Valued Constraint Satisfaction in Structures with an Oligomorphic Automorphism Group

D I S S E R T A T I O N

zur Erlangung des akademischen Grades

Doctor of Philosophy
(Ph. D.)

vorgelegt

dem Bereich Mathematik und Naturwissenschaften
der Technischen Universität Dresden

von

Mgr. Žaneta Semanišinová

Eingereicht am 24. Januar 2025

Gutachter:
Prof. Dr.rer.nat. Manuel Bodirsky
Dr hab. Marcin Kozik

Die Dissertation wurde in der Zeit von Oktober 2021 bis
Januar 2025 am Institut für Algebra angefertigt.

"Du wirst sehen, dass alles gutgehen wird."
"Ich wüsste nicht, wie," meinte Atréju.
"Ich auch nicht," erwiderte der Drache,
"aber das ist gerade das Schöne."

— Michael Ende, *Die unendliche Geschichte*

"Everything will turn out all right. You'll see."
"I can't imagine how," said Atreyu.
"Neither can I," said the luckdragon.
"But that's the best part of it."

— Michael Ende, *The Neverending Story*

Acknowledgements

First and foremost, I want to express my deepest gratitude to my PhD supervisor Manuel Bodirsky for his guidance and support throughout all of my doctoral studies. He shared his knowledge and advice with me, discussed mathematical questions with me as with a colleague from the very beginning, and always pushed my ambitions a bit further than I would push myself. I greatly appreciate that he took me as a PhD student and dedicated his time to my questions and progress.

I am also very grateful to David Stanovský, who supervised my bachelor and master theses. He showed me the beauty of algebra and research and took me to my first mathematical conference. Thanks to him I was captivated by the idea of doing research and pursuing a PhD. I also want to thank Víťa Kala who didn't take no for an answer when we discussed me doing a PhD abroad.

I want to thank my coauthors and collaborators, Antoine, Barny, Carsten, Colin, Édouard, George, Peter, Jakub, Michał, Santiago and of course Manuel for working with me on exciting CSP problems, exchanging our ideas and making something out of them.

I am also greatly thankful to all my current and former colleagues at the TU Dresden. My deepest gratitude goes to everybody who made me feel at home in Dresden when I came, especially to Philip for being the best German teacher I could ever think of and to Edoardo for suffering through this with me, to Albert for being unconditionally nice to me even when I was mean to him, to Jakub for reminding me that 'sky is the limit' (both in mathematics and bouldering) and for swimming to that beach in Molveno, and to Richard and Verona for our bouldering and chatting sessions. I also want to thank Philipp for helping me to become a person of culture and for Polish jokes, and, together with Andrew, Colin and Santiago for being my amazing office mates. Thanks to Aaron and Piotr for being my favorite office duo, for table tennis and Bockbieranstich sessions. I hope that Dresden lunch discussions remain legendary.

Huge thanks goes to the whole POCOCOP and CSP community for creating a wonderful research network and expanding my horizons in so many ways, from minor conditions to ferratas, from Ramsey theory to karaoke. In addition to those who have already been mentioned, thanks to Dima, Johanna, both Michaels, Paolo, Silvia, and Tomáš for making this journey unforgettable. I certainly hope to continue to be part of this wonderland of reflections.

I want to thank my friends in Prague who have not given up on me when I moved away, especially to Kristín and Dano for our fine dining sessions and lovely holidays together. I am also very grateful to my friends and colleagues from the association STROM, where I shaped my love for mathematics and education.

Finally, my sincerest gratitude goes to my mum and dad, who showed me the beauty of mathematics when I was little and who have greatly supported me on this journey and gave me a lot of important advice on academia and life. I also want to thank my sister Deniska and my brother Riško for believing in me and supporting me. And most importantly, I am deeply grateful to my fiancé Jakub, who convinced me to get this adventure started despite all the troubles it brought us and who has been there for me in every moment of this difficult journey.

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 author 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.

Contents

Introduction

Constraint satisfaction problems (CSPs) are computational problems that appear in many areas of computer science, for example in temporal and spatial reasoning in artificial intelligence [15] or in database theory [7, 55]. Let $\mathfrak{B}$ be a fixed relational structure over a finite signature τ, sometimes called the *template*. The CSP of $\mathfrak{B}$ (denoted $\mathrm{CSP}(\mathfrak{B})$) is the problem of deciding whether a given conjunction of atomic τ-formulas is satisfiable in $\mathfrak{B}$. We give a natural and well-understood example of a CSP.

Example 0.1. *For $(i, j, k) \in \{0, 1\}^3$, let $R_{ijk} = \{0, 1\}^3 \setminus \{(i, j, k)\}$. Consider the relational structure $\mathfrak{B} = (\{0, 1\}; R_{000}, R_{001}, R_{011}, R_{111})$. Note that every input to $\mathrm{CSP}(\mathfrak{B})$ corresponds to a conjunction of 3-clauses (disjunction of three literals, each of which is either a Boolean variable or its negation): for example, the input*

$$R_{001}(x_1, x_3, x_2) \wedge R_{011}(x_4, x_3, x_2)$$

corresponds to

$$(x_1 \vee x_3 \vee \neg x_2) \wedge (x_4 \vee \neg x_3 \vee \neg x_2).$$

It follows that $\mathrm{CSP}(\mathfrak{B})$ is the same problem as the problem of satisfiability of Boolean conjunctions of 3-clauses, known as 3-SAT, which is one of the most significant and intensely studied NP-complete problems [48].

The computational complexity of CSPs is of central interest, and a general research goal is to obtain systematic complexity classification results, in particular about CSPs that are in P and CSPs that are NP-hard. The following breakthrough result was obtained independently by Bulatov [32] and by Zhuk [82,83], which confirmed the famous Feder-Vardi conjecture [42].

Theorem 0.2. *Let $\mathfrak{A}$ be a relational structure with a finite domain. Then $\mathrm{CSP}(\mathfrak{A})$ is in P or NP-complete.*

Note that by the theorem of Ladner [61], if P$\neq$NP, then there are problems in NP which are neither in P nor NP-complete, so-called *NP-intermediate problems.* Theorem 0.2 implies that there are no such problems in the class of finite-domain CSPs. A stronger formulation of Theorem 0.2 which includes an algebraic dichotomy that aligns with the complexity dichotomy will be given in Theorem 3.4. A key component for the algebraic dichotomy and the proof of Theorem 0.2 is the algebraic theory that was developed for CSPs, and, mainly, the concepts of polymorphisms. Polymorphisms are homomorphisms from a finite power of the template to the template itself and they can be viewed as operations on the template. In some sense, polymorphisms capture the symmetries of a relational structure and they may provide a polynomial-time algorithm for the CSP of the structure.

