

The Complexity of Resilience for Digraph Queries

Manuel Bodirsky  

Institut für Algebra, TU Dresden, Germany

Žaneta Semaništinová  

Institute of Discrete Mathematics and Geometry, TU Wien, Austria

Abstract

We prove a complexity dichotomy for the resilience problem for unions of conjunctive digraph queries (i.e., for existential positive sentences over the signature $\{R\}$ of directed graphs). Specifically, for every union μ of conjunctive digraph queries, the following problem is in P or NP-complete: given a directed multigraph G and a natural number u , can we remove u edges from G so that $G \models \neg\mu$? In fact, we verify a more general dichotomy conjecture from [6] for all resilience problems in the special case of directed graphs, and show that for such unions of queries μ there exists a countably infinite (‘dual’) valued structure Δ_μ which either primitively positively constructs 1-in-3-3-SAT, and hence the resilience problem for μ is NP-complete by general principles, or has a pseudo cyclic canonical fractional polymorphism, and the resilience problem for μ is in P.

2012 ACM Subject Classification Theory of computation \rightarrow Problems, reductions and completeness; Theory of computation \rightarrow Complexity theory and logic; Theory of computation \rightarrow Database query processing and optimization (theory)

Keywords and phrases valued constraints, unions of conjunctive queries, resilience, computational complexity, pp-constructions

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.15

Related Version *Full Version:* <https://arxiv.org/abs/2601.05346> [7]

Funding *Manuel Bodirsky:* The author has 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.

Žaneta Semaništinová: The author has been funded by the European Research Council (Project POCOCOP, ERC Synergy Grant 101071674). 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. This research was funded in whole or in part by the Austrian Science Fund (FWF) 10.55776/ESP6949724.

1 Introduction

The *resilience problem* for a fixed conjunctive query, or more generally for a union of conjunctive queries μ , is the problem of deciding for a given database \mathfrak{A} and $u \in \mathbb{N}$ whether it is possible to remove at most u tuples from \mathfrak{A} so that \mathfrak{A} does not satisfy μ . The resilience problem lies at the core of algorithmic challenges in various forms of reverse data management, where an action is required on the input data to achieve a desired outcome in the output data [18]. The computational complexity of this problem depends on the query μ . The resilience problem is always in NP, and often NP-complete, but for some queries μ the problem can be solved in polynomial time; see, e.g., [12, 13, 17, 6] for some partial classification results.



© Manuel Bodirsky and Žaneta Semaništinová;

licensed under Creative Commons License CC-BY 4.0

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).

Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thăng

Article No. 15; pp. 15:1–15:20



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



The computational complexity of the problem also depends on whether we view the database \mathfrak{A} under set semantics (i.e., \mathfrak{A} is treated as a relational structure) or under bag semantics (i.e., \mathfrak{A} is a structure where each tuple appears with some multiplicity), and both settings have been studied in the literature (in particular, see [6, 17] for results in bag semantics). The importance of bag semantics stems from applications: bag databases represent SQL databases more faithfully. There are examples of conjunctive queries μ for which the resilience problem in bag and set semantics have different complexities [17].

Recently, a connection between the resilience problem under bag semantics and valued constraint satisfaction has been established [6]. The connection is based on the fact that for every union of connected conjunctive queries μ , the resilience problem in bag semantics is equal to a valued constraint satisfaction problem (VCSP) for some template dependent on μ , and therefore the algebraic tools developed for describing the complexity of VCSPs can be utilized. To do so, we focus in this paper on the resilience problem exclusively under bag semantics and from now on, always implicitly assume this semantics for resilience problems.

It has been conjectured that resilience problems exhibit a complexity dichotomy in the sense that all problems are NP-complete or in P [17]. This conjecture has been verified in some special cases, for instance if μ is a conjunctive query which is self-join-free [17], or if μ is a union of conjunctive queries that are Berge-acyclic [6]. The proof of the latter is based on a connection to finite-domain VCSPs, which also covers resilience problems for regular path queries (RPQs), and even two-way RPQs. Resilience of (one-way) RPQs has also been studied recently in [1] where the authors present language-theoretic conditions for computational hardness. However, the conjecture in full generality remains open.

In this article, we confirm the complexity dichotomy conjecture in the special case where the signature of the database consists of a single binary relation symbol R , that is, we prove the following:

► **Theorem 1.** *If μ is a union of conjunctive queries over a binary signature $\{R\}$, then the resilience problem for μ is in P or NP-complete.*

The class of unions of conjunctive queries over $\{R\}$ is incomparable with the class of queries studied in [13] (in set semantics), since they study queries with arbitrary signatures, but with a single repetition of a single binary relation symbol. In the case that the signature is equal to $\{R\}$, the query expresses a directed graph property and the resilience problem can be phrased as follows: given a directed multigraph G and a natural number u , can we remove u edges from G so that $G \models \neg\mu$? Edge-removal problems have been studied from a computational complexity perspective in the graph theory community as well, especially for concrete properties [15, Section A1.2]. In [11] the authors study edge-removal problems for first-order logic properties in general; however, they only consider simple undirected graphs and study the problem from the perspective of *parametrized complexity*, where the number of edges that is removed is the parameter.

The scope of our contribution extends beyond verifying the complexity dichotomy conjecture for digraph resilience problems: we also verify a variant of a stronger conjecture (from [6]) which provides a precise mathematical condition aiming at predicting the border between NP-hardness and polynomial-time tractability, based on simulations of a hard Boolean constraint satisfaction problem (CSP) using so-called *pp-constructions*. This condition is one-sided correct in the sense that if it applies, the corresponding resilience problem is NP-hard. The authors of [6] conjectured that if the condition does not apply, the resilience problem is in P.

Several results in the present paper are relevant for the larger research goal of classifying the complexity of all resilience problems in bag semantics by modeling them as VCSPs and applying methods and results from the VCSP literature. For instance, our result that the two

conditions of the dichotomy statement are disjoint (Corollary 33) holds for resilience problems in general (without the assumption that $P \neq NP$). Another result that holds for resilience in general is Theorem 40, which provides pp-constructions (and, therefore, polynomial-time reductions) based on the idea of *self-join variations* from [13]. We believe that this paper is an important step towards classifying the complexity for resilience problems of queries with self joins and understanding reductions between resilience problems by algebraic and logic tools.

2 Preliminaries

In this section, we provide preliminaries that cover the notions appearing in Section 4, where the main theorem of the article (Theorem 34) is stated. Since the theorem provides not only a complexity dichotomy, but also an algebraic one, this requires several notions from the theory of VCSPs. For readers mostly interested in the complexity of resilience problems on its own, we recommend reading only Sections 2.1–2.4 and skipping Sections 2.5–2.8, and coming back to them when they are needed in the proofs in the article.

The set $\{0, 1, 2, \dots\}$ of natural numbers is denoted by \mathbb{N} . For $k \in \mathbb{N}$, the set $\{1, \dots, k\}$ will be denoted by $[k]$. The set of rational numbers is denoted by \mathbb{Q} and the standard strict linear order on \mathbb{Q} by $<$. The set of real numbers is denoted by \mathbb{R} . We also need an additional value ∞ ; all we need to know about ∞ is that

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

Let A be a set and $k \in \mathbb{N}$. If $t \in A^k$, then we implicitly assume that $t = (t_1, \dots, t_k)$, where $t_1, \dots, t_k \in A$. If $\ell \in \mathbb{N}$ and $f: A^\ell \rightarrow A$ is an operation on A and $t^1, \dots, t^\ell \in A^k$, then we denote $(f(t_1^1, t_1^2, \dots, t_1^\ell), \dots, f(t_k^1, t_k^2, \dots, t_k^\ell))$ by $f(t^1, \dots, t^\ell)$ and say that f is *applied componentwise*.

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 $R: C^k \rightarrow \mathbb{Q} \cup \{\infty\}$. We write $\mathcal{R}_C^{(k)}$ for the set of all valued relations over C of arity k , and define

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

A valued relation is called *finite-valued* if it takes values only in \mathbb{Q} . Usual relations will also be called *crisp* relations. A valued relation $R \in \mathcal{R}_C^{(k)}$ that only takes values from $\{0, \infty\}$ will be identified with the crisp relation $\{t \in C^k \mid R(t) = 0\}$. The unary empty relation, where every element evaluates to ∞ , is denoted by \perp . The crisp equality relation, where a pair of elements evaluates to 0 if they are equal and evaluates to ∞ otherwise, is denoted by $(=)^\infty$. For $R \in \mathcal{R}_C^{(k)}$ the *feasibility relation of R* is defined as $\text{Feas}(R) := \{t \in C^k \mid R(t) < \infty\}$.

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 $R^\Gamma \in \mathcal{R}_C^{(k)}$ for each relation symbol $R \in \tau$ of arity k . All valued structures in this article have countable domains. We often write R instead of R^Γ if the valued structure is clear from the context. A valued τ -structure where all valued relations only take values from $\{0, \infty\}$ may be viewed as a *relational* or *crisp* τ -structure in the classical sense. When not specified, we assume that the domains of relational structures $\mathfrak{A}, \mathfrak{B}, \dots$ are denoted A, B, \dots , respectively, and the domains of valued structures Γ, Δ, \dots are denoted C, D, \dots , respectively.

► **Example 2.** Let R be a binary relation symbol. Then Γ_{MC} with the domain $\{0, 1\}$ and the signature $\{R\}$ where $R^{\Gamma_{\text{MC}}}(x, y) = 0$ if $x = 0$ and $y = 1$, and $R^{\Gamma_{\text{MC}}}(x, y) = 1$ otherwise, is a valued structure.

