \dots, t_\beta \leq t'_\beta\}, \end{aligned}$$

where s_i, s'_i, t_j, t'_j are τ_0 -terms for all i, j , and $\Psi_1 \cup \Psi_2 \cup \Psi_3$ is satisfiable in positive numbers. Clearly the space of positive solutions of $\Psi_1 \cup \Psi_2$ must be contained in that of Ψ_3 . In fact, by construction, they intersect: consider any straight line segment connecting a solution of $\Psi_1 \cup \Psi_2 \cup \Psi_3$ and a solution of $\Psi_1 \cup \Psi_2$ not satisfying Ψ_3 , on this segment there must be a solution of $\Psi_1 \cup \Psi_2 \cup \Psi_3$ lying on the boundary of one of the inequalities of Ψ_3 , contradicting the maximality of γ . By the last observation it suffices to prove that there is a solution of $\Psi_1 \cup \Psi_2$ taking values in $C_{\Phi,d}$. Put an edge between two variables x_i and x_j when they appear in the same formula of $\Psi_1 \cup \Psi_2$. For each connected component of the graph thus defined, either it contains at least one variable v_i such that there is a constraint of the form $h \cdot v_i = k \cdot 1$, or all constraints are of the form $h \cdot v_i = h' \cdot v_j$. In the first

case assign $v_i = \frac{k}{h}$, in the second assign one of the variables v_i arbitrarily to 1, then, in any case, since the diameter of the connected component is smaller than d , all variables in this connected component are forced to take values in $C_{\Phi,d}$ by simple propagation of v_i . \square

Lemma 4.8. *Let Φ be a finite set of atomic τ_0 -formulas having free variables in $\{v_1, \dots, v_d\}$. Assume that the formulas in Φ are simultaneously satisfiable in \mathbb{Q} . Then they are simultaneously satisfiable in*

$$D_{\Phi,d} := -C_{\Phi,d}^* \cup \{0\} \cup C_{\Phi,d}^* \subseteq \mathbb{Q}^*$$

where

$$C_{\Phi,d}^* = \{x + nx\epsilon \mid x \in C_{\Phi,d}, n \in \mathbb{Z}, -d \leq n \leq d\} \subseteq \mathbb{Q}^*$$

$C_{\Phi,d}$ is defined as in Lemma 4.7, and $-C_{\Phi,d}^*$ denotes the set $\{-x \mid x \in C_{\Phi,d}^*\}$.

Proof. First we fix a solution $v_i = a_i$ for $i = 1 \dots d$ of Φ . In general, some of the values a_i will be positive, some 0, and some negative: we look for a new solution $z_1, \dots, z_d \in D_{\Phi,d}$ such that z_i is positive, respectively 0 or negative, if and only if a_i is. To this aim we rewrite the formulas in Φ replacing each variable v_i with either y_i , or 0 (formally $0 \cdot 1$), or $-y_i$ (formally $-1 \cdot y_i$). We call Φ^+ the new set of formulas, which, by construction, is satisfiable in positive numbers $y_i = b_i$. To establish the lemma, it suffices to find a solution of Φ^+ taking values in $C_{\Phi,d}^*$.

By Lemma 4.7, we have an assignment $y_i = c_i$ of values c_1, \dots, c_d in $C_{\Phi^+,d} \subseteq C_{\Phi,d}$ that satisfies simultaneously all formulas $\bar{\phi}$ with $\phi \in \Phi^+$. Let $n_1, \dots, n_d \in \{n \in \mathbb{Z} \mid -d \leq n \leq d\}$ be such that for all i, j

$$\begin{aligned} n_i < n_j & \quad \text{if and only if} \quad \frac{b_i}{c_i} < \frac{b_j}{c_j} \\ 0 < n_i & \quad \text{if and only if} \quad 1 < \frac{b_i}{c_i} \\ n_i < 0 & \quad \text{if and only if} \quad \frac{b_i}{c_i} < 1. \end{aligned}$$

Such numbers exist: simply sort the set $\{1\} \cup \{\frac{b_i}{c_i} \mid i = 1, \dots, d\}$ and consider the positions in the sorted sequence counting from that of 1. We claim that the assignment $y_i = c_i + n_i c_i \epsilon \in \mathbb{Q}^*$ satisfies all formulas of Φ^+ . To check this, we consider the different cases for atomic formulas

- $k \cdot y_i < h \cdot y_j$: if $kc_i < hc_j$ this is obviously satisfied. Otherwise $kc_i = hc_j$, in this case k and h are positive and the constraint

$$kc_i + kn_i c_i \epsilon < hc_j + hn_j c_j \epsilon$$

is equivalent to $n_i < n_j$. This, in turn, is equivalent by construction to $\frac{b_i}{c_i} < \frac{b_j}{c_j}$ which we get by observing that $b_i hc_j = b_i kc_i < b_j hc_i$.

- $k \cdot y_i = h \cdot y_j$: obviously $kb_i = hb_j$ and $kc_i = hc_j$, therefore $\frac{b_i}{c_i} = \frac{b_j}{c_j}$, and, as a consequence, also $n_i = n_j$ from which the statement.
- $k \cdot 1 < h \cdot y_j$: similarly to the first case, if $k < hc_j$ this is immediate. Otherwise $k = hc_j$, so k and h are positive, the constraint

$$k \cdot 1 < hc_j + hn_j c_j \epsilon$$

is equivalent to $0 < n_j$, in other words $1 < \frac{b_j}{c_j}$, which follows observing that $hc_j = k < hb_j$.

- $k \cdot y_i < h \cdot 1$: as the case above.
- $k \cdot 1 = h \cdot y_j$: obviously $k \cdot 1 = hb_j = hc_j$, therefore $\frac{b_j}{c_j} = 1$, so $n_j = 0$ and the case follows.
- $k \cdot y_i = h \cdot 1$: as the case above.

□

of Theorem 4.6. On input d , the sampling algorithm produces the finite substructure $\mathfrak{A}_{\text{At}(\tau),d}^*$ of \mathfrak{A}^* having domain $D_{\text{At}(\tau),d}$ where $\text{At}(\tau) := \bigcup_{R \in \tau} \text{At}(\phi_R)$, that is, the τ -structure with domain $D_{\text{At}(\tau),d}$ in which each relation symbol $R \in \tau$ denotes the restriction of $R^{\mathfrak{A}^*}$ to $D_{\text{At}(\tau),d}$. It is immediate to observe that this structure can be computed in polynomial time in d . Since $\mathfrak{A}_{\text{At}(\tau),d}^*$ is a substructure of \mathfrak{A}^* , it is clear that if an instance is satisfiable in $\mathfrak{A}_{\text{At}(\tau),d}^*$, then it is a fortiori satisfiable in \mathfrak{A}^* .

The vice versa follows from Lemma 4.8. In fact, consider a set Ψ of atomic τ -formulas having free variables x_1, \dots, x_d . Assume that Ψ is satisfied in \mathfrak{A}^* by one assignment $x_i = a_i$ for $i \in \{1, \dots, d\}$. For each $R \in \Psi$ let $\Phi_R \subset \text{At}(\phi_R)$ be the set of atomic subformulas of ϕ_R which are satisfied by our assignment a_i . Clearly the atomic τ_0 -formulas $\Phi := \bigcup_{R \in \Psi} \Phi_R$ are simultaneously satisfiable. Remembering that the formulas ϕ_R have no negations by construction, it is obvious that any simultaneous solution of Φ must also satisfy Ψ . By Lemma 4.8, Φ has a solution in the set $D_{\Phi,d}$ defined therein. We can observe that $C_{\Phi,d} \subset C_{\text{At}(\tau),d}$, hence $D_{\Phi,d} \subset D_{\text{At}(\tau),d}$ and the claim follows. □

From Theorem 4.6 it easily follows that there is an efficient sampling algorithm for \mathfrak{A} .

Corollary 4.9. *There exists an efficient sampling algorithm for \mathfrak{A} .*

Proof. On input d , the sampling algorithm produces the τ_0 -reduct of the sample for \mathfrak{A}^* (on input d) described in the proof of Theorem 4.6. By Corollary 4.5 and Theorem 4.6, this is an efficient sampling algorithm for \mathfrak{A} . □

Definition 4.10. A k -ary operation $g: \mathbb{Q}^k \rightarrow \mathbb{Q}$ is called a *polymorphism of \mathfrak{A}* if for all $R \in \tau$ we have that $g(x^1, \dots, x^k) \in R^{\mathfrak{A}}$ for all $x^1, \dots, x^k \in R^{\mathfrak{A}}$, namely $R^{\mathfrak{A}}$ is *preserved* by g (where g is applied component-wise).

As an application of Theorem 4.6 we prove the following result:

Theorem 4.11. *Let \mathfrak{A} be a PLH relational structure that is preserved by \max . Then $\text{CSP}(\mathfrak{A})$ is polynomial-time solvable.*

This result is incomparable to known results about semilinear relations preserved by \max [4]. In particular, there, the weaker bound $\text{NP} \cap \text{co-NP}$ has been shown for a larger class, and polynomial tractability only for a smaller class (which does not contain many PLH relations preserved by \max , for instance $x \geq \max(y, z)$). To prove Theorem 4.11 we use the notion of totally symmetric polymorphism and a result from [3].

Definition 4.12. A k -ary operation g is *totally symmetric* if

$$g(x^1, \dots, x^k) = g(y^1, \dots, y^k) \text{ whenever } \{x^1, \dots, x^k\} = \{y^1, \dots, y^k\}.$$

Example 4.13. For every $k \geq 1$, the k -ary totally symmetric operation $\max^{(k)}$ defined by $(x^1, \dots, x^k) \mapsto \max(x^1, \dots, x^k)$ is totally symmetric. \triangle

Theorem 4.14 ([3], Theorem 2.5). *Let \mathfrak{A} be a structure over a finite relational signature with totally symmetric polymorphisms of all arities. If there exists an efficient sampling algorithm for \mathfrak{A} then $\text{CSP}(\mathfrak{A})$ is in P .*

of Theorem 4.11. Since \mathfrak{A} is preserved by \max , for every $k \geq 1$, the k -ary totally symmetric operation $\max^{(k)}$ (see Example 4.13) preserves \mathfrak{A} . To see this, observe that for every $x^1, \dots, x^k \in \mathbb{Q}$, it holds that

$$\max^{(k)}(x^1, \dots, x^k) = \max(x^1, \max(x^2, \dots, \max(x^{k-1}, x^k) \dots)).$$

Therefore, as by Corollary 4.9 there exists an efficient sampling algorithm for \mathfrak{A} , from Theorem 4.14 it follows that $\text{CSP}(\mathfrak{A})$ can be solved in polynomial time. \square

5 Efficient Sampling for PLH Valued Structures

In this section we introduce the notion of a sampling algorithm for valued structures and we exhibit an efficient sampling algorithm for PLH valued structures. Analogously to the case of relational structures, before defining the notion of a sampling algorithm for valued structures, we define a generalisation of the universal algebraic concept of a homomorphism to the valued constraint satisfaction framework.

Definition 5.1. Let Δ and Γ be valued structures with the same signature τ with domain D and C respectively. Let C^D denote the set of all functions $g: D \rightarrow C$. A *fractional homomorphism* [20] from Δ to Γ is a function $\omega: C^D \rightarrow Q_{\geq 0}$ with finite support, $\text{Supp}(\omega) := \{g \in C^D \mid \omega(g) > 0\}$, such that $\sum_{g \in C^D} \omega(g) = 1$, and such that for every function symbol $f \in \tau$ and tuple $a \in D^{\text{ar}(f)}$, it holds that

$$\sum_{g \in C^D} \omega(g) f^\Gamma(g(a)) \leq f^\Delta(a),$$

where the functions g are applied componentwise. We say that Δ is *fractionally homomorphic* to Γ and write $\Delta \rightarrow_f \Gamma$ to indicate the existence of a fractional homomorphism from Δ to Γ .