The hardness condition for finite-domain CSPs builds on the simple observation that primitively positively definable relations can be added to the template without changing the complexity of the CSP, which inspired the concepts of pp-interpretations [9, 33] and pp-constructions [6]. These concepts generalize primitive positive definability and provide a framework for reductions between CSPs on different domains. In fact, it turns out that, unless P=NP, the only source of hardness for finite-domain CSPs is a pp-construction of some (equivalently: all) relational structure on a finite domain with an NP-complete CSP.

CSPs on infinite domains

As a natural generalization, the research direction of CSPs on countably infinite domains emerged. The motivation for this research is two-fold: to see how far the algebraic and complexity results on finite-domain CSPs can be pushed and at the same time, to study numerous natural computational problems that can be modeled as CSPs on an infinite domain only. This is the case for most of the CSPs in temporal and spatial reasoning and for many of the CSPs that appear in database theory (e.g., most of the CSPs in the logic MMSNP, which is a fragment of existential second-order logic introduced by Feder and Vardi [42] important for database theory [7], cannot be formulated as CSPs with a finite template [64]). For CSPs with countably infinite templates we may not hope for general classification results [14]; however, we may hope for such results if we restrict our attention to classes of templates that are model-theoretically well-behaved. An example of such a class is the class of all relational structures with the domain $\mathbb{Q}$ where all relations are definable with a first-order formula over the structure $(\mathbb{Q}; <)$, so-called *temporal* structures.

For CSPs of relational structures with a rich automorphism group, a powerful algebraic machinery, inspired by the tools for finite-domain CSPs, was developed [6, 23], which has led to classification results for many concrete automorphism groups: we list [10, 18, 22, 26, 68] as a representative sample. In 2011, Bodirsky and Pinsker identified a class of infinite-domain CSPs which are all in NP and conjectured a complexity dichotomy for it.

Conjecture 0.3 ([27]). *Let $\mathfrak{A}$ be a relational structure with a countable domain. If $\mathfrak{A}$ is a reduct of a finitely bounded homogeneous relational structure $\mathfrak{B}$, then* $\mathrm{CSP}(\mathfrak{A})$ *is in P or NP-complete.*

All relational structures in the scope of Conjecture 0.3 have a rich automorphism group in the sense that it is *oligomorphic*: for every $k \in \mathbb{N}$, there are only finitely many orbits of k-tuples under the action of this group. Relational structures with an oligomorphic automorphism group are in some sense finite-like: in many contexts, instead of working with the elements of the domain of the structure, one may work with orbits of k-tuples for a sufficiently large k, of which there are finitely many. Thanks to this property, many essential algebraic properties of polymorphisms and pp-constructions that were key for the proof of Theorem 0.2 generalize to this setting.

Valued CSPs

Simultaneously with the progress in the area of infinite-domain CSPs, another influential generalization of the finite-domain CSP setting emerged: a framework that allows to model optimization problems, so-called *valued constraint satisfaction problems* (VCSPs). A VCSP is parameterized by a *valued structure* Γ (the *template*), which consists of a *domain* C and *cost functions*, each defined on C^k for some k. The input to the VCSP consists of a finite set of variables, a finite sum of cost functions applied to 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. In analogy to the CSP setting, we refer to the cost functions as *valued relations*. The computational complexity of VCSPs has been studied depending on the valued structure that parameterizes the problem.

CSPs can be viewed as a variant of VCSPs with costs from the set $\{0, \infty\}$: every constraint is either satisfied or surpasses every finite threshold. VCSPs also generalize min-CSPs, which are the natural variant of CSPs where, instead of asking whether all constraints can be satisfied at once, we search for an assignment that minimizes the number of unsatisfied constraints. Such problems can be modeled as VCSPs with costs from the set $\{0, 1\}$.

Example 0.4. *Recall the* 3*-SAT problem, its representation as* $\mathrm{CSP}(\mathfrak{B})$ *and relations* R_{ijk} *from Example 0.1. We show how to model Min-CSP*$(\mathfrak{B})$ *as a VCSP; this problem is essentially the same problem as* Max-3-SAT. *For* $(i, j, k) \in \{0, 1\}^3$*, let*

$$(R_{ijk})_0^1(x, y, z) = \begin{cases} 0 & (x, y, z) \in R_{ijk}; \\ 1 & (x, y, z) \notin R_{ijk}. \end{cases}$$

Consider the valued structure $\Gamma = (\{0, 1\}; (R_{000})_0^1, (R_{001})_0^1, (R_{011})_0^1, (R_{111})_0^1)$. *Given an instance of* $\mathrm{VCSP}(\Gamma)$ *with a threshold* u*, we can view the instance as a conjunction of* 3*-clauses and see* $\mathrm{VCSP}(\Gamma)$ *as a question if there is an assignment to* ϕ *which satisfies all but at most* u *clauses. Note that since the summands in the instance may repeat, it allows to give weights to the constraints in the input and hence prefer satisfying one constraint over another. Therefore, to be precise,* $\mathrm{VCSP}(\Gamma)$ *models Min-CSP*$(\mathfrak{B})$ *if we restrict the inputs only to sums with non-repeating summands.*

A major achievement of the field of finite-domain VCSPs is the following analogue of Theorem 0.2.

Theorem 0.5. *Let* Γ *be a valued structure with a finite signature and a finite domain. Then* $\mathrm{VCSP}(\Gamma)$ *is in P or NP-complete.*

The result above has an intriguing history. The classification task was first considered in [41] with important first results that indicated that we might expect a good systematic theory for such VCSPs. A milestone was reached by Thapper and Živný with the proof of a complexity dichotomy for the case where the valued relations never take value ∞ [79]. On the hardness side, Kozik and Ochremiak [58] formulated a condition that implies hardness for VCSP(Γ) and found equivalent characterisations that suggested that this condition characterises NP-hardness (unless P=NP, of course). Kolmogorov, Krokhin, and Rolínek [57] then showed that if the hardness condition from [58] does not apply, linear programming relaxation in combination with algorithms for classical CSPs can be used to solve VCSP(Γ), conditional on the tractability conjecture for (classical) CSPs. Finally, this conjecture about CSPs has been confirmed [82], thus completing the complexity dichotomy for VCSP(Γ) for finite-domain templates Γ.

