The Complexity of Resilience Problems via Valued Constraint Satisfaction Problems

Manuel Bodirsky
Žaneta Semanišinová
manuel.bodirsky@tu-dresden.de
zaneta.semanisinova@tu-dresden.de
Institut für Algebra, TU Dresden
Dresden, Germany

Carsten Lutz
Institut für Informatik, Leipzig University
Leipzig, Germany
clu@informatik.uni-leipzig.de

ABSTRACT

Valued constraint satisfaction problems (VCSPs) constitute a large class of computational optimisation problems. It was shown recently that, over finite domains, every VCSP is in P or NP-complete, depending on the admitted cost functions. In this article, we study cost functions over countably infinite domains whose automorphisms form an oligomorphic permutation group. Our results include a hardness condition based on a generalisation of pp-constructability as known from classical CSPs and a polynomial-time tractability condition based on the concept of fractional polymorphisms. We then observe that the resilience problem for unions of conjunctive queries (UCQs) studied in database theory, under bag semantics, may be viewed as a special case of the VCSPs that we consider. We obtain a complexity dichotomy for the case of incidence-acyclic UCQs and exemplarily use our methods to determine the complexity of a query that had remained open in the literature. Further, we conjecture that our hardness and tractability conditions match for resilience problems for UCQs.

CCS CONCEPTS

• **Theory of computation → Problems, reductions and completeness**; **Complexity theory and logic**; **Database query processing and optimization (theory)**.

KEYWORDS

valued constraints, conjunctive queries, resilience, oligomorphic automorphism groups, computational complexity, pp-constructions, fractional polymorphisms, polynomial-time tractability

ACM Reference Format:
Manuel Bodirsky, Žaneta Semanišinová, and Carsten Lutz. 2024. The Complexity of Resilience Problems via Valued Constraint Satisfaction Problems. In *39th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '24), July 8–11, 2024, Tallinn, Estonia.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3661814.3662071

LICS '24, July 8–11, 2024, Tallinn, Estonia

ACM ISBN 979-8-4007-0660-8/24/07
https://doi.org/10.1145/3661814.3662071

1 INTRODUCTION

In data management, the *resilience* of a query μ in a relational database $\mathfrak{A}$ is the minimum number of tuples that need to be removed from $\mathfrak{A}$ to achieve that μ is false in $\mathfrak{A}$. This associates every fixed query μ with a natural decision problem: given a database $\mathfrak{A}$ and a candidate resilience $n \in \mathbb{N}$, decide whether the resilience of μ in $\mathfrak{A}$ is at most n. Significant efforts have been invested into classifying the complexity of such resilience problems depending on the query μ, concentrating on the case that μ is a conjunctive query [21, 22, 36]. Notably, research has identified several classes of conjunctive queries for which the resilience problem is in polynomial time and others for which it is NP-complete. A general classification, however, has remained open.

Resilience problems have been considered on set databases [21, 22] and, more recently, also on bag databases [36]. The latter means that every fact in the database is associated with a multiplicity, that is, a tuple of constants may have multiple occurrences in the same relation. Bag databases are of importance because they represent SQL databases more faithfully than set databases [14]. Note that if the resilience problem of a query μ can be solved in polynomial time on bag databases, then it can be solved in polynomial time also on set databases. Regarding the converse, Makhija and Gatterbauer [36] identify a conjunctive query for which the resilience problem on bag databases is NP-hard whereas the resilience problem on set databases is in P.

In this paper, we present a surprising link between the resilience problem for (unions of) conjunctive queries under bag semantics and *valued constraint satisfaction problems (VCSPs)*, which constitute a large class of computational optimisation problems. In a VCSP, we are given a finite set of variables, a finite sum of cost functions on these variables, and a threshold u, and the task is to find an assignment to the variables so that the sum of the costs is at most u. The computational complexity of such problems has been studied depending on the admitted cost functions, which we may view as a *valued structure*. A complete classification has been obtained for valued structures with a finite domain, showing that the corresponding VCSPs are in P or NP-hard [12, 31, 33, 44, 45]. There are also some results about VCSPs of valued structures with infinite domains [7, 43].

We show that the resilience problem for every union of connected conjunctive queries can be formulated as a VCSP for a valued structure with an *oligomorphic automorphism group*, i.e., a structure with a countable domain that, for every fixed k, has only finitely many orbits of k-tuples under the action of the automorphism group. This property is important for classical CSPs (which can be seen as VCSPs where all cost functions take values in $\{0, \infty\}$) since it enables the use and extension of some tools from finite-domain CSPs (see, e.g., [3]). The complexity classification for general, not necessarily connected conjunctive queries can be reduced to the connected case. In the important special case that the conjunctive query is incidence-acyclic (meaning that it has an acyclic incidence graph), we even obtain a VCSP for a valued structure with a finite domain and consequently obtain a P versus NP-complete dichotomy from the known dichotomy for such VCSPs.

The above results actually hold in the more general setting where some relations or tuples may be declared to be *exogenous*, meaning that they may *not* be removed from the database to make the query false. This is useful when the data in the database stems from different sources; it was also considered in [21, 22, 36].

As a main contribution of this paper, the novel connection between resilience and VCSPs leads us to initiating the systematic study of VCSPs of countably infinite valued structures whose automorphisms form an oligomorphic permutation group. In particular, we develop a notion of *expressive power* which is based on *primitive positive definitions* and other complexity-preserving operators, inspired by the techniques known from VCSPs over finite domains. We use the expressive power to obtain polynomial-time reductions between VCSPs and use them as the basis for formulating a hardness condition for infinite domain VCSPs. We conjecture that for VCSPs that stem from resilience problems, this hardness condition is not only sufficient but also necessary, unless P = NP. More precisely, we conjecture that this is the case when the automorphism group of the valued structure is identical to that of a structure which is a reduct of a countable finitely bounded homogeneous structure.

We also present an algebraic condition for (infinite-domain) valued structures which implies that the induced VCSP is in P, based on the concept of *fractional polymorphisms* which generalise classical polymorphisms, a common tool for proving tractability of CSPs. To prove membership in P, we use a reduction to finite-domain VCSPs which can be solved by a linear programming relaxation technique. We conjecture that the resulting algorithm solves all resilience problems that are in P. We demonstrate the utility of our tractability condition by applying it to a concrete conjunctive query for which the computational complexity of resilience has been stated as an open problem in the literature [22] (Section 8.5).

Related Work. The study of resilience problems was initiated in [21] under set semantics. The authors study the class of self-join-free conjunctive queries, i.e., queries in which each relation symbol occurs at most once, and obtain a P versus NP-complete dichotomy for this class. In a subsequent paper [22], several results are obtained for conjunctive queries with self-joins of a specific form, while the authors also state a few open problems of similar nature that cannot be handled by their methods. In the latest article [36], Gatterbauer and Makhija present a unified approach to resilience problems based on integer linear programming that works also for queries with self-joins, both under bag semantics and under set semantics. The new complexity results in [36] again concern self-join-free queries. Our approach is independent from self-joins and hence allows to study conjunctive queries that were difficult to treat before.

We stress that VCSPs of countable valued structures with an oligomorphic automorphism group greatly surpass resilience problems. For example, many problems in the recently very active area of graph separation problems [28, 29] such as *directed feedback edge set problem* and *directed symmetric multicut problem* can be formulated as VCSPs of appropriate countable valued structures with an oligomorphic automorphism group. Several of these problems such as the *multicut problem* and the *coupled min cut problem* can even be formulated as VCSPs over finite domains.

Outline. The article is organised from the general to the specific, starting with VCSPs in full generality (Section 2), then focusing on valued structures with an oligomorphic automorphism group (Section 3), for which our notion of expressive power (Section 4) leads to polynomial-time reductions. Our general hardness condition, which also builds upon the notion of expressive power, is presented in Section 5. To study the expressive power and to formulate general polynomial-time tractability results, we introduce the concept of *fractional polymorphisms* in Section 6 (they are probability distributions over operations on the valued structure). We take inspiration from the theory of VCSPs for finite-domain valued structures, but apply some non-trivial modifications that are specific to the infinite-domain setting (because the considered probability distributions are over uncountable sets) making sure that our definitions specialize to the standard ones over finite domains. We then present a general polynomial-time tractability result (Theorem 7.13) which is phrased in terms of fractional polymorphisms. Section 8 applies the general theory to resilience problems. We illustrate the power of our approach by settling the computational complexity of a resilience problem for a concrete conjunctive query from the literature (Section 8.5). Section 9 closes with open problems for future research.

Many of statements in the article have long and technical proofs, therefore the proofs are omitted and can be found in the long version of the article available on arXiv [11].

2 PRELIMINARIES

The set $\{0, 1, 2, \dots\}$ of natural numbers is denoted by $\mathbb{N}$, the set of rational numbers is denoted by $\mathbb{Q}$, the set of non-negative rational numbers by $\mathbb{Q}_{\geq 0}$ and the set of positive rational numbers by $\mathbb{Q}_{>0}$. We use analogous notation for the set of real numbers $\mathbb{R}$ and the set of integers $\mathbb{Z}$. We also need an additional value ∞; all we need to know about ∞ is that

- $a < \infty$ for every $a \in \mathbb{Q}$,
- $a + \infty = \infty + a = \infty$ for all $a \in \mathbb{Q} \cup \{\infty\}$, and
- $0 \cdot \infty = \infty \cdot 0 = 0$ and $a \cdot \infty = \infty \cdot a = \infty$ for $a > 0$.

2.1 Valued Structures

Let C be a set and let $k \in \mathbb{N}$. A *valued relation of arity k over C* is a function $R\colon C^k \to \mathbb{Q} \cup \{\infty\}$. We write $\mathscr{R}_C^{(k)}$ for the set of all valued relations over C of arity k, and define

$$\mathscr{R}_C := \bigcup_{k \in \mathbb{N}} \mathscr{R}_C^{(k)}.$$

A valued relation is called *finite-valued* if it takes values only in $\mathbb{Q}$.

Example 2.1. *The* valued equality relation $R_=$ *is the binary valued relation defined over* C *by* $R_=(x,y) = 0$ *if* $x = y$ *and* $R_=(x,y) = \infty$ *otherwise. The* empty relation $R_\emptyset$ *is the unary valued relation defined over* C *by* $R_\emptyset(x) = \infty$ *for all* $x \in C$.

A valued relation $R \in \mathscr{R}_C^{(k)}$ that only takes values from $\{0, \infty\}$ will be identified with the 'crisp' relation $\{a \in C^k \mid R(a) = 0\}$. For $R \in \mathscr{R}_C^{(k)}$ the *feasibility relation of* R is defined as

$$\mathrm{Feas}(R) := \{a \in C^k \mid R(a) < \infty\}.$$

A *(relational) signature* τ is a set of *relation symbols*, each of them equipped with an arity from $\mathbb{N}$. A *valued* τ*-structure* Γ consists of a set C, which is also called the *domain* of Γ, and a valued relation $R^\Gamma \in \mathscr{R}_C^{(k)}$ for each relation symbol $R \in \tau$ of arity k. A τ*-structure* in the usual sense may then be identified with a valued τ-structure where all valued relations only take values from $\{0, \infty\}$.

Example 2.2. *Let* $\tau = \{<\}$ *be a relational signature with a single binary relation symbol* $<$. *Let* $\Gamma_<$ *be the valued* τ*-structure with domain* $\{0, 1\}$ *and where* $<(x,y) = 0$ *if* $x < y$, *and* $<(x,y) = 1$ *otherwise.*

An *atomic* τ*-expression* is an expression of the form $R(x_1, \ldots, x_k)$ for $R \in \tau$ and (not necessarily distinct) variable symbols $x_1, \ldots, x_k$. A τ*-expression* is an expression ϕ of the form $\sum_{i \leq m} \phi_i$ where $m \in \mathbb{N}$ and ϕ_i for $i \in \{1, \ldots, m\}$ is an atomic τ-expression. Note that the same atomic τ-expression might appear several times in the sum. We write $\phi(x_1, \ldots, x_n)$ for a τ-expression where all the variables are from the set $\{x_1, \ldots, x_n\}$. If Γ is a valued τ-structure, then a τ-expression $\phi(x_1, \ldots, x_n)$ defines over Γ a member of $\mathscr{R}_C^{(n)}$, which we denote by ϕ^Γ. If ϕ is the empty sum then ϕ^Γ is constant 0.