Let Γ be a valued structure with domain C and signature τ . A *valued substructure* Δ of Γ is a valued structure with domain $D \subseteq C$, with signature τ , and such that the interpretation of a function symbol f from τ is the restriction of f^Γ to the domain D . Trivially, if Γ is a valued structure and Δ is a valued substructure of Γ , then Δ is fractionally homomorphic to Γ .

Definition 5.2. Let Γ be a valued structure with domain C and finite signature τ . A *sampling algorithm* for Γ takes as input a positive integer d and computes a finite-domain valued τ -structure Δ fractionally homomorphic to Γ such that, for every finite sum ϕ of τ -terms having at most d distinct variables, $V = \{x_1, \dots, x_d\}$, and every $u \in \mathbb{Q}$, there exists a solution $h: V \rightarrow C$ with $\phi^\Gamma(h(x_1), \dots, h(x_d)) \leq u$ if, and only if, there exists a solution $h': V \rightarrow D$ with $\phi^\Delta(h'(x_1), \dots, h'(x_d)) \leq u$. A sampling algorithm is called *efficient* if its running time is bounded by a polynomial in d .

The valued finite structure computed by a sampling algorithm is called a *sample*.

Remark 5.3. Observe that the output Δ_d of a sampling algorithm for a given valued structure Γ with finite signature τ does not depend on the rational threshold u . Therefore, the given sampling algorithm has the property that given a finite sum ϕ of function symbols from τ with variables $V := \{x_1, \dots, x_n\}$, it holds that

$$\inf_{h: V \rightarrow \text{dom}(\Gamma)} \phi^\Gamma(h(x_1), \dots, h(x_n)) = \inf_{h': V \rightarrow \text{dom}(\Delta_d)} \phi^{\Delta_d}(h'(x_1), \dots, h'(x_n)).$$

5.1 The Ring of Formal Laurent Power Series

In the present section we extend the method developed in Section 4 to the treatment of VCSPs. To better highlight the analogy with Section 4, so that the reader already familiar with it may quickly get an intuition of the arguments here, we use identical notation to represent corresponding objects.

This choice has the drawback that some symbols, notably \mathbb{Q}^* , need to be re-defined (the new \mathbb{Q}^* will contain the old one). In this section, we will sometimes skip details that can be borrowed unchanged from Section 4.

Definition 5.4. Let \mathbb{Q}^* denote the ring $\mathbb{Q}((\epsilon))$ of formal Laurent power series in the indeterminate ϵ , that is, \mathbb{Q}^* is the set of formal expressions

$$\sum_{i=-\infty}^{+\infty} a_i \epsilon^i$$

where $a_i \neq 0$ for only finitely many *negative* values of i . Clearly, \mathbb{Q} is embedded in \mathbb{Q}^* (the embedding is defined similarly to its corresponding in Section 4). The ring operations on \mathbb{Q}^* are defined as usual

$$\begin{aligned} \sum_{i=-\infty}^{+\infty} a_i \epsilon^i + \sum_{i=-\infty}^{+\infty} b_i \epsilon^i &= \sum_{i=-\infty}^{+\infty} (a_i + b_i) \epsilon^i \\ \sum_{i=-\infty}^{+\infty} a_i \epsilon^i \cdot \sum_{i=-\infty}^{+\infty} b_i \epsilon^i &= \sum_{i=-\infty}^{+\infty} \left(\sum_{j=-\infty}^{+\infty} a_j b_{i-j} \right) \epsilon^i \end{aligned}$$

where the sum in the product definition is always finite by the hypothesis on a_i, b_i with negative index i . The order is the lexicographical order induced by $0 < \epsilon \ll 1$, i.e.,

$$\sum_{i=-\infty}^{+\infty} a_i \epsilon^i < \sum_{i=-\infty}^{+\infty} b_i \epsilon^i \quad \text{iff} \quad \exists i \, a_i < b_i \wedge \forall j < i \, a_j = b_j.$$

It is well known that \mathbb{Q}^* is an ordered field, that is, all non-zero elements have a multiplicative inverse and the order is compatible with the field operations. We define the following subsets of \mathbb{Q}^* for $m \leq n$

$$\mathbb{Q}_{m,n}^* := \left\{ \sum_{i=m}^n \epsilon^i a_i \mid a_i \in \mathbb{Q} \right\} \subset \mathbb{Q}^*.$$

Definition 5.5. We define a new structure \mathfrak{L}^* , which is an extension of an expansion (or an expansion of an extension) of \mathfrak{L} , having \mathbb{Q}^* as domain and $\tau_1 := \tau_0 \cup \{k\}_{k \in \mathbb{Q}_{-1,1}^*}$ as signature, where the interpretation of symbols in τ_0 is formally the same as for \mathfrak{L} and the symbols $k \in \mathbb{Q}_{-1,1}^*$ denote constants (zero-ary functions).

Notice that, for technical reasons, we allow only constants from $\mathbb{Q}_{-1,1}^*$. In the remainder of this section, τ_1 -formulas will be interpreted in the structure \mathfrak{L}^* . We make for τ_1 -formulas the same assumptions as in Section 4 (that atomic subformulas are of the form either $c_1 \cdot 1 \leq x_i$, or $x_i \leq c_2 \cdot 1$, or $c_1 \cdot x_i \leq c_2 \cdot x_j$ with constants c_1 and c_2 not both negative and where function symbols c_i .

are never composed). Also $H(\Phi)$ and $K(\Phi)$, where Φ is a set of atomic τ_1 -formulas, are defined similarly to Section 4, and the only difference is that now Φ is a set of atomic τ_1 -formulas rather than τ_0 -formulas. Observe that the reduct of \mathfrak{L}^* obtained by restricting the signature to τ_0 is *elementarily equivalent* to \mathfrak{L} , namely it satisfies the same first-order sentences.

Similarly as in Section 4, for every valued PLH structure Γ and for every positive integer number d , we explicitly give a τ_1 -structure with finite domain D^* such that every instance of $\text{VCSP}(\Gamma)$ with at most d distinct free variables has a solution with values in \mathbb{Q} if, and only if, it has a solution with values in $D^* \subseteq \mathbb{Q}^*$ (Lemma 5.8). We will need two preliminary results: Lemma 5.6 and Lemma 5.7, which are analogues of Lemma 4.7 and Lemma 4.8 from Section 4. More specifically, in Lemma 5.6 we consider the positive solutions to the closures of finitely τ_1 -formulas, and in Lemma 5.7 we consider the positive solutions to finitely many τ_1 -formulas.

Lemma 5.6. *Let Φ be a finite set of atomic τ_1 -formulas having free variables in $\{v_1, \dots, v_d\}$ and let $\bar{\Phi}$ be the set $\bigcup_{\phi \in \Phi} \bar{\phi}$. Suppose that there is $0 < r \in \mathbb{Q}^*$ such that all satisfying assignments of $\bar{\Phi}$ with values (x_1, \dots, x_d) in \mathbb{Q}^* also satisfy $0 < x_i \leq r$ for all i . Let $u, \alpha_1, \dots, \alpha_d$ be elements of \mathbb{Q}^* . Assume that the formulas in Φ are simultaneously satisfiable by a point $(x_1, \dots, x_d) \in (\mathbb{Q}^*)^d$ such that $\sum_i \alpha_i x_i < u$. Let us define the set*

$$C_{\Phi,d} = \left\{ |k| \prod_{i=1}^s |h_i|^{e_i} \mid k \in K(\Phi), e_1, \dots, e_s \in \mathbb{Z}, \sum_{r=1}^s |e_r| < d \right\} \subseteq \mathbb{Q}_{-1,1}^*$$

where h_1, \dots, h_s is an enumeration of the (finitely many) elements of $H(\Phi)$. Then there is a point in $(x'_1, \dots, x'_d) \in (C_{\Phi,d})^d \subseteq (\mathbb{Q}^*)^d$ with $\sum_i \alpha_i x'_i < u$ that satisfies simultaneously all $\bar{\phi}$, for $\phi \in \Phi$.

Proof. As in the proof of Lemma 4.7 (to which we direct the reader for many details) we take a maximal $\gamma \leq \beta$ such that there are Ψ_1, Ψ_2, Ψ_3 with

$$\begin{aligned} \bar{\Phi} &= \{s_1 = s'_1, \dots, s_\alpha = s'_\alpha\} \cup \{t_1 \leq t'_1, \dots, t_\beta \leq t'_\beta\} \\ \Psi_1 &= \{s_1 = s'_1, \dots, s_\alpha = s'_\alpha\} \\ \Psi_2 &= \{t_1 = t'_1, \dots, t_\gamma = t'_\gamma\} \\ \Psi_3 &= \{t_{\gamma+1} \leq t'_{\gamma+1}, \dots, t_\beta \leq t'_\beta\} \end{aligned}$$

and $\Psi_1 \cup \Psi_2 \cup \Psi_3$ is satisfiable by an assignment (x_1, \dots, x_d) with $\sum_i \alpha_i x_i < u$. As in the proof of Lemma 4.7 the set of solutions of $\Psi_1 \cup \Psi_2$ satisfying $\sum_i \alpha_i x_i < u$ is contained in the solutions of Ψ_3 . So, here too, it suffices to show that there is a solution of $\Psi_1 \cup \Psi_2$ with $\sum_i \alpha_i v_i < u$ taking values in $C_{\Phi,d}$. The proof of Lemma 4.7 shows that there is a solution of $\Psi_1 \cup \Psi_2$ taking values (x_1, \dots, x_d) in $C_{\Phi,d}$ without necessarily meeting the requirement that $\sum_i \alpha_i x_i < u$. We will prove that, in fact, any such solution meets the additional constraint.

Let $x_i = a_i, b_i$ be two distinct satisfying assignments for $\Psi_1 \cup \Psi_2$ such that $\sum_i \alpha_i a_i < u$ and $\sum_i \alpha_i b_i \geq u$. We know that the first exists, and we assume the second towards a contradiction. The two assignments must differ, so, without loss of generality $a_1 \neq b_1$. For $t \in \mathbb{Q}^*$, with $t \geq 0$, define the assignment $x_i(t) = (1+t)a_i - tb_i$. Since all constraints in $\Psi_1 \cup \Psi_2$ are equalities, it is clear that the new assignment $x_i(t)$ satisfies $\Psi_1 \cup \Psi_2$ for all $t \in \mathbb{Q}^*$. Moreover, if $t \geq 0$

$$\sum_i \alpha_i x_i(t) \leq \sum_i \alpha_i a_i - t \left(\sum_i \alpha_i b_i - \sum_i \alpha_i a_i \right) < u$$

Let $t = \frac{2r}{|b_1 - a_1|}$. Then

$$x_1(t) = a_1 + \frac{2r}{|b - a|}(a - b)$$

is either not smaller than $2r$ or smaller than 0 depending on the sign of $(a-b)$. In either case we have a solution $x_i = x_i(t)$ of $\Psi_1 \cup \Psi_2$ satisfying $\sum_i \alpha_i x_i(t) < u$, which must therefore be a solution of Φ that does not satisfy $0 < x_i \leq r$. \square

Lemma 5.7. *Let Φ be a finite set of atomic τ_1 -formulas having free variables in $\{v_1, \dots, v_d\}$. Suppose that there are $0 < l < r \in \mathbb{Q}^*$ such that all satisfying assignments (x_1, \dots, x_d) of Φ in the domain \mathbb{Q}^* also satisfy $l < x_i < r$ for all i . Let $\alpha_1, \dots, \alpha_d$ be rational numbers and $u \in \mathbb{Q}_{-1,1}^*$. Assume that the formulas in Φ are simultaneously satisfiable by a point $(x_1, \dots, x_d) \in (\mathbb{Q}^*)^d$ such that $\sum_i \alpha_i x_i \leq u$. Then the same formulas are simultaneously satisfiable by a point $(x'_1, \dots, x'_d) \in (C_{\Phi,d}^*)^d \subseteq (\mathbb{Q}^*)^d$ such that $\sum_i \alpha_i x'_i \leq u$ where*

$$C_{\Phi,d}^* = \{x + nx\epsilon^3 \mid x \in C_{\Phi,d}, n \in \mathbb{Z}, -d \leq n \leq d\} \subseteq \mathbb{Q}_{-1,4}^*.$$

Proof. We consider two cases: either all satisfying assignments (x_1, \dots, x_d) satisfy the inequality $\sum_i \alpha_i x_i \geq u$, or there exists a satisfying assignment (x_1, \dots, x_d) for Φ such that $\sum_i \alpha_i x_i < u$.