If $\sigma \subseteq \tau$ and Γ' is a valued σ -structure such that $R^{\Gamma'} = R^{\Gamma}$ for every $R \in \sigma$, then we call Γ' a *reduct* of Γ and Γ an *expansion* of Γ' .

Let τ be a relational signature. A first-order formula is called atomic if it is of the form $R(x_1, \dots, x_k)$ for some $R \in \tau$ of arity k , $x = y$, or \perp . We introduce a generalization of conjunctions of atomic formulas to the valued setting. An *atomic τ -expression* is an expression of the form $R(x_1, \dots, x_k)$ for $R \in \tau \cup \{(\text{=})_0^\infty, \perp\}$ and (not necessarily distinct) variable symbols x_1, \dots, x_k . A *τ -expression* is an expression ϕ of the form $\sum_{i=1}^m \phi_i$ where $m \in \mathbb{N}$ and ϕ_i for $i \in \{1, \dots, m\}$ is an atomic τ -expression. Note that the same atomic τ -expression might appear several times in the sum. We write $\phi(x_1, \dots, x_n)$ for a τ -expression where all the variables are from the set $\{x_1, \dots, x_n\}$. If Γ is a valued τ -structure, then a τ -expression $\phi(x_1, \dots, x_n)$ defines over Γ a member of $\mathcal{R}_C^{(n)}$ in a natural way, which we denote by ϕ^Γ . If ϕ is the empty sum then ϕ^Γ is constant 0.

► **Definition 3.** Let $k \in \mathbb{N}$, let $R \in \mathcal{R}_C^{(k)}$, and let α be a permutation of C . Then α preserves R if for all $t \in C^k$ we have $R(\alpha(t)) = R(t)$. 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 Γ .

The set of all automorphisms of Γ is denoted by $\text{Aut}(\Gamma)$, and forms a group with respect to composition.

Let A be a set and $R \subseteq A^k$. An operation $f: A^\ell \rightarrow A$ on the set A preserves R if $f(t^1, \dots, t^\ell) \in R$ for every $t^1, \dots, t^\ell \in R$. If \mathfrak{A} is a relational structure and f preserves all relations of \mathfrak{A} , then f is called a *polymorphism* of \mathfrak{A} . The set of all polymorphisms of \mathfrak{A} is denoted by $\text{Pol}(\mathfrak{A})$ and is closed under composition. We write $\text{Pol}^{(\ell)}(\mathfrak{A})$ for the set of ℓ -ary operations in $\text{Pol}(\mathfrak{A})$, $\ell \in \mathbb{N}$. Unary polymorphisms are called *endomorphisms* and $\text{Pol}^{(1)}(\mathfrak{A})$ is also denoted by $\text{End}(\mathfrak{A})$.

Let τ be a relational signature and let \mathfrak{A} and \mathfrak{B} be relational τ -structures. A map $h: A \rightarrow B$ is called a *homomorphism* from \mathfrak{A} to \mathfrak{B} if for every $R \in \tau$ of arity k and every $t \in R^{\mathfrak{A}}$, $h(t) \in R^{\mathfrak{B}}$. \mathfrak{A} and \mathfrak{B} are called *homomorphically equivalent* if there is a homomorphism from \mathfrak{A} to \mathfrak{B} and from \mathfrak{B} to \mathfrak{A} , and they are called *homomorphically incomparable* if there is no homomorphism from \mathfrak{A} to \mathfrak{B} or from \mathfrak{B} to \mathfrak{A} . The generalizations of the notions of polymorphisms and homomorphisms to valued structures will be defined in Sections 2.7 and 2.8.

2.2 Valued constraint satisfaction problems

In this section we assume that Γ is a fixed valued τ -structure for a *finite* signature τ . We first define the valued constraint satisfaction problem of a relational structure and then explain the connection to the less general constraint satisfaction problem in Remark 5.

► **Definition 4.** The valued constraint satisfaction problem for Γ , denoted by $\text{VCSP}(\Gamma)$, is the computational problem to decide for a given τ -expression $\phi(x_1, \dots, x_n)$ and a given $u \in \mathbb{Q}$ whether there exists $t \in C^n$ such that $\phi^\Gamma(t) \leq u$. We refer to $\phi(x_1, \dots, x_n)$ as an instance of $\text{VCSP}(\Gamma)$, and to u as the threshold. Tuples $t \in C^n$ such that $\phi^\Gamma(t) \leq u$ are called a solution for (ϕ, u) . The cost of ϕ (with respect to Γ) is defined to be

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

In some contexts, it is beneficial to consider only a given τ -expression ϕ to be the input of $\text{VCSP}(\Gamma)$ (rather than ϕ and the threshold u) and a tuple $t \in C^n$ is then called a *solution for ϕ* if the cost of ϕ equals $\phi^\Gamma(t)$. Note that in general there might not be any solution; however, this is never the case for VCSPs considered in this paper as they stem from resilience problems. If there exists a tuple $t \in C^n$ such that $\phi^\Gamma(t) < \infty$ then ϕ is called *satisfiable*.

For relational structures, VCSPs specialize to CSPs, as explained below.

► **Remark 5.** If \mathfrak{A} is a relational τ -structure, then $\text{CSP}(\mathfrak{A})$ is the problem of deciding satisfiability of conjunctions of atomic formulas over τ in \mathfrak{A} . Note that for every τ -expression $\phi(x_1, \dots, x_n)$, $\phi^{\mathfrak{A}}$ defines a crisp relation and can be viewed as a conjunction of atomic formulas, which defines the same relation. Minimizing $\phi^{\mathfrak{A}}$ then corresponds to finding $t \in A^n$ such that $\phi^{\mathfrak{A}}(t) = 0$, i.e. t that satisfies all atomic formulas in the conjunction. Therefore, $\text{VCSP}(\mathfrak{A})$ and $\text{CSP}(\mathfrak{A})$ are essentially the same problem.

► **Example 6.** The problem $\text{VCSP}(\Gamma_{\text{MC}})$ for the valued structure Γ_{MC} from Example 2 models the *directed max-cut* problem: given a finite directed multigraph (V, E) , find a partition of 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 answer corresponds to finding the evaluation map $f: V \rightarrow \{0, 1\}$ that minimises the value

$$\sum_{(x,y) \in E} R^{\Gamma_{\text{MC}}}(f(x), f(y)),$$

which can be formulated as an instance of $\text{VCSP}(\Gamma_{\text{MC}})$. Conversely, every instance of $\text{VCSP}(\Gamma_{\text{MC}})$ corresponds to a directed max-cut instance.

► **Example 7.** Consider the relation $\text{OIT} = \{(0, 0, 1), (0, 1, 0), (1, 0, 0)\}$. $\text{CSP}(\{0, 1\}; \text{OIT})$ is the so called 1-in-3-SAT problem, which is known to be NP-complete (see, e.g., [4, Example 1.2.2]).

2.3 Conjunctive queries and resilience

A first-order formula is called *primitive positive* if it is an existentially quantified conjunction of atomic formulas. A *conjunctive query* over a (relational) signature τ 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.

If \mathfrak{A} is a relational τ -structure and μ is a union of conjunctive queries over τ with a quantifier-free part $\mu'(v_1, \dots, v_n)$, we say that $\alpha: \{v_1, \dots, v_n\} \rightarrow A$ *witnesses that $\mathfrak{A} \models \mu$* if $\mathfrak{A} \models \mu'(\alpha(v_1), \dots, \alpha(v_n))$. Given conjunctive queries μ_1 and μ_2 over τ , we say that μ_1 is *equivalent* to μ_2 if $\mathfrak{A} \models \mu_1$ if and only if $\mathfrak{A} \models \mu_2$ for every finite relational τ -structure \mathfrak{A} . We say a conjunctive query μ is *minimal* if every conjunctive query which is equivalent to μ has at least as many atoms as μ . For every conjunctive query μ , there exists a minimal equivalent query μ' that can be obtained from μ by removing zero or more atoms [9].

A *multiset relation* on a set A of arity k is a multiset with elements from A^k and a *bag database* \mathfrak{A} over a relational signature τ consists of a finite domain A and for every $R \in \tau$ of arity k , a multiset relation $R^{\mathfrak{A}}$ of arity k . A bag database \mathfrak{A} satisfies a union of conjunctive queries μ if the relational structure obtained from \mathfrak{A} by forgetting the multiplicities of tuples in its relations satisfies μ . In the present paper, we study the resilience problem for unions of conjunctive queries in bag semantics; from now on we will refer to this problem just as the

resilience problem. Let τ be a finite relational signature and μ a union of conjunctive queries over τ . The input to the *resilience problem for μ* consists of a bag database \mathfrak{A} over τ , and the task is to compute the number of tuples that have to be removed from relations of \mathfrak{A} so 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. It is easy to see that the resilience problem for any union of conjunctive queries is in NP.

The *canonical database* of a conjunctive query μ with relational signature τ is the relational τ -structure \mathfrak{A} whose domain are the variables of μ and where $(x_1, \dots, x_k) \in R^{\mathfrak{A}}$ for $R \in \tau$ of arity k if and only if μ contains the conjunct $R(x_1, \dots, x_k)$; we denote the canonical database by \mathfrak{D}_μ .

► **Remark 8.** All terminology introduced for τ -structures also applies to conjunctive queries over τ : by definition, a query has the property if its canonical database has the property.