2.2 Valued Constraint Satisfaction

In this section we assume that Γ is a fixed valued τ-structure for a *finite* signature τ. The valued relations of Γ are also called *cost functions*. The *valued constraint satisfaction problem for* Γ, denoted by VCSP(Γ), is the computational problem to decide for a given τ-expression $\phi(x_1, \ldots, x_n)$ and a given $u \in \mathbb{Q}$ whether there exists $a \in C^n$ such that $\phi^\Gamma(a) \leq u$. We refer to $\phi(x_1, \ldots, x_n)$ as an *instance* of VCSP(Γ), and to u as the *threshold*. Tuples $a \in C^n$ such that $\phi^\Gamma(a) \leq u$ are called a *solution for* (ϕ, u). The *cost* of ϕ (with respect to Γ) is defined to be

$$\inf_{a \in C^n} \phi^\Gamma(a).$$

In some contexts, it will be beneficial to consider only a given τ-expression ϕ to be the input of VCSP(Γ) (rather than ϕ and the threshold u) and a tuple $a \in C^n$ will then be called a *solution for* ϕ if the cost of ϕ equals $\phi^\Gamma(a)$. Note that in general there might not be any solution. If there exists a tuple $a \in C^n$ such that $\phi^\Gamma(a) < \infty$ then ϕ is called *satisfiable*.

Note that our setting also captures classical CSPs, which can be viewed as the VCSPs for valued structures Γ that only contain cost functions that take value 0 or ∞. In this case, we will sometimes write CSP(Γ) for VCSP(Γ). Below we give an example of a known optimisation problem that can be formulated as a valued constraint satisfaction problem.

Example 2.3. *The problem* VCSP($\Gamma_<$) *for the valued structure* $\Gamma_<$ *from Example 2.2 models the* directed max-cut *problem: given a finite directed graph* (V, E) *(we do allow loops and multiple edges), partition the vertices* V *into two classes* A *and* B *such that the number of edges from* A *to* B *is maximal. Maximising the number of edges from* A *to* B *amounts to minimising the number* e *of edges within* A*, within* B*, and from* B *to* A*. So when we associate* A *to the preimage of* 0 *and* B *to the preimage of* 1*, computing the number* e *corresponds to finding the evaluation map* $s \colon V \to \{0, 1\}$ *that minimises the value* $\sum_{(x,y) \in E} <(s(x), s(y))$*, which can be formulated as an instance of* VCSP($\Gamma_<$). *Conversely, every instance of* VCSP($\Gamma_<$) *corresponds to a directed max-cut instance. It is known that* VCSP($\Gamma_<$) *is NP-complete (even if we do not allow loops and multiple edges in the input) [24]. We mention that this problem can be viewed as a resilience problem in database theory as explained in Section 8.*

3 OLIGOMORPHICITY

Many facts about VCSPs for valued structures with a finite domain can be generalised to a large class of valued structures over an infinite domain, defined in terms of automorphisms. We define automorphisms of valued structures as follows.

Definition 3.1. *Let* $k \in \mathbb{N}$, *let* $R \in \mathscr{R}_C^{(k)}$, *and let* α *be a permutation of* C. *Then* α preserves R *if for all* $a \in C^k$ *we have* $R(\alpha(a)) = R(a)$. *If* Γ *is a valued structure with domain* C, *then an* automorphism *of* Γ *is a permutation of* C *that preserves all valued relations of* R.

The set of all automorphisms of Γ is denoted by Aut(Γ), and forms a group with respect to composition. Let $k \in \mathbb{N}$. An *orbit of* k*-tuples* of a permutation group G is a set of the form $\{\alpha(a) \mid \alpha \in G\}$ for some $a \in C^k$. A permutation group G on a countable set is called *oligomorphic* if for every $k \in \mathbb{N}$ there are finitely many orbits of k-tuples in G [13]. From now on, whenever we write that a structure has an oligomorphic automorphism group, we also imply that its domain is countable. Clearly, every valued structure with a finite domain has an oligomorphic automorphism group. A countable structure has an oligomorphic automorphism group if and only if it is ω*-categorical*, i.e., if all countable models of its first-order theory are isomorphic [25].

Example 3.2. *Let* $\tau = \{E, N\}$ *be a relational signature with two binary relation symbols* E *and* N. *Let* Γ_{LCC} *be the valued* τ*-structure with domain* $\mathbb{N}$ *where* $E(x,y) = 0$ *if* $x = y$ *and* $E(x,y) = 1$ *otherwise, and where* $N(x,y) = 0$ *if* $x \neq y$ *and* $N(x,y) = 1$ *otherwise. Note that* Aut(Γ_{LCC}) *is the full symmetric group on* $\mathbb{N}$. *This group is oligomorphic; for example, there are five orbits of triples represented by the tuples* $(1,2,3)$, $(1,1,2)$, $(1,2,1)$, $(2,1,1)$ *and* $(1,1,1)$.

The problem of least correlation clustering with partial information *[42, Example 5] is equal to* VCSP(Γ_{LCC}). *It is a variant of the min-correlation clustering problem [1] that does not require precisely one constraint between any two variables. The problem is NP-complete in both settings [24, 42].*

The following lemma shows that valued τ-structures always realize infima of τ-expressions.

Lemma 3.3. *Let* Γ *be a valued structure with a countable domain* C *and an oligomorphic automorphism group. Then for every instance*

$\phi(x_1, \dots, x_n)$ of VCSP(Γ) *there exists $a \in C^n$ such that the cost of ϕ equals $\phi^\Gamma(a)$.*

A first-order sentence is called *universal* if it is of the form $\forall x_1, \dots, x_l.\ \psi$ where ψ is quantifier-free. Every quantifier-free formula is equivalent to a formula in conjunctive normal form, so we generally assume that quantifier-free formulas are of this form.

Recall that a τ-structure $\mathfrak{A}$ *embeds* into a τ structure $\mathfrak{B}$ if there is an injective map from A to B that preserves all relations of $\mathfrak{A}$ and their complements; the corresponding map is called an *embedding*. The *age* of a τ-structure is the class of all finite τ-structures that embed into it. A structure $\mathfrak{B}$ with a finite relational signature τ is called

- *finitely bounded* if there exists a universal τ-sentence ϕ such that a finite structure $\mathfrak{A}$ is in the age of $\mathfrak{B}$ iff $\mathfrak{A} \models \phi$.
- *homogeneous* if every isomorphism between finite substructures of $\mathfrak{B}$ can be extended to an automorphism of $\mathfrak{B}$.

If $\tau' \subseteq \tau$, then a τ'-structure $\mathfrak{B}'$ is called the *reduct* of $\mathfrak{B}$ if $\mathfrak{B}$ and $\mathfrak{B}'$ have the same domain and $R^{\mathfrak{B}'} = R^{\mathfrak{B}}$ for every $R \in \tau'$.

Note that for every structure $\mathfrak{B}$ with a finite relational signature, for every n there are only finitely many non-isomorphic substructures of $\mathfrak{B}$ of size n. Therefore, all countable homogeneous structures with a finite relational signature and all of their reducts have finitely many orbits of k-tuples for all $k \in \mathbb{N}$, and hence an oligomorphic automorphism group.

Theorem 3.4. *Let Γ be a countable valued structure with finite signature such that there exists a finitely bounded homogeneous structure $\mathfrak{B}$ with* Aut($\mathfrak{B}$) ⊆ Aut(Γ)*. Then* VCSP(Γ) *is in NP.*

4 EXPRESSIVE POWER

One of the fundamental concepts in the theory of constraint satisfaction is the concept of *primitive positive definitions*, which is the fragment of first-order logic where only equality, existential quantification, and conjunction are allowed (in other words, negation, universal quantification, and disjunction are forbidden). The motivation for this concept is that relations with such a definition can be added to the structure without changing the complexity of the respective CSP. The natural generalisation to *valued* constraint satisfaction is the following notion of expressibility.

Definition 4.1. *Let Γ be a valued τ-structure. We say that $R \in \mathscr{R}_C^{(k)}$ can be* expressed *by Γ if there exists a τ-expression*

$$\phi(x_1, \dots, x_k, y_1, \dots, y_n)$$

such that for all $a \in C^k$ we have

$$R(a) = \inf_{b \in C^n} \phi^\Gamma(a, b).$$

Note that $\inf_{b \in C^n} \phi^\Gamma(a, b)$ might be irrational or $-\infty$. If this is the case in Definition 4.1, then ϕ does not witness that R can be expressed in Γ since valued relations must have weights from $\mathbb{Q} \cup \{\infty\}$. If C has an oligomorphic permutation group, however, then Lemma 3.3 guarantees existence. We will further see in Lemma 4.6 that if Γ has an oligomorphic automorphism group, then the addition of valued relations that are expressible by Γ does not change the computational complexity of VCSP(Γ).

More ways to derive new relations from existing ones that preserves the computational complexity of the original VCSP are introduced in the following definition.

Definition 4.2. *Let $R, R' \in \mathscr{R}_C$. We say that R' can be obtained from R by*

- non-negative scaling *if there exists $r \in \mathbb{Q}_{\geq 0}$ such that $R = rR'$;*
- shifting *if there exists $s \in \mathbb{Q}$ such that $R = R' + s$.*

If R is of arity k, then the relation that contains all minimal-value tuples of R is

$$\text{Opt}(R) := \{a \in \text{Feas}(R) \mid R(a) \leq R(b) \text{ for every } b \in C^k\}.$$

Definition 4.3 (valued relational clone). *A* valued relational clone *(over C) is a subset of $\mathscr{R}_C$ that contains $R_=$ and $R_\emptyset$ (from Example 2.1), and is closed under expressibility, shifting, and non-negative scaling,* Feas*, and* Opt*. For a valued structure Γ with domain C, we write ⟨Γ⟩ for the smallest relational clone that contains the valued relations of Γ.*

The following example shows that neither the operator Opt nor the operator Feas is redundant in the definition above.

Example 4.4. *Consider the domain $C = \{0, 1, 2\}$ and the unary valued relation R on C defined by $R(0) = 0$, $R(1) = 1$ and $R(2) = \infty$. Then the relation* Feas(R) *cannot be obtained from R by expressing, shifting, non-negative scaling and use of* Opt*. Similarly, the relation* Opt(R) *cannot be obtained from R by expressing, shifting, non-negative scaling and use of* Feas*.*

Remark 4.5. *Note that for every valued structure Γ and $R \in$ ⟨Γ⟩, every automorphism of Γ is an automorphism of R.*

The motivation for Definition 4.3 for valued CSPs stems from the following lemma, which shows that adding relations in ⟨Γ⟩ does not change the complexity of the VCSP up to polynomial-time reductions. For finite-domain valued structures this is proved in [16], except for the operator Opt, for which a proof can be found in [23, Theorem 5.13]. Parts of the proof have been generalised to infinite-domain valued structures without further assumptions; see, e.g. [40] and [42, Lemma 7.1.4]. However, in these works the definition of VCSPs was changed to ask whether there is a solution of a cost strictly less than u, to circumvent problems about infima that are not realised. Moreover, in [40] the authors restrict themselves to finite-valued relations and hence do not consider the operator Opt. It is visible from Example 4.4 that neither the operator Opt nor the operator Feas can be simulated by the other ones already on finite domains, which is why they both appear in [23] (Feas was included implicitly by allowing to scale by 0 and defining $0 \cdot \infty = \infty$). In this article we work with valued structures with an oligomorphic automorphism group so that the values of expressions are attained and hence we can adapt the proof from the finite-domain case to show that the complexity is preserved.

Lemma 4.6. *Let Γ be a valued structure with an oligomorphic automorphism group and a finite signature. Suppose that Δ is a valued structure with a finite signature over the same domain C such that every cost function of Δ is from ⟨Γ⟩. Then there is a polynomial-time reduction from* VCSP(Δ) *to* VCSP(Γ)*.*

The next two examples illustrate the use of Lemma 4.6 for obtaining hardness results.