In the first case, we claim that all satisfying assignments, in fact, satisfy $\sum_i \alpha_i x_i = u$. In fact, assume that $x_i = a_i, b_i$ are two satisfying assignments such that $\sum_i \alpha_i a_i = u$ and $v := \sum_i \alpha_i b_i > u$. As in the proof of Lemma 5.6, consider assignments of the form $x_i(t) = (1+t)a_i - tb_i$ for $t \in \mathbb{Q}^*$. Clearly $\sum_i \alpha_i x_i(t) = u - t(v-u) < u$ for all $t > 0$. As in Lemma 5.6, the new assignment must satisfy all equality constraints in Φ . Each inequality constraint implies a strict inequality on t (remember that Φ only has strict inequalities). Since all of these must be satisfied by $t = 0$, there is an open interval of acceptable values of t around 0 , and, in particular, an acceptable $t > 0$. Our claim is thus established. Therefore, in this case, it suffices to find any satisfying assignment for Φ taking values in $C_{\Phi,d}^*$. The assignment is now

constructed as in the proof of Lemma 4.8, replacing the formal symbol ϵ in that proof by ϵ^3 . Namely take a satisfying assignment $x_i = b_i$ for Φ , and, by Lemma 5.6, one satisfying assignment $x_i = c_i$ for $\bar{\Phi}$ taking values in $C_{\Phi,d}$. Observe that the hypothesis that all solutions of Φ satisfy $l < x_i$ for all i is used here to ensure that all solutions of $\bar{\Phi}$ assign positive values to the variables, which is required by Lemma 5.6. Let $-d \leq n_1, \dots, n_d \leq d$ be integers such that for all i, j

$$\begin{aligned} n_i < n_j & \text{ if and only if } \frac{b_i}{c_i} < \frac{b_j}{c_j} \\ 0 < n_i & \text{ if and only if } 1 < \frac{b_i}{c_i} \\ n_i < 0 & \text{ if and only if } \frac{b_i}{c_i} < 1 \end{aligned}$$

The assignment $y_i = c_i + n_i c_i \epsilon^3$ can be seen to satisfy all formulas of Φ by the same check as in the proof of Lemma 4.8. Observe that we have to replace ϵ in Lemma 4.8 by ϵ^3 here, so that $\mathbb{Q}_{-1,1}^* \cap \epsilon^3 \mathbb{Q}_{-1,1}^* = \emptyset$.

For the second case, fix a satisfying assignment $x_i = b_i$. By Lemma 5.6 there is an assignment $x_i = c_i \in C_{\Phi,d}$ such that $\sum_i \alpha_i c_i < u$ and this assignment satisfies $\bar{\phi}$ for all $\phi \in \Phi$. From these two assignments construct the numbers n_i and then the assignment $y_i = c_i + n_i c_i \epsilon^3$ as before. For the same reason it is clear that the new assignment satisfies Φ . To conclude that $\sum_i \alpha_i y_i < u$ we write

$$\sum_i \alpha_i y_i = \sum_i \alpha_i c_i + \epsilon^3 \sum_i \alpha_i n_i c_i < u$$

because the first summand is in $\mathbb{Q}_{-1,1}^*$ and smaller than u , therefore the second summand is neglected in the lexicographical order. \square

Lemma 5.8. *Let Φ be a finite set of atomic τ_0 -formulas having free variables in $\{v_1, \dots, v_d\}$. Let $u, \alpha_1, \dots, \alpha_d$ be rational numbers. Then the following are equivalent:*

1. *The formulas in Φ are simultaneously satisfiable in \mathbb{Q} , by a point $(x_1, \dots, x_d) \in \mathbb{Q}^d$ such that $\sum_i \alpha_i x_i \leq u$.*
2. *The formulas in Φ are simultaneously satisfiable in $D_{\Phi,d} \subseteq \mathbb{Q}^*$, by a point $(x'_1, \dots, x'_d) \in D_{\Phi,d}^d$ such that $\sum_i \alpha_i x'_i \leq u$, where the set $D_{\Phi,d}$ is defined as follows*

$$\begin{aligned} D_{\Phi,d} &:= -C_{\Phi',d}^* \cup \{0\} \cup C_{\Phi',d}^* \subseteq \mathbb{Q}_{-1,4}^* \\ \Phi' &:= \Phi \cup \{x > \epsilon, x < -\epsilon, x > -\epsilon^{-1}, x < \epsilon^{-1}\}. \end{aligned}$$

Remark 5.9. Observe that the set $D_{\Phi,d}$ depends neither on the α_i 's nor on u . In fact, it only depends on the set of formulas Φ and on the number d of free variables.

Proof. The implication $2 \rightarrow 1$ is immediate observing that the conditions Φ and $\sum_i \alpha_i x_i \leq u$ are first-order definable in \mathfrak{S} . In fact, any assignments with values in $D_{\Phi,d}$ satisfying the conditions is, in particular, an assignment in \mathbb{Q}^* , and, by the completeness of the first order theory of ordered \mathbb{Q} -vector spaces, we have an assignment taking values in \mathbb{Q} .

For the vice versa, fix any assignment $x_i = a_i$ with $a_i \in \mathbb{Q}$ for $i \in \{1, \dots, d\}$. We pre-process the formulas in Φ producing a new set of atomic formulas Φ' as follows. We replace every variable x_i such that $a_i = 0$ with the constant $0 = 0 \cdot 1$. Then we replace each of the remaining variables x_i with either y_i or $-y_i$ according to the sign of a_i . Finally, we add the constraints $\epsilon < y_i$ and $y_i < \epsilon^{-1}$ for each of these variables. Similarly we produce new coefficients $\alpha'_i = \text{sign}(a_i)\alpha_i$. It is clear that the new set of formulas Φ' has a satisfying assignment in positive rational numbers with $\sum_i \alpha'_i y_i \leq u$. Observing that a positive rational x always satisfies $\epsilon < x < \epsilon^{-1}$, we see that Φ' satisfies the hypothesis of Lemma 5.7 with $l = \epsilon$ and $r = \epsilon^{-1}$. Hence the statement. \square

Remark 5.10. Lemma 5.8 provides a polynomial-time many-one reduction of the VCSP for a PLH valued structure Γ with finite signature to the VCSP for a valued structure Δ^* with finite signature having as domain a finite subset of \mathbb{Q}^* . We want to point out that, however, Lemma 5.8 does not give rise to a sampling algorithm: firstly, the computed finite domain is a subset of \mathbb{Q}^* rather than \mathbb{Q} , secondly, the signature of Δ^* is strictly larger than the signature of Γ . In the next section we show how to obtain an efficient sampling algorithm for PLH valued structures using Lemma 5.8.

5.2 Reduction to a VCSP over a Finite Rational Domain

In this section we use the results achieved in Section 5.1 to provide an efficient sampling algorithm for PLH valued structures.

Let Φ be a finite set of τ_0 -formulas having at most d distinct free variables. By Lemma 5.8, for every $\alpha_1, \dots, \alpha_d, u \in \mathbb{Q}$ the formulas in Φ are simultaneously satisfiable in \mathbb{Q} by a point (x_1, \dots, x_d) such that $\sum_{i=1}^d \alpha_i x_i \leq u$ if, and only if, they are simultaneously satisfiable in $D_{\Phi,d} \subseteq \mathbb{Q}_{-1,4}^*$ by a point (x'_1, \dots, x'_d) such that $\sum_{i=1}^d \alpha_i x'_i \leq u$. The elements in $D_{\Phi,d} \subseteq \mathbb{Q}_{-1,4}^*$ are of the form

$$\sum_{i=-1}^4 x_i \epsilon^i \text{ where } x_i \in -C_{\Phi,d} \cup \{0\} \cup C_{\Phi,d}$$

where $C_{\Phi,d}$ is as in Lemma 4.7. In \mathfrak{L}^* it holds that

$$\begin{aligned} -\epsilon^{-1} < x < -\epsilon & \quad \text{for every } x \in -C_{\Phi,d}, \\ \epsilon < x < \epsilon^{-1} & \quad \text{for every } x \in C_{\Phi,d}. \end{aligned}$$

(See Lemmas 5.6, 5.7, 5.8, and Lemmas 4.7, 4.8.)

Lemma 5.11. *Let Φ be a finite set of τ_0 -formulas having at most d distinct free variables. Let $D_{\Phi,d}$ be defined as in Lemma 5.8 (2.) and let \mathfrak{D}^* be the finite substructure of \mathfrak{L}^* with domain $D_{\Phi,d} \subseteq \mathbb{Q}_{-1,4}^*$. Then there exists a positive rational ε with $-\varepsilon^{-1} < x < -\varepsilon$ for every $x \in -C_{\Phi,d}$ and $\varepsilon < x < \varepsilon^{-1}$ for every $x \in C_{\Phi,d}$ such that the map $\eta: D_{\Phi,d} \rightarrow \mathbb{Q}$ defined by*

$$\sum_{j=1}^4 x_j \epsilon^j \mapsto \sum_{j=1}^4 x_j \varepsilon^j$$

is a homomorphism from the τ_0 -reduct of \mathfrak{D}^ to \mathfrak{L} . Moreover, ε and η are computable in time polynomial in d .*

Proof. Let $C_{\Phi,d} \subseteq \mathbb{Q}$ be as in Lemma 4.7 and let $-C_{\Phi,d} := \{-x \mid x \in C_{\Phi,d}\}$. We define $C := \{x - y \mid x, y \in C_{\Phi,d} \text{ and } x - y > 0\}$. Observe that $C \supseteq C_{\Phi,d}$. Define

$$\varepsilon := \frac{1}{6} \frac{\min C}{\max C}.$$

The number ε is positive and rational and it can be computed in time polynomial in d . Furthermore,

$$\begin{aligned} 0 < \varepsilon < 1, \\ \varepsilon < \min C < \min C_{\Phi,d} & \quad \text{implying } \max(-C_{\Phi,d}) < -\varepsilon, \text{ and} \\ \max C = \max C_{\Phi,d} < \varepsilon^{-1} & \quad \text{implying } -\varepsilon^{-1} < \min(-C_{\Phi,d}). \end{aligned}$$

It follows that $-\varepsilon^{-1} < x < -\varepsilon$ for every $x \in -C_{\Phi,d}$, and $\varepsilon < x < \varepsilon^{-1}$ for every $x \in C_{\Phi,d}$.

It is easy to see that η preserves the scalar multiplication by rational elements and the identity in \mathbb{Q} . We prove now that η is order preserving (and therefore injective). Let us consider

$$x := \sum_{j=-1}^4 x_j \epsilon^j, \text{ and } y := \sum_{j=-1}^4 y_j \epsilon^j \in D_{\Phi,d}$$

such that x is smaller than y in the lexicographic order induced by $\mathbb{Q}_{-1,4}^*$. This means that there exists an index $i \in \{-1, \dots, 4\}$ such that $x_j = y_j$ for $-1 \leq j < i$ and $x_i < y_i$. Since $\varepsilon > 0$, it holds that $x_i \varepsilon^i < y_i \varepsilon^i$ and, consequently,

$$\sum_{j=-1}^i x_j \varepsilon^j < \sum_{j=-1}^i y_j \varepsilon^j.$$

Moreover, if $i \neq 4$ then for all $j \in \{i+1, \dots, 4\}$

$$(x_j - y_j) \varepsilon^{j-i} \leq (x_j - y_j) \varepsilon \leq (\max C) \varepsilon \leq \frac{\min C}{6} \leq \frac{y_i - x_i}{6}$$

because $\varepsilon < 1$, and $x_j - y_j \leq \max C$ (indeed, if $x_j - y_j > 0$ then $x_j - y_j \in C$, otherwise $x_j - y_j$ is smaller than any element in C). Therefore we get $(x_j - y_j)\varepsilon^j \leq \frac{y_i - x_i}{6}\varepsilon^i$, and

$$\sum_{j=i+1}^4 (x_j - y_j)\varepsilon^j \leq \sum_{j=i+1}^4 \frac{y_i - x_i}{6}\varepsilon^i \leq \frac{5}{6}(y_i - x_i)\varepsilon^i < (y_i - x_i)\varepsilon^i.$$

It follows that

$$\sum_{j=i}^4 x_j \varepsilon^j < \sum_{j=i}^4 y_j \varepsilon^j,$$

and, because $x_j = y_j$ for $-1 \leq j \leq i-1$,

$$\sum_{j=-1}^4 x_j \varepsilon^j < \sum_{j=-1}^4 y_j \varepsilon^j,$$

i.e., η preserves the order. \square

Let Φ be a finite set of τ_0 -formulas with at most d variables, and let η be the map introduced in Lemma 5.11, we set $E_{\Phi,d} := \eta(D_{\Phi,d}) \subseteq \mathbb{Q}$. Lemma 5.11 implies the following corollary.