Note that by the above remark, we can talk about homomorphisms between queries and queries being homomorphically incomparable. Observe that if two queries are non-equivalent and minimal, they must be homomorphically incomparable (see, e.g., [9]).

A relational τ -structure is *connected* if it cannot be written as the disjoint union of two relational τ -structures with non-empty domains. We show that when classifying the resilience problem for conjunctive queries, it suffices to consider queries that are connected.

► **Lemma 9** ([6, Lemma 8.5]). *Let ν_1, \dots, ν_k be conjunctive queries such that ν_i does not imply ν_j if $i \neq j$. Let $\nu = (\nu_1 \wedge \dots \wedge \nu_k)$ 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 for each μ_i , then the resilience problem for μ is in P as well.*

By applying Lemma 9 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.

2.4 Connection between resilience and VCSPs

In this section we summarize the key points of the connection between resilience problems and VCSPs, originally introduced in [6].

► **Definition 10.** *Let \mathfrak{B} be a relational τ -structure. Define \mathfrak{B}_0^1 to be the valued τ -structure on the same domain as \mathfrak{B} such that for each $R \in \tau$, $R^{\mathfrak{B}_0^1}(a) = 0$ if $a \in R^{\mathfrak{B}}$ and $R^{\mathfrak{B}_0^1}(a) = 1$ otherwise.*

If μ is a union of conjunctive queries with signature τ , then a *dual* of μ is a relational τ -structure \mathfrak{B} with the property that a finite relational τ -structure \mathfrak{A} has a homomorphism to \mathfrak{B} if and only if \mathfrak{A} does not satisfy μ . If \mathfrak{B} and \mathfrak{B}' are both duals of μ , then they are homomorphically equivalent by compactness [4, Lemma 4.1.7].

► **Proposition 11** ([6, Proposition 8.14]). *Let μ be a union of connected conjunctive queries with signature τ . Then the resilience problem for μ is polynomial-time equivalent to VCSP(\mathfrak{B}_0^1) for any dual \mathfrak{B} of μ .*

Let $k \in \mathbb{N}$, let C be a set and G a permutation group on C . An *orbit of k -tuples* of G is a set of the form $\{\alpha(t) \mid \alpha \in G\}$ for some $t \in C^k$. 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 [8]. 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 relational 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 [16].

A relational τ -structure \mathfrak{A} *embeds* into a relational 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 relational τ -structure is the class of all finite relational τ -structures that embed into it. A relational structure \mathfrak{B} with a relational signature τ is called

- *finitely bounded* if τ is finite and there exists a universal τ -sentence ϕ such that a finite relational 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 \mathfrak{B} is finitely bounded and homogeneous, then $\text{Aut}(\mathfrak{B})$ is oligomorphic.

► **Theorem 12** ([6, Theorem 8.12]). *For every union μ of connected conjunctive queries over a finite relational signature τ there exists a τ -structure \mathfrak{B}_μ such that the following statements hold:*

1. \mathfrak{B}_μ is a reduct of a finitely bounded and homogeneous structure \mathfrak{B} .
2. A countable τ -structure \mathfrak{A} satisfies $\neg\mu$ if and only if it embeds into \mathfrak{B}_μ .
3. \mathfrak{B}_μ is finitely bounded.
4. $\text{Aut}(\mathfrak{B})$ and $\text{Aut}(\mathfrak{B}_\mu)$ are oligomorphic.

The existence of the dual \mathfrak{B}_μ for a union of connected conjunctive queries μ is the key to obtaining another dual \mathfrak{C}_μ , which has a strong model-theoretic property introduced in the following definition. If G is a permutation group on a set C , then \overline{G} denotes the closure of G in the space C^C with respect to the topology of pointwise convergence. This is the unique topology such that the closed subsets of C^C are precisely the endomorphism monoids of relational structures; see, e.g., [4, Proposition 4.4.2]. Note that \overline{G} might contain some operations that are not surjective, but if $G = \text{Aut}(\mathfrak{B})$ for some relational 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 13.** *A relational structure \mathfrak{B} with an oligomorphic automorphism group is a model-complete core if $\overline{\text{Aut}(\mathfrak{B})} = \text{End}(\mathfrak{B})$.*

For every relational structure \mathfrak{B} with an oligomorphic automorphism group, there exists a model-complete core \mathfrak{C} homomorphically equivalent to \mathfrak{B} , which is unique up to isomorphism called *the model-complete core of \mathfrak{B}* [3, Theorem 16], [4, Proposition 4.7.7]. Intuitively, the model-complete core of \mathfrak{B} is in a sense a “minimal” structure with the same CSP as \mathfrak{B} . If the domain of \mathfrak{B} is finite, then the domain of its model-complete core (usually just called *core*) is also finite.

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.

The following is an analogue of Theorem 12 for the model-complete core of \mathfrak{B}_μ . The statements in the theorem should be considered to be previously known; we provide a proof with references to the literature for the convenience of the reader.

► **Theorem 14.** *Let μ be a union of connected conjunctive queries over a finite relational signature τ . Then the model-complete core \mathfrak{C}_μ^1 of the structure \mathfrak{B}_μ from Theorem 12 satisfies the following:*

1. \mathfrak{C}_μ is a reduct of a finitely bounded and homogeneous structure \mathfrak{B} .
2. A countable τ -structure \mathfrak{A} satisfies $\neg\mu$ if and only if there is a homomorphism from \mathfrak{A} to \mathfrak{C}_μ .
3. If for each query ν in μ , the Gaifman graph of ν is complete, then \mathfrak{C}_μ is homogeneous.
4. $\text{Aut}(\mathfrak{B})$ and $\text{Aut}(\mathfrak{C}_\mu)$ are oligomorphic.

Proof. Item (1) follows from results in [19]; see [5, Corollary 7.5.15] for an explicit reference.

Item (2) is a consequence of \mathfrak{C}_μ being homomorphically equivalent to \mathfrak{B}_μ .

To prove (3), suppose that for each query ν in μ , the Gaifman graph of ν is complete. By [6, Theorem 8.13], there exists a dual \mathfrak{H} of μ , which is homogeneous. By [4, Proposition 4.7.7], the model-complete core of \mathfrak{H} is also homogeneous. Note that \mathfrak{C}_μ is homomorphically equivalent to \mathfrak{H} as they are both duals of μ and hence, by uniqueness, it is the model-complete core of \mathfrak{H} .

For item (4), note that the automorphism group of \mathfrak{B} is oligomorphic since it is homogeneous with finite relational signature. The automorphism group of \mathfrak{C}_μ is oligomorphic, since this property is clearly preserved under taking reducts. ◀

Note that since \mathfrak{C}_μ is unique up to isomorphism and homomorphic equivalence is transitive, the structure \mathfrak{C}_μ does not depend on the concrete choice of \mathfrak{B}_μ . For a union of connected conjunctive queries μ , let $\Delta_\mu := (\mathfrak{C}_\mu)_0^1$. In most results, this will be the valued structure to which we apply results about \mathfrak{B}_0^1 for a dual \mathfrak{B} of μ .

2.5 Expressive power

The concept of *expressive power* introduced in this section is a basis for polynomial-time gadget reductions between VCSPs.

► **Definition 15.** *Let A be a set and $R, R' \in \mathcal{R}_A$. We say that R' can be obtained from R by*

- projecting if R' is of arity k , R is of arity $k + n$ and for all $s \in A^k$

$$R'(s) = \inf_{t \in A^n} R(s, t).$$

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

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

$$\text{Opt}(R) := \{t \in \text{Feas}(R) \mid R(t) \leq R(s) \text{ for every } s \in A^k\}.$$

Note that $\inf_{t \in A^n} R(s, t)$ in item (1) might be irrational or $-\infty$. If this is the case, then $\inf_{t \in A^n} R(s, t)$ does not express a valued relation because valued relations must have weights from $\mathbb{Q} \cup \{\infty\}$. However, if R is preserved by all permutations of an oligomorphic automorphism group, then R attains only finitely many values and therefore this is never the case.

¹ In [6], the notation \mathfrak{C}_μ was used for a different dual of μ , which we do not need in this paper.

If $\mathcal{S} \subseteq \mathcal{R}_A$, then an atomic expression over \mathcal{S} is an atomic τ -expression where $\tau = \mathcal{S}$. We say that \mathcal{S} is *closed under forming sums of atomic expressions* if it contains all valued relations defined by sums of atomic expressions over \mathcal{S} .

► **Definition 16** (valued relational clone). *A valued relational clone (over a set C) is a subset of \mathcal{R}_C that is closed under forming sums of atomic expressions, projecting, shifting, non-negative scaling, Feas, and Opt; we refer to expressions formed this way as pp-expressions. For a valued structure Γ with the domain C , we write $\langle \Gamma \rangle$ for the smallest relational clone that contains the valued relations of Γ . If $R \in \langle \Gamma \rangle$, we say that Γ pp-expresses R .*

The acronym “pp” stands for primitive positive, since the concept of pp-expressions for valued structures is a generalization of primitive positive definitions used for reductions between CSPs.

2.6 Fractional maps

Let A and B be sets. We equip the space A^B of functions from B to A with the topology of pointwise convergence, where A is taken to be discrete. In this topology, a basis of open sets is given by $\mathcal{S}_{s,t} := \{f \in A^B \mid f(s) = t\}$ for $s \in B^k$ and $t \in A^k$ for some $k \in \mathbb{N}$. For any topological space T , we denote by $\mathcal{B}(T)$ 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 17** (fractional map). *Let A and B be sets. A fractional map from B to A is a probability distribution $(A^B, \mathcal{B}(A^B), \omega: \mathcal{B}(A^B) \rightarrow [0, 1])$, that is, $\omega(A^B) = 1$ and ω is countably additive: if $S_1, S_2, \dots \in \mathcal{B}(A^B)$ are disjoint, then*

$$\omega\left(\bigcup_{i \in \mathbb{N}} S_i\right) = \sum_{i \in \mathbb{N}} \omega(S_i).$$

We often use ω for both the entire fractional map and for the map $\omega: \mathcal{B}(A^B) \rightarrow [0, 1]$.