Example 4.7. *Recall the structure $\Gamma_<$ from Example 2.2. We have seen in Example 2.3 that* $\text{VCSP}(\Gamma_<)$ *is the directed max-cut problem. Consider the classical relation* NAE *on* $\{0,1\}$ *defined by*

$$\text{NAE} := \{0,1\}^3 \setminus \{(0,0,0),(1,1,1)\}.$$

Note that

$$\text{NAE}(x,y,z) = \text{Opt}\left(<(x,y) + <(y,z) + <(z,x)\right).$$

Since $\text{CSP}(\{0,1\},\text{NAE})$ *is known to be NP-hard (see, e.g., [3]), this provides an alternative proof of the NP-hardness of the directed max-cut problem via Lemma 4.6.*

Example 4.8. *We revisit the countably infinite valued structure* Γ_{LCC} *from Example 3.2. Recall that* $\text{VCSP}(\Gamma_{LCC})$ *is the least correlation clustering problem with partial information and that* $\text{Aut}(\Gamma_{LCC})$ *is oligomorphic. Let* Γ_{EC} *be the relational structure with the same domain as* Γ_{LCC} *and the relation* $R := \{(x,y,z) \mid (x = y \wedge y \neq z) \vee (x \neq y \wedge y = z)\}$ *(attaining values* 0 *and* ∞*). Note that*

$$R(x,y,z) = \text{Opt}(N(x,z) + N(x,z) + E(x,y) + E(y,z)).$$

This provides an alternative proof of NP-hardness of the least correlation clustering with partial information via Lemma 4.6, because $\text{CSP}(\Gamma_{EC})$ *is known to be NP-hard [5].*

Note that we can replace $N(x,z) + N(x,z)$ *in the definition of* R *by* $\text{Opt}(N)(x,z)$ *and that* $\text{Opt}(N)$ *is equal to the classical relation* $\neq$ *(attaining values* 0 *and* ∞*). This shows that even* $\text{VCSP}(\mathbb{N};E,\neq)$ *is NP-hard.*

5 HARDNESS FROM PP-CONSTRUCTIONS

A universal-algebraic theory of VCSPs for finite valued structures has been developed in [33], following the classical approach to CSPs which is based on the concepts of cores, addition of constants, and primitive positive interpretations. Subsequently, an important conceptual insight has been made for classical CSPs which states that every structure that can be interpreted in the expansion of the core of the structure by constants can also be obtained by taking a pp-power if we then consider structures up to homomorphic equivalence [2]. We are not aware of any published reference that adapts this perspective to the algebraic theory of VCSPs, so we develop (parts of) this approach here. As in [2], we immediately step from valued structures with a finite domain to the more general case of valued structures with an oligomorphic automorphism group. The results in this section are adaptations of the known results for classical relational structures or finite-domain valued structures.

Definition 5.1 (pp-power). *Let* Γ *be a valued structure with domain* C *and let* $d \in \mathbb{N}$*. Then a (d-th)* pp-power *of* Γ *is a valued structure* Δ *with domain* C^d *such that for every valued relation* R *of* Δ *of arity* k *there exists a valued relation* S *of arity* kd *in* $\langle\Gamma\rangle$ *such that*

$$R((a^1_1,\dots,a^1_d),\dots,(a^k_1,\dots,a^k_d)) = S(a^1_1,\dots,a^1_d,\dots,a^k_1,\dots,a^k_d).$$

The name 'pp-power' comes from 'primitive positive power', since for relational structures expressibility is captured by primitive positive formulas. The following proposition shows that the VCSP of a pp-power reduces to the VCSP of the original structure.

Proposition 5.2. *Let* Γ *and* Δ *be valued structures such that* $\text{Aut}(\Gamma)$ *is oligomorphic and* Δ *is a pp-power of* Γ*. Then* $\text{Aut}(\Delta)$ *is oligomorphic and there is a polynomial-time reduction from* $\text{VCSP}(\Delta)$ *to* $\text{VCSP}(\Gamma)$*.*

If C and D are sets, then we equip the space C^D of functions from D to C with the topology of pointwise convergence, where C is taken to be discrete. In this topology, a basis of open sets is given by

$$\mathcal{S}_{a,b} := \{f \in C^D \mid f(a) = b\}$$

for $a \in D^k$ and $b \in C^k$ for some $k \in \mathbb{N}$, and f is applied componentwise. For any topological space T, we denote by $B(T)$ the Borel σ-algebra on T, i.e., the smallest subset of the powerset $\mathcal{P}(T)$ which contains all open sets and is closed under countable intersection and complement. We write $[0,1]$ for the set $\{x \in \mathbb{R} \mid 0 \leq x \leq 1\}$.

Definition 5.3 (fractional map). *Let* C *and* D *be sets. A* fractional map *from* D *to* C *is a probability distribution*

$$(C^D, B(C^D), \omega\colon B(C^D) \to [0,1]),$$

that is, $\omega(C^D) = 1$ *and* ω *is countably additive: if* $A_1, A_2, \dots \in B(C^D)$ *are disjoint, then*

$$\omega(\bigcup_{i\in\mathbb{N}} A_i) = \sum_{i\in\mathbb{N}} \omega(A_i).$$

If $f \in C^D$, we often write $\omega(f)$ instead of $\omega(\{f\})$. Note that $\{f\} \in B(C^D)$ for every f. The set $[0,1]$ carries the topology inherited from the standard topology on $\mathbb{R}$. We also view $\mathbb{R} \cup \{\infty\}$ as a topological space with a basis of open sets given by all open intervals (a,b) for $a,b \in \mathbb{R}$, $a < b$ and additionally all sets of the form $\{x \in \mathbb{R} \mid x > a\} \cup \{\infty\}$.

A *(real-valued) random variable* is a *measurable function* $X\colon T \to \mathbb{R} \cup \{\infty\}$, i.e., pre-images of elements of $B(\mathbb{R} \cup \{\infty\})$ under X are in $B(T)$. If X is a real-valued random variable, then the *expected value of* X *(with respect to a probability distribution* ω*)* is denoted by $E_\omega[X]$ and is defined via the Lebesgue integral

$$E_\omega[X] := \int_T X d\omega.$$

Recall that the Lebesgue integral $\int_T X d\omega$ need not exist, in which case $E_\omega[X]$ is undefined; otherwise, the integral equals a real number, ∞, or $-\infty$. The definition and some properties of the Lebesgue integral, specialised to our setting, can be found in the arXiv version of the article [11]. Also recall that the expected value is

- *linear*, i.e., for every $a,b \in \mathbb{R}$ and random variables X, Y such that $E_\omega[X]$ and $E_\omega[Y]$ exist and $aE_\omega[X] + bE_\omega[Y]$ is defined we have
$$E_\omega[aX + bY] = aE_\omega[X] + bE_\omega[Y];$$
- *monotone*, i.e., if X, Y are random variables such that $E_\omega[X]$ and $E_\omega[Y]$ exist and $X(f) \leq Y(f)$ for all $f \in T$, then $E_\omega[X] \leq E_\omega[Y]$.

Let C and D be sets. In the rest of the paper, we will work exclusively on a topological space C^D of maps $f\colon D \to C$ and the special case $\mathscr{O}_C^{(\ell)}$ for some $\ell \in \mathbb{N}$ (i.e., $D = C^\ell$). Note that if C and D are infinite, then these spaces are uncountable and hence there are probability distributions ω such that $\omega(A) = 0$ for every 1-element set A. Therefore, in these cases, $E_\omega[X]$ for a random variable X might not be expressible as a sum.

Definition 5.4 (fractional homomorphism). *Let* Γ *and* Δ *be valued* τ*-structures with domains* C *and* D*, respectively. A* fractional homomorphism *from* Δ *to* Γ *is a fractional map from* D *to* C *such that for every* $R \in \tau$ *of arity* k *and every tuple* $a \in D^k$ *it holds for the random variable* $X: C^D \to \mathbb{R} \cup \{\infty\}$ *given by*

$$f \mapsto R^\Gamma(f(a))$$

that $E_\omega[X]$ *exists and that*

$$E_\omega[X] \leq R^\Delta(a).$$

The following lemma shows that if $\mathrm{Aut}(\Gamma)$ is oligomorphic, then the expected value from Definition 5.4 always exists.

Lemma 5.5. *Let* C *and* D *be sets,* $a \in D^k$*,* $R \in \mathscr{R}_C^{(k)}$*. Let* $X: C^D \to \mathbb{R} \cup \{\infty\}$ *be the random variable given by*

$$f \mapsto R(f(a)).$$

If $\mathrm{Aut}(C;R)$ *is oligomorphic, then* $E_\omega[X]$ *exists and* $E_\omega[X] > -\infty$*.*

Lemma 5.6. *Let* $\Gamma_1, \Gamma_2, \Gamma_3$ *be countable valued* τ*-structures such that there exists a fractional homomorphism* ω_1 *from* Γ_1 *to* Γ_2 *and a fractional homomorphism* ω_2 *from* Γ_2 *to* Γ_3*. Then there exists a fractional homomorphism* $\omega_3 := \omega_2 \circ \omega_1$ *from* Γ_1 *to* Γ_3*.*

We say that two valued τ-structures Γ and Δ are *fractionally homomorphically equivalent* if there exists a fractional homomorphisms from Γ to Δ and from Δ to Γ. Clearly, fractional homomorphic equivalence is indeed an equivalence relation on valued structures of the same signature.

Proposition 5.7. *Let* Γ *and* Δ *be valued* τ*-structures with oligomorphic automorphism groups that are fractionally homomorphically equivalent. Then* $\mathrm{VCSP}(\Gamma)$ *and* $\mathrm{VCSP}(\Delta)$ *are polynomial-time equivalent.*

Remark 5.8. *If* Γ *and* Δ *are classical relational* τ*-structures that are homomorphically equivalent in the classical sense, then they are fractionally homomorphically equivalent when we view them as valued structures: if* h_1 *is the homomorphism from* Γ *to* Δ *and* h_2 *is the homomorphism from* Δ *to* Γ*, then this is witnessed by the fractional homomorphisms* ω_1 *and* ω_2 *such that* $\omega_1(h_1) = \omega_2(h_2) = 1$*.*

Definition 5.9 (pp-construction). *Let* Γ, Δ *be valued structures. Then* Δ *has a* pp-construction *in* Γ *if* Δ *is fractionally homomorphically equivalent to a structure* Δ' *which is a pp-power of* Γ*.*

Combining Proposition 5.2 and Proposition 5.7 gives the following.

Corollary 5.10. *Let* Γ *and* Δ *be valued structures with finite signatures and oligomorphic automorphism groups such that* Δ *has a pp-construction in* Γ*. Then there is a polynomial-time reduction from* $\mathrm{VCSP}(\Delta)$ *to* $\mathrm{VCSP}(\Gamma)$*.*

Note that the hardness proofs in Examples 4.7 and 4.8 are special cases of Corollary 5.10. A more involved example uses the relational structure with a hard CSP introduced below and is shown in Example 8.18.

Let OIT be the following relation

$$\mathrm{OIT} = \{(0,0,1), (0,1,0), (1,0,0)\}.$$

It is well-known (see, e.g., [3]) that $\mathrm{CSP}(\{0,1\};\mathrm{OIT})$ is NP-complete. Via Corollary 5.10, the NP-hardness of $\mathrm{CSP}(\{0,1\};\mathrm{OIT})$ yields:

Corollary 5.11. *Let* Γ *be a valued structure with a finite signature and oligomorphic automorphism group such that* $(\{0,1\};\mathrm{OIT})$ *has a pp-construction in* Γ*. Then* $\mathrm{VCSP}(\Gamma)$ *is NP-hard.*

We finish this section by the following useful lemma.

Lemma 5.12. *The relation of pp-constructibility on the class of countable valued structures is transitive.*

6 FRACTIONAL POLYMORPHISMS

In this section we introduce *fractional polymorphisms* of valued structures; they are an important tool for formulating tractability results and complexity classifications of VCSPs. For valued structures with a finite domain, our definition specialises to the established notion of a fractional polymorphism which has been used to study the complexity of VCSPs for valued structures over finite domains (see, e.g. [41]). Our approach is different from the one of Schneider and Viola [40, 42] and Viola and Živný [43] in that we work with arbitrary probability spaces instead of distributions with finite support or countable additivity property. As we will see in Section 7, fractional polymorphisms can be used to give sufficient conditions for tractability of VCSPs of certain valued structures with oligomorphic automorphism groups. This justifies our more general notion of a fractional polymorphism, as it might provide a tractability proof for more problems.

Let $\mathscr{O}_C^{(\ell)}$ be the set of all operations $f: C^\ell \to C$ on a set C of arity ℓ. We equip $\mathscr{O}_C^{(\ell)}$ with the topology of pointwise convergence, where C is taken to be discrete. That is, the basic open sets are of the form

$$\mathscr{S}_{a^1,\dots,a^\ell,b} := \{f \in \mathscr{O}_C^{(\ell)} \mid f(a^1,\dots,a^\ell) = b\} \quad (1)$$

where $a^1,\dots,a^\ell, b \in C^k$, for some $k \in \mathbb{N}$, and f is applied componentwise. Let

$$\mathscr{O}_C := \bigcup_{\ell \in \mathbb{N}} \mathscr{O}_C^{(\ell)}.$$

Definition 6.1 (fractional operation). *Let* $\ell \in \mathbb{N}$*. A* fractional operation on C of arity ℓ *is a probability distribution*

$$\left(\mathscr{O}_C^{(\ell)}, B(\mathscr{O}_C^{(\ell)}), \omega: B(\mathscr{O}_C^{(\ell)}) \to [0,1]\right).$$

The set of all fractional operations on C *of arity* ℓ *is denoted by* $\mathscr{F}_C^{(\ell)}$*, and* $\mathscr{F}_C := \bigcup_{\ell \in \mathbb{N}} \mathscr{F}_C^{(\ell)}$*.*

If the reference to C is clear, we occasionally omit the subscript C. We often use ω for both the entire fractional operation and for the map $\omega: B(\mathscr{O}_C^{(\ell)}) \to [0,1]$.