Corollary 5.12. *Let Γ be a PLH valued structure with finite signature τ . Then there exists an efficient sampling algorithm for Γ .*

Proof. For every cost function f in τ let us consider the quantifier-free τ_0 -formula ϕ_f defining $R_f^\Gamma = \text{dom}(f)$ over \mathbb{Q} . Let $\text{At}(\phi_f)$ denote the set of atomic subformulas of ϕ_f and let $\text{At}(\tau) := \bigcup_{f \in \tau} \text{At}(\phi_f)$.

On input d , the algorithm produces the valued finite substructure Δ of Γ having domain $\eta(D_{\text{At}(\tau),d})$, where η is defined as in Lemma 5.11. It is immediate to see that the valued structure Δ has size polynomial in d and can be computed in time polynomial in d , because $D_{\text{At}(\tau),d}$ and $\eta(D_{\text{At}(\tau),d})$ can be computed in polynomial time in d (see Remark 5.9, and Lemma 5.11). Let ϕ be a finite sum of function symbols from τ with at most d variables from $V := \{v_1, \dots, v_d\}$, and let u be a rational number. By Lemma 5.8 there exists an assignment $h: V \rightarrow \mathbb{Q}$ such that $\phi^\Gamma(h) \leq u$ if, and only if, there exists an assignment $h': V \rightarrow D_{\text{At}(\tau),d}$ such that $\phi^{\Gamma^*}(h) \leq u$. Furthermore, by Lemma 5.11 there exists a positive rational number ε such that the map $\eta: D_{\text{At}(\tau),d} \rightarrow \mathbb{Q}$ is a homomorphism of τ_0 -structures from the τ_0 -reduct of \mathfrak{D}^* to \mathfrak{L} . Since η is injective, the assignment $\eta \circ h': V \rightarrow \eta(D_{\text{At}(\tau),d})$ has cost $\phi^\Delta(\eta \circ h') \leq u$. \square

6 Fully Symmetric Fractional Polymorphisms

In this section we give sufficient algebraic conditions under which valued structures with an infinite domain that admit an efficient sampling algorithm can be solved in polynomial time.

6.1 Universal Algebraic Tools

In this section we survey the concepts from universal algebra that we need and use in the remainder of the article.

Definition 6.1. Let Δ and Γ be valued structures with the same signature τ with domain D and C respectively. Let C^D denote the set of all functions $g: D \rightarrow C$. A *fractional homomorphism* [20] from Δ to Γ is a function $\omega: C^D \rightarrow \mathbb{Q}_{\geq 0}$ with finite support, $\text{Supp}(\omega) := \{g \in C^D \mid \omega(g) > 0\}$, such that $\sum_{g \in C^D} \omega(g) = 1$, and such that for every function symbol $f \in \tau$ and tuple $a \in D^{\text{ar}(f)}$, it holds that

$$\sum_{g \in C^D} \omega(g) f^\Gamma(g(a)) \leq f^\Delta(a),$$

where the functions g are applied componentwise. We write $\Delta \rightarrow_f \Gamma$ to indicate the existence of a fractional homomorphism from Δ to Γ .

The following proposition is adapted from [20], Proposition 2.1 where it is stated for valued structures with finite domains. In fact, we need not explicitly prove it, as the proof of the proposition in the finite-domain case can be easily modified to work in the infinite-domain case.

Proposition 6.2. *Let Δ and Γ be valued structures with the same signature τ and with domain D and C , respectively. Assume that $\Delta \rightarrow_f \Gamma$. Let I be an instance of $\text{VCSP}(\Gamma)$ having variables $V_I = \{x_1, \dots, x_n\}$, objective function $\phi_I(x_1, \dots, x_n)$, and threshold $u_I \in \mathbb{Q} \cup \{+\infty\}$. Suppose that there exists an assignment $h: V_I \rightarrow D$ such that $\phi_I^\Delta(h(x_1), \dots, h(x_n)) \leq u_I$. Then there exists an assignment $h': V_I \rightarrow C$ such that $\phi_I^\Gamma(h'(x_1), \dots, h'(x_n)) \leq u_I$. In particular, it holds that*

$$\inf_{c \in C^n} \phi_I^\Gamma(c) \leq \inf_{d \in D^n} \phi_I^\Delta(d).$$

Let Γ be a valued structure with countable (possibly infinite) domain C and signature τ . An m -ary operation on C is a function $g: C^m \rightarrow C$. Let $\mathcal{O}_C^{(m)}$ denote the set of all m -ary operations on C .

Definition 6.3. An m -ary *fractional operation* is a function $\omega: \mathcal{O}_C^{(m)} \rightarrow \mathbb{Q}_{\geq 0}$. The set $\{g \in \mathcal{O}_C^{(m)} \mid \omega(g) > 0\}$ is called the *support* of ω and it is denoted by $\text{Supp}(\omega)$. An m -ary fractional operation with finite support is a *fractional polymorphism* if $\sum_{g \in \text{Supp}(\omega)} \omega(g) = 1$ and for every $f \in \tau$ and tuples $a^1, \dots, a^m \in C^{\text{ar}(f)}$ it holds that

$$\sum_{g \in \mathcal{O}_C^{(m)}} \omega(g) f^\Gamma(g(a^1, \dots, a^m)) \leq \frac{1}{m} \sum_{i=1}^m f^\Gamma(a^i)$$

(where g is applied component-wise). The set of all fractional polymorphisms of a valued structure Γ is denoted by $\text{fPol}(\Gamma)$. If ω is a fractional polymorphism of Γ we say that Γ is *improved* by ω .

Let S_m be the symmetric group on $\{1, \dots, m\}$. An m -ary operation g is *fully symmetric* if for every permutation $\pi \in S_m$, we have

$$g(x_1, \dots, x_m) = g(x_{\pi(1)}, \dots, x_{\pi(m)}).$$

Note that every totally symmetric operation (see Definition 4.12) is also fully symmetric. A fractional operation is *fully symmetric* (*totally symmetric*, respectively) if every operation in its support is fully symmetric (totally symmetric, respectively). The fractional operations in the next example are totally symmetric.

Example 6.4. Let D be a totally ordered set. Let \min and \max be, respectively, the binary operations giving the smallest and the largest among two arguments, respectively. The fractional operation $\omega_{\max}: \mathcal{O}_D^{(2)} \rightarrow [0, 1]$,

$$\omega_{\max}(g) := \begin{cases} 1 & \text{if } g = \max \\ 0 & \text{otherwise} \end{cases}$$

is a binary totally symmetric fractional operation with finite support. The fractional operation $\omega_{\min}: \mathcal{O}_D^{(2)} \rightarrow [0, 1]$, is defined dually and it is a binary totally symmetric fractional operation with finite support, too. \triangle

The next example shows a family of fractional operations that are fully symmetric but not totally symmetric.

Example 6.5. Let $\text{avg}^{(k)}: D^k \rightarrow D$ be the k -ary arithmetic average operation defined, for every $(x^1, \dots, x^k) \in D^k$, by

$$\text{avg}(x^1, \dots, x^k) = \frac{1}{k} \sum_{i=1}^k x^i.$$

Let us define, for every $k \geq 2$, the fractional operation $\omega_{\text{conv}}^{(k)}: \mathcal{O}_D^{(k)} \rightarrow \mathbb{Q}_{\geq 0}$ such that

$$\omega_{\text{conv}}^{(k)}(g) = \begin{cases} 1 & \text{if } g = \text{avg}^{(k)} \\ 0 & \text{otherwise.} \end{cases}$$

The fractional operations $\omega_{\text{conv}}^{(k)}$ are fully symmetric, but for $k \geq 3$ they are not totally symmetric, because, e.g., $\{1, 1, 2\} = \{1, 2, 2\}$ but $\frac{1+1+2}{3} \neq \frac{1+2+2}{3}$. \triangle

6.2 The Basic Linear Programming Relaxation

Every VCSP over a finite domain has a natural linear programming relaxation. Let Γ be a valued structure with finite domain D and let I be an instance of VCSP(Γ) with set of free variables $V_I = \{x_1, \dots, x_d\}$, and objective function

$$\phi_I(x_1, \dots, x_d) = \sum_{j \in J} f_j(x_1^j, \dots, x_{n_j}^j),$$

with $f_j \in \Gamma$, $x^j = (x_1^j, \dots, x_{n_j}^j) \in V^{n_j}$, for all $j \in J$ (the set J is finite and indexing the cost functions that are summands of ϕ_I). Define the sets W_1 , W_2 , and W of variables $\lambda_j(t)$, $\mu_{x_i}(a)$, for $j \in J$, $t \in D^{n_j}$, $x_i \in V$, and $a \in D$, as follows.

$$\begin{aligned} W_1 &:= \{\lambda_j(t) \mid j \in J \text{ and } t = (t_1, \dots, t_{n_j}) \in D^{n_j}\}, \\ W_2 &:= \{\mu_{x_i}(a) \mid x_i \in V \text{ and } a \in D\}, \\ W &:= W_1 \cup W_2. \end{aligned}$$

Then the *basic linear programming, BLP, relaxation associated to I* (see [20], [16], and references therein) is a linear program with variables W and is defined as follows

$$\text{BLP}(I, \Gamma) := \min_{\lambda_j(t) \in \mathbb{Q}} \sum_{j \in J} \sum_{t \in D^{n_j}} \lambda_j(t) f_j^\Gamma(t)$$

such that

$$\begin{aligned} \sum_{t \in D^{n_j} : t_l = a} \lambda_j(t) &= \mu_{x_l^j}(a) && \text{for all } j \in J, l \in \{1, \dots, n_j\}, a \in D \\ \sum_{a \in D} \mu_{x_i}(a) &= 1 && \text{for all } x_i \in V \\ \lambda_j(t) &= 0 && \text{for all } j \in J, t \notin \text{dom}(f_j) \\ 0 \leq \lambda_j(t), \mu_{x_i}(a) &\leq 1 && \text{for all } \lambda_j(t) \in W_1, \mu_{x_i}(a) \in W_2. \end{aligned}$$

If there is no feasible solution to this linear program, then $\text{BLP}(I, \Gamma) = +\infty$. We say that the BLP relaxation *solves I* if $\text{BLP}(I, \Gamma) = \min_{x \in D^d} \phi_I(x)$. We say that the BLP relaxation solves VCSP(Γ) if it solves all instances of VCSP(Γ). For a given instance of a finite-domain VCSP, the corresponding BLP relaxation can be computed in polynomial time. Therefore, if the VCSP for a valued structure Γ is solved by the BLP relaxation, then VCSP(Γ) can be solved in polynomial time.

Let Γ be a valued structure that admits an efficient sampling algorithm. We may solve VCSP(Γ) by using the following algorithm that computes the BLP relaxation of the sample, and then solves the BLP relaxation.

ALGORITHM 1: Sampling + BLP Algorithm

Input: $I := (V_I, \phi_I, u_I)$.