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 $\{x \in \mathbb{R} \mid x > a\} \cup \{\infty\}$ (thus, $\mathbb{R} \cup \{\infty\}$ is equipped with its order topology when ordered in the natural way).

A (real-valued) random variable is a measurable function $X: T \rightarrow \mathbb{R} \cup \{\infty\}$, i.e., pre-images of elements of $\mathcal{B}(\mathbb{R} \cup \{\infty\})$ under X are in $\mathcal{B}(T)$. 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

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

In the rest of the paper, we will work exclusively with topological spaces T of the form A^B for some sets A and B .

2.7 Pp-constructions

In this section, we introduce a concept of pp-constructions which generalize pp-expressions and provide polynomial-time reductions between VCSPs. We first define fractional homomorphisms.

► **Definition 18** (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 ω from D to C such that for every $R \in \tau$ of arity k and every tuple $t \in D^k$ it holds for the random variable $X: C^D \rightarrow \mathbb{R} \cup \{\infty\}$ given by $f \mapsto R^\Gamma(f(t))$ that $E_\omega[X]$ exists and that $E_\omega[X] \leq R^\Delta(t)$.*

We refer to [6] for a detailed introduction to fractional homomorphisms. Two valued τ -structures Γ and Δ are said to be *fractionally homomorphically equivalent*, if there is a fractional homomorphism from Γ to Δ and from Δ to Γ .

► **Remark 19.** If μ is a union of conjunctive queries with duals \mathfrak{B} and \mathfrak{C} , then \mathfrak{B} and \mathfrak{C} are homomorphically equivalent. Hence, \mathfrak{B}_0^1 and \mathfrak{C}_0^1 are fractionally homomorphically equivalent witnessed by fractional maps where the respective homomorphisms have probability 1.

As a next step towards the definition of a pp-construction, we define pp-powers.

► **Definition 20** (pp-power). *Let Γ be a valued structure with a domain C and let $d \in \mathbb{N}$. Then a (d -th) pp-power of Γ is a valued structure Δ with the 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*

$$R((x_1^1, \dots, x_d^1), \dots, (x_1^k, \dots, x_d^k)) = S(x_1^1, \dots, x_d^1, \dots, x_1^k, \dots, x_d^k).$$

We can now define the notion of a pp-construction.

► **Definition 21** (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 .*

The relation of pp-constructability is transitive: if Γ_1, Γ_2 , and Γ_3 are valued structures such that Γ_1 pp-constructs Γ_2 and Γ_2 pp-constructs Γ_3 , then Γ_1 pp-constructs Γ_3 [6, Lemma 5.12]. Note that whenever μ is a union of connected conjunctive queries and Δ_μ pp-constructs a valued structure Γ , then for every dual \mathfrak{B} of μ , the valued structure \mathfrak{B}_0^1 pp-constructs Γ , because Δ_μ and \mathfrak{B}_0^1 are fractionally homomorphically equivalent (Remark 19).

The motivation for introducing pp-constructions stems from the following lemma: pp-constructions give rise to polynomial-time reductions.

► **Lemma 22** ([6, Corollary 5.10 and 5.11]). *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 $\text{VCSP}(\Delta)$ to $\text{VCSP}(\Gamma)$. In particular, if $\Delta = (\{0, 1\}; \text{OIT})$, then $\text{VCSP}(\Gamma)$ is NP-hard.*

► **Example 23.** Recall the valued structure Γ_{MC} from Example 2. It is known that Γ_{MC} pp-constructs $(\{0, 1\}, \text{OIT})$ [25, Example 2.18] and by Lemma 22, $\text{VCSP}(\Gamma_{\text{MC}})$ is NP-hard.

2.8 Fractional polymorphisms

We now introduce *fractional polymorphisms* of valued structures, which generalize polymorphisms of relational structures. 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. [22]); it is known that fractional polymorphisms of a finite-domain valued structure capture the complexity of its VCSP up to polynomial-time reductions [10, 14]. Our definition is taken from [6] and allows arbitrary probability distributions in contrast to [23, 21, 24].

Let $\ell \in \mathbb{N}$. A *fractional operation on A of arity ℓ* is a fractional map from A^ℓ to A . The set of all fractional operations on a set A of arity ℓ is denoted by $\mathcal{F}_A^{(\ell)}$.

► **Definition 24.** A fractional operation $\omega \in \mathcal{F}_A^{(\ell)}$ improves a valued relation $R \in \mathcal{R}_A^{(k)}$ if for all $t^1, \dots, t^\ell \in A^k$

$$E := E_\omega[f \mapsto R(f(t^1, \dots, t^\ell))] \text{ exists, and } E \leq \frac{1}{\ell} \sum_{j=1}^{\ell} R(t^j). \quad (1)$$

Note that (1) has the interpretation that the expected value of $R(f(t^1, \dots, t^\ell))$ is at most the average of the values $R(t^1), \dots, R(t^\ell)$.

► **Definition 25** (fractional polymorphism). If a fractional operation ω improves every valued relation in a valued structure Γ , then ω is called a fractional polymorphism of Γ ; the set of all fractional polymorphisms of Γ is denoted by $\text{fPol}(\Gamma)$.

► **Remark 26.** A fractional polymorphism of arity ℓ of a valued τ -structure Γ might also be viewed as a fractional homomorphism from a specific ℓ -th pp-power of Γ , which we denote by Γ^ℓ , to Γ : the domain of Γ^ℓ is C^ℓ , and for every $R \in \tau$ of arity k we have

$$R^{\Gamma^\ell}((x_1^1, \dots, x_\ell^1), \dots, (x_1^k, \dots, x_\ell^k)) := \frac{1}{\ell} \sum_{i=1}^{\ell} R^\Gamma(x_i^1, \dots, x_i^k).$$

► **Example 27.** Let A be a set and $\pi_i^\ell \in \mathcal{O}_A^{(\ell)}$ be the i -th projection of arity ℓ , which is given by $\pi_i^\ell(x_1, \dots, x_\ell) = x_i$. The fractional operation Id_ℓ of arity ℓ such that $\text{Id}_\ell(\pi_i^\ell) = \frac{1}{\ell}$ for every $i \in \{1, \dots, \ell\}$ is a fractional polymorphism of every valued structure with domain A .

► **Lemma 28** (Lemma 6.8 in [6]). Let Γ be a valued τ -structure Γ over a countable domain C . Then every valued relation $R \in \langle \Gamma \rangle$ is improved by all fractional polymorphisms of Γ .

Let \mathfrak{A} be a relational structure and G a permutation group on the domain A of \mathfrak{A} . Let $\ell \geq 2$ and $f: A^\ell \rightarrow A$. The operation f is *pseudo cyclic with respect to G* if there exist $e_1, \dots, e_\ell \in \overline{G}$ such that for all $x_1, \dots, x_\ell \in A$,

$$e_1 f(x_1, x_2, \dots, x_\ell) = e_2 f(x_2, \dots, x_\ell, x_1) = \dots = e_\ell f(x_\ell, x_1, \dots, x_{\ell-1}).$$

The operation f is *canonical with respect to G* if for all $k \in \mathbb{N}$ and $t^1, \dots, t^\ell \in A^k$, the orbit of the k -tuple $f(t^1, \dots, t^\ell)$ with respect to G only depends on the orbits of t^1, \dots, t^ℓ with respect to G . A fractional operation ω on C of arity ℓ is called *pseudo cyclic with respect to G* if $\omega(S) = 1$ for the set S of all pseudo cyclic operations with respect to G of arity ℓ . *Canonicity* for fractional operations is defined analogously. The following theorem is a special case of [6, Theorem 7.13].

► **Theorem 29.** Let μ be a union of connected conjunctive queries and let \mathfrak{A} be a finitely bounded and homogenous expansion of \mathfrak{C}_μ (it exists by Theorem 14). If Δ_μ has a canonical pseudo cyclic fractional polymorphism with respect to $\text{Aut}(\mathfrak{A})$, then $\text{VCSP}(\Delta_\mu)$ and the resilience problem for μ is in P .

We formulate an adaptation of [6, Conjecture 8.17] for the valued structure Δ_μ , which replaces the structure Γ_μ used in [6] (and without considering so-called *exogenous relations*, which we do not introduce in this paper).

► **Conjecture 30.** Let μ be a union of connected conjunctive queries over the signature τ and let \mathfrak{A} be a finitely bounded homogeneous expansion of \mathfrak{C}_μ . Then exactly one of the following holds:

- $(\{0, 1\}; \text{OIT})$ has a pp-construction in Δ_μ , and $\text{VCSP}(\Delta_\mu)$ is NP-complete.
- Δ_μ has a fractional polymorphism of arity $\ell \geq 2$ which is canonical and pseudo cyclic with respect to $\text{Aut}(\mathfrak{A})$, and $\text{VCSP}(\Delta_\mu)$ is in P .