Definition 6.2. *A fractional operation* $\omega \in \mathscr{F}_C^{(\ell)}$ improves *a* k*-ary valued relation* $R \in \mathscr{R}_C^{(k)}$ *if for all* $a^1,\dots,a^\ell \in C^k$

$$E := E_\omega[f \mapsto R(f(a^1,\dots,a^\ell))]$$

exists and

$$E \leq \frac{1}{\ell} \sum_{j=1}^{\ell} R(a^j). \quad (2)$$

Note that (2) has the interpretation that the expected value of $R(f(a^1,\dots,a^\ell))$ is at most the average of the values $R(a^1),\dots,R(a^\ell)$. Also note that if R is a classical relation improved by a fractional operation ω and $\omega(f) > 0$ for $f \in \mathscr{O}^{(\ell)}$, then f must preserve R in the usual sense. It follows from Lemma 5.5 that if $\mathrm{Aut}(C; R)$ is oligomorphic, then $E_\omega[f \mapsto R(f(a^1, \dots, a^\ell))]$ always exists and is greater than $-\infty$.

Definition 6.3 (fractional polymorphism). *If ω improves every valued relation in Γ, then ω is called a* fractional polymorphism *of Γ; the set of all fractional polymorphisms of Γ is denoted by* $\mathrm{fPol}(\Gamma)$.

Remark 6.4. *Our notion of fractional polymorphism coincides with the previously used notions of fractional polymorphisms with finite support [40, 42] or the countable additivity property [43], since in this case the expected value on the left-hand side of (2) is equal to the weighted arithmetic mean.*

Remark 6.5. *A fractional polymorphism of arity ℓ of a valued structure Γ might also be viewed as a fractional homomorphism from a specific ℓ-th pp-power of Γ to Γ, which we denote by Γ^ℓ: if C is the domain and τ the signature of Γ, then the domain of Γ^ℓ is C^ℓ, and for every $R \in \tau$ of arity k we have*

$$R^{\Gamma^\ell}((a^1_1, \dots, a^1_\ell), \dots, (a^k_1, \dots, a^k_\ell)) := \frac{1}{\ell} \sum_{i=1}^{\ell} R^\Gamma(a^1_i, \dots, a^k_i).$$

Example 6.6. *Let $\pi^\ell_i \in \mathscr{O}^{(\ell)}_C$ be the i-th projection of arity ℓ, which is given by $\pi^\ell_i(x_1, \dots, x_\ell) = x_i$. The fractional operation Id_ℓ of arity ℓ such that $\mathrm{Id}_\ell(\pi^\ell_i) = \frac{1}{\ell}$ for every $i \in \{1, \dots, \ell\}$ is a fractional polymorphism of every valued structure with domain C.*

Let $\mathscr{C} \subseteq \mathscr{F}_C$. We write $\mathscr{C}^{(\ell)}$ for $\mathscr{C} \cap \mathscr{F}^{(\ell)}_C$ and $\mathrm{Imp}(\mathscr{C})$ for the set of valued relations that are improved by every fractional operation in $\mathscr{C}$.

Lemma 6.7. *Let $R \in \mathscr{R}^{(k)}_C$ and let Γ be a valued structure with domain C and an automorphism $\alpha \in \mathrm{Aut}(\Gamma)$ which does not preserve R. Then $R \notin \mathrm{Imp}(\mathrm{fPol}(\Gamma)^{(1)})$.*

Parts of the arguments in the proof of the following lemma can be found in the proof of [42, Lemma 7.2.1]; note that the author works with a more restrictive notion of fractional operation, so we cannot reuse her result. However, the arguments can be generalized to our notion of fractional polymorphism for all countable valued structures.

Lemma 6.8. *For every valued τ-structure Γ over a countable domain C we have*

$$\langle \Gamma \rangle \subseteq \mathrm{Imp}(\mathrm{fPol}(\Gamma)).$$

The following example shows an application of Lemma 6.8.

Example 6.9. *Let $<$ be the binary relation on $\{0, 1\}$ and $\Gamma_<$ the valued structure from Example 2.2. By definition, $\mathrm{Opt}(<) \in \langle \Gamma_< \rangle$. Denote the minimum operation on $\{0, 1\}$ by $\min$ and let ω be a binary fractional operation defined by $\omega(\min) = 1$. Note that $\omega \in \mathrm{fPol}(\{0, 1\}; \mathrm{Opt}(<))$. However,*

$$< \left(\min\left(\begin{pmatrix}0\\1\end{pmatrix}, \begin{pmatrix}0\\0\end{pmatrix}\right)\right) = <(0, 0) = 1,$$

while $(1/2) \cdot <(0, 1) + (1/2) \cdot <(0, 0) = 1/2$. This shows that ω does not improve $<$ and hence $< \notin \langle (\{0, 1\}; \mathrm{Opt}(<)) \rangle$ by Lemma 6.8.

7 TRACTABILITY VIA CANONICAL FRACTIONAL POLYMORPHISMS

In this section we make use of a tractability result for finite-domain VCSPs of Kolmogorov, Krokhin, and Rolinek [31], building on earlier work of Kolmogorov, Thapper, and Živný [32, 41]. To exploit this result, the key ingredient is a polynomial-time reduction for VCSPs of valued structures with an oligomorphic automorphism group satisfying certain assumptions to VCSPs of finite-domain structures. This reduction is inspired by a similar reduction in the classical relational setting [8]. We then generalize and adapt several statements for finite-domain valued structures to be able to use the complexity classification for finite-domain VCSPs.

Definition 7.1. *An operation $f : C^\ell \to C$ for $\ell \geq 2$ is called* cyclic *if*

$$f(x_1, \dots, x_\ell) = f(x_2, \dots, x_\ell, x_1)$$

for all $x_1, \dots, x_\ell \in C$. Let $\mathrm{Cyc}^{(\ell)}_C \subseteq \mathscr{O}^{(\ell)}_C$ be the set of all operations on C of arity ℓ that are cyclic.

If G is a permutation group on a set C, then $\overline{G}$ denotes the closure of G in the space of functions from $C \to C$ with respect to the topology of pointwise convergence. Note that $\overline{G}$ might contain some operations that are not surjective, but if $G = \mathrm{Aut}(\mathfrak{B})$ for some structure $\mathfrak{B}$, then all operations in $\overline{G}$ are still embeddings of $\mathfrak{B}$ into $\mathfrak{B}$ that preserve all first-order formulas.

Definition 7.2. *Let G be a permutation group on the set C. An operation $f : C^\ell \to C$ is called* pseudo cyclic with respect to G *if there are $e_1, e_2 \in \overline{G}$ such that for all $x_1, \dots, x_\ell \in C$*

$$e_1(f(x_1, \dots, x_\ell)) = e_2(f(x_2, \dots, x_\ell, x_1)).$$

Let $\mathrm{PC}^{(\ell)}_G \subseteq \mathscr{O}^{(\ell)}_C$ be the set of all operations on C of arity ℓ that are pseudo cyclic with respect to G.

Note that $\mathrm{PC}^{(\ell)}_G \in B(\mathscr{O}^{(\ell)}_C)$. Indeed, the complement can be written as a countable union of sets of the form $\mathscr{S}_{a^1,\dots,a^\ell,b}$ where for all $f \in \mathscr{O}^{(\ell)}_C$ the tuples $f(a^1, \dots, a^\ell)$ and $f(a^2, \dots, a^\ell, a^1)$ lie in different orbits with respect to G.

Definition 7.3. *Let G be a permutation group with domain C. An operation $f : C^\ell \to C$ for $\ell \geq 2$ is called* canonical with respect to G *if for all $k \in \mathbb{N}$ and $a^1, \dots, a^\ell \in C^k$ the orbit of the k-tuple $f(a^1, \dots, a^\ell)$ only depends on the orbits of $a^1, \dots, a^\ell$ with respect to G. Let $\mathrm{Can}^{(\ell)}_G \subseteq \mathscr{O}^{(\ell)}_C$ be the set of all operations on C of arity ℓ that are canonical with respect to G.*

Remark 7.4. *Note that if h is an operation over C of arity ℓ which is canonical with respect to G, then h induces for every $k \in \mathbb{N}$ an operation h^* of arity ℓ on the orbits of k-tuples of G. Note that if h is pseudo cyclic with respect to G, then h^* is cyclic.*

Note that $\mathrm{Can}^{(\ell)}_G \in B(\mathscr{O}^{(\ell)}_C)$, since the complement is a countable union of sets of the form $\mathscr{S}_{a^1,\dots,a^\ell,b} \cap \mathscr{S}_{c^1,\dots,c^\ell,d}$ where for all $i \in \{1, \dots, \ell\}$ the tuples a^i and c^i lie in the same orbit with respect to G, but b and d do not.

Definition 7.5. *A fractional operation ω is called* pseudo cyclic with respect to G *if for every $A \in B(\mathscr{O}^{(\ell)}_C)$ we have $\omega(A) = \omega(A \cap \mathrm{PC}^{(\ell)}_G)$. Canonicity with respect to G and cyclicity for fractional operations are defined analogously.*

We refer to Section 8.3 for examples of concrete fractional polymorphisms of valued structures Γ that are cyclic and canonical with respect to $\mathrm{Aut}(\Gamma)$. We may omit the specification 'with respect to G' when G is clear from the context.

We prove below that canonical pseudo cyclic fractional polymorphisms imply polynomial-time tractability of the corresponding VCSP, by reducing to a tractable VCSP over a finite domain. Motivated by Theorem 3.4 and the infinite-domain tractability conjecture from [10], we state these results for valued structures related to finitely bounded homogeneous structures.

Definition 7.6 (Γ_m^*). *Let Γ be a valued structure with signature τ such that $\mathrm{Aut}(\Gamma)$ contains the automorphism group of a homogeneous structure $\mathfrak{B}$ with a finite relational signature. Let m be at least as large as the maximal arity of the relations of $\mathfrak{B}$. Let Γ_m^* be the following valued structure.*

- *The domain of Γ_m^* is the set of orbits of m-tuples of $\mathrm{Aut}(\Gamma)$.*
- *For every $R \in \tau$ of arity $k \leq m$ the signature of Γ_m^* contains a unary relation symbol R^*, which denotes in Γ_m^* the unary valued relation that returns on the orbit of an m-tuple $t = (t_1, \dots, t_m)$ the value of $R^\Gamma(t_1, \dots, t_k)$ (this is well-defined as the value is the same for all representatives t of the orbit).*
- *For every $p \in \{1, \dots, m\}$ and $i, j \colon \{1, \dots, p\} \to \{1, \dots, m\}$ there exists a binary relation $C_{i,j}$ which returns 0 for two orbits of m-tuples O_1 and O_2 if for every $s \in O_1$ and $t \in O_2$ we have that $(s_{i(1)}, \dots, s_{i(p)})$ and $(t_{j(1)}, \dots, t_{j(p)})$ lie in the same orbit of p-tuples of $\mathrm{Aut}(\Gamma)$, and returns ∞ otherwise.*

Note that $\mathrm{Aut}(\mathfrak{B})$ and hence $\mathrm{Aut}(\Gamma)$ has finitely many orbits of k-tuples for every $k \in \mathbb{N}$ and therefore Γ_m^* has a finite domain. The following generalises a known reduction for CSPs from [8].

Theorem 7.7. *Let Γ be a valued structure such that $\mathrm{Aut}(\Gamma)$ equals the automorphism group of a finitely bounded homogeneous structure $\mathfrak{B}$. Let r be the maximal arity of the relations of $\mathfrak{B}$ and the valued relations in Γ, let v be the maximal number of variables that appear in a single conjunct of the universal sentence ψ that describes the age of $\mathfrak{B}$, and let $m \geq \max(r+1, v, 3)$. Then there is a polynomial-time reduction from $\mathrm{VCSP}(\Gamma)$ to $\mathrm{VCSP}(\Gamma_m^*)$.*

For studying canonical operations, we can use known results about operations on finite domains.