Output: accepts if there exists an assignment $h: V_I \rightarrow \text{dom}(\Gamma)$ such that $\phi_I(h(x_1, \dots, x_{|V_I|})) \leq u_I$.

$\Delta := \text{Sampling}_\Gamma(|V_I|)$;

$\text{BLP}(I, \Delta)$;

if $\text{BLP}(I, \Delta) \leq u_I$ **then**

 accept;

else

 reject;

end

Note that Algorithm 1 runs in polynomial time in $|V_I|$, and that if it rejects, then indeed the answer to $\text{VCSP}(\Gamma)$ is no, without further assumptions.

In the following we present a sufficient condition under which Algorithm 1 correctly solves $\text{VCSP}(\Gamma)$. Valued structures over a finite domain that can be solved by the BLP relaxation have been characterised by Kolmogorov, Thapper, and Živný.

Theorem 6.6 ([16], Theorem 1). *Let Δ be a valued structure with finite signature and finite domain. Then the following are equivalent:*

1. *the BLP relaxation solves $\text{VCSP}(\Delta)$;*
2. *Δ has fully symmetric fractional polymorphisms of all arities.*

One of the main results of this section, Theorem 6.10, states that if Γ is improved by fully symmetric fractional operations of all arities, then Algorithm 1 correctly solves $\text{VCSP}(\Gamma)$ in polynomial time. Note that there are examples of PLH valued structures which have fully symmetric fractional polymorphisms, but the computed sample does not have fully symmetric fractional polymorphisms (Example 6.12).

Definition 6.7. Let Δ be a valued τ -structure with domain D and let $m \geq 1$. The *multiset-structure* $\mathcal{P}^m(\Delta)$ [20] is the valued structure with domain $\binom{D}{m}$ i.e., the set of multisets of elements from D of size m , and for every k -ary function symbol $f \in \tau$, and $\alpha_1, \dots, \alpha_k \in \binom{D}{m}$ the function $f^{\mathcal{P}^m(\Delta)}$ is defined as follows

$$f^{\mathcal{P}^m(\Delta)}(\alpha_1, \dots, \alpha_k) := \frac{1}{m} \min_{t^1, \dots, t^k \in D^m: \{t^l\} = \alpha_l} \sum_{i=1}^m f^\Delta(t_i^1, \dots, t_i^k).$$

(Here we denote by $\{t^l\}$ the multiset whose elements are the coordinates of t^l .)

Lemma 6.8 ([20], Lemma 2.2). *Let Δ be a valued structure with finite domain, and $m \geq 2$. Then $\mathcal{P}^m(\Delta) \rightarrow_f \Delta$ if, and only if, Δ has an m -ary fully symmetric fractional polymorphism.*

Lemma 6.9. *Let Γ be a valued structure (with finite signature τ), and let Δ be a valued structures with finite domain such that $\Delta \rightarrow_f \Gamma$, and let $m \geq 2$ be an integer. If Γ has an m -ary fully symmetric fractional polymorphism, then $\mathcal{P}^m(\Delta) \rightarrow_f \Gamma$.*

Lemma 6.9 is a generalisation of Lemma 6.8 to valued structures with arbitrary domain. However, while Lemma 6.8 follows directly from the definition of $\mathcal{P}^m(\Delta)$, our proof of Lemma 6.9 is inspired by the proof of [3]. In fact, Lemma 6.9 is a generalisation of [3], Lemma 2.4 to valued structures, and to the weaker assumption that the polymorphisms are fully symmetric rather than totally symmetric.

of Lemma 6.9. Let C and D be, respectively, the domain of Γ and the domain of Δ , respectively. Let χ be a fractional homomorphism from Δ to Γ and let ω be an m -ary fully symmetric fractional polymorphism of Γ . For every $g \in \text{Supp}(\omega) \subseteq \mathcal{O}_C^{(m)}$ and every $h \in \text{Supp}(\chi) \subseteq C^D$ we define

$$(g \circ h): \left(\binom{D}{m} \right) \rightarrow C$$

by setting

$$(g \circ h)(\alpha) = g(h(a^1), \dots, h(a^m))$$

for every $\alpha = \{a^1, \dots, a^m\} \in \left(\binom{D}{m} \right)$. Observe that $(g \circ h)$ is well defined as g is fully symmetric (the order of $h(a^1), \dots, h(a^m)$ does not matter). We define the function $\omega': C^{\left(\binom{D}{m} \right)} \rightarrow \mathbb{Q}_{\geq 0}$ as follows, for every $g \in C^{\left(\binom{D}{m} \right)}$,

$$\omega'(g') = \sum_{g \in \text{Supp}(\omega)} \sum_{h \in \text{Supp}(\chi): g \circ h = g'} \omega(g) \chi(h).$$

We claim that ω' is a fractional homomorphism from $\mathcal{P}^m(\Delta)$ to Γ . Indeed, the support $\text{Supp}(\omega') = \{(g \circ h) \mid g \in \text{Supp}(\omega), h \in \text{Supp}(\chi)\}$, is finite as the support of ω and the support of χ are finite. It also holds that

$$\sum_{g' \in \text{Supp}(\omega')} \omega'(g') = \sum_{g \in \text{Supp}(\omega)} \omega(g) \sum_{h \in \text{Supp}(\chi)} \chi(h) = \sum_{g \in \text{Supp}(\omega)} \omega(g) = 1.$$

Furthermore, for every k -ary $f \in \tau$ and tuple $(\alpha_1, \dots, \alpha_k) \in \left(\binom{D}{m} \right)^k$, with

$\alpha_i = \{\alpha_i^1, \dots, \alpha_i^m\}$, it holds that

$$\begin{aligned}
& \sum_{g' \in C\left(\binom{D}{m}\right)} \omega'(g') f^\Gamma(g'(\alpha_1, \dots, \alpha_k)) \\
&= \sum_{g \in \text{Supp}(\omega)} \sum_{h \in \text{Supp}(\chi)} \omega(g) \chi(h) f^\Gamma(g(h(\alpha_1^1), \dots, h(\alpha_1^m)), \dots, g(h(\alpha_k^1), \dots, h(\alpha_k^m))) \\
&\leq \sum_{h \in \text{Supp}(\chi)} \chi(h) \left(\frac{1}{m} \sum_{i=1}^m f^\Gamma\left(h(\alpha_1^{\pi_1(i)}), \dots, h(\alpha_1^{\pi_1(i)})\right) \right) \tag{3}
\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{m} \sum_{i=1}^m \sum_{h \in \text{Supp}(\chi)} \chi(h) f^\Gamma\left(h(\alpha_1^{\pi_1(i)}), \dots, h(\alpha_1^{\pi_1(i)})\right) \\
&\leq \frac{1}{m} \sum_{i=1}^m f^\Delta(\alpha_1^{\pi_1(i)}, \dots, \alpha_k^{\pi_k(i)}) \tag{4}
\end{aligned}$$

for every $\pi_1, \dots, \pi_k \in S_m$. Inequality (3) holds because ω is a fully symmetric fractional polymorphism of Γ , and Inequality (4) holds because χ is a fractional homomorphism from Δ to Γ . Then, in particular it holds that

$$\begin{aligned}
& \sum_{g' \in C\left(\binom{D}{m}\right)} \omega'(g') f^\Gamma(g'(\alpha_1, \dots, \alpha_k)) \leq \frac{1}{m} \min_{t^1, \dots, t^k \in D^m: \{t^l\} = \alpha_l} \sum_{i=1}^m f^\Delta(t_i^1, \dots, t_i^k) \\
&= f^{\mathcal{P}^m(\Delta)}(\alpha_1, \dots, \alpha_k).
\end{aligned}$$

□

Theorem 6.10. *Let Γ be a valued structure with finite signature having fully symmetric fractional polymorphisms of all arities. If there exists an efficient sampling algorithm for Γ , then Algorithm 1 correctly solves $\text{VCSP}(\Gamma)$ (in polynomial time).*

Proof. Let I be an instance of $\text{VCSP}(\Gamma)$ with variables $V_I = \{x_1, \dots, x_n\}$, objective function $\phi_I(x_1, \dots, x_n) = \sum_{j \in J} \gamma_j(x^j)$ where J is a finite set of indices, $\gamma_j \in \Gamma$, and $x^j \in V_I^{\text{ar}(f_j)}$, and threshold u_I . Let Δ be the finite-domain valued structure computed by the sampling algorithm for Γ on input $|V_I|$.

Let C be the (possibly infinite) domain of Γ and D the finite domain of Δ . Note that if $\text{BLP}(I, \Delta) \not\leq u_I$ (this also includes the case $\text{BLP}(I, \Delta) = +\infty$, i.e., the case that I is not feasible) then $\inf_{d \in D^n} \phi_I(d) \not\leq u_I$ which implies $\inf_{c \in C^n} \phi_I(c) \not\leq u_I$ since D was produced by the sampling algorithm for Γ on input $|V_I|$. We may therefore safely reject. Otherwise, if $\text{BLP}(I, \Delta) \leq u_I$, then $\inf_{\alpha \in \left(\binom{D}{m}\right)^n} \phi_I(\alpha) \leq \text{BLP}(I, \Delta)$. The proof of this last statement is contained in the first part of the proof of Theorem 3.2 in [20]; we report it

here for completeness. Let (λ^*, μ^*) be an optimal solution to $\text{BLP}(I, \Delta)$ and let M be a positive integer such that $M \cdot \lambda^*$, and $M \cdot \mu^*$ are both integral. Let $\nu: V_I \rightarrow \binom{D}{M}$ be defined by mapping the variable x_i to the multiset in which the elements are distributed accordingly to $\mu_{x_i}^*$, i.e., for every $a \in D$ the number of occurrences of a in $\nu(x_i)$ is equal to $M\mu_{x_i}^*(a)$. Let f_j be a k -ary function symbol in τ that occurs in a term $f_j(x^j)$ of the objective function ϕ_I . Now we write

$$M \cdot \sum_{t \in D^k} \lambda_j^*(t) f_j^\Delta(t) = f_j^\Delta(\alpha^1) + \dots + f_j^\Delta(\alpha^M),$$

where the $\alpha^i \in D^k$ are such that $\lambda_j^*(t)$ -fractions are equal to t . Let us define $\alpha'_i := (\alpha_i^1, \dots, \alpha_i^M)$ for $1 \leq i \leq k$. We get

$$\begin{aligned} \sum_{t \in D^k} \lambda_j^*(t) f_j^\Delta(t) &= \frac{1}{M} \sum_{i=1}^M f_j^\Delta(\alpha^i) = \frac{1}{M} \sum_{i=1}^M f_j^\Delta(\alpha_1^i, \dots, \alpha_k^i) \\ &\geq \frac{1}{M} \min_{t^1, \dots, t^k \in D^M: \{t^l\} = \{\alpha'_l\}} \sum_{i=1}^M f_j^\Delta(t_i^1, \dots, t_i^k) = f_j^{\mathcal{P}^M(\Delta)}(\alpha'_1, \dots, \alpha'_k) \\ &= f_j^{\mathcal{P}^M(\Delta)}(\nu(x)), \end{aligned}$$

where the last equality follows as the number of a 's in α'_i is

$$M \cdot \sum_{t \in D^k: t^i = a} \lambda_j^*(t) = M \cdot \mu_{x_i}^*(a).$$

Then

$$\begin{aligned} \text{BLP}(I, \Delta) &= \sum_{j \in J} \sum_{t \in D^{\text{ar}(f_j)}} \lambda_j^*(t) f_j^\Delta(t) = \sum_{j \in J} \left(\sum_{t \in D^{\text{ar}(f_j)}} \lambda_j^*(t) f_j^\Delta(t) \right) \\ &\geq \sum_{j \in J} \left(f_j^{\mathcal{P}^M(\Delta)}(\nu(x)) \right) \geq \inf_{\alpha \in \left(\binom{D}{M}\right)^n} \phi_I(\alpha). \end{aligned}$$