15:12 The Complexity of Resilience for Digraph Queries

The main reason to use Δ_μ instead of $\Gamma_\mu := (\mathfrak{B}_\mu)_0^1$ in this conjecture is Corollary 33, which shows that for Δ_μ the converse of the implication in Conjecture 30 is true: if Δ_μ has a canonical and pseudo cyclic fractional polymorphism, then it does not pp-construct $(\{0, 1\}; \text{OIT})$; see also the discussion in Section 3. The relationship between the two conjectures will be a subject of further investigation; at the moment we cannot prove that if Δ_μ has a canonical and pseudo cyclic fractional polymorphism, then so does Γ_μ , or vice versa.

3 Disjointness of the two cases of Conjecture 30

In this section we prove that the two cases in the complexity dichotomy of Conjecture 30 are disjoint. For a valued structure Γ , we denote by Γ^* the relational structure on the same domain whose relations are all relations from $\langle \Gamma \rangle$ that attain only values 0 and ∞ . Observe that by Lemma 28, $\text{Aut}(\Gamma) \subseteq \text{Aut}(\Gamma^*)$.

► **Observation 31.** *Let μ be a union of conjunctive queries. Then Δ_μ^* is a model-complete core.*

Proof. Note that for every $R \in \tau$, the structure Δ_μ^* contains $R^{\mathfrak{C}_\mu} = \text{Opt}(R^{\Delta_\mu})$. In particular, $\text{End}(\Delta_\mu^*) \subseteq \text{End}(\mathfrak{C}_\mu)$ by Lemma 28.

$$\begin{aligned} \text{End}(\Delta_\mu^*) \subseteq \text{End}(\mathfrak{C}_\mu) &= \overline{\text{Aut}(\mathfrak{C}_\mu)} && (\mathfrak{C}_\mu \text{ is a model-compl. core}) \\ &= \overline{\text{Aut}(\Delta_\mu)} \subseteq \overline{\text{Aut}(\Delta_\mu^*)} \subseteq \text{End}(\Delta_\mu^*). \end{aligned}$$

Therefore, Δ_μ^* is a model-complete core. ◀

Let G be a permutation group on a set C . An operation $f: C^\ell \rightarrow C$ on a set C is called *pseudo Taylor with respect to G* if for every $i \in \{1, \dots, \ell\}$ there exist $e_1, e_2 \in \overline{G}$ and variables $z_1, \dots, z_\ell, z'_1, \dots, z'_\ell \in \{x, y\}$ such that $z_i \neq z'_i$ and for all $x, y \in C$, $e_1(f(z_1, \dots, z_n)) = e_2(f(z'_1, \dots, z'_n))$. A fractional operation ω on C of arity ℓ is called *pseudo Taylor with respect to G* if $\omega(T) = 1$ for the set T of all pseudo Taylor operations with respect to G on C of arity ℓ . Note that every pseudo cyclic operation with respect to G is pseudo Taylor with respect to G ; similarly, pseudo Taylor fractional operations generalize pseudo cyclic fractional operations. The following result is not specific to resilience problems, but holds for VCSPs of valued structures with an oligomorphic automorphism group in general.

► **Theorem 32.** *Let Γ be a valued structure with an oligomorphic automorphism group such that Γ^* is a model-complete core and such that Γ has a pseudo cyclic (or, more generally, a pseudo Taylor) fractional polymorphism ω with respect to $\text{Aut}(\Gamma)$. Then Γ does not pp-construct K_3 .*

Proof. Suppose for contradiction that Γ pp-constructs K_3 . By Proposition 2.22 in [25], Γ^* pp-constructs K_3 as well. By results in [2] (see, e.g., Theorem 10.3.5 in [4]), Γ^* cannot have a pseudo Taylor polymorphism with respect to $\text{Aut}(\Gamma)$, and in particular, it cannot have a pseudo cyclic polymorphism with respect to $\text{Aut}(\Gamma)$.

By the definition of a pseudo cyclic fractional operation, there is a set S of pseudo cyclic operations of arity ℓ on C such that $\omega(S) = 1$. By Lemma 28, ω is also a fractional polymorphism of Γ^* . By Proposition 3.22 in [25], $\omega(S \cap \text{Pol}^{(\ell)}(\Gamma^*)) = 1$. In particular, $S \cap \text{Pol}^{(\ell)}(\Gamma^*)$ is non-empty. This is in contradiction to $\text{Pol}(\Gamma^*)$ not containing any pseudo cyclic operations. The proof in the case that ω is just a pseudo Taylor operation is analogous. ◀

► **Corollary 33.** *Let μ be a union of conjunctive queries such that Δ_μ has a pseudo cyclic, or, more generally, a pseudo Taylor fractional polymorphism ω with respect to $\text{Aut}(\Delta_\mu)$. Then Δ_μ does not pp-construct K_3 .*

Proof. By Observation 31, the structure Δ_μ^* is a model-complete core. Now the statement follows from Theorem 32. ◀

Observe that to prove Conjecture 30 it suffices to show that whenever Δ_μ does not pp-construct $(\{0, 1\}; \text{OIT})$, it has a canonical pseudo cyclic fractional polymorphism: this follows from Corollary 33 (the two cases are known to be disjoint), Theorem 29 (the tractability result for canonical pseudo cyclic fractional polymorphisms) and Lemma 22 (the hardness condition based on pp-constructions).

The main reason to work with the dual \mathfrak{C}_μ in this paper, instead of the dual \mathfrak{B}_μ that was used in [6], comes from the proof of Theorem 32 above: we need the property that \mathfrak{C}_μ is a model-complete core to get that Δ_μ^* is a model-complete core and hence to be able to apply the results from [2].

4 Complexity Dichotomy for Digraph Resilience Problems

From now on, R denotes a binary relational symbol. We will often view $\{R\}$ -structures as directed graphs. Let

$$\begin{aligned}\mu_\ell &:= \exists x R(x, x), \\ \mu_e &:= \exists x, y R(x, y), \text{ and} \\ \mu_c &:= \exists x, y (R(x, y) \wedge R(y, x)).\end{aligned}$$

The main result of the present article is the following theorem, which is a stronger version of Theorem 1 presented in Section 1.

► **Theorem 34.** *Let μ be a union of conjunctive queries over the signature $\{R\}$. Then the resilience problem of μ is in P or NP-complete. If all conjunctive queries in μ are minimal, connected, and pairwise non-equivalent, then exactly one of the following holds:*

1. μ is equal to μ_ℓ , μ_e , or μ_c , and the resilience of μ is in P. In this case, Δ_μ has a fractional polymorphism, which is canonical and pseudo cyclic with respect to $\text{Aut}(\mathfrak{C}_\mu)$.
2. Δ_μ pp-constructs $(\{0, 1\}; \text{OIT})$ and the resilience problem of μ is NP-complete.

We first sketch the proof strategy for Theorem 34. First observe that one may assume without loss of generality that all queries in μ are minimal, connected, and pairwise non-equivalent. If μ is equal to μ_ℓ , μ_e , μ_c , then the properties from item 1 are proven in Lemma 35. Otherwise, we prove that either μ contains a query μ_0 that contains a cycle of length ≥ 3 , or it has a finite dual without directed cycles. In both of these cases we show that item 2 holds.

It is easy to see that the resilience problem for μ_ℓ , μ_e or μ_c is in P. In Lemma 35 we give a stronger algebraic statement which corresponds to item 1 in Theorem 34; this was essentially known before, and we prove it in the long version of this article [7] for the convenience of the reader.

► **Lemma 35.** *For every $\mu \in \{\mu_\ell, \mu_e, \mu_c\}$, the valued structure Δ_μ has a canonical pseudo cyclic fractional polymorphism with respect to \mathfrak{C}_μ . In particular, the resilience problem for μ is in P.*

5 Self-join-free queries and self-join variations

Self-join-freeness is a fundamental and frequently used concept in database theory.

► **Definition 36** (self-join-free queries). *A union of conjunctive queries μ is called self-join-free if every relation symbol appears at most once in μ .*

Note that this is a more restrictive notion than a union of self-join-free conjunctive queries.

► **Lemma 37.** *Let μ be a self-join-free union of conjunctive queries over the signature τ containing a conjunctive query ν with signature $\rho \subseteq \tau$. Let \mathfrak{B} be the ρ -reduct of \mathfrak{C}_μ . Then \mathfrak{B} is a dual of ν , and the ρ -reduct of Δ_μ is equal to \mathfrak{B}_0^1 .*

Proof. Clearly, $\mathfrak{B} \models \neg\nu$. Let \mathfrak{A} be a finite relational ρ -structure such that $\mathfrak{A} \models \neg\nu$. Let \mathfrak{A}' be a τ -expansion of \mathfrak{A} where $R^{\mathfrak{A}'}$ for every $R \in \tau \setminus \rho$ is empty. Then $\mathfrak{A}' \models \neg\mu$ and hence has a homomorphism to \mathfrak{C}_μ . The same map is a homomorphism from \mathfrak{A} to \mathfrak{B} . It follows that \mathfrak{B} is a dual of ν . The last statement is clear from the definitions. ◀

We introduce a construction for obtaining queries with self joins from self-join-free queries, which will be crucial in our hardness proofs.

► **Definition 38.** *Let ν be a self-join-free union of conjunctive queries over the signature τ and let $f: \tau \rightarrow \tau$ be a map that preserves the arities. Then the union of queries resulting from ν by replacing each atom $R(x_1, \dots, x_k)$ by $f(R)(x_1, \dots, x_k)$ is denoted by $f(\nu)$. We say that f is ν -injective if for all $R, S \in \tau$ of the same arity k such that ν contains a query with atoms $R(x_1, \dots, x_k)$ and $S(x_1, \dots, x_k)$ for some variables x_1, \dots, x_k , $f(R) \neq f(S)$.*


A union of queries of the form $f(\nu)$ for some self-join-free ν and arity-preserving f is often called a *self-join variation* of ν in the literature [13].