Definition 7.8. *Let ω be a fractional operation of arity ℓ on a finite domain C. Then the* support of ω *is the set*

$$\mathrm{Supp}(\omega) = \{f \in \mathscr{O}_C^{(\ell)} \mid \omega(f) > 0\}.$$

If $\mathscr{F}$ is a set of fractional operations, then

$$\mathrm{Supp}(\mathscr{F}) := \bigcup_{\omega \in \mathscr{F}} \mathrm{Supp}(\omega).$$

Note that, a fractional operation ω on a finite domain is determined by the values $\omega(f)$, $f \in \mathrm{Supp}(\omega)$, in contrast to fractional operations on infinite domains. Moreover, a fractional polymorphism ω of a valued structure with a finite domain is cyclic if and only if all operations in its support are cyclic, in accordance to the definitions from [33]. An operation $f \colon C^4 \to C$ is called *Siggers* if $f(a,r,e,a) = f(r,a,r,e)$ for all $a, r, e \in C$.

Lemma 7.9. *Let Γ and Δ be valued structures with finite domains that are fractionally homomorphically equivalent.*

- *If Γ has a cyclic fractional polymorphism, then Δ has a cyclic fractional polymorphism of the same arity.*
- *If the set $\mathrm{Supp}(\mathrm{fPol}(\Gamma))$ contains a cyclic operation, then the set $\mathrm{Supp}(\mathrm{fPol}(\Delta))$ contains a cyclic operation of the same arity.*

The following lemma is a variation of Proposition 39 from [33], which is phrased there only for valued structures Γ that are cores and for idempotent cyclic operations.

Lemma 7.10. *Let Γ be a valued structure over a finite domain. Then Γ has a cyclic fractional polymorphism if and only if $\mathrm{Supp}(\mathrm{fPol}(\Gamma))$ contains a cyclic operation.*

The following outstanding result classifies the computational complexity of VCSPs for valued structures over finite domains; it does not appear in this form in the literature, but it can be derived from results in [12, 31, 33, 44, 45].

Theorem 7.11. *Let Γ be a valued structure with a finite signature and a finite domain. If $(\{0,1\}; \mathrm{OIT})$ does not have a pp-construction in Γ, then Γ has a fractional cyclic polymorphism, and $\mathrm{VCSP}(\Gamma)$ is in P, and it is NP-hard otherwise.*

The problem of deciding for a given valued structure Γ with finite domain and finite signature whether Γ satisfies the condition given in the previous theorem can be solved in exponential time [30]. We now state consequences of this result for certain valued structures with an infinite domain.

Proposition 7.12. *Let $\mathfrak{B}$ be a finitely bounded homogeneous structure and let Γ be a valued structure with finite relational signature such that $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\mathfrak{B})$. Let m be as in Theorem 7.7. Then the following are equivalent.*

1. $\mathrm{fPol}(\Gamma)$ *contains a fractional operation which is canonical and pseudo cyclic with respect to* $\mathrm{Aut}(\mathfrak{B})$*;*
2. $\mathrm{fPol}(\Gamma_m^*)$ *contains a cyclic fractional operation;*
3. $\mathrm{Supp}(\mathrm{fPol}(\Gamma_m^*))$ *contains a cyclic operation.*
4. $\mathrm{Supp}(\mathrm{fPol}(\Gamma_m^*))$ *contains a Siggers operation.*

Note that item (4) in the previous proposition can be decided algorithmically for a given valued structure Γ_m^* (which has a finite domain and finite signature) by testing all 4-ary operations on Γ_m^*.

Theorem 7.13. *If the conditions from Proposition 7.12 hold, then $\mathrm{VCSP}(\Gamma)$ is in P.*

8 APPLICATION: RESILIENCE

We introduce the resilience problem for conjunctive queries and, more generally, unions of conjunctive queries. We generally work with Boolean queries, i.e., queries without free variables. A *conjunctive query* is a primitive positive τ-sentence and a *union of conjunctive queries* is a (finite) disjunction of conjunctive queries. Note that every existential positive sentence can be written as a union of conjunctive queries.

Let τ be a finite relational signature and μ a conjunctive query over τ. The input to the *resilience problem for μ* consists of a finite τ-structure $\mathfrak{A}$, called a *database*[1], and the task is to compute the number of tuples that have to be removed from relations of $\mathfrak{A}$ so

[1] To be precise, a finite relational structure is not exactly the same as a database because the latter may not contain elements that are not contained in any relation. This difference, however, is inessential for the problems studied in this paper.

Fixed: a relational signature τ, a subset $\sigma \subseteq \tau$, and a union μ of conjuctive queries over τ.
Input: A bag database $\mathfrak{A}$ in signature τ and $u \in \mathbb{N}$.
$m :=$ minimal number of tuples to be removed from the relations in $\{R^{\mathfrak{A}} \mid R \in \tau \setminus \sigma\}$ so that $\mathfrak{A} \not\models \mu$.
Output: Is $m \leq u$?

Figure 1: The resilience problem considered in this paper.

that $\mathfrak{A}$ does *not* satisfy μ. This number is called the *resilience* of $\mathfrak{A}$ (with respect to μ). As usual, this can be turned into a decision problem where the input also contains a natural number $u \in \mathbb{N}$ and the question is whether the resilience is at most u. Clearly, $\mathfrak{A}$ does not satisfy μ if and only if its resilience is 0.

A natural variation of the problem is that the input database is a *bag database*, meaning that it may contain tuples with *multiplicities*. Formally, a bag database is a valued structure with all weights (which represent multiplicities) taken from $\mathbb{N}$. In this paper, we focus on bag databases whose relevance has already been discussed in the introduction.

The basic resilience problem defined above can be generalized by admitting the decoration of databases with a subsignature $\sigma \subseteq \tau$, in this way declaring all tuples in $R^{\mathfrak{A}}$, $R \in \sigma$, to be *exogenous*. This means that we are not allowed to remove such tuples from $\mathfrak{A}$ to make μ false; the tuples in the other relations are then called *endogenous*. For brevity, we also refer to the relations in σ as being exogenous and those in $\tau \setminus \sigma$ as being endogenous. If not specified, then $\sigma = \emptyset$, i.e., all tuples are endogenous. As an alternative, one may also declare individual tuples as being endogenous or exogenous. Under bag semantics, however, this case can be reduced to the one studied here (see Remark 8.15). The resilience problem that we study is summarized in Figure 1.

The *canonical database* of a conjunctive query μ with relational signature τ is the τ-structure $\mathfrak{A}$ whose domain are the variables of μ and where $a \in R^{\mathfrak{A}}$ for $R \in \tau$ if and only if μ contains the conjunct $R(a)$. Conversely, the *canonical query* of a relational τ-structure $\mathfrak{A}$ is the conjunctive query whose variable set is the domain A of $\mathfrak{A}$, and which contains for every $R \in \tau$ and $\bar{a} \in R^{\mathfrak{A}}$ the conjunct $R(\bar{a})$.

Remark 8.1. *For every conjunctive query μ, the resilience problem for μ parameterized by the threshold u is fixed-parameter tractable (FPT, which we refrain from defining here, see [18]). Indeed, there is a parameterized reduction to k-Hitting Set parametrized by the threshold u, which is known to be FPT [18].[2] The reduction is as follows. Let τ be the signature of μ, k the number of conjuncts of μ, and $\mathfrak{A}$ an input database. Construct a k-uniform hypergraph H where*

- *the vertices take the form $\langle R(a), i\rangle$ with $R \in \tau$, $a \in R^{\mathfrak{A}}$, and $i \geq 1$ bounded by the multiplicity of a in $R^{\mathfrak{A}}$ and*
- *every homomorphism h from the canonical database of μ to $\mathfrak{A}$ gives rise to an hyperedge $\{\langle R_1(h(x^1)), i_1\rangle, \dots, \langle R_\ell(h(x^\ell)), i_\ell\rangle\}$ where $R_1(x^1), \dots, R_\ell(x^\ell)$ are the atoms in μ and each i_j is bounded by the multiplicity of $h(x^j)$ in $R^{\mathfrak{A}}$.*

[2]We thank Peter Jonsson and George Osipov for suggesting this remark to us and letting us include it in this paper.

Figure 2: The query μ from Example 8.2 (on the left) and the corresponding structure $\mathfrak{B}$ (on the right).

Then μ has resilience at most u with respect to $\mathfrak{A}$ if and only if there is a set of vertices of H of size at most u that intersects every hyperedge.

We next explain how to represent resilience problems as VCSPs using appropriately chosen valued structures with oligomorphic automorphism groups.

Example 8.2. *The following query is taken from Meliou, Gatterbauer, Moore, and Suciu [37]; they show how to solve its resilience problem without multiplicities in polynomial time by a reduction to a max-flow problem. Let μ be the query*

$$\exists x, y, z\big(R(x, y) \wedge S(y, z)\big).$$

Observe that a finite τ-structure satisfies μ if and only if it does not have a homomorphism to the τ-structure $\mathfrak{B}$ with domain $B = \{0, 1\}$ and the relations $R^{\mathfrak{B}} = \{(0, 1), (1, 1)\}$ and $S^{\mathfrak{B}} = \{(0, 0), (0, 1)\}$ (see Figure 2). We turn $\mathfrak{B}$ into the valued structure Γ with domain $\{0, 1\}$ where $R^{\Gamma}(0, 1) = R^{\Gamma}(1, 1) = 0 = S^{\Gamma}(0, 0) = S^{\Gamma}(0, 1)$ and R^{Γ} and S^{Γ} take value 1 otherwise. Then $\mathrm{VCSP}(\Gamma)$ is precisely the resilience problem for μ (with multiplicities). Our results reprove the result from [36] that even with multiplicities, the problem can be solved in polynomial time (see Theorem 7.13, Proposition 8.14, and Example 8.11).

Example 8.3. *Let μ be the conjunctive query*

$$\exists x, y, z(R(x, y) \wedge S(x, y, z)).$$

This query is linear *in the sense of Freire, Gatterbauer, Immerman, and Meliou and thus its resilience problem without multiplicities can be solved in polynomial time (Theorem 4.5 in [37]; also see Fact 3.18 in [20]). Our results reprove the result from [36] that this problem remains polynomial-time solvable with multiplicities (see Theorem 7.13, Proposition 8.14 and Example 8.16).*

8.1 Connectivity

We show that when classifying the resilience problem for conjunctive queries, it suffices to consider queries that are connected.

A τ-structure is *connected* if it cannot be written as the disjoint union of two τ-structures with non-empty domains.

Remark 8.4. *All terminology introduced for τ-structures also applies to conjunctive queries with signature τ: by definition, the query has the property if the canonical database has the property.*

Lemma 8.5. *Let $\nu_1, \dots, \nu_k$ be conjunctive queries such that ν_i does not imply ν_j if $i \neq j$. Let $\nu = \nu_1 \wedge \cdots \wedge \nu_k$ and suppose that ν occurs in a union μ of conjunctive queries. For $i \in \{1, \dots, k\}$, let μ_i be the union of queries obtained by replacing ν by ν_i in μ. Then the resilience problem for μ is NP-hard if the resilience problem for one of the μ_i is NP-hard. Conversely, if the resilience problem is in P (in NP) for each μ_i, then the resilience problem for μ is in P as well (in NP, respectively). The same is true in the setting without multiplicities and/or exogeneous relations.*

By applying Lemma 8.5 finitely many times, we obtain that, when classifying the complexity of the resilience problem for unions of conjunctive queries, we may restrict our attention to unions of connected conjunctive queries.

8.2 Finite Duals

If μ is a union of conjunctive queries with signature τ, then a *dual* of μ is a τ-structure $\mathfrak{B}$ with the property that a finite τ-structure $\mathfrak{A}$ has a homomorphism to $\mathfrak{B}$ if and only if $\mathfrak{A}$ does not satisfy μ. The conjunctive query in Example 8.2, for instance, even has a *finite* dual. There is an elegant characterisation of the (unions of) conjunctive queries that have a finite dual. To state it, we need some basic terminology from database theory.

Definition 8.6. *The* incidence graph *of a relational τ-structure $\mathfrak{A}$ is the bipartite undirected multigraph whose first colour class is A, and whose second colour class consists of expressions of the form $R(b)$ where $R \in \tau$ has arity k, $b \in A^k$, and $\mathfrak{A} \models R(b)$. An edge $e_{a,i,R(b)}$ joins $a \in A$ with $R(b)$ if $b_i = a$. A structure is called* incidence-acyclic *(also known as Berge-acyclic) if its incidence graph is acyclic, i.e., it contains no cycles (if two vertices are linked by two different edges, then they establish a cycle). A structure is called a* tree *if it is incidence-acyclic and connected in the sense defined in Section 8.1.*

What follows is due to Nešetřil and Tardif [39]; see also [19, 35].