Since we assumed $\text{BLP}(I, \Delta) \leq u_I$, we obtain $\inf_{\alpha \in \left(\binom{D}{M}\right)^n} \phi_I(\alpha) \leq u_I$. Moreover, since Γ has fully symmetric fractional polymorphisms of all arities, Lemma 6.9 implies the existence of a fractional homomorphism $\omega: C^{\left(\binom{D}{M}\right)} \rightarrow \mathbb{Q}_{\geq 0}$. From Proposition 6.2 it follows that

$$\inf_{c \in C^n} \phi_I(c) \leq \inf_{\alpha \in \left(\binom{D}{M}\right)^n} \phi_I(\alpha) \leq \text{BLP}(I, \Delta) \leq u_I.$$

□

We now give examples of valued structures that satisfy the hypothesis of Theorem 6.10. A set $S \subseteq \mathbb{Q}^n$ is said to be *convex* if for any two points $x, y \in S$ every point between them is still in S , i.e., for any $\lambda \in [0, 1] \cap \mathbb{Q}$, $\lambda x + (1 - \lambda)y \in S$. A function $f: \mathbb{Q}^n \rightarrow \mathbb{Q} \cup \{+\infty\}$ is said to be *convex* if for any two points $x, y \in \mathbb{Q}^n$ and for any $\lambda \in [0, 1] \cap \mathbb{Q}$, it holds

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y).$$

Proposition 6.11. *Let $D \subseteq \mathbb{Q}$ be a convex set. Let Γ be a valued structure with domain D such that every cost function in Γ is convex. Then, for every $k \geq 2$, the valued structure Γ is improved by the fully symmetric fractional operation $\omega_{conv}^{(k)}: \mathcal{O}_D^{(k)} \rightarrow \mathbb{Q}_{\geq 0}$ (see Example 6.5) such that*

$$\omega_{conv}^{(k)}(g) = \begin{cases} 1 & \text{if } g = \text{avg}^{(k)} \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Let $f: D^n \rightarrow \mathbb{Q} \cup \{+\infty\}$ be a convex cost function. Jensen's Inequality (cf. [13]) implies that for all $k \geq 2$ and for all $x^1, \dots, x^k \in D^n$

$$f\left(\frac{1}{k} \sum_{i=1}^k x^i\right) \leq \frac{1}{k} \sum_{i=1}^k f(x^i).$$

Therefore, for every $k \geq 2$, the function f is improved by the fully symmetric fractional operation $\omega_{conv}^{(k)}$. \square

Example 6.12. Let Γ be a convex valued structure with domain \mathbb{Q} that admits an efficient sampling algorithm. Then, for all $k \geq 2$, the valued structure Γ is improved by the k -ary fully symmetric fractional operation $\omega_{conv}^{(k)}$ defined in Example 6.5 (see Proposition 6.11), and therefore, by Theorem 6.10, Algorithm 1 solves VCSP(Γ). Note that the finite-domain valued structure computed by the sampling algorithm might not have fully symmetric fractional polymorphisms of all arities. In fact, the fractional polymorphisms $\omega_{conv}^{(k)}$ are not even inherited by valued finite substructures of Γ whose domain contains more than one element. \triangle

If the VCSP for a valued structure is solvable by Algorithm 1, then it is possible to find a solution in polynomial time, by applying the self-reduction algorithm (see [16]) to the instance I of VCSP(Δ). Observe that the assignment obtained by self-reduction has values in D and therefore it is also an assignment with values in C .

Proposition 6.13 ([16], Proposition 8). *Let Γ be an arbitrary valued structure with finite domain. Let $I := (V_I, \phi_I, u_I)$ be an instance of VCSP(Γ). If the BLP relaxation solves VCSP(Γ), then an assignment to the variables in V_I with cost at most u_I can be found in polynomial time.*

We remark that Theorem 6.10 generalises the following known result (that we used in Section 4 to prove the polynomial-time complexity of the CSP for PLH relational structures preserved by \max).

Theorem 4.14 ([3], Theorem 2.5). *Let \mathfrak{A} be a structure over a finite relational signature with totally symmetric polymorphisms of all arities. If there exists an efficient sampling algorithm for \mathfrak{A} then $\text{CSP}(\mathfrak{A})$ is in P .*

More precisely, we extended Theorem 4.14 to valued structures and, at the same time, to the weaker assumption of having fully symmetric polymorphisms of all arities rather than totally symmetric polymorphisms of all arities. The following example is adapted from [18] (Example 99), and it exhibits a PL valued structure having fully symmetric polymorphisms of all arities but having no totally symmetric polymorphism of arity 3.

Example 6.14. Let us consider the PL valued structure Γ with domain \mathbb{Q} , signature $\{f_+, f_-\}$, and such that $f_+^\Gamma, f_-^\Gamma: \mathbb{Q}^3 \rightarrow \mathbb{Q} \cup \{+\infty\}$ are defined by

$$f_+^\Gamma(x_1, x_2, x_3) := \begin{cases} x_1 + x_2 + x_3 & \text{if } x_1 + x_2 + x_3 \geq 1 \\ +\infty & \text{otherwise,} \end{cases}$$

and

$$f_-^\Gamma(x_1, x_2, x_3) := \begin{cases} x_1 + x_2 + x_3 & \text{if } x_1 + x_2 + x_3 \leq -1 \\ +\infty & \text{otherwise.} \end{cases}$$

Clearly, the cost functions f_+^Γ , and f_-^Γ are PL and it is easy to see that they are convex. As all cost functions in Γ are PL and convex, by Proposition 6.11, the valued structure Γ is improved by the fully symmetric fractional operations $\omega_{\text{conv}}^{(k)}$, for every $k \geq 2$, i.e., Γ has fully symmetric fractional polymorphisms of all arities. We already observed that the fractional operations $\omega_{\text{conv}}^{(k)}$ are not totally symmetric for $k \geq 3$ (see Example 6.5).

Assume now that ω is a ternary totally symmetric fractional polymorphism of Γ and let $t: \mathbb{Q}^3 \rightarrow \mathbb{Q}$ be a totally symmetric operation in $\text{Supp}(\omega)$, then, in particular, t is a polymorphism of $\text{Feas}(\Gamma)$, i.e., t preserves the feasibility relations

$$\begin{aligned} \text{dom}(f_+^\Gamma) &= \{(x_1, x_2, x_3) \in \mathbb{Q}^3 \mid x_1 + x_2 + x_3 \geq 1\}, \text{ and} \\ \text{dom}(f_-^\Gamma) &= \{(x_1, x_2, x_3) \in \mathbb{Q}^3 \mid x_1 + x_2 + x_3 \leq -1\}. \end{aligned}$$

By the total symmetry of t , there exists $a \in \mathbb{Q}$ such that

$$\begin{aligned} t(1, 1, -1) &= t(1, -1, 1) = t(-1, 1, 1) \\ &= t(-1, -1, 1) = t(-1, 1, -1) = t(1, -1, -1) = a. \end{aligned}$$

Observe that $(1, 1, -1), (1, -1, 1), (-1, 1, 1) \in \text{dom}(f_+^\Gamma)$, then, by applying t componentwise we get $(a, a, a) \in \text{dom}(f_+^\Gamma)$, i.e., $a \geq \frac{1}{3}$; on the other hand, $(-1, -1, 1), (-1, 1, -1), (1, -1, -1) \in \text{dom}(f_-^\Gamma)$, then, by applying t componentwise we get $(a, a, a) \in \text{dom}(f_-^\Gamma)$, i.e., $a \leq -\frac{1}{3}$, that is a contradiction. \triangle

In the next section we will apply Theorem 6.10 to PLH valued structures.

7 Polynomial-Time PLH VCSPs

We apply the results of Sections 5 and 6 to state the polynomial-time tractability of the VCSP for PLH valued structures with finite signature that are improved by fully symmetric operations of all arities.

Proposition 7.1. *Let Γ be a valued PLH constraint language with finite signature that is improved by fully symmetric fractional polymorphisms of all arities. Then $\text{VCSP}(\Gamma)$ can be solved in polynomial time.*

Proof. The statement is an immediate consequence of Theorem 6.10 and the fact that PLH valued structures with finite signature admit an efficient sampling algorithm (Corollary 5.12). \square

7.1 Submodular PLH Valued Structures

We have already mentioned the importance of submodularity in the introduction to this article. In this section we apply Proposition 7.1 to prove that the VCSP of submodular PLH valued structures with finite signature can be solved in polynomial time. Moreover, we show that submodularity defines a *maximally tractable* class of valued structures within the class of PLH valued structures.

Corollary 7.2. *Let Γ be a PLH valued structure with finite signature such that the cost functions in Γ are submodular. Then $\text{VCSP}(\Gamma)$ can be solved in polynomial time.*

We already observed in Section 7.1 that a function over a totally ordered set D is submodular if, and only if, it is improved by the binary fully symmetric fractional operation $\omega_{\text{sub}}: \mathcal{O}_D^2 \rightarrow \mathbb{Q}_{\geq 0}$ such that

$$\omega_{\text{sub}}(g) = \begin{cases} \frac{1}{2} & \text{if } g = \max \\ \frac{1}{2} & \text{if } g = \min \\ 0 & \text{otherwise.} \end{cases}$$

In fact, there is another equivalent characterisation of submodularity based on fractional polymorphisms. For every $k \geq 2$ and every $i \in \{1, \dots, k\}$ we define $s_i^{(k)}: D^k \rightarrow D$ to be the operation returning the i -th smallest of its arguments with respect to the total order in D . Observe that for $k = 2$ we get $s_1^{(2)} = \min$ and $s_2^{(2)} = \max$. We define for every $k \geq 2$ the k -ary fractional operation $\omega_{\text{sub}}^{(k)}: \mathcal{O}_D^{(k)} \rightarrow \mathbb{Q}_{\geq 0}$ having support

$$\text{Supp}(\omega_{\text{sub}}^{(k)}) = \{s_i^{(k)} \mid 1 \leq i \leq k\}$$

by setting

$$\omega_{\text{sub}}^{(k)}(g) := \begin{cases} \frac{1}{k} & \text{if } g \in \text{Supp}(\omega_{\text{sub}}^{(k)}) \\ 0 & \text{otherwise.} \end{cases}$$

The operations $s_i^{(k)}(x^1, \dots, x^k)$ are fully symmetric for all $i \in \{1, \dots, k\}$ and all $k \geq 2$, therefore the fractional operations $\omega_{\text{sub}}^{(k)}$ are fully symmetric for all $k \geq 2$. It is an easy observation that, for $k = 2$, the fractional operation $\omega_{\text{sub}}^{(2)}$ is exactly the fractional operation ω_{sub} characterising submodular functions.