A stronger formulation of Theorem 0.5 that includes the algebraic conditions characterizing the respective complexities can be found in Theorem 3.36. The hardness condition is based on the notion of *expressibility*, which is a generalization of primitive positive definability for VCSPs, and can be phrased using the generalization of pp-constructions to the VCSP setting (see Chapter 2). For the tractability condition, a generalization of polymorphisms is utilized: instead of operations, we consider probability distributions on operations with certain preservation properties, so-called *fractional polymorphisms.*

Valued CSPs on infinite domains

Motivated by the progress in the area of infinite-domain CSPs, we focus in this thesis on VCSPs on infinite domains. Many important optimization problems in the literature cannot be modeled as VCSPs if we restrict to templates on a finite domain; VCSPs that require an infinite domain are, for example, the min-correlation-clustering problem with partial information [4, 80], ordering min-CSPs [52], phylogeny min-CSPs [36], VCSPs with semilinear constraints [20], and the class of resilience problems from database theory [30, 45, 46, 65]. Since VCSPs generalize CSPs, we cannot achieve a general classification for VCSPs with countably infinite templates [14]; however, similar structural restrictions as for infinite-domain CSPs can restrict the complexity to the class NP. In this thesis, we focus on the class of VCSPs whose templates have an oligomorphic automorphism group; an automorphism of a valued relation R is a permutation of the domain of R that preserves the values of R when applied componentwise.

The systematic study of infinite-domain VCSPs was initiated in the dissertation of Viola [80]. The focus of [80] is on *piecewise linear* and *piecewise linear homogeneous* valued structures. These are valued structures over the domain $\mathbb{Q}$ where the cost functions are first-order definable as partial functions in $(\mathbb{Q}; +, 1, \leq)$, or in $(\mathbb{Q}; <, 1, (x \mapsto cx)_{c\in\mathbb{Q}})$, respectively (if the function is undefined, the cost is considered to be ∞). For example, linear programming can be modeled as a VCSP of a piecewise linear valued structure, but not as a VCSP of a valued structure with an oligomorphic automorphism group; this can be seen from [9, Corollary 4.6.2] applied to the feasibility problem (see Section 1.1) of a valued structure that models linear programming as a VCSP (see, e.g. [80, Section

1.7]). On the other hand, it is easy to see that $+$ and $x \mapsto 2x$ have only the trivial automorphism when viewed as valued relations: for example, an automorphism of $+$ would have to satisfy $\alpha(x) + \alpha(y) = x + y$ for all $x, y \in \mathbb{Q}$. [1] This signals that piecewise linear and piecewise linear homogeneous valued structures typically do not have an oligomorphic automorphism group.

VCSPs stemming from resilience problems (see Chapter 5) provide ample examples of VCSP templates with an oligomorphic automorphism group that are neither piecewise linear nor piecewise linear homogeneous. For a concrete one, we may consider the valued structure Γ from Example 5.23. Its cost function R^Γ viewed as a function $R^\Gamma \colon \mathbb{Q}^2 \to \{0,1\}$ is first-order definable in neither $(\mathbb{Q}; +, 1, \leq)$ nor $(\mathbb{Q}; <, 1, (x \mapsto cx)_{c \in \mathbb{Q}})$. To see this, first note that $(\mathbb{Q}; <, 1, (x \mapsto cx)_{c \in \mathbb{Q}})$ is first-order definable in $(\mathbb{Q}; +, 1, \leq)$ so it is enough to show the claim for $(\mathbb{Q}; +, 1, \leq)$. The theory of $(\mathbb{Q}; +, 1, \leq)$ is NIP ('not the independence property', see [78] for a definition), because the theory of real closed fields is NIP [78]. On the other hand, the structure $(\mathbb{Q}; R^\Gamma, 1)$ has the IP by [39, Proposition 5.2], because the formula $R^\Gamma(x, y) \neq 1$ encodes every undirected bipartite graph (see [39, Section 5] for a formal definition). This implies that R^Γ cannot be first-order definable in $(\mathbb{Q}; +, 1, \leq)$. A similar argument could be used for numerous structures in the scope of this thesis.[2]

At the starting point of the research work presented in this thesis, there were only a few articles studying classes of VCSPs on infinite domains [20, 77, 81] and, as in [80], the valued structures considered in these papers rarely have an oligomorphic automorphism group. Since the oligomorphicity assumption on templates proved to be crucial for CSPs, in this thesis, we focus on valued constraint satisfaction problems with countable templates with an oligomorphic automorphism group and present results of the author in this research area.

Resilience problems

As an application of the complexity results for VCSPs obtained in this thesis, we study the computational complexity of resilience problems from database theory. A *resilience problem* is parameterized by a query q. The input is a finite database $\mathfrak{A}$ and the question is what is the minimum number of tuples to be removed from the database relations so that $\mathfrak{A}$ does not satisfy q. The notion of resilience captures how 'robust' an answer to a query is, which is particularly important if the facts in the input database may be incorrect [2]. 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 [67].

The systematic study of complexity of resilience problems was initiated in [45], focus-

[1] Note that there are two different notions of automorphisms one can apply to an operation $f \colon \mathbb{Q}^n \to \mathbb{Q}$: if we view f as a valued relation, an automorphism α must satisfy $f(\alpha(x_1), \dots, \alpha(x_n)) = f(x_1, \dots, x_n)$, and if we view f as a function on $\mathbb{Q}$, an automorphism α must satisfy $f(\alpha(x_1), \dots, \alpha(x_n)) = \alpha(f(x_1, \dots, x_n))$. We will always consider the former notion in this thesis.

[2] We thank Paolo Marimon for suggesting this elegant argument.

ing on conjunctive queries that are self-join-free, i.e., they contain every relation symbol at most once. Since then, the problem continued to be intensively studied in various settings [2, 45, 65]. One of the variations of the problem that lately gained attention is to consider *bag databases*, i.e., where the database relations are multisets, as opposed to the standard setting with *set databases*. Despite the substantial effort dedicated to classifying the complexity of resilience problems depending on the query μ, the complete classification still remains open even in the case where μ is a conjunctive query, both in bag and set semantics.