Theorem 8.7. *A conjunctive query μ has a finite dual if and only if the canonical database of μ is homomorphically equivalent to a tree. A union of conjunctive queries has a finite dual if and only if the canonical database for each of the conjunctive queries is homomorphically equivalent to a tree.*

The theorem shows that in particular Example 8.3 does not have a finite dual, since the query given there is not incidence-acyclic and hence cannot be homomorphically equivalent to a tree. To construct valued structures from duals, we introduce the following notation.

Definition 8.8. *Let $\mathfrak{B}$ be a τ-structure and $\sigma \subseteq \tau$. Define $\Gamma(\mathfrak{B}, \sigma)$ to be the valued τ-structure on the same domain as $\mathfrak{B}$ such that*

- *for each $R \in \tau \setminus \sigma$, $R^{\Gamma(\mathfrak{B},\sigma)}(a) := 0$ if $a \in R^{\mathfrak{B}}$ and $R^{\Gamma(\mathfrak{B},\sigma)}(a) := 1$ otherwise, and*
- *for each $R \in \sigma$, $R^{\Gamma(\mathfrak{B},\sigma)}(a) := 0$ if $a \in R^{\mathfrak{B}}$ and $R^{\Gamma(\mathfrak{B},\sigma)}(a) := \infty$ otherwise.*

Note that $\mathrm{Aut}(\mathfrak{B}) = \mathrm{Aut}(\Gamma(\mathfrak{B}, \sigma))$ for any τ-structure $\mathfrak{B}$ and any σ. In the following result we use a correspondence between resilience problems for incidence-acyclic conjunctive queries and valued CSPs. The result then follows from the P versus NP-complete dichotomy theorem for valued CSPs over finite domains stated in Theorem 7.11.

Theorem 8.9. *Let μ be a union of incidence-acyclic conjunctive queries with relational signature τ and let $\sigma \subseteq \tau$. Then the resilience problem for μ with exogenous relations from σ is in P or NP-complete. Moreover, it is decidable whether the resilience problem for a given union of incidence-acyclic conjunctive queries is in P. If μ is a union of queries each of which is homomorphically equivalent to a tree and $\mathfrak{B}$ is the finite dual of μ (which exists by Theorem 8.7), then* $\mathrm{VCSP}(\Gamma(\mathfrak{B}, \sigma))$ *is polynomial-time equivalent to the resilience problem for μ with exogenous relations from σ.*

Remark 8.10. *We mention that Theorem 8.9 also applies to (2-way) regular path queries, which can be shown to always have a finite dual, more details can be found in [11].*

Theorem 8.9 can be combined with the tractability results for VCSPs from Section 7 that use fractional polymorphisms. To illustrate fractional polymorphisms and how to find them, we revisit a known tractable resilience problem from [20–22, 37] and show that it has a fractional canonical pseudo cyclic polymorphism.

Example 8.11. *We revisit Example 8.2. Consider again the conjunctive query*

$$\exists x, y, z(R(x, y) \wedge S(y, z)).$$

There is a finite dual $\mathfrak{B}$ of μ with domain $\{0, 1\}$ which is finitely bounded homogeneous, as described in Example 8.2. That example also describes a valued structure Γ which is actually $\Gamma(\mathfrak{B}, \emptyset)$. Let ω be the fractional cyclic operation given by $\omega(\min) = \omega(\max) = \frac{1}{2}$. Since $\mathrm{Aut}(\Gamma)$ *is trivial, ω is canonical. The fractional operation ω improves both valued relations R and S (they are* submodular*; see, e.g., [34]) and hence is a canonical cyclic fractional polymorphism of Γ.*

Combining Theorem 7.13 and 8.9, Example 8.11 reproves the results from [21] (without multiplicities) and [36] (with multiplicities) that the resilience problem for this query is in P.

8.3 Infinite Duals

Conjunctive queries might not have a finite dual (see Example 8.3), but unions of connected conjunctive queries always have a countably infinite dual. Cherlin, Shelah and Shi [15] showed that in this case we may even find a dual with an oligomorphic automorphism group (see Theorem 8.12 below). This is the key insight to phrase resilience problems as VCSPs for valued structures with oligomorphic automorphism groups. The not necessarily connected case again reduces to the connected case by Lemma 8.5.

In Theorem 8.12 below we state a variant of a theorem of Cherlin, Shelah, and Shi [15] (also see [3, 26, 27]). If $\mathfrak{B}$ is a structure, we write $\mathfrak{B}_{\mathrm{pp}(m)}$ for the expansion of $\mathfrak{B}$ by all relations that can be defined with a connected primitive positive formula (see Remark 8.4) with at most m variables, at least one free variable, and without equality. For a union of conjunctive queries μ over the signature τ, we write $|\mu|$ for the maximum of the number of variables of each conjunctive query in μ, the maximal arity of τ, and 2.

Theorem 8.12. *For every union μ of connected conjunctive queries over a finite relational signature τ there exists a τ-structure $\mathfrak{B}_\mu$ such that the following statements hold:*

1. *$(\mathfrak{B}_\mu)_{pp(|\mu|)}$ is homogeneous.*
2. *$\mathrm{Age}(\mathfrak{B}_{pp(|\mu|)})$ is the class of all substructures of structures of the form $\mathfrak{A}_{pp(|\mu|)}$ for a finite structure $\mathfrak{A}$ that satisfies $\neg\mu$.*
3. *A countable τ-structure $\mathfrak{A}$ satisfies $\neg\mu$ if and only if it embeds into $\mathfrak{B}_\mu$.*
4. *$\mathfrak{B}_\mu$ is finitely bounded.*
5. *$\mathrm{Aut}(\mathfrak{B}_\mu)$ is oligomorphic.*
6. *$(\mathfrak{B}_\mu)_{pp(|\mu|)}$ is finitely bounded.*

By Properties (1) and (6) of Theorem 8.12, $\mathfrak{B}_\mu$ is always a reduct of a finitely bounded homogeneous structure. For short, we write Γ_μ for $\Gamma(\mathfrak{B}_\mu, \emptyset)$ and $\Gamma_{\mu,\sigma}$ for $\Gamma(\mathfrak{B}_\mu, \sigma)$, see Definition 8.8. For some queries μ, the structure $\mathfrak{B}_\mu$ can be replaced by a simpler structure

$\mathfrak{C}_\mu$. This will be convenient for some examples that we consider later, because the structure $\mathfrak{C}_\mu$ is finitely bounded and homogeneous itself and hence admits the application of Theorem 7.13. To define the respective class of queries, we need the following definition. The *Gaifman graph* of a relational structure $\mathfrak{A}$ is the undirected graph with vertex set A where $a, b \in A$ are adjacent if and only if $a \neq b$ and there exists a tuple in a relation of $\mathfrak{A}$ that contains both a and b. The Gaifman graph of a conjunctive query is the Gaifman graph of the canonical database of that query.

Theorem 8.13. *For every union μ of connected conjunctive queries over a finite relational signature τ such that the Gaifman graph of each of the conjunctive queries in μ is complete, there exists a countable τ-structure $\mathfrak{C}_\mu$ such that the following statements hold:*

(1) *$\mathfrak{C}_\mu$ is homogeneous.*
(2) *$\mathrm{Age}(\mathfrak{C}_\mu)$ is the class of all finite structures $\mathfrak{A}$ that satisfy $\neg\mu$.*

Moreover, $\mathfrak{C}_\mu$ is finitely bounded, $\mathrm{Aut}(\mathfrak{C}_\mu)$ is oligomorphic, and a countable τ-structure satisfies $\neg\mu$ if and only if it embeds into $\mathfrak{C}_\mu$.

Note that $\mathfrak{C}_\mu$ is homomorphically equivalent to $\mathfrak{B}_\mu$ by [3, Lemma 4.1.7]. Therefore, $\Gamma(\mathfrak{C}_\mu, \sigma)$ is homomorphically equivalent to $\Gamma_{\mu,\sigma}$ for any $\sigma \subseteq \tau$.

The following proposition follows straightforwardly from the definitions and provides a valued constraint satisfaction problem that is polynomial-time equivalent to the resilience problem for μ, similar to Theorem 8.9.

Proposition 8.14. *The resilience problem for a union of connected conjunctive queries μ where the relations from $\sigma \subseteq \tau$ are exogenous is polynomial-time equivalent to $\mathrm{VCSP}(\Gamma(\mathfrak{B}, \sigma))$ for any dual $\mathfrak{B}$ of μ; in particular, to $\mathrm{VCSP}(\Gamma_{\mu,\sigma})$.*

In [36] one may find a seemingly more general notion of exogenous tuples, where in a single relation there might be both endogenous and exogenous tuples. However, one can show that classifying the complexity of resilience problems according to our original definition also entails a classification of this variant, using the operator Opt and a similar reduction as in Proposition 8.14.

Remark 8.15. *Consider a union μ of conjunctive queries with the signature τ, let $\sigma \subseteq \tau$, and let $\rho \subseteq \tau \setminus \sigma$. Suppose we would like to model the resilience problem for μ where the relations in σ are exogenous and the relations in ρ might contain both endogenous and exogenous tuples. Let $\mathfrak{B}$ be a dual of μ and Γ be the expansion of $\Gamma(\mathfrak{B}, \sigma)$ where for every relational symbol $R \in \rho$, there is also a relation $(R^x)^\Gamma = R^{\mathfrak{B}}$, i.e., a classical relation that takes values 0 and ∞. The resilience problem for μ with exogenous tuples specified as above is polynomial-time equivalent to $\mathrm{VCSP}(\Gamma)$ by analogous reductions as in Proposition 8.14. Note that $(R^x)^\Gamma = \mathrm{Opt}\left(R^{\Gamma(\mathfrak{B},\sigma)}\right)$ for every $R \in \rho$, and therefore by Lemma 4.6, $\mathrm{VCSP}(\Gamma)$ is polynomial-time equivalent to $\mathrm{VCSP}(\Gamma(\mathfrak{B}, \sigma))$ and thus to the resilience problem for μ where the relations in σ are exogeneous and the relations in $\tau \setminus \sigma$ are purely endogeneous. This justifies the restriction to our setting for exogenous tuples. Moreover, the same argument shows that if resilience of μ with all tuples endogenous is in P, then all variants of resilience of μ with exogenous tuples are in P as well.*

Similarly as in Example 8.11, Proposition 8.14 can be combined with the tractability results for VCSPs from Section 7 that use fractional polymorphisms to prove tractability of resilience problems.

Example 8.16. *We revisit Example 8.3. Consider the conjunctive query $\exists x, y, z\, (R(x, y) \wedge S(x, y, z))$ over the signature $\tau = \{R, S\}$. Note that the Gaifman graph of μ is complete; let $\mathfrak{C}_\mu$ be the structure from Theorem 8.13. We construct a binary pseudo cyclic canonical fractional polymorphism of $\Gamma(\mathfrak{C}_\mu, \emptyset)$. Let $\mathfrak{M}$ be the τ-structure with domain $(C_\mu)^2$ and where*

- *$((b_1^1, b_1^2), (b_2^1, b_2^2)) \in R^{\mathfrak{M}}$ if $(b_1^1, b_2^1) \in R^{\mathfrak{C}_\mu}$ and $(b_1^2, b_2^2) \in R^{\mathfrak{C}_\mu}$*
- *$((b_1^1, b_1^2), (b_2^1, b_2^2), (b_3^1, b_3^2)) \in S^{\mathfrak{M}}$ if $(b_1^1, b_2^1, b_3^1) \in S^{\mathfrak{C}_\mu}$ or $(b_1^2, b_2^2, b_3^2) \in S^{\mathfrak{C}_\mu}$.*

Similarly, let $\mathfrak{N}$ be the τ-structure with domain $(C_\mu)^2$ and where

- *$((b_1^1, b_1^2), (b_2^1, b_2^2)) \in R^{\mathfrak{N}}$ if $(b_1^1, b_2^1) \in R^{\mathfrak{C}_\mu}$ or $(b_1^2, b_2^2) \in R^{\mathfrak{C}_\mu}$,*
- *$((b_1^1, b_1^2), (b_2^1, b_2^2), (b_3^1, b_3^2)) \in S^{\mathfrak{N}}$ if $(b_1^1, b_2^1, b_3^1) \in S^{\mathfrak{C}_\mu}$ and $(b_1^2, b_2^2, b_3^2) \in S^{\mathfrak{C}_\mu}$.*

Note that $\mathfrak{M} \not\models \mu$ and $\mathfrak{N} \not\models \mu$ and hence there are embeddings $f: \mathfrak{M} \to \mathfrak{C}_\mu$ and $g: \mathfrak{N} \to \mathfrak{C}_\mu$. Clearly, both f and g regarded as operations on the set C_μ are pseudo cyclic (but in general not cyclic) and canonical with respect to $\mathrm{Aut}(\mathfrak{C}_\mu)$ (see Claim 6 in Proposition 8.19 for a detailed argument of this type). Let ω be the fractional operation given by $\omega(f) = \frac{1}{2}$ and $\omega(g) = \frac{1}{2}$. Then ω is a binary fractional polymorphism of $\Gamma := \Gamma(\mathfrak{C}_\mu, \emptyset)$: for $b^1, b^2 \in (C_\mu)^2$ we have

$$\sum_{h \in \mathscr{O}^{(2)}} \omega(h) R^\Gamma(h(b^1, b^2)) = \frac{1}{2} R^\Gamma(f(b^1, b^2)) + \frac{1}{2} R^\Gamma(g(b^1, b^2))$$
$$= \frac{1}{2} \sum_{j=1}^{2} R^\Gamma(b^j). \qquad (3)$$

so ω improves R, and similarly we see that ω improves S.