Proposition 7.3. *Let D be a totally ordered set and let f be a submodular function over D . Then for all $k \geq 2$, the fractional operation $\omega_{\text{sub}}^{(k)}$ improves the function f .*

Proof. Clearly,

$$\sum_{g \in \text{Supp}(\omega_{\text{sub}}^{(k)})} \omega_{\text{sub}}^{(k)}(g) = 1.$$

We want to prove that for all $k \geq 2$ and for all $x^1, \dots, x^k \in D^n$ it holds that

$$\frac{1}{k} \sum_{i=1}^k f(s_i^{(k)}(x^1, \dots, x^k)) \leq \frac{1}{k} (f(x^1) + \dots + f(x^k)). \quad (5)$$

By using the submodularity of f we can write

$$\begin{aligned} f(x^1) + \dots + f(x^k) &= \frac{1}{k-1} \sum_{1 \leq i < j \leq k} (f(x^i) + f(x^j)) \\ &\geq \frac{1}{k-1} \sum_{1 \leq i < j \leq k} (f(\min(x^i, x^j)) + f(\max(x^i, x^j))) \\ &= \sum_{1 \leq i < j \leq k} \frac{f(\min(x_1^i, x_1^j), \dots, \min(x_n^i, x_n^j)) + f(\max(x_1^i, x_1^j), \dots, \max(x_n^i, x_n^j))}{k-1} \\ &\geq \sum_{i=1}^k f(s_i^{(k)}(x_1^1, \dots, x_1^k), \dots, s_i^{(k)}(x_n^1, \dots, x_n^k)), \end{aligned} \quad (6)$$

from which Inequality (5) follows. We prove Inequality (6) by induction on n . Observe that for every coordinate $1 \leq l \leq n$, the following equality between multisets holds:

$$\begin{aligned} &\left\{ \min(x_l^i, x_l^j) \mid 1 \leq i < j \leq k \right\} \cup \left\{ \max(x_l^i, x_l^j) \mid 1 \leq i < j \leq k \right\} \\ &= \underbrace{\left\{ s_1^{(k)}(x_l^1, \dots, x_l^k) \right\}}_{k-1 \text{ occurrences}} \underbrace{\left\{ s_2^{(k)}(x_l^1, \dots, x_l^k) \right\}}_{k-1 \text{ occurrences}} \dots \underbrace{\left\{ s_k^{(k)}(x_l^1, \dots, x_l^k) \right\}}_{k-1 \text{ occurrences}}. \end{aligned} \quad (7)$$

If f has arity $n = 1$, then Inequality (6) immediately follows from Equality (7). Let $n \geq 2$, assume that Inequality (6) is true for submodular functions

of arity at most $n - 1$, and let us prove it for submodular functions of arity n . From Equality (7) and from the inductive hypothesis it follows that there exist $(k - 1)$ -many permutations $\pi_1, \dots, \pi_{k-1} \in S_k$ such that

$$\begin{aligned} & \sum_{1 \leq i < j \leq k} \left(f(\min(x_1^i, x_1^j), \dots, \min(x_n^i, x_n^j)) + f(\max(x_1^i, x_1^j), \dots, \max(x_n^i, x_n^j)) \right) \\ & \geq \sum_{p=1}^{k-1} \sum_{i=1}^k f \left(s_i^{(k)}(x_1^1, \dots, x_1^k), \dots, s_i^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_{\pi_p(i)}^{(k)}(x_n^1, \dots, x_n^k) \right). \end{aligned} \quad (8)$$

We claim that for every $p \in \{1, \dots, k - 1\}$ it holds that

$$\begin{aligned} & \sum_{i=1}^k f \left(s_i^{(k)}(x_1^1, \dots, x_1^k), \dots, s_i^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_{\pi_p(i)}^{(k)}(x_n^1, \dots, x_n^k) \right) \\ & \geq \sum_{i=1}^k f \left(s_i^{(k)}(x_1^1, \dots, x_1^k), \dots, s_i^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_i^{(k)}(x_n^1, \dots, x_n^k) \right). \end{aligned} \quad (9)$$

To prove Inequality (9), let $j := \max \{i \in \{1, \dots, k\} \mid \pi_p(i) \neq i\}$. Then there exists $l \in \{1, \dots, j - 1\}$ such that $\pi_p(l) = j$. By the submodularity of f we have that

$$\begin{aligned} & f \left(s_j^{(k)}(x_1^1, \dots, x_1^k), \dots, s_j^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_{\pi_p(j)}^{(k)}(x_n^1, \dots, x_n^k) \right) \\ & + f \left(s_l^{(k)}(x_1^1, \dots, x_1^k), \dots, s_l^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_j^{(k)}(x_n^1, \dots, x_n^k) \right) \\ & \geq f \left(s_j^{(k)}(x_1^1, \dots, x_1^k), \dots, s_j^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_j^{(k)}(x_n^1, \dots, x_n^k) \right) \\ & + f \left(s_l^{(k)}(x_1^1, \dots, x_1^k), \dots, s_l^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_{\pi_p(j)}^{(k)}(x_n^1, \dots, x_n^k) \right). \end{aligned}$$

After this step

$$\begin{aligned} & \sum_{i=1}^k f \left(s_i^{(k)}(x_1^1, \dots, x_1^k), \dots, s_i^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_{\pi_p(i)}^{(k)}(x_n^1, \dots, x_n^k) \right) \\ & \geq \sum_{i=j}^k f \left(s_i^{(k)}(x_1^1, \dots, x_1^k), \dots, s_i^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_i^{(k)}(x_n^1, \dots, x_n^k) \right) \\ & \quad + \sum_{i=1}^{j-1} f \left(s_i^{(k)}(x_1^1, \dots, x_1^k), \dots, s_i^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_{\pi'_p(i)}^{(k)}(x_n^1, \dots, x_n^k) \right) \end{aligned}$$

where $\pi'_p \in S_k$ is the permutation defined by

$$\pi'_p(i) = \begin{cases} \pi_p(j) & \text{if } i = l \\ \pi_p(i) & \text{otherwise.} \end{cases}$$

By reiterating the procedure for at most $j - 1 \leq k$ times for every $p \in \{1, \dots, k - 1\}$, we get the claim. Then, using Inequality (9), we can rewrite Inequality (8) as

$$\begin{aligned}
& \sum_{1 \leq i < j \leq k} \left(f(\min(x_1^i, x_1^j), \dots, \min(x_n^i, x_n^j)) + f(\max(x_1^i, x_1^j), \dots, \max(x_n^i, x_n^j)) \right) \\
& \geq \sum_{p=1}^{k-1} \sum_{i=1}^k f \left(s_i^{(k)}(x_1^1, \dots, x_1^k), \dots, s_i^{(k)}(x_{n-1}^1, \dots, x_{n-1}^k), s_i^{(k)}(x_n^1, \dots, x_n^k) \right) \\
& = \sum_{p=1}^{k-1} \sum_{i=1}^k f \left(s_i^{(k)}(x^1, \dots, x^k) \right) = (k-1) \sum_{i=1}^k f \left(s_i^{(k)}(x^1, \dots, x^k) \right),
\end{aligned}$$

that is, Inequality (6) holds and this concludes the proof. \square

The next corollary immediately follows from the full symmetry of fractional operations $\omega_{\text{sub}}^{(k)}$ and Proposition 7.3.

Corollary 7.4. *Let Γ be a valued PL (or PLH) submodular constraint language. Then Γ has fully symmetric fractional polymorphisms of all arities.*

of Corollary 7.2. The proof follows directly from Proposition 7.1, since every valued submodular PLH constraint language has fully symmetric fractional polymorphisms of all arities (Corollary 7.4). \square

We show that submodular PLH valued structures are *maximal tractable* within the class of PLH valued structures. Let Γ be a valued structure with signature τ . A *valued reduct* Γ' of Γ is a valued structure with domain $\text{dom}(\Gamma)$ and such that the signature τ' can be obtained from τ by dropping some of the function symbols. The interpretation of function symbols from τ' in Γ' is as in Γ . In this case we also say that Γ is an *expansion* of Δ . A valued reduct of a valued structure is said *finite* if its signature is finite.

Definition 7.5. Let \mathcal{V} be a class of valued structures with fixed domain D and let Γ be a valued structure in \mathcal{V} . We say that Γ is *maximal tractable within \mathcal{V}* if

- $\text{VCSP}(\Gamma')$ is polynomial time solvable for every valued finite reduct Γ' of Γ ; and
- for every valued structure Δ in \mathcal{V} that is an expansion of Γ , there exists a valued finite reduct Δ' of Δ such that $\text{VCSP}(\Delta')$ is NP-hard.

We will make use of the following result.

Theorem 7.6 ([8], Theorem 6.7). *Let D be a finite totally ordered set. Then the valued structure containing all submodular cost functions over D is maximal tractable within the class of all valued structures with domain D .*

We show that the class of submodular PLH valued structures is maximal tractable within the class of PLH valued structures.

Definition 7.7. Given a finite domain $D \subset \mathbb{Q}$ and a partial function $f: D^n \rightarrow \mathbb{Q}$ we define the *canonical extension* of f as $\hat{f}: \mathbb{Q}^n \rightarrow \mathbb{Q}$, by

$$\hat{f}(x) = \begin{cases} f(x) & x \in D^n \\ +\infty & \text{otherwise.} \end{cases}$$

Note that the canonical extension of a submodular function over a finite domain is submodular and PLH.

Theorem 7.8. *The valued structure containing all submodular PLH cost functions is maximal tractable within the class of PLH valued structures.*

Proof. Polynomial-time tractability of the VCSP for submodular PLH valued structures with finite signature has been stated in Theorem 7.2.

Let $f^{\mathbb{Q}}$ be an m -ary PLH cost function over \mathbb{Q} that is not submodular, i.e., there exists a couple of points, $a := (a_1, \dots, a_m)$, $b := (b_1, \dots, b_m) \in \mathbb{Q}^m$ such that

$$f^{\mathbb{Q}}(a) + f^{\mathbb{Q}}(b) < f^{\mathbb{Q}}(\min(a, b)) + f^{\mathbb{Q}}(\max(a, b)).$$

Let

$$D := \{a_1, \dots, a_m, b_1, \dots, b_m\} \subset \mathbb{Q},$$

and let Δ be the valued structure with domain D and such that its signature τ contains a function symbol for every submodular cost functions on D . Notice that the restriction $f^{\mathbb{Q}}|_D$ is not submodular, for our choice of D . Therefore, by Theorem 7.6, there exists a valued structure Δ' having domain D and signature $\tau' \cup \{f\}$, where $\tau' \subseteq \tau$ is finite and the cost function $f^{\Delta'} = f^{\mathbb{Q}}|_D$, such that $\text{VCSP}(\Delta')$ NP-hard.

We define Γ' to be the (submodular PLH) valued structure with domain D , signature $\tau' \cup \{f, \chi_D\}$, and such that the interpretation of functions symbols in the signature is as it follows:

- for every $g \in \tau'$, the cost function $g^{\Gamma'}$ is the canonical extension of $g^{\Delta'}$,
- the cost function $f^{\Gamma'}$ is $f^{\mathbb{Q}}$, and
- the unary cost function $\chi_D^{\Gamma'}: \mathbb{Q} \rightarrow \mathbb{Q} \cup \{+\infty\}$ is defined, for every $x \in \mathbb{Q}$ as

$$\chi_D^{\Gamma'}(x) = \begin{cases} 0 & \text{if } x \in D \\ +\infty & \text{if } x \in \mathbb{Q} \setminus D. \end{cases}$$

Observe that $\chi_D^{\Gamma'}$ is submodular and PLH.

The valued structure Γ' is an extension of an expansion (or an expansion of an extension) of Δ' . We claim that $\text{VCSP}(\Gamma')$ is NP-hard. Indeed, for every instance $I := (V, \phi_I, u)$ of $\text{VCSP}(\Delta')$, with $V := \{v_1, \dots, v_n\}$, we define the instance $J := (V, \phi_J, u)$ of $\text{VCSP}(\Gamma')$ such that

$$\phi_J(v_1, \dots, v_n) := \phi_I(v_1, \dots, v_n) + \sum_{i=1}^n \chi_D(v_i).$$

Because of the terms involving χ_D , an assignment $h: V \rightarrow \mathbb{Q}$ is such that $\phi_J^{\Gamma'}(h(v_1), \dots, h(v_n))$ is smaller than $+\infty$ if, and only if, $h(v_i) \in D$ for all $v_i \in V$. In this case,

$$\phi_J^{\Gamma'}(h(v_1), \dots, h(v_n)) = \phi_I^{\Delta'}(h(v_1), \dots, h(v_n))$$

and therefore, deciding whether there exists an assignment $h: V \rightarrow \mathbb{Q}$ such that

$$\phi_J^{\Gamma'}(h(v_1), \dots, h(v_n)) \leq u$$

is equivalent to decide whether there exists an assignment $h': V \rightarrow D$ such that

$$\phi_I^{\Delta'}(h'(v_1), \dots, h'(v_n)) \leq u.$$

Since J is computable in polynomial-time from I , the NP-hardness of the problem $\text{VCSP}(\Gamma')$ follows from the NP-hardness of $\text{VCSP}(\Delta')$. \square

7.2 Convex PLH Valued Structures

We have already seen that if Γ is a valued structure such that all cost functions in Γ are convex, then Γ has fully symmetric fractional polymorphisms of all arities (see Proposition 6.11). Therefore, the next corollary follows directly from Proposition 6.11 and Corollary 7.1.