Contributions and Structure of the Thesis

In this thesis, we present results on valued constraint satisfaction in valued structures with a countable domain and an oligomorphic automorphism group obtained during the doctoral studies of the author. Most of the results appeared in the conference paper of Manuel Bodirsky, Carsten Lutz and the author [30], and a preprint written jointly with Manuel Bodirsky and Édouard Bonnet [12].

In the following we outline the structure of the thesis, highlighting the key content and contributions of each chapter. In Chapter 1, we give the basic definitions for our setting and several general results, accompanied by examples.

Motivated by the success of the pp-constructability framework for CSPs of structures with an oligomorphic automorphism group [6], we develop parts of this approach for VCSPs of such structures in Chapter 2. Pp-constructions give rise to polynomial-time reductions between (V)CSPs and they enable elegant phrasing of hardness conditions for these problems. To define pp-constructability, we introduce fractional homomorphisms (Section 2.2), which are probability distributions on potentially uncountable sets. For the convenience of the reader, we also provide an overview of the notions from topology and measure theory that are needed in the thesis, specialized to our setting.

In Chapter 3 we generalize the influential concept of fractional polymorphisms to the most permissive setting on infinite domains considered so far. We study the properties of fractional polymorphisms of valued structures with an oligomorphic automorphism group. In Section 3.6, we reprove the complexity dichotomy for finite-domain VCSPs using the newly introduced notion of pp-constructions and variations of results from [58]. The most important contribution of this chapter is presented in Section 3.7, where we generalize the polynomial-time reduction from [21] based on canonical polymorphisms to the VCSP setting and hence enable reducing infinite-domain VCSPs to finite-domain VCSPs under some assumptions. We utilize this reduction to prove a sufficient condition for tractability for VCSPs.

The focus of Chapter 4 is a complexity classification of the class of VCSPs on the domain $\mathbb{Q}$ whose templates are preserved by all order-preserving permutations on $\mathbb{Q}$. In analogy to the terminology for CSPs, we call such VCSPs *temporal*. This is the first complete complexity classification of a class of infinite-domain VCSPs preserved by a fixed automorphism group; classifying complexity in such classes is standard in the research on CSPs on infinite domains, where this corresponds to studying a class of structures first-order definable in some fixed relational structure (such as $(\mathbb{Q};<)$). The class of temporal VCSPs contains important optimization problems such as *least correlation clustering* (see Example 1.11) and the *minimum feedback arc set problem* (see Example 1.16). Moreover, temporal CSPs proved to be one of the fundamental test cases for shaping the theory of infinite-domain CSPs and we believe that the class of temporal VCSPs plays a similar role for infinite-domain VCSPs.

The final Chapter 5 is concerned with an application of the theory of VCSPs of valued structures with an oligomorphic automorphism group to the computational complexity of resilience problems in bag semantics. Using the notion of homomorphism duality, we translate the resilience problem for a query q to a VCSP of a template Γ_q; the domain of Γ_q might need to be countably infinite. We exploit the tools introduced in Chapter 2 and 3 to obtain a sufficient hardness condition and a sufficient tractability condition for resilience problems.

The author published two more research articles based on the results obtained during her doctoral studies [16,76]. Since the results concern different CSP frameworks (classical CSPs and quantified CSPs), the results are not presented in this thesis.

Chapter 1

Preliminaries

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 $\mathbb{Q}$, the set of non-negative rational numbers by $\mathbb{Q}_{\geq 0}$ and the set of positive rational numbers by $\mathbb{Q}_{>0}$. The standard strict linear order on $\mathbb{Q}$ is denoted by $<$. 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 $a \in \mathbb{R}$,
- $a + \infty = \infty + a = \infty$ 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 $a > 0$.

If A is a set, then $\mathrm{Sym}(A)$ denotes the group of all permutations of A. If $t \in A^k$, then we implicitly assume that $t = (t_1, \dots, t_k)$, where $t_1, \dots, t_k \in A$. A tuple $t \in A^k$ is called *injective* if it has pairwise distinct entries. If $f\colon A^\ell \to A$ is an operation on A and $t^1, \dots, t^\ell \in A^k$, then we denote

$$(f(t^1_1, t^2_1, \dots, t^\ell_1), \dots, f(t^1_k, t^2_k, \dots, t^\ell_k))$$

by $f(t^1, \dots, t^\ell)$ and say that f *is applied componentwise.*

1.1 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 $\mathbb{Q}$.

Usual relations will also be called *crisp* relations. A valued relation $R \in \mathscr{R}_C^{(k)}$ that only takes values from $\{0, \infty\}$ will be identified with the crisp relation $\{t \in C^k \mid R(t) = 0\}$. Valued relations that take at most one finite value will be called *essentially crisp.*

For $R \in \mathscr{R}_C^{(k)}$ the *feasibility relation of* R is defined as

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

For $S \subseteq C^k$ and $a, b \in \mathbb{Q} \cup \{\infty\}$, we denote by S_a^b the valued relation such that $S_a^b(t) = a$ if $t \in R$, and $R_a^b(t) = b$ otherwise. We often write S_0^∞ to stress that S is a crisp relation viewed as a valued relation.

Example 1.1. *On the domain C, the valued relation $(=)_0^\infty$ denotes the crisp equality relation and $(\emptyset)_0^\infty$ is the unary empty relation (where every $c \in C$ evaluates to ∞). If $C = \mathbb{Q}$, then $(<)_0^1$ denotes the valued relation*

$$(<)_0^1(x, y) = \begin{cases} 0 & x < y; \\ 1 & x \geq y. \end{cases}$$

A *(relational) signature* τ is a set of *relation symbols*, each of them equipped with an arity from $\mathbb{N}$. A *relational τ-structure* $\mathfrak{C}$ consists of a set C, which is also called the *domain* of $\mathfrak{C}$, and a relation $R^{\mathfrak{C}} \subseteq C^k$ for each relation symbol $R \in \tau$ of arity k. Relational structure will also be called *crisp structures.* A *valued τ-structure* Γ consists of a *domain* C and a valued relation $R^\Gamma \in \mathscr{R}_C^{(k)}$ for each relation symbol $R \in \tau$ of arity k. A relational τ-structure may be identified with a valued τ-structure where all valued relations only take values from $\{0, \infty\}$. 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 $\Gamma, \Delta, \dots$ are denoted $C, D, \dots$, respectively. If $\mathcal{R}$ is a set of valued relations over a common domain C, we write $(C; \mathcal{R})$ for the valued structure Γ whose relations are precisely the relations from $\mathcal{R}$; we only use this notation if the precise choice of the signature does not matter.