We proved that the corresponding valued structure has a binary canonical pseudo cyclic fractional polymorphism. By Theorem 7.13 and 8.14, this reproves the results from [21] (without multiplicities) and [36] (with multiplicities) that the resilience problem for this query is in P.

8.4 The Resilience Tractability Conjecture

In this section we present a conjecture which implies, together with Corollary 5.11 and Lemma 8.5, a P versus NP-complete dichotomy for resilience problems for finite unions of conjunctive queries.

Conjecture 8.17. *Let μ be a union of connected conjunctive queries over the signature τ, and let $\sigma \subseteq \tau$. If the structure $(\{0, 1\}; \mathrm{OIT})$ has no pp-construction in $\Gamma := \Gamma_{\mu,\sigma}$, then Γ has a fractional polymorphism of arity $\ell \geq 2$ which is canonical and pseudo cyclic with respect to $\mathrm{Aut}(\Gamma)$ (and in this case, $\mathrm{VCSP}(\Gamma)$ is in P by Theorem 7.13).*

The conjecture is intentionally only formulated for VCSPs that stem from resilience problems, because it is known to be false for the more general situation of VCSPs for valued structures Γ that have the same automorphisms as a reduct of a finitely bounded homogeneous structure [3] (Section 12.9.1; the counterexample is even a CSP). However, the structures $\mathfrak{B}_\mu$ from Theorem 8.12 that we need to formulate resilience problems as VCSPs are particularly well-behaved for the universal-algebraic approach and more specifically, for canonical operations (see, e.g., [4, 6, 38]), which is why we believe in the strong formulation of Conjecture 8.17. See Conjecture 9.3 for a conjecture that could hold for VCSPs in the more general setting of reducts of finitely bounded homogeneous structures.

For the following conjunctive query μ, the NP-hardness of the resilience problem without multiplicities was shown in [21]; to illustrate our condition, we verify that the structure $(\{0,1\};\mathrm{OIT})$ has a pp-construction in Γ_μ and thus prove in a different way that the resilience problem (with multiplicities) for μ is NP-hard.

Example 8.18 (Triangle query). *Let τ be the signature that consists of three binary relation symbols R, S, and T, and let μ be the conjunctive query*

$$\exists x,y,z\big(R(x,y) \wedge S(y,z) \wedge T(z,x)\big).$$

The resilience problem without multiplicities for μ is NP-complete [21], and hence $\mathrm{VCSP}(\Gamma_\mu)$ is NP-hard (Proposition 8.14). Since the Gaifman graph of μ is NP-complete, the structure $\mathfrak{C}_\mu$ from Theorem 8.13 exists. Let $\Gamma := \Gamma(\mathfrak{C}_\mu, \emptyset)$. We provide a pp-construction of $(\{0,1\};\mathrm{OIT})$ in Γ, which also proves NP-hardness of $\mathrm{VCSP}(\Gamma)$ and hence the resilience problem of μ with multiplicities by Corollary 5.11. Since Γ is homomorphically equivalent to Γ_μ, this also provides a pp-construction of $(\{0,1\};\mathrm{OIT})$ in Γ_μ (see Lemma 5.12).

Let C be the domain of Γ. Denote the relations $\mathrm{Opt}(R)$, $\mathrm{Opt}(S)$, $\mathrm{Opt}(T)$ by R^*, S^*, T^*, respectively. In the following, for $U \in \{R,S,T\}$ and variables x,y we write $2U(x,y)$ for short instead of $U(x,y) + U(x,y)$. Let $\phi(a,b,c,d,e,f,g,h,i)$ be the τ-expression

$$R(a,b) + 2S(b,c) + 2T(c,d) + 2R(d,e) + 2S(e,f) + 2T(f,g) + 2R(g,h) + S(h,i) \quad (4)$$

$$+ T^*(i,g) + S^*(h,f) + R^*(g,e) + T^*(f,d) + S^*(e,c) + R^*(d,b) + T^*(c,a). \quad (5)$$

For an illustration of μ and ϕ, see Figure 3. Note that ϕ can be viewed as 7 non-overlapping copies of μ (if we consider the doubled constraints as two separate constraints) with some constraints forbidden to violate.

In what follows, we say that an atomic τ-expression holds if it evaluates to 0 and an atomic τ-expression is violated if it does not hold. Since there are 7 non-overlapping copies of μ in ϕ, the cost of ϕ is at least 7. Every assignment where

- *all atoms in (5) hold, and*
- *either every atom at even position or every atom at odd position in (4) holds,*

evaluates ϕ to 7 and hence is a solution to ϕ.

Let $RT \in \langle\Gamma\rangle$ be given by

$$RT(a,b,f,g) := \mathrm{Opt} \inf_{c,d,e,h,i \in C} \phi.$$

Note that $RT(a,b,f,g)$ holds if and only if

- *$R(a,b)$ holds and $T(f,g)$ does not hold, or*
- *$T(f,g)$ holds and $R(a,b)$ does not hold,*

where the reverse implication uses that $\mathfrak{C}_\mu$ is homogeneous and embeds all finite structures that do not satisfy μ. Define $RS \in \langle\Gamma\rangle$ by

$$RS(a,b,h,i) := \mathrm{Opt} \inf_{c,d,e,f,g \in C} \phi.$$

Note that $RS(a,b,h,i)$ holds if and only if

- *$R(a,b)$ holds and $S(h,i)$ does not hold, or*
- *$S(h,i)$ holds and $R(a,b)$ does not hold.*

Next, we define the auxiliary relation $\widetilde{RS}(a,b,e,f)$ to be

$$\mathrm{Opt} \inf_{c,d,g,h,i \in C} \phi.$$

Note that $\widetilde{RS}(a,b,e,f)$ holds if and only if

- *both $R(a,b)$ and $S(e,f)$ hold, or*
- *neither $R(a,b)$ and nor $S(e,f)$ holds.*

This allows us to define the relation

$$RR(u,v,x,y) := \inf_{w,z \in C} RS(u,v,w,z) + \widetilde{RS}(x,y,w,z)$$

which holds if and only if

- *$R(u,v)$ holds and $R(x,y)$ does not hold, or*
- *$R(x,y)$ holds and $R(u,v)$ does not hold.*

Define $M \in \langle\Gamma\rangle$ as

$$M(u,v,u',v',u'',v'') := \mathrm{Opt} \inf_{x,y,z \in C} \big(R(x,y) + S(y,z) + T(z,x) + RR(u,v,x,y) + RS(u',v',y,z) + RT(u'',v'',z,x)\big).$$

Note that $R(x,y)$, $S(y,z)$ and $T(z,x)$ cannot hold at the same time and therefore $(u,v,u',v',u'',v'') \in M$ if and only if exactly one of of $R(u,v)$, $R(u',v')$, and $R(u'',v'')$ holds. Let Δ be the pp-power of $(C;M)$ of dimension two with signature $\{\mathrm{OIT}\}$ such that

$$\mathrm{OIT}^\Delta((u,v),(u',v'),(u'',v'')) := M(u,v,u',v',u'',v'').$$

Then Δ is homomorphically equivalent to $(\{0,1\};\mathrm{OIT})$, witnessed by the homomorphism from Δ to $(\{0,1\};\mathrm{OIT})$ that maps (u,v) to 1 if $R(u,v)$ and to 0 otherwise, and the homomorphism $(\{0,1\};\mathrm{OIT}) \to \Delta$ that maps 1 to any pair of vertices $(u,v) \in R$ and 0 to any pair of vertices $(u,v) \notin R$. Therefore, Γ pp-constructs $(\{0,1\};\mathrm{OIT})$.

We mention that another conjecture concerning a P vs. NP-complete complexity dichotomy for resilience problems appears in [36, Conjecture 7.7]. The conjecture has a similar form as Conjecture 8.17 in the sense that it states that a sufficient hardness condition for resilience is also necessary. The relationship between our hardness condition from Corollary 5.11 and the condition from [36] remains open.

8.5 An example of formerly open complexity

We use our approach to settle the complexity of the resilience problem for a conjunctive query that was mentioned as an open problem in [22] (Section 8.5):

$$\mu := \exists x,y (S(x) \wedge R(x,y) \wedge R(y,x) \wedge R(y,y)) \quad (6)$$

Proposition 8.19. *There is a finitely bounded homogeneous dual $\mathfrak{B}$ of μ such that the valued τ-structure $\Gamma := \Gamma(\mathfrak{B}, \emptyset)$ has a binary fractional polymorphism which is canonical and pseudo cyclic with respect to $\mathrm{Aut}(\Gamma)$. Hence, $\mathrm{VCSP}(\Gamma)$ and the resilience problem for μ are in P. The polynomial-time tractability result even holds for resilience of μ with exogeneous relations from any $\sigma \subseteq \tau$.*

9 CONCLUSION AND FUTURE WORK

We formulated a general hardness condition for VCSPs of valued structures with an oligomorphic automorphism group and a new polynomial-time tractability result. We use the latter to resolve a resilience problem whose complexity was left open in the literature and conjecture that our conditions exactly capture the hard and easy

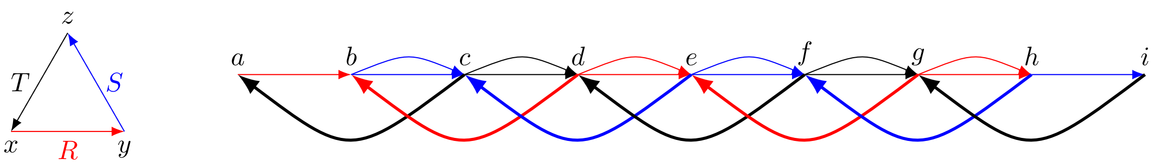

Figure 3: Example 8.18, visualisation of μ and ϕ. The thick edges cannot be removed.

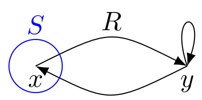

Figure 4: Visualisation of the query μ from (6).

resilience problems for conjunctive queries (under bag semantics), respectively. In fact, a full classification of resilience problems for conjunctive queries based on our approach seems feasible, but requires further research, as discussed in the following.

We have proved that if Γ is a valued structure with an oligomorphic automorphism group and R is a valued relation in the smallest valued relational clone that contains the valued relations of Γ, then R is preserved by all fractional polymorphisms of Γ (Lemma 6.8). We do not know whether the converse is true. It is known to hold for the special cases of finite-domain valued structures [17, 23] and for classical relational structures with 0-∞ valued relations (CSP setting) having an oligomorphic automorphism group [9].

Question 9.1. *Let Γ be a valued structure with an oligomorphic automorphism group. Is it true that $R \in \langle\Gamma\rangle$ if and only if $R \in \mathrm{Imp}(\mathrm{fPol}(\Gamma))$?*

A natural attempt to positively answer Question 9.1 would be to combine the proof strategy for finite-domain valued structures from [17, 23] with the one for relational structures with oligomorphic automorphism group from [9]. However, since non-improving of R is not a closed condition, the compactness argument from [9] cannot be used to construct an operation from $\mathrm{fPol}(\Gamma)$ that does not improve R. A positive answer to Question 9.1 would imply that the computational complexity of VCSPs for valued structures Γ with an oligomorphic automorphism group, and in particular the complexity of resilience problems, is fully determined by the fractional polymorphisms of Γ.