Corollary 7.9. *Let Γ be a PLH valued structure with finite signature such that the cost functions in Γ are convex. Then $\text{VCSP}(\Gamma)$ can be solved in polynomial time.*

7.3 Componentwise Increasing PLH Valued Structures

Let $f: \mathbb{Q}^n \rightarrow \mathbb{Q} \cup \{+\infty\}$ be an n -ary function. We say that

- f is *componentwise increasing* if

$$f(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_n) \leq f(x_1, \dots, x_{i-1}, z_i, x_{i+1}, \dots, x_n),$$

for every $y_i < z_i$, $1 \leq i \leq k$, and $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n \in \mathbb{Q}$;

- f is *componentwise decreasing* if

$$f(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_n) \geq f(x_1, \dots, x_{i-1}, z_i, x_{i+1}, \dots, x_n),$$

for every $y_i < z_i$, $1 \leq i \leq k$, and $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n \in \mathbb{Q}$;

In [8], componentwise increasing functions, and componentwise decreasing functions are respectively referred to as *monotone* functions and *antitone* functions.

Let us define, for every $k \geq 2$, the fractional operations $\omega_{\min}^{(k)}: \mathcal{O}_D^{(k)} \rightarrow \mathbb{Q}_{\geq 0}$, and $\omega_{\max}^{(k)}: \mathcal{O}_D^{(k)} \rightarrow \mathbb{Q}_{\geq 0}$ by setting, respectively,

$$\omega_{\min}^{(k)}(g) := \begin{cases} 1 & \text{if } g = \min^{(k)} \\ 0 & \text{otherwise,} \end{cases}$$

and $\omega_{\max}^{(k)}(g)$ is dually defined. The fractional operations $\omega_{\min}^{(k)}$ and $\omega_{\max}^{(k)}$ are fully symmetric (in fact they are totally symmetric).

A valued structure Γ is said to be *componentwise increasing* (*componentwise decreasing*, respectively) if all the cost functions in Γ are componentwise increasing (componentwise decreasing, respectively).

Corollary 7.10. *Let Γ be a componentwise increasing PLH valued structure with finite signature. Then $\text{VCSP}(\Gamma)$ can be solved in polynomial time.*

Proof. The proof follows from Proposition 7.1, since for every $k \geq 2$ the fully symmetric fractional operation $\omega_{\min}^{(k)}$ improves Γ as the next lemma shows. \square

Lemma 7.11. *A function $f: \mathbb{Q}^n \rightarrow \mathbb{Q} \cup \{+\infty\}$ is componentwise increasing if, and only if, f is improved by $\omega_{\min}^{(k)}$, for every $k \geq 2$.*

Proof. Let $k \geq 2$ and let us first consider f componentwise increasing with arity 1. Let us consider $x_1, \dots, x_k \in \mathbb{Q}$ and let us assume without loss of generality that $x_1 = \min(x_1, \dots, x_k)$. Then we have that

$$f(\min(x_1, \dots, x_k)) = f(x_1) = \frac{\overbrace{f(x_1) + \dots + f(x_1)}^{k \text{ times}}}{k} \leq \frac{f(x_1) + \dots + f(x_k)}{k},$$

i.e., $\omega_{\min}^{(k)}$ is a fractional polymorphism of f . Assume now that, for $n \geq 2$, every $(n-1)$ -ary componentwise increasing function is improved by $\omega_{\min}^{(k)}$, and let us prove that for every $n \in \mathbb{N}$, an n -ary componentwise increasing function f is improved by $\omega_{\min}^{(k)}$. Let us fix $x^1, \dots, x^k \in \mathbb{Q}^n$. The restricted function

$$f(\cdot, \min(x_n^1, \dots, x_n^k)): \mathbb{Q}^{n-1} \rightarrow \mathbb{Q} \cup \{+\infty\},$$

which maps every $(z_1, \dots, z_{n-1}) \in \mathbb{Q}^{n-1}$ to $f(z_1, \dots, z_{n-1}, \min(x_n^1, \dots, x_n^k))$, is clearly componentwise increasing and therefore, by the inductive hypothesis, it is improved by $\omega_{\min}^{(k)}$, that is,

$$\begin{aligned} & f(\min(x_1^1, \dots, x_1^k), \dots, \min(x_{n-1}^1, \dots, x_{n-1}^k), \min(x_n^1, \dots, x_n^k)) \\ & \leq \frac{f(x_1^1, \dots, x_{n-1}^1, \min(x_n^1, \dots, x_n^k)) + \dots + f(x_1^k, \dots, x_{n-1}^k, \min(x_n^1, \dots, x_n^k))}{k}. \end{aligned}$$

Again by the fact that f is componentwise increasing we get that the right-hand side of the last inequality is

$$\leq \frac{1}{k}(f(x_1^1, \dots, x_{n-1}^1, x_n^1) + \dots + f(x_1^k, \dots, x_{n-1}^k, x_n^k)),$$

i.e., f is improved by $\omega_{\min}^{(k)}$.

Conversely, if $f: \mathbb{Q}^n \rightarrow \mathbb{Q} \cup \{+\infty\}$ is improved by $\omega_{\min}^{(k)}$ for all $k \geq 2$, then in particular it is improved by $\omega_{\min} = \omega_{\min}^{(2)}$. For all $i \in \{1, \dots, n\}$, and $x_1, \dots, x_{i-1}, y_i, z_i, x_{i+1}, \dots, x_n \in \mathbb{Q}$ such that $y_i < z_i$ it holds that

$$\begin{aligned} & f(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_n) \\ & \leq \frac{f(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_n)}{2} + \frac{f(x_1, \dots, x_{i-1}, z_i, x_{i+1}, \dots, x_n)}{2}. \end{aligned}$$

It follows that

$$f(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_n) \leq f(x_1, \dots, x_{i-1}, z_i, x_{i+1}, \dots, x_n),$$

i.e., f is componentwise increasing. \square

Lemma 7.11 and Corollary 7.10 may be shown to have a dual form that holds in the case of componentwise decreasing cost functions.

The valued structure containing all componentwise decreasing PLH functions (all componentwise increasing PLH cost functions, respectively) is maximal tractable within the class of PLH valued structures. The proof is similar to the proof of the maximal tractability of submodular PLH valued structures (Theorem 7.8) and uses a result by Cohen, Cooper, Jeavons, and Krokhin stating the maximal tractability of componentwise decreasing finite-domain valued structures (see [8], Theorem 6.15).

8 Conclusion and Outlook

We have presented two main results: an efficient sampling algorithm for PLH valued structures with finite signature, and an algebraic condition for the VCSP of infinite-domain valued structures that admit an efficient sampling algorithm to be solved in polynomial time by the BLP relaxation. By combining these two results we obtained a polynomial-time algorithm for PLH

valued structures over the rationals that are improved by fully symmetric fractional operations of all arities. Finally, we showed that this polynomial-time algorithm solves the VCSP for PLH cost functions that are submodular, convex, or piecewise increasing. In fact, our algorithm not only decides whether there exists a solution of cost at most u_I , but it can also be adapted to efficiently compute the infimum of the cost of all solutions (which might be $-\infty$), and decides whether the infimum is attained. The modification is straightforward observing that the computed sample does not depend on the threshold u_I .

We also showed that submodular PLH cost functions are *maximally tractable* within the class of PLH cost functions. Such maximal tractability results are of particular importance for the more ambitious goal of classifying the complexity of the VCSP for all classes of PLH cost functions: to prove a complexity dichotomy it suffices to identify *all* maximally tractable classes. In the same direction, the study of the algebraic properties that make a PLH valued structure NP-hard, namely an algebraic theory of hardness, is an interesting field of future research. Another challenge is to understand which algebraic properties of fractional polymorphisms of an infinite-domain valued structure are necessary the polynomial-time tractability of the VCSP under the assumption that the valued structure admits an efficient sampling algorithm. Finally, we would like to extend our tractability result to the class of all submodular *piecewise linear* VCSPs. We believe that submodular piecewise linear VCSPs are in P, too. But note that already the structure $(\mathbb{Q}; 0, S, D)$ where $S := \{(x, y) \mid y = x+1\}$ and $D := \{(x, y) \mid y = 2x\}$ (which has both min and max as polymorphisms) does *not* admit an efficient sampling algorithm (it is easy to see that for every $d \in \mathbb{N}$ every sample computed on input d must have exponentially many vertices in d), so a different approach than the approach in this article is needed.

Acknowledgement

We would like to thank Jakub Opršal for helpful discussions and comments on a draft of this article.

References

- [1] N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Machine Learning*, 56(1-3):89–113, 2004.
- [2] M. Bodirsky and M. Grohe. Non-dichotomies in constraint satisfaction complexity. In L. Aceto, I. Damgård, L. A. Goldberg, M. M. Halldórsson, A. Ingólfssdóttir, and I. Walukiewicz, editors, *Proceedings of the International Colloquium on Automata, Languages and Program-*

- ming*, *ICALP'08*, Lecture Notes in Computer Science, pages 184–196. Springer Verlag, July 2008.
- [3] M. Bodirsky, D. Macpherson, and J. Thapper. Constraint satisfaction tractability from semi-lattice operations on infinite sets. *Transaction of Computational Logic (ACM-TOCL)*, 14(4):1–30, 2013.
 - [4] M. Bodirsky and M. Mamino. Tropically convex constraint satisfaction. *Theory of Computing Systems*, 62(3):481–509, 2018.
 - [5] M. Bodirsky, M. Mamino, and C. Viola. Submodular functions and valued constraint satisfaction problems over infinite domains. In *27th EACSL Annual Conference on Computer Science Logic, CSL'18, September 4-7, 2018, Birmingham, UK*, pages 12:1–12:22, 2018.
 - [6] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
 - [7] A. A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS'17, Berkeley, CA, USA, October 15-17, 2017*, pages 319–330, 2017.
 - [8] 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.
 - [9] J. Ferrante and C. Rackoff. A decision procedure for the first order theory of real addition with order. *SIAM Journal on Computing*, 4(1):69–76, 1975.
 - [10] S. Fujishige. *Submodular Functions and Optimization*. Volume 58 of Annals of Discrete Mathematics. North-Holland, Amsterdam, 2005. 2nd edition.
 - [11] M. Grötschel, L. Lovász, and L. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer, Heidelberg, 1994. 2nd edition.
 - [12] W. Hodges. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.
 - [13] J. L. W. V. Jensen. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30(1):175–193, Dec 1906.
 - [14] P. Jonsson, F. Kuivinen, and J. Thapper. Min CSP on four elements: Moving beyond submodularity. In *Principles and Practice of Constraint Programming - 17th International Conference, CP'11, Perugia, Italy, September 12-16, 2011. Proceedings*, pages 438–453, 2011.

- [15] V. Kolmogorov, A. A. Krokhin, and M. Rolinek. The complexity of general-valued CSPs. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS'15, Berkeley, CA, USA, 17-20 October, 2015*, pages 1246–1258, 2015.
- [16] V. Kolmogorov, J. Thapper, and S. Zivny. The power of linear programming for general-valued CSPs. *SIAM J. Comput.*, 44(1):1–36, 2015.
- [17] M. Kozik and J. Ochremiak. Algebraic properties of valued constraint satisfaction problem. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP'15, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 846–858, 2015.
- [18] G. Kun and M. Szegedy. A new line of attack on the dichotomy conjecture. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing, STOC'09*, pages 725–734, New York, NY, USA, 2009. ACM.
- [19] A. Schrijver. *Theory of Linear and Integer Programming*. Wiley - Interscience Series in Discrete Mathematics and Optimization, 1998.
- [20] J. Thapper and S. Živný. The power of linear programming for valued CSPs. In *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, FOCS'12*, pages 669–678, Washington, DC, USA, 2012. IEEE Computer Society.
- [21] J. Thapper and S. Zivny. The complexity of finite-valued CSPs. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 695–704, 2013.
- [22] L. van den Dries. *Tame topology and o-minimal structures*, volume 248 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1998.
- [23] D. Zhuk. A proof of CSP dichotomy conjecture. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS'17, Berkeley, CA, USA, October 15-17, 2017*, pages 331–342, 2017.