If Γ is a valued τ-structure on the domain C and Δ is a valued τ-structure on the domain $D \subseteq C$ such that for every $R \in \tau$ of arity k, R^Δ is the restriction of R^Γ on D^k, then we call Δ a *substructure* of Γ. Every $C' \subseteq C$ induces a substructure Γ' of Γ by setting $R^{\Gamma'}$ to be the restriction of R^Γ on $(C')^k$ for every $R \in \tau$ of arity k.

A valued structure is called *essentially crisp* if all of its valued relations are essentially crisp. If Γ is a valued τ-structure on the domain C, then $\mathrm{Feas}(\Gamma)$ denotes the relational τ-structure $\mathfrak{C}$ on the domain C where $R^{\mathfrak{C}} = \mathrm{Feas}(R^\Gamma)$ for every $R \in \tau$. 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'.

We give three simple examples of valued structures.

Example 1.2. *Let $<_2$ denote the standard strict order on the set $\{0, 1\}$. Then $\Gamma_{max} = (\{0, 1\}; (<_2)_0^1)$ is a valued structure.*

Example 1.3. *Let $\tau = \{E\}$, where E is a binary relation symbol. Then define $K_3 := (\{0,1,2\}; E^{K_3})$ where $E^{K_3} = \{0,1,2\}^2 \setminus \{(0,0),(1,1),(2,2)\}$ is a relational structure that represents the complete graph on 3 vertices. We also view K_3 as a valued structure where for every $x, y \in \{0,1,2\}$*

$$E^{K_3}(x,y) = \begin{cases} 0 \text{ if } x \neq y, \\ \infty \text{ if } x = y. \end{cases}$$

Example 1.4. *Let $\tau = \{E\}$, where E is a binary relation symbol and let Γ be a valued τ-structure on the domain $\{0,1,2\}$ where*

$$E^{\Gamma}(x,y) = \begin{cases} 42 \text{ if } x \neq y, \\ \infty \text{ if } x = y. \end{cases}$$

Then Γ is essentially crisp, but not crisp, and $\mathrm{Feas}(\Gamma) = K_3$.

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 now 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 $R \in \tau \cup \{(=)_0^\infty, (\emptyset)_0^\infty\}$ and (not necessarily distinct) variable symbols $x_1, \dots, x_k$. 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 $i \in \{1, \dots, m\}$ 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 $\mathscr{R}_C^{(n)}$ in a natural way, which we denote by ϕ^Γ. If ϕ is the empty sum then ϕ^Γ is constant 0.

1.2 Valued Constraint Satisfaction

In this section we assume that Γ is a fixed valued τ-structure for a *finite* signature τ. We first define constraint satisfaction problem of a relational structure and then give a definition of a more general valued constraint satisfaction problem. These problems are closely related, see Remark 1.8.

Definition 1.5. *Let $\mathfrak{A}$ be a relational structure over a finite signature τ. The constraint satisfaction problem for $\mathfrak{A}$, denoted by* $\mathrm{CSP}(\mathfrak{A})$*, is the computational problem to decide for a given conjunction $\phi(x_1, \dots, x_n)$ of atomic τ-formulas, whether there exists $t \in A^n$ such that $\mathfrak{A} \models \phi(t)$.*

Definition 1.6. *The* valued constraint satisfaction problem for Γ*, denoted by* $\mathrm{VCSP}(\Gamma)$*, 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* $\mathrm{VCSP}(\Gamma)$*, and to u as the* threshold*. We also refer to the pair (ϕ, u) as* *a (positive or negative)* instance *of* VCSP(Γ). *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 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 $t \in C^n$ will then be called a *solution for* ϕ if the cost of ϕ equals $\phi^\Gamma(t)$. Note that in general there might not be any solution. If there exists a tuple $t \in C^n$ such that $\phi^\Gamma(t) < \infty$ then ϕ is called *satisfiable.*

Example 1.7. *The problem* VCSP(Γ_{max}) *for the valued structure* Γ_{max} *from Example 1.2 models the* directed max-cut *problem: given a finite directed graph* (V, E) *(we do allow loops and multiple edges), 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* $f \colon V \to \{0, 1\}$ *that minimises the value*

$$\sum_{(x,y) \in E} (<_2)_0^1(f(x), f(y)),$$

which can be formulated as an instance of VCSP(Γ_{max}). *Conversely, every instance of* VCSP(Γ_{max}) *corresponds to a directed max-cut instance. It is known that* VCSP(Γ_{max}) *is NP-complete (even if we do not allow loops and multiple edges in the input) [48]. We mention that this problem can be viewed as a resilience problem as explained in Chapter 5, Example 5.16.*

For relational structures, VCSPs specialize to CSPs.

Remark 1.8. *If* $\mathfrak{A}$ *be a relational* τ*-structure, then* 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* $\phi(x_1, \ldots, x_n)$*,* $\phi^\mathfrak{A}$ *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,* VCSP($\mathfrak{A}$) *and* CSP($\mathfrak{A}$) *are essentially the same problem.*

Example 1.9. *Recall the structure* K_3 *from Example 1.3.* CSP(K_3) *is the problem of deciding whether a given conjunction of formulas of the form* $E(x, y)$ *and* $x = y$ *is satisfiable in* K_3*. Viewing variables as vertices of a given graph with the edges given by the input, it is clear that* CSP(K_3) *is the* 3*-coloring problem for graphs, which is known to be an NP-complete problem [48].*

1.3 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 1.10. *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 $t \in C^k$ 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 R.*

The set of all automorphisms of Γ is denoted by $\mathrm{Aut}(\Gamma)$, 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 $\{\alpha(t) \mid \alpha \in G\}$ 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 [35]. 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 [51]. In this thesis, however, we stick to the oligomorphicity notion which naturally generalizes to valued structures.