► **Lemma 39.** *Every union of minimal conjunctive queries μ over a signature σ can be written as $f(\nu)$ for some self-join-free union of conjunctive queries ν with signature τ containing σ and some ν -injective $f: \tau \rightarrow \sigma$.*

Proof. For every $R \in \sigma$, let n_R be the number of occurrences of R in μ and let $R_0 := R$. Let $\tau = \bigcup_{R \in \sigma} \{R_0, R_1, \dots, R_{n_R-1}\}$, where all symbols R_i , $i \geq 1$, are fresh and of the same arity as R . Define ν to be the union of conjunctive queries obtained from μ by replacing the occurrences of R in μ by R_0, \dots, R_{n_R-1} (each of the symbols is used once) for every $R \in \sigma$; observe that ν is self-join-free. Let $f: \tau \rightarrow \sigma \subseteq \tau$ be defined by $f(R_0) = \dots = f(R_{n_R-1}) = R$, $R \in \sigma$. Then $f(\nu) = \mu$. Moreover, f is ν -injective, because queries in μ are minimal and therefore contain each atom at most once. ◀

We proceed to present the main result of this section – Theorem 40. The theorem and its proof is inspired by [13, Lemma 21]; their result is a special case of Theorem 40, because it only applies to conjunctive queries rather than unions of conjunctive queries, and because it only states a polynomial-time reduction, whereas our result even provides a pp-construction (which implies a polynomial-time reduction via Lemma 22).

► **Theorem 40.** *Let ν be a self-join-free union of connected conjunctive queries over the signature τ and let $f: \tau \rightarrow \tau$ be a ν -injective map that preserves arities. If all queries in $f(\nu)$ are minimal and pairwise non-equivalent, then $\Delta_{f(\nu)}$ pp-constructs Δ_ν . In particular, the resilience problem for ν reduces in polynomial time to the resilience problem for $f(\nu)$.*



■ **Figure 1** An illustration of the relationship between the structures \mathfrak{C}_μ and \mathfrak{B} from the proof of Theorem 40 for $\nu = \exists x, y(R(x, y) \wedge S(y, x))$ and $\mu = \exists x, y(R(x, y) \wedge R(y, x))$. The tuples in R are depicted in red and tuples in S in blue.

Proof. Let V be the finite set of variables of ν , which is also the set of variables of $\mu := f(\nu)$ and let Q be the set of conjunctive queries that form the union ν . Note that since all queries in ν are connected, the same is true for μ . Also, since all queries in μ are pairwise non-equivalent and minimal, they are pairwise homomorphically incomparable (see Section 2.3).

Let D_μ be the domain of Δ_μ . Let $D := (D_\mu)^{V \times Q}$, i.e., D is a finite power of D_μ ; it will be more convenient to use $V \times Q$ as an indexing set rather than the set $\{1, \dots, |V \times Q|\}$. We define a pp-power Δ of Δ_μ on the domain D with the signature τ . For every $R \in \tau$ of arity k and $(d^1, \dots, d^k) \in D^k$, if $R(x_1, \dots, x_k)$ is an atom in a query ν_0 in ν , then

$$R^\Delta(d^1, \dots, d^k) := f(R)^{\Delta_\mu}(d_{x_1, \nu_0}^1, \dots, d_{x_k, \nu_0}^k).$$

The idea is that the combination of query, relation symbol, and variables uniquely identifies an atom in μ and therefore encodes the difference between relation symbols R and R' from τ such that $f(R) = f(R')$.

Note that since relations in Δ_μ are 0-1-valued, the same is true for the relations in Δ and hence $\Delta = \mathfrak{B}_0^1$ for a relational τ -structure \mathfrak{B} on the domain D where for every $R \in \tau$ and $(d^1, \dots, d^k) \in D^k$, if $R(x_1, \dots, x_k)$ is an atom in a query ν_0 in ν , then

$$\mathfrak{B} \models R(d^1, \dots, d^k) \Leftrightarrow \mathfrak{C}_\mu \models f(R)(d_{x_1, \nu_0}^1, \dots, d_{x_k, \nu_0}^k)$$

(see Figure 1 for an illustration of the relationship between \mathfrak{C}_μ and \mathfrak{B}).

▷ **Claim.** \mathfrak{B} is a dual of ν .

To see this, we first argue that $\mathfrak{B} \not\models \nu$. If $\alpha: V \rightarrow D$ witnesses that $\mathfrak{B} \models \nu_0$ where ν_0 is a query in ν , then it is straightforward to verify that the map $\alpha': V \rightarrow D_\mu$ defined by $v \mapsto \alpha(v)_{v, \nu_0}$ witnesses that $\mathfrak{C}_\mu \models f(\nu_0)$ and hence $\mathfrak{C}_\mu \models \mu$, a contradiction with \mathfrak{C}_μ being a dual for μ . Therefore, $\mathfrak{B} \not\models \nu$.

It remains to show that if \mathfrak{A} is a relational τ -structure on a finite domain A such that $\mathfrak{A} \not\models \nu$, then \mathfrak{A} maps homomorphically into \mathfrak{B} . To this end, we construct an $f(\tau)$ -structure \mathfrak{A}' on the domain $A' = \{a_{v, \nu_0} \mid a \in A, v \in V, \nu_0 \in Q\}$. For every $R \in \tau$, if $R(x_1, \dots, x_k)$ is an atom in a query ν_0 in ν and $(a_1, \dots, a_k) \in R^{\mathfrak{A}}$, we put the tuple $((a_1)_{x_1, \nu_0}, \dots, (a_k)_{x_k, \nu_0})$ in $f(R)^{\mathfrak{A}'}$. No other tuples are in the relations of \mathfrak{A}' .

We argue that $\mathfrak{A}' \not\models \mu$. Suppose for contradiction that there exists a conjunctive query μ_0 in μ over a variable set $V_0 \subseteq V$ and a map $\beta: V_0 \rightarrow A'$ witnessing that $\mathfrak{A}' \models \mu_0$. Then for every atom $f(R)(x_1, \dots, x_k)$ in μ_0 we have $(\beta(x_1), \dots, \beta(x_k)) \in f(R)^{\mathfrak{A}'}$. We define maps

15:16 The Complexity of Resilience for Digraph Queries

$\beta_Q: V_0 \rightarrow Q$ and $\beta_V: V_0 \rightarrow V$ by setting $\beta_Q(x) := \nu_0$ and $\beta_V(x) := y$ where $\nu_0 \in Q$ and $y \in V$ are such that $\beta(x) = a_{y,\nu_0}$ for some $a \in A$. Recall that μ_0 is connected. Therefore, by the construction, β_Q is constant; let $\nu_0 \in Q$ be the only element of the image of β_Q .

Let $f(R)(x_1, \dots, x_k)$ be an atom in μ_0 . Since the tuple $(\beta(x_1), \dots, \beta(x_k))$ has been put in $f(R)^{\mathfrak{A}'}$, ν_0 contains an atom $R'(\beta_V(x_1), \dots, \beta_V(x_k))$ where $R' \in \tau$ is such that $f(R') = f(R)$. Therefore, there is an atom $f(R)(\beta_V(x_1), \dots, \beta_V(x_k))$ in $f(\nu_0)$. Hence, β_V defines a homomorphism from μ_0 to $f(\nu_0)$. Since the queries in μ are pairwise homomorphically incomparable, we must have $f(\nu_0) = \mu_0$. Moreover, since $f(R)(x_1, \dots, x_k)$ is an atom in μ_0 , we must have the atoms $f(R)(\beta_V(x_1), \dots, \beta_V(x_k))$, $f(R)(\beta_V^2(x_1), \dots, \beta_V^2(x_k))$, \dots , $f(R)(\beta_V^p(x_1), \dots, \beta_V^p(x_k))$ in $f(\nu_0) = \mu_0$, for all $p \in \mathbb{N}$.

Since β_V is a homomorphism from μ_0 to μ_0 and μ_0 is minimal, the image of β_V is equal to V_0 . Therefore, $\beta_V: V_0 \rightarrow V_0$ is surjective. Since V_0 is a finite set, this implies that β_V is a permutation of V_0 with an inverse $\beta_V^{-1} = \beta_V^p$ for some $p \in \mathbb{N}$. By the previous paragraph $f(R)(\beta_V^{-1}(x_1), \dots, \beta_V^{-1}(x_k))$ is an atom in μ_0 .

Let $\beta': V \rightarrow A$ be any map satisfying for every $x \in V_0$ that $\beta'(x) = a$ for the $a \in A$ such that $\beta(\beta_V^{-1}(x)) = a_{x,\nu_0}$. Then for every atom $R(x_1, \dots, x_k)$ in ν_0 , we have that $f(R)(x_1, \dots, x_k)$ is an atom of μ_0 , and therefore $f(R)(\beta_V^{-1}(x_1), \dots, \beta_V^{-1}(x_k))$ is an atom of μ_0 as well. Since β witnesses that μ_0 holds in \mathfrak{A}' ,

$$(\beta'(x_1)_{x_1,\nu_0}, \dots, \beta'(x_k)_{x_k,\nu_0}) = (\beta(\beta_V^{-1}(x_1)), \dots, \beta(\beta_V^{-1}(x_k))) \in f(R)^{\mathfrak{A}'}$$

By the ν -injectivity of f , there is no atom $R'(x_1, \dots, x_k)$ in ν_0 with $f(R') = f(R)$, so we must have $(\beta'(x_1), \dots, \beta'(x_k)) \in R^{\mathfrak{A}'}$ by the definition of \mathfrak{A}' . Thus, β' witnesses that $\mathfrak{A} \models \nu_0$ and hence, $\mathfrak{A} \models \nu$, a contradiction. It follows that $\mathfrak{A}' \not\models \mu$.