Note that in all examples that arise from resilience problems that we considered so far, it was sufficient to work with fractional polymorphisms ω that are *finitary*, i.e., there are finitely many operations $f_1, \dots, f_k \in \mathscr{O}_C$ such that $\sum_{i\in\{1,\dots,k\}} \omega(f_i) = 1$. It is therefore possible that all fractional polymorphisms relevant for resilience have discrete probability distributions. This motivates the following question.

Question 9.2.

- *Does our notion of pp-constructability change if we restrict to finitary fractional homomorphisms ω?*
- *Is there a valued structure Γ with an oligomorphic automorphism group and a valued relation R such that R is not improved by all fractional polymorphism of Γ, but* is *improved by all finitary fractional polymorphisms ω?*
- *In particular, are these statements true if we restrict to valued τ-structures Γ that arise from resilience problems as described in Proposition 8.14?*

In the following, we formulate a common generalisation of the complexity-theoretic implications of Conjecture 8.17 and the infinite-domain tractability conjecture from [10] that concerns a full complexity classification of VCSPs for valued structures from reducts of finitely bounded homogeneous structures.

Conjecture 9.3. *Let Γ be a valued structure with finite signature such that $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\mathfrak{B})$ for some reduct $\mathfrak{B}$ of a countable finitely bounded homogeneous structure. If $(\{0,1\}; \mathrm{OIT})$ has no pp-construction in Γ, then $\mathrm{VCSP}(\Gamma)$ is in P (otherwise, we already know that $\mathrm{VCSP}(\Gamma)$ is NP-complete by Theorem 3.4 and Corollary 5.11).*

One might hope to prove this conjecture under the assumption of the infinite-domain tractability conjecture. Recall that also the finite-domain VCSP classification was first proven conditionally on the finite-domain tractability conjecture [31, 33], which was only confirmed later [12, 45].

We also believe that the 'meta-problem' of deciding whether for a given conjunctive query the resilience problem is in P is decidable. This would follow from a positive answer to Conjecture 8.17 because Γ_m^* can be computed and Item 4 of Proposition 7.12 for the finite-domain valued structure Γ_m^* can be decided algorithmically using linear programming [30].

ACKNOWLEDGMENTS

The first two authors have been funded by the European Research Council (Project POCOCOP, ERC Synergy Grant 101071674) and by the DFG (Project FinHom, Grant 467967530). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. The third author was supported by the DFG project LU 1417/3-1 QTEC.

REFERENCES

[1] Nikhil Bansal, Avrim Blum, and Shuchi Chawla. 2004. Correlation clustering. *Machine Learning* 56, 1-3 (2004), 89–113.

[2] Libor Barto, Jakub Opršal, and Michael Pinsker. 2018. The wonderland of reflections. *Israel Journal of Mathematics* 223, 1 (2018), 363–398.

[3] Manuel Bodirsky. 2021. *Complexity of Infinite-Domain Constraint Satisfaction.* Cambridge University Press. https://doi.org/10.1017/9781107337534

[4] Manuel Bodirsky and Bertalan Bodor. 2024. A Complexity Dichotomy in Spatial Reasoning via Ramsey Theory. 16, 2, Article 10 (jun 2024), 39 pages. https://doi.org/10.1145/3649445

[5] Manuel Bodirsky and Jan Kára. 2008. The Complexity of Equality Constraint Languages. *Theory of Computing Systems* 3, 2 (2008), 136–158. https://doi.org/10.1007/s00224-007-9083-9 A conference version appeared in the proceedings of Computer Science Russia (CSR'06).

[6] Manuel Bodirsky, Florent Madelaine, and Antoine Mottet. 2018. A universal-algebraic proof of the complexity dichotomy for Monotone Monadic SNP. In *Proceedings of the Symposium on Logic in Computer Science (LICS)*. 105–114. Preprint arXiv:1802.03255.

[7] Manuel Bodirsky, Marcello Mamino, and Caterina Viola. 2022. Piecewise Linear Valued CSPs Solvable by Linear Programming Relaxation. *ACM Transactions of Computational Logic* 23, 1 (2022), 1–35. https://doi.org/10.1145/3488721 Preprint arXiv:1912.09298.

[8] Manuel Bodirsky and Antoine Mottet. 2018. A Dichotomy for First-Order Reducts of Unary Structures. *Logical Methods in Computer Science* 14, 2 (2018). https://doi.org/10.23638/LMCS-14(2:13)2018

[9] Manuel Bodirsky and Jaroslav Nešetřil. 2003. Constraint Satisfaction with Countable Homogeneous Templates. In *Proceedings of Computer Science Logic (CSL)*. Vienna, 44–57.

[10] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. 2021. Projective clone homomorphisms. *Journal of Symbolic Logic* 86, 1 (2021), 148–161.

[11] Manuel Bodirsky, Žaneta Semanišinová, and Carsten Lutz. 2023. The Complexity of Resilience Problems via Valued Constraint Satisfaction Problems. arXiv:2309.15654 [math.LO]

[12] Andrei A. Bulatov. 2017. A Dichotomy Theorem for Nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17*. 319–330.

[13] Peter J. Cameron. 1990. *Oligomorphic permutation groups*. Cambridge University Press, Cambridge.

[14] Surajit Chaudhuri and Moshe Y. Vardi. 1993. Optimization of Real Conjunctive Queries. In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems* (Washington, D.C., USA) *(PODS '93)*. Association for Computing Machinery, New York, NY, USA, 59–70. https://doi.org/10.1145/153850.153856

[15] Gregory Cherlin, Saharon Shelah, and Niangdong Shi. 1999. Universal graphs with forbidden subgraphs and algebraic closure. *Advances in Applied Mathematics* 22 (1999), 454–491.

[16] David A. Cohen, Martin C. Cooper, Páidí Creed, Peter G. Jeavons, and Stanislav Živný. 2013. An Algebraic Theory of Complexity for Discrete Optimization. *SIAM J. Comput.* 42, 5 (2013), 1915–1939.

[17] David A. Cohen, Martin C. Cooper, and Peter Jeavons. 2006. An Algebraic Characterisation of Complexity for Valued Constraints. In *Principles and Practice of Constraint Programming - CP 2006, 12th International Conference, CP 2006, Nantes, France, September 25-29, 2006, Proceedings*. 107–121. https://doi.org/10.1007/11889205_10

[18] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. 2015. *Parameterized Algorithms* (1st ed.). Springer Publishing Company, Incorporated.

[19] Jan Foniok. 2007. Homomorphisms and Structural Properties of Relational Systems. Doctoral Thesis, Charles University Prague, Faculty of Mathematics and Physics.

[20] Cibele Freire, Wolfgang Gatterbauer, Neil Immerman, and Alexandra Meliou. 2015. A Characterization of the Complexity of Resilience and Responsibility for Self-join-free Conjunctive Queries. *CoRR* abs/1507.00674 (2015). arXiv:1507.00674 http://arxiv.org/abs/1507.00674

[21] Cibele Freire, Wolfgang Gatterbauer, Neil Immerman, and Alexandra Meliou. 2015. The Complexity of Resilience and Responsibility for Self-Join-Free Conjunctive Queries. *Proc. VLDB Endow.* 9, 3 (2015), 180–191. https://doi.org/10.14778/2850583.2850592

[22] Cibele Freire, Wolfgang Gatterbauer, Neil Immerman, and Alexandra Meliou. 2020. New Results for the Complexity of Resilience for Binary Conjunctive Queries with Self-Joins. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2020, Portland, OR, USA, June 14-19, 2020*. 271–284. https://doi.org/10.1145/3375395.3387647

[23] Peter Fulla and Stanislav Živný. 2016. A Galois Connection for Weighted (Relational) Clones of Infinite Size. *TOCT* 8, 3 (2016), 9:1–9:21.

[24] Michael Garey and David Johnson. 1978. *A guide to NP-completeness*. CSLI Press, Stanford.

[25] Wilfrid Hodges. 1997. *A shorter model theory*. Cambridge University Press, Cambridge.

[26] Jan Hubička and Jaroslav Nešetřil. 2015. Universal Structures with Forbidden Homomorphisms. In *Logic Without Borders - Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics*. De Gruyter, 241–264. arXiv:0907.4079.

[27] Jan Hubička and Jaroslav Nešetřil. 2016. Homomorphism and Embedding Universal Structures for Restricted Classes. *Multiple-Valued Logic and Soft Computing* 27, 2-3 (2016), 229–253. arXiv:0909.4939.

[28] Eun Jung Kim, Stefan Kratsch, Marcin Pilipczuk, and Magnus Wahlström. 2021. Solving hard cut problems via flow-augmentation. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*. 149–168.

[29] Eun Jung Kim, Stefan Kratsch, Marcin Pilipczuk, and Magnus Wahlström. 2022. Directed flow-augmentation. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*. 938–947.

[30] Vladimir Kolmogorov. 2019. Testing the Complexity of a Valued CSP Language. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 132)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 77:1–77:12. https://doi.org/10.4230/LIPIcs.ICALP.2019.77

[31] Vladimir Kolmogorov, Andrei A. Krokhin, and Michal Rolínek. 2017. The Complexity of General-Valued CSPs. *SIAM J. Comput.* 46, 3 (2017), 1087–1110. https://doi.org/10.1137/16M1091836

[32] Vladimir Kolmogorov, Johan Thapper, and Stanislav Živný. 2015. The Power of Linear Programming for General-Valued CSPs. *SIAM J. Comput.* 44, 1 (2015), 1–36.

[33] Marcin Kozik and Joanna Ochremiak. 2015. Algebraic Properties of Valued Constraint Satisfaction Problem. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*. 846–858.

[34] Andrei Krokhin and Stanislav Živný. 2017. The Complexity of Valued CSPs. In *The Constraint Satisfaction Problem: Complexity and Approximability*, Andrei Krokhin and Stanislav Zivny (Eds.). Dagstuhl Follow-Ups, Vol. 7. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 233–266. https://doi.org/10.4230/DFU.Vol7.15301.233

[35] Benoit Larose, Cynthia Loten, and Claude Tardif. 2007. A Characterisation of First-Order Constraint Satisfaction Problems. *Logical Methods in Computer Science* 3, 4:6 (2007).

[36] Neha Makhija and Wolfgang Gatterbauer. 2023. A Unified Approach for Resilience and Causal Responsibility with Integer Linear Programming (ILP) and LP Relaxations. *Proc. ACM Manag. Data* 1, 4, Article 228 (dec 2023), 27 pages. https://doi.org/10.1145/3626715

[37] Alexandra Meliou, Wolfgang Gatterbauer, Katherine F. Moore, and Dan Suciu. 2010. The Complexity of Causality and Responsibility for Query Answers and non-Answers. *Proc. VLDB Endow.* 4, 1 (2010), 34–45. https://doi.org/10.14778/1880172.1880176

[38] Antoine Mottet and Michael Pinsker. 2020. Smooth Approximations and CSPs over finitely bounded homogeneous structures. (2020). Preprint arXiv:2011.03978.

[39] Jaroslav Nešetřil and Claude Tardif. 2000. Duality theorems for finite structures. *Journal of Combinatorial Theory, Series B* 80 (2000), 80–97.

[40] Friedrich Martin Schneider and Caterina Viola. 2022. An Application of Farkas' Lemma to Finite-Valued Constraint Satisfaction Problems over Infinite Domains. *CoRR* abs/2208.04912 (2022). https://doi.org/10.48550/arXiv.2208.04912 arXiv:2208.04912

[41] Johan Thapper and Stanislav Živný. 2013. The complexity of finite-valued CSPs. In *Proceedings of the Symposium on Theory of Computing Conference (STOC), Palo Alto, CA, USA, June 1-4, 2013*. 695–704.

[42] Caterina Viola. 2020. *Valued Constraint Satisfaction Problems over Infinite Domains*. Ph. D. Dissertation. TU Dresden.

[43] Caterina Viola and Stanislav Živný. 2021. The Combined Basic LP and Affine IP Relaxation for Promise VCSPs on Infinite Domains. *ACM Trans. Algorithms* 17, 3 (2021), 21:1–21:23. https://doi.org/10.1145/3458041

[44] Dmitriy Zhuk. 2020. A Proof of the CSP Dichotomy Conjecture. *J. ACM* 67, 5 (2020), 30:1–30:78. https://doi.org/10.1145/3402029

[45] Dmitriy N. Zhuk. 2017. A Proof of CSP Dichotomy Conjecture. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17*. 331–342. https://arxiv.org/abs/1704.01914..

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009