Example 1.11. *Let Γ_{LCC} be the valued structure $(\mathbb{N}; (=)_0^1, (\neq)_0^1)$. Note that $\mathrm{Aut}(\Gamma_{LCC})$ 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 *[80, Example 5] is equal to* VCSP(Γ_{LCC}). *It is a variant of the min correlation clustering problem [4] that does not require precisely one constraint between any two variables. The problem is NP-complete in both settings [48, 80].*

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

Lemma 1.12. *Let τ be a relational signature. Let Γ be a valued τ-structure with a countable domain C and an oligomorphic automorphism group. Then for every τ-expression $\phi(x_1, \dots, x_n)$ there exists $t \in C^n$ such that $\inf_{s \in C^n} \phi^\Gamma(s) = \phi^\Gamma(t)$.*

Proof. By the assumption, there are only finitely many orbits of n-tuples of $\mathrm{Aut}(\Gamma)$. Therefore, there are only finitely many possible values from $\mathbb{Q} \cup \{\infty\}$ for $\phi^\Gamma(s)$, which implies the statement. □

Let Γ be a valued τ-structure and $\mathfrak{B}$ a relational structure. Suppose that $\mathrm{Aut}(\mathfrak{B})$ is oligomorphic and $\mathrm{Aut}(\mathfrak{B}) \subseteq \mathrm{Aut}(\Gamma)$ (and hence $\mathrm{Aut}(\Gamma)$ is oligomorphic). Then $S \subseteq B^k$ is first-order definable over $\mathfrak{B}$ if and only if S is preserved by $\mathrm{Aut}(\mathfrak{B})$ (see, e.g., [9, Theorem 4.2.9]). Let $R \in \tau$ be of arity k. Then R^Γ attains only finitely many values by the oligomorphicity of $\mathrm{Aut}(\Gamma)$. Moreover, if for some $s, t \in C^k$ we have $R^\Gamma(s) \neq R^\Gamma(t)$, then s and t lie in a different orbit of $\mathrm{Aut}(\mathfrak{B})$. Therefore, for every value $a \in \mathbb{Q} \cup \{\infty\}$, there is a union U_a of orbits of k-tuples under the action of $\mathrm{Aut}(\mathfrak{B})$ such that $R^\Gamma(t) = a$ if and only if $t \in U_a$. Since U_a is preserved by $\mathrm{Aut}(\mathfrak{B})$, it is first-order definable over $\mathfrak{B}$

by a formula ϕ_a. Hence, R can be given by a list of values a in the range of R and first-order formulas ϕ_a over $\mathfrak{B}$. Such a collection

$$((R, a, \phi_a) \mid R \in \tau, \exists t \in C^k(R(t) = a))$$

will be called a *first-order definition* of Γ in $\mathfrak{B}$. Clearly, if a valued structure Γ has a first-order definition in a relational structure $\mathfrak{B}$, then $\mathrm{Aut}(\mathfrak{B}) \subseteq \mathrm{Aut}(\Gamma)$. We will use first-order definitions of valued structures to be able to give valued structures as an input to decision problems (see Remark 4.9 and Proposition 4.35).

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.

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 $\mathfrak{B}$.

Note that for every relational structure $\mathfrak{B}$ with a finite signature, for every n there are only finitely many non-isomorphic substructures of $\mathfrak{B}$ of size n. In a homogeneous relational structure $\mathfrak{B}$, the orbit of $t \in B^k$ under the action of $\mathrm{Aut}(\mathfrak{B})$ is determined by the atomic formulas that hold on entries of t, equivalently, the substructure of $\mathfrak{B}$ induced by the entries of t. Therefore, all countable homogeneous relational structures with a finite signature have finitely many orbits of k-tuples for all $k \in \mathbb{N}$, and hence an oligomorphic automorphism group. In particular, any supergroup of the automorphism group of a finitely bounded homogeneous structure on a countable domain is oligomorphic.

Let $\mathfrak{A}$ be a relational structure with a finite signature and a countable domain whose automorphism group contains the automorphism group of a finitely bounded homogeneous structure, equivalently, whose relations are first-order-definable in this finitely bounded homogeneous structure. It is well-known (see, e.g., [9, Proposition 2.3.16]) that this condition is equivalent to $\mathfrak{A}$ being a reduct of a (possibly different) finitely bounded homogeneous relational structure. Note that this condition implies that $\mathrm{Aut}(\mathfrak{A})$ is oligomorphic. Recall that the class of reducts of finitely bounded homogeneous structures is precisely the class for which a CSP complexity dichotomy is conjectured in Conjecture 0.3.

A homogeneous relational structure $\mathfrak{B}$ over a finite signature has quantifier elimination [51]. Therefore, whenever a valued structure Γ has a first-order definition over $\mathfrak{B}$, then the defining formulas ϕ_a can be chosen to be quantifier-free, and hence disjunctions of conjunctions of atomic formulas over $\mathfrak{B}$. We will then refer to a *quantifier-free definition of* Γ *in* $\mathfrak{B}$.

Example 1.13. *The relational structure* $(\mathbb{Q};<)$ *is finitely bounded and homogeneous. A finite relational* $\{<\}$*-structure* $\mathfrak{A}$ *embeds into* $(\mathbb{Q};<)$ *if and only if it satisfies the universal sentence*

$$\forall x,y,z\ \big(\neg(x<x)\wedge(x<y\vee y<x\vee x=y)\wedge\neg(x<y\wedge y<x)\\ \wedge(\neg(x<y)\vee\neg(y<z)\vee x<z)\big).$$

Moreover, it is well-known that every isomorphism between finite substructures of $(\mathbb{Q};<)$ *can be extended to an automorphism of* $(\mathbb{Q};<)$*; this is a standard back-and-forth argument, see, e.g. [9, Remark 4.1.2]. Since* $(\mathbb{Q};<)$ *is homogeneous and over a finite signature,* $\mathrm{Aut}(\mathbb{Q};<)$ *is oligomorphic.*

We finish this section by a theorem that explains why finitely bounded homogeneous structures are important in the context of complexity classification of VCSPs.