Since \mathfrak{C}_μ is a dual of μ , there is a homomorphism $h': \mathfrak{A}' \rightarrow \mathfrak{C}_\mu$. Let $h: \mathfrak{A} \rightarrow \mathfrak{B}$ be defined by $h(a) := (h'(a_{v,\nu_0}))_{v \in V, \nu_0 \in Q}$. We claim that h is a homomorphism from \mathfrak{A} to \mathfrak{B} . To see this, let $R \in \tau$ be of arity k and $(a_1, \dots, a_k) \in R^{\mathfrak{A}}$. Let ν_0 be a query in ν with an atom $R(x_1, \dots, x_k)$. Then $((a_1)_{x_1,\nu_0}, \dots, (a_k)_{x_k,\nu_0}) \in f(R)^{\mathfrak{A}'}$ and since h' is a homomorphism, $(h'((a_1)_{x_1,\nu_0}), \dots, h'((a_k)_{x_k,\nu_0})) \in f(R)^{\mathfrak{C}_\mu}$. Then, by the definition of \mathfrak{B} ,

$$(h(a_1), \dots, h(a_k)) = ((h'((a_1)_{v,\nu_0}))_{v \in V, \nu_0 \in Q}, \dots, (h'((a_k)_{v,\nu_0}))_{v \in V, \nu_0 \in Q}) \in R^{\mathfrak{B}}$$

It follows that \mathfrak{B} is a dual of ν .

Since \mathfrak{B} and \mathfrak{C}_ν are duals of ν , they are homomorphically equivalent and hence Δ and Δ_ν are fractionally homomorphically equivalent (see Remark 19). Since Δ is a pp-power of Δ_μ , it follows that Δ_μ pp-constructs Δ_ν . The final statement of the theorem follows from Lemma 22 and Proposition 11. \blacktriangleleft

6 Hardness proofs

The goal of this section is to present several hardness results that will be used in the proof of Theorem 34. First we have to define several graph-theoretical notions that will be useful in this section. Let $\mathfrak{G} = (V; E)$ be a directed multigraph and $k \in \mathbb{N}$. A *directed walk* in \mathfrak{G} of length k is a sequence $W = (v_0, v_1, \dots, v_k)$ of elements of V such that $(v_i, v_{i+1}) \in E$ for every $i \in \{0, \dots, k-1\}$. The walk W is *closed* if $v_0 = v_k$. A *directed path* in \mathfrak{G} of length k is a directed walk (v_0, \dots, v_k) such that $v_i \neq v_j$ for all distinct $i, j \in \{0, \dots, k\}$. A *directed cycle* in \mathfrak{G} of length k is a closed directed walk (v_0, \dots, v_k) such that $v_i \neq v_j$ for all distinct $i, j \in \{0, \dots, k-1\}$. An *oriented cycle* in \mathfrak{G} of length k is a sequence (v_0, v_1, \dots, v_k) of elements of V such that $v_k = v_0$, for every $i \in \{0, \dots, k-1\}$, $(v_i, v_{i+1}) \in E$ or $(v_{i+1}, v_i) \in E$, and for every $j \in \{0, \dots, k-1\}$, $j \neq i$, $v_i \neq v_j$.

Suppose now that \mathfrak{G} is undirected. A *cycle* in \mathfrak{G} is any sequence that forms an oriented cycle in \mathfrak{G} when viewed as a directed multigraph. We say that \mathfrak{G} is a *tree* if it does not contain any cycles and if it is *connected* in the sense that the graph obtained from \mathfrak{G} by replacing multiple edges by single edges is connected (see Section 2.3).

6.1 Hardness for queries with orientations of cycles

A signature τ is called *binary* if all relation symbols in τ are binary. In this section, we work with binary signatures in general rather than just the signature $\{R\}$. For any conjunctive query μ over a binary signature τ , let $\text{Multigraph}(\mu)$ denote the undirected multigraph whose edge relation is the union (as a multiset) of all the relations of the canonical database \mathfrak{D}_μ .

In this section we prove hardness for the resilience problem for minimal connected conjunctive queries μ over a binary signature such that $\text{Multigraph}(\mu)$ contains a cycle of length at least 3, and, more generally, for unions that contain such a query. To this end, we start with a result about self-join-free conjunctive queries, which together with Theorem 40 will yield a hardness proof for any query over a binary signature.

► **Theorem 41.** *Let μ be a connected self-join-free conjunctive query over a binary signature τ . If $\text{Multigraph}(\mu)$ contains a cycle of length ≥ 3 , then Δ_μ pp-constructs $(\{0, 1\}, \text{OIT})$.*

The proof of Theorem 41 is inspired by a much simpler pp-construction presented in [6, Example 8.18] for the query $\mu_\Delta := \exists x, y, z(R(x, y) \wedge S(y, z) \wedge T(z, x))$, which is the simplest query in the scope of Theorem 41. We recommend having a look at this example as a warm-up for this proof; the proof of Theorem 41 can be found in the long version of the article [7]. The main difference from the general proof is that μ_Δ is a query with a complete Gaifman graph and therefore has a homogeneous dual², which significantly simplifies the second part of the construction. The following corollary generalizes Theorem 41.

► **Corollary 42.** *Let μ be a union of minimal connected pairwise non-equivalent conjunctive queries over a binary signature containing a conjunctive query μ_0 . If $\text{Multigraph}(\mu_0)$ contains a cycle of length ≥ 3 , then Δ_μ pp-constructs $(\{0, 1\}; \text{OIT})$ and the resilience problem for μ is NP-complete.*

Proof. Let μ' be a self-join-free union of connected conjunctive queries over a binary signature τ' such that $\mu = f(\mu')$ for some μ' -injective $f: \tau' \rightarrow \tau$; it exists by Lemma 39. Let μ'_0 be a conjunctive query from μ' such that $\mu_0 = f(\mu'_0)$ and let $\tau'_0 \subseteq \tau'$ be the signature of μ'_0 .

By Theorem 40, Δ_μ pp-constructs $\Delta_{\mu'}$. By Lemma 37, the τ'_0 -reduct of $\Delta_{\mu'}$ is equal to \mathfrak{B}_0^1 for some dual \mathfrak{B} of μ'_0 . Since \mathfrak{B} and $\mathfrak{C}_{\mu'_0}$ are both duals of μ'_0 , the valued structure \mathfrak{B}_0^1 is fractionally homomorphically equivalent to $\Delta_{\mu'_0} = (\mathfrak{C}_{\mu'_0})_0^1$ (Remark 19). Therefore, $\Delta_{\mu'}$ pp-constructs $\Delta_{\mu'_0}$. By Theorem 41, $\Delta_{\mu'_0}$ pp-constructs $(\{0, 1\}; \text{OIT})$. By the transitivity of pp-constructability, Δ_μ pp-constructs $(\{0, 1\}; \text{OIT})$. By Lemma 22, $\text{VCSP}(\Delta_\mu)$ is NP-hard. By Proposition 11, the resilience problem for μ is NP-hard, and thus NP-complete. ◀

² We remark that the dual used in [6, Example 8.18] is a different dual from $\mathfrak{C}_{\mu_\Delta}$. The dual used there embeds every finite relational $\{R, S, T\}$ -structure \mathfrak{A} that does not satisfy μ_Δ , whereas our dual is a model-complete core; so, for example, the empty structure on two vertices maps homomorphically into $\mathfrak{C}_{\mu_\Delta}$, but does not have an embedding.

6.2 Hardness for queries with finite acyclic duals

In this section we prove that the resilience problem for queries μ that have a non-trivial finite dual without directed cycles is NP-hard. We stress that this lemma crucially relies on our approach to analyse the complexity of the resilience problem for μ using the dual structure \mathfrak{C}_μ and Δ_μ .

► **Lemma 43.** *Let μ be a union of conjunctive queries over the signature $\{R\}$ such that the domain of \mathfrak{C}_μ is finite. Assume that \mathfrak{C}_μ contains at least one edge and does not contain any directed cycles. Then Δ_μ pp-constructs $(\{0, 1\}; \text{OIT})$.*

Proof. Let C be the domain of \mathfrak{C}_μ . Let k be the length of the longest directed path in \mathfrak{C}_μ ; it exists, because \mathfrak{C}_μ is finite and does not contain any directed cycles. Let $\phi(x_0, x_1)$ be the pp-expression

$$\inf_{x_2, \dots, x_k} (R(x_0, x_1) + \text{Opt}(R)(x_1, x_2) \cdots + \text{Opt}(R)(x_{k-1}, x_k)).$$

Let $(x, y) \in C^2$. Then $\phi^{\Delta_\mu}(x, y) = 0$ if and only if there is a directed path in \mathfrak{C}_μ of length k starting with the edge (x, y) . If $\phi^{\Delta_\mu}(x, y) \neq 0$, then $\phi^{\Delta_\mu}(x, y) = 1$ if and only if there is a directed path in \mathfrak{C}_μ of length $k - 1$ starting in y . Finally, if $\phi^{\Delta_\mu}(x, y) \notin \{0, 1\}$, then $\phi^{\Delta_\mu}(x, y) = \infty$. Let Γ be the valued $\{R\}$ -structure on the domain C where $R^\Gamma(x, y) = \phi^{\Delta_\mu}(x, y)$ for all $(x, y) \in C^2$. Note that $R^\Gamma \in \langle \Delta_\mu \rangle$.