Theorem 1.14. *Let* Γ *be a valued structure with a finite signature and a countable domain such that there exists a finitely bounded homogeneous relational structure* $\mathfrak{B}$ *with* $\mathrm{Aut}(\mathfrak{B})\subseteq\mathrm{Aut}(\Gamma)$*. Then* $\mathrm{VCSP}(\Gamma)$ *is in NP.*

Proof. Let (ϕ,u) be an instance of $\mathrm{VCSP}(\Gamma)$ with n variables. Note that two tuples $(a_1,\dots,a_n)$ and $(b_1,\dots,b_n)$ lie in the same orbit of $\mathrm{Aut}(\mathfrak{B})$ if and only if the map that maps a_i to b_i for $i\in\{1,\dots,n\}$ is an isomorphism between the substructures induced by $\mathfrak{B}$ on $\{a_1,\dots,a_n\}$ and on $\{b_1,\dots,b_n\}$. Since $\mathrm{Aut}(\mathfrak{B})\subseteq\mathrm{Aut}(\Gamma)$, every orbit of n-tuples of $\mathrm{Aut}(\Gamma)$ is a union of orbits of $\mathrm{Aut}(\mathfrak{B})$ and hence determined by the substructure induced by $\mathfrak{B}$ on the elements of some tuple from the orbit. Whether a given finite relational structure $\mathfrak{A}$ is in the age of a fixed finitely bounded structure $\mathfrak{B}$ can be decided in polynomial time: if ϕ is the universal τ-sentence which describes the age of $\mathfrak{B}$, it suffices to exhaustively check all possible instantiations of the variables of ϕ with elements of A and verify whether ϕ is true in $\mathfrak{A}$ under the instantiation. Hence, we may non-deterministically generate a relational structure $\mathfrak{A}$ with domain $\{1,\dots,n\}$ from the age of $\mathfrak{B}$ and then verify in polynomial time whether the value $\phi^\Gamma(b_1,\dots,b_n)$ is at most u for any tuple $(b_1,\dots,b_n)\in B^n$ such that $i\mapsto b_i$ is an embedding of $\mathfrak{A}$ into $\mathfrak{B}$. $\square$

All concrete VCSPs that we study in this thesis satisfy the assumptions of Theorem 1.14 and hence are contained in the complexity class NP.

1.4 Examples

In this section we give several examples to illustrate the variety of templates that a VCSP can have and various problems they capture. We start with the a problem that is dual to the max-cut problem introduced in Example 1.7.

Example 1.15. *Consider the valued structure* $\Gamma_{min}=(\{0,1\};(\geq_2)^1_0)$*, where* $\geq_2$ *is the complement of the relation* $<_2$*. Similarly to Example 1.7,* $\mathrm{VCSP}(\Gamma_{min})$ *models the* directed min-cut problem*, i.e., given a finite directed graph* (V,E)*, partition the vertices* V *into two classes* A *and* B *such that the number of edges from* A *to* B *is minimal. The min-cut problem is solvable in polynomial time; see, e.g., [49].*

We also give two examples of VCSPs with infinite templates.

Example 1.16. *Let $\tau = \{E\}$ and let Γ be a valued τ-structure on the domain $\mathbb{Q}$ where $E^\Gamma = (<)^1_0$. Then every τ-expression can be interpreted as a (not necessarily simple) digraph with the edge relation E and every digraph corresponds to a τ-expression. Therefore,* VCSP(Γ) *is the* minimum feedback arc set problem*, i.e., the problem of finding the minimum number of edges to be removed from a digraph to make it acyclic. This problem is known to be NP-complete [48]. An analogous argument shows that* CSP($\mathbb{Q}; <$) *is the* digraph acyclicity problem*, which is polynomial-time tractable, for example, by performing depth-first search. Note that* $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\mathbb{Q}; <)$ *is an oligomorphic permutation group (Example 1.13).*

Example 1.17. *Let* $(\mathbb{V}; E)$ *be the* countable random graph*, that is,* $\mathbb{V}$ *is a countably infinite set,* E *is a binary irreflexive symmetric relation and every finite simple, undirected graph* $\mathfrak{G}$ *embeds into* $(\mathbb{V}; E)$ *(this property is sometimes called* universality*). It is folklore that* $(\mathbb{V}; E)$ *is homogeneous, see, e.g., [9, Theorem 2.3.8, Example 2.3.9]. It is also finitely bounded: a finite* $\{E\}$*-structure embeds into* $(\mathbb{V}; E)$ *if and only if it satisfies*

$$\forall x, y\, \neg E(x, x) \wedge (\neg E(x, y) \vee E(y, x)).$$

Let Γ_{graph} *denote the valued structure* $(\mathbb{V}; R)$ *where*

$$R(x, y) := \begin{cases} 0 & E(x, y), \\ 1 & \neg E(x, y) \wedge x \neq y, \\ 2 & x = y, \end{cases}$$

where E *is the edge relation of the random graph. Note that* $\mathrm{Aut}(\Gamma_{graph}) = \mathrm{Aut}(\mathbb{V}; E)$*. We will show in Example 3.54 that* VCSP(Γ_{graph}) *is in P.*

We continue with several examples of crisp structures, where we can equivalently consider CSPs instead of VCSPs. We will often use reductions from NP-complete CSPs to VCSPs to show their hardness in later chapters.

Example 1.18. *Let* $\oplus$ *denote the addition on* $\{0, 1\}$ *modulo 2. Let*

$$\Gamma = (\{0, 1\}; \{0\}, \{1\}, \{(x, y, z) \in \{0, 1\}^3 \mid x \oplus y \oplus z = 0\}).$$

Then Γ *is a relational structure and every instance of* CSP(Γ) *can be viewed as a system of linear equations modulo 2. Therefore,* CSP(Γ) *is solvable in polynomial time by Gaussian elimination.*

The next two examples are variations of the SAT problem, i.e., the problem of satisfiability of boolean formulas. Recall that we have already seen 3-SAT in Example 0.1.

Example 1.19. *Let* OIT *be the following relation*

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

CSP($\{0, 1\}$; OIT) *is the so called 1-in-3-3-SAT problem, which is known to be be NP-complete (see, e.g., [9, Example 1.2.2]).*

Example 1.20. *Consider the relation*

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

$\mathrm{CSP}(\{0,1\};\mathrm{NAE})$ *is the so called Not-All-Equal-3-SAT problem, which is a well-known NP-complete problem (see, e.g., [9, Example 1.2.2]).*

We finish this section with an example coming from so-called *equality CSPs* (see Section 4.2).