Recall the valued structure Γ_{MC} from Example 2. Let $a, b \in C$ be such that there is a directed path in \mathfrak{C}_μ of length k starting with the edge (a, b) . Let $f: \{0, 1\} \rightarrow C$ be defined by $f(0) := a$ and $f(1) := b$. It is straightforward to verify that ω_f defined by $\omega_f(f) = 1$ is a fractional homomorphism from Γ_{MC} to Γ . Let $g: C \rightarrow \{0, 1\}$ be defined by $g(x) = 0$ for every $x \in C$ such that there is a directed path of length k starting in x and $g(x) = 1$ otherwise. We argue that ω_g defined by $\omega_g(g) = 1$ is a fractional homomorphism from Γ to Γ_{MC} . Let $(x, y) \in C^2$. Note that if $R^\Gamma(x, y) \geq 1$, then trivially $R^{\Gamma_{\text{MC}}}(g(x), g(y)) \leq R^\Gamma(x, y)$. Suppose therefore that $R^\Gamma(x, y) = 0$. Then by the definition of ϕ , there is a directed path in \mathfrak{B} of length k starting with the edge (x, y) . By the definition of g , we have $g(x) = 0$. Since there is no directed path of length $k + 1$ in \mathfrak{C}_μ and \mathfrak{C}_μ does not contain directed cycles, there is no directed path of length k starting in y and therefore $g(y) = 1$. Hence, $R^{\Gamma_{\text{MC}}}(g(x), g(y)) \leq R^\Gamma(x, y)$. It follows that ω_g is a fractional homomorphism from Γ to Γ_{MC} .

By the previous paragraph, Γ is fractionally homomorphically equivalent to Γ_{MC} . Since $R^\Gamma \in \langle \Delta_\mu \rangle$, we have that Δ_μ pp-constructs Γ_{MC} . We have already mentioned in Example 23 that Γ_{MC} pp-constructs $(\{0, 1\}, \text{OIT})$. By the transitivity of pp-constructability, Δ_μ pp-constructs $(\{0, 1\}, \text{OIT})$. ◀

7 Proof of Theorem 34

We are now ready to prove the main result of the paper.

Proof of Theorem 34. Since $\text{Aut}(\mathfrak{C}_\mu) = \text{Aut}(\Delta_\mu)$, items (1) and (2) are mutually exclusive by Corollary 33. Hence it is enough to prove that item (1) or item (2) holds. Without loss of generality, we may assume that all queries in μ are pairwise non-equivalent, minimal and connected (see Lemma 9). In particular, the queries in μ are pairwise homomorphically incomparable (see Section 2.3). By this assumption, if μ contains μ_ℓ or μ_e , then μ is equal to μ_ℓ or to μ_e , respectively, in which case item (1) holds by Lemma 35. We may therefore assume that μ contains neither μ_ℓ nor μ_e .

If μ contains a conjunctive query μ_0 such that $\text{Multigraph}(\mu_0)$ contains a cycle of length ≥ 3 , then item (2) holds by Corollary 42. Suppose that this is not the case. Then for every query ν in μ , $\text{Multigraph}(\nu)$ is a tree, or ν contains the atoms $R(x, y)$ and $R(y, x)$ for some variables $x \neq y$, in which case $\nu = \mu_c$ by the minimality of ν . Note that every ν such that $\text{Multigraph}(\nu)$ is a tree has a homomorphism to μ_c , so $\mu = \mu_c$ whenever μ contains μ_c . In this case, item (1) holds by Lemma 35. Suppose therefore that $\mu \neq \mu_c$ and hence, $\text{Multigraph}(\nu)$ is a tree for every query ν in μ . Then μ has a finite dual by [20] (see also [6, Theorem 8.7]). Since \mathfrak{C}_μ is the model-complete core of this dual, it also has a finite domain, so it is a finite directed graph with an edge relation R . Note that \mathfrak{C}_μ contains at least one edge, because $\mu \neq \mu_e$. It is easy to see that every orientation of a tree (in particular, every ν in μ) maps homomorphically to every directed cycle; thus, \mathfrak{C}_μ does not contain directed cycles. By Lemma 43, Δ_μ pp-constructs $(\{0, 1\}; \text{OIT})$. By Lemma 22 and Proposition 11, the resilience problem for μ is NP-complete. Therefore, item (2) holds. ◀

References

- 1 Antoine Amarilli, Wolfgang Gatterbauer, Neha Makhija, and Mikaël Monet. Resilience for regular path queries: Towards a complexity classification. *Proc. ACM Manag. Data*, 3(2), June 2025. doi:10.1145/3725245.
- 2 Libor Barto, Michael Kompatscher, Miroslav Olšák, Trung Van Pham, and Michael Pisker. Equations in oligomorphic clones and the constraint satisfaction problem for ω -categorical structures. *Journal of Mathematical Logic*, 19(2):#1950010, 2019. doi:10.1142/S0219061319500107.
- 3 Manuel Bodirsky. Cores of countably categorical structures. *Logical Methods in Computer Science (LMCS)*, 3(1):1–16, 2007. doi:10.2168/LMCS-3(1:2)2007.
- 4 Manuel Bodirsky. *Complexity of Infinite-Domain Constraint Satisfaction*. Lecture Notes in Logic (52). Cambridge University Press, 2021. doi:10.1017/9781107337534.
- 5 Manuel Bodirsky. Automorphism groups, 2023. Course Notes, TU Dresden, <https://wwwpub.zih.tu-dresden.de/~bodirsky/Automorphism-Groups.pdf>.
- 6 Manuel Bodirsky, Žaneta Semanišínová, and Carsten Lutz. The complexity of resilience problems via valued constraint satisfaction problems. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '24*, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3661814.3662071.
- 7 Manuel Bodirsky and Žaneta Semanišínová. The complexity of resilience for digraph queries, 2026. arXiv:2601.05346.
- 8 Peter J. Cameron. *Oligomorphic permutation groups*. Cambridge University Press, Cambridge, 1990.
- 9 Ashok K. Chandra and Philip M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 77–90, 1977. doi:10.1145/800105.803397.
- 10 David A. Cohen, Martin C. Cooper, Páidí Creed, Peter G. Jeavons, and Stanislav Živný. An algebraic theory of complexity for discrete optimization. *SIAM J. Comput.*, 42(5):1915–1939, 2013. doi:10.1137/130906398.
- 11 Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. On the parameterized complexity of graph modification to first-order logic properties. *Theory of Computing Systems*, 64(2):251–271, 2020. doi:10.1007/s00224-019-09938-8.
- 12 Cibele Freire, Wolfgang Gatterbauer, Neil Immerman, and Alexandra Meliou. The complexity of resilience and responsibility for self-join-free conjunctive queries. *Proc. VLDB Endow.*, 9(3):180–191, 2015. doi:10.14778/2850583.2850592.

- 13 Cibele Freire, Wolfgang Gatterbauer, Neil Immerman, and Alexandra Meliou. 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*, pages 271–284, 2020. doi:10.1145/3375395.3387647.
- 14 Peter Fulla and Stanislav Živný. A Galois connection for weighted (relational) clones of infinite size. *TOCT*, 8(3):9:1–9:21, 2016. doi:10.1145/2898438.
- 15 Michael Garey and David Johnson. *A guide to NP-completeness*. CSLI Press, Stanford, 1978.
- 16 Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.
- 17 Neha Makhija and Wolfgang Gatterbauer. A unified approach for resilience and causal responsibility with integer linear programming (ilp) and lp relaxations. *Proc. ACM Manag. Data*, 1(4), December 2023. doi:10.1145/3626715.
- 18 Alexandra Meliou, Wolfgang Gatterbauer, and Dan Suciu. Reverse data management. *Proc. VLDB Endow.*, 4(12):1490–1493, August 2011. doi:10.14778/3402755.3402803.
- 19 Antoine Mottet and Michael Pinsker. Cores over Ramsey structures. *Journal of Symbolic Logic*, 86(1):352–361, 2021. doi:10.1017/jsl.2021.6.
- 20 J. Nešetřil and C. Tardif. Duality theorems for finite structures. *Journal of Combinatorial Theory, Series B*, 80:80–97, 2000. doi:10.1006/jctb.2000.1970.
- 21 Friedrich Martin Schneider and Caterina Viola. An application of Farkas’ lemma to finite-valued constraint satisfaction problems over infinite domains. *Journal of Mathematical Analysis and Applications*, 517(1):126591, 2023. doi:10.1016/j.jmaa.2022.126591.
- 22 Johan Thapper and Stanislav Živný. 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*, pages 695–704, 2013. doi:10.1145/2488608.2488697.
- 23 Caterina Viola. *Valued Constraint Satisfaction Problems over Infinite Domains*. PhD thesis, TU Dresden, 2020.
- 24 Caterina Viola and Stanislav Živný. The combined basic LP and affine IP relaxation for promise VCSPs on infinite domains. *ACM Trans. Algorithms*, 17(3):21:1–21:23, 2021. doi:10.1145/3458041.
- 25 Žaneta Semanišinová. *Valued Constraint Satisfaction in Structures with an Oligomorphic Automorphism Group*. Dissertation, TU Dresden, Dresden, Germany, 2025. Available at <https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa2-974737>.