Example 1.21. *Consider a valued structure* $(\mathbb{Q};\mathrm{Dis})$*, where* Dis *is the ternary relation*

$$\{(x,y,z) \in \mathbb{Q}^3 \mid (x = y \wedge y \neq z) \vee (x \neq y \wedge y = z)\}.$$

$\mathrm{CSP}(\mathbb{Q};\mathrm{Dis})$ *is known to be NP-complete, [9, Theorem 7.4.1]; we will exploit this fact in Section 4.2.*

1.5 Expressive power

A first-order formula is called *primitive positive* if it is an existentially quantified conjunction of atomic formulas. A relation on a set A is called *primitively positively definable* over a relational structure $\mathfrak{A}$, if it is definable by a primitive positive formula over $\mathfrak{A}$. If $\mathscr{S}$ is a set of relations on A, then a relation is primitively positively definable from $\mathscr{S}$ if it is primitively positively definable over $(A;\mathscr{S})$. A *relational clone* on A is a set of relations on A closed under primitive positive definability. The main reason to consider primitive positive definability is that relations with a primitive positive definition can be added to the structure without changing the complexity of the respective CSP. We now generalize the definition of primitive positive definability and relational clones to the VCSP setting.

Definition 1.22. *Let A be a set and $R, R' \in \mathscr{R}_A$. 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

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

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.

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

Definition 1.23 (valued relational clone)**.** *A* valued relational clone (over A) *is a subset of $\mathscr{R}_A$ that is closed under forming sums of atomic expressions, projecting, shifting, non-negative scaling,* Feas*, and* Opt*. For a valued structure $\mathfrak{A}$ with the domain A, we write $\langle \mathfrak{A} \rangle$ for the smallest relational clone that contains the valued relations of $\mathfrak{A}$. If $R \in \langle \mathfrak{A} \rangle$, we say that $\mathfrak{A}$* expresses *R.*

Remark 1.24. *Note that if a valued relational clone $\mathscr{C}$ contains a set $\mathscr{S} \subseteq \mathscr{R}_A$ of crisp relations, then every relation which is primitively positively definable from $\mathscr{S}$ is in $\mathscr{C}$ by forming a sum of the corresponding atomic expressions and projecting on the variables that are not existentially quantified. Therefore, valued relational clones are a generalization of relational clones.*

Moreover, if $\mathfrak{A}$ is a relational structure and $R \in \langle \mathfrak{A} \rangle$, then R is essentially crisp and $\mathrm{Feas}(R)$ *is primitively positively definable from $\mathfrak{A}$; this is easily verified by induction.*

The following example shows that neither the operator Opt nor the operator Feas is redundant in Definition 1.23.

Example 1.25. *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* $\mathrm{Feas}(R)$ *cannot be obtained from R by expressing, shifting, non-negative scaling and use of* Opt*. Similarly, the relation* $\mathrm{Opt}(R)$ *cannot be obtained from R by expressing, shifting, non-negative scaling and use of* Feas*.*

Remark 1.26. *Note that for every valued structure Γ and $R \in \langle \Gamma \rangle$, every automorphism of Γ is an automorphism of R.*

The motivation for Definition 1.23 for valued CSPs stems from the following lemma, which shows that adding relations in $\langle \Gamma \rangle$ does not change the complexity of the VCSP up to polynomial-time reductions. For finite-domain valued structures this is proved in [40], except for the operator Opt, for which a proof can be found in [47, Theorem 5.13]. Parts of the proof have been generalised to infinite-domain valued structures without further assumptions; see, e.g. [77] and [80, 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 [77] the authors restrict themselves to finite-valued relations and hence do not consider the operator Opt. It is visible from Example 1.25 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 [47] (Feas was included implicitly by allowing to scale by 0 and defining $0 \cdot \infty = \infty$). In this thesis we work with valued structures with an oligomorphic automorphism group; in this setting, the cost of an expression is always realized by some tuple and therefore we can adapt the proof from the finite-domain case to show that the complexity is preserved. The result below originally appeared in [30].

Lemma 1.27. *Let $\mathfrak{A}$ be a valued structure on a countable domain with an oligomorphic automorphism group and a finite signature. Suppose that $\mathfrak{B}$ is a valued structure with a finite signature over the same domain A such that every valued relation of $\mathfrak{B}$ is from $\langle \mathfrak{A} \rangle$. Then there is a polynomial-time reduction from* VCSP($\mathfrak{B}$) *to* VCSP($\mathfrak{A}$).

Proof. Let τ be the signature of Γ. It suffices to prove the statement for expansions of Γ to signatures $\tau \cup \{R\}$ that extend τ with a single relation R, $R^\Delta \in \langle \Gamma \rangle$.

If $R^\Delta = (\emptyset)_0^\infty$, then an instance ϕ of VCSP(Δ) with threshold $u \in \mathbb{Q}$ is unsatisfiable if and only if ϕ contains the symbol R or if it does not contain R and is unsatisfiable viewed as an instance of VCSP(Γ). In the former case, choose a k-ary relation symbol $S \in \tau$ and note that S^Γ attains only finitely many values, by the oligomorphicity of Aut(Γ). Let $u' \in \mathbb{Q}$ be smaller than all of them. Then $S(x_1, \ldots, x_k)$ is an instance of VCSP(Γ) that never meets the threshold u', so this provides a correct reduction. In the latter case, for every $t \in C^n$ we have that $\phi^\Delta(t) = \phi^\Gamma(t)$; this provides a polynomial-time reduction.

Now suppose that R^Δ is equal to $(=)_0^\infty$. Let $\psi(x_{i_1}, \ldots, x_{i_k})$ be obtained from an instance $\phi(x_1, \ldots, x_n)$ of VCSP(Δ) by identifying all variables x_i and x_j such that ϕ contains the summand $R(x_i, x_j)$. Then ϕ is satisfiable if and only if the instance ψ is satisfiable, and $\inf_{t \in C^n} \phi^\Delta(t) = \inf_{s \in C^k} \psi^\Gamma(s)$; Again, this provides a polynomial-time reduction.

Next, consider that for some τ-expression $\delta(y_1, \ldots, y_l, z_1, \ldots, z_k)$ we have

$$R^\Delta(y_1, \ldots, y_l) = \inf_{t \in C^k} \delta^\Gamma(y_1, \ldots, y_l, t_1, \ldots, t_k).$$

Let $\phi(x_1, \ldots, x_n)$ be an instance of VCSP(Δ). We replace each summand $R(y_1, \ldots, y_l)$ in ϕ by $\delta(y_1, \ldots, y_l, z_1, \ldots, z_k)$ where $z_1, \ldots, z_k$ are new variables (different for each summand). After doing this for all summands that involve R, let $\theta(x_1, \ldots, x_n, w_1, \ldots, w_t)$ be the resulting τ-expression. For any $t \in C^n$ we have that

$$\phi(t_1, \ldots, t_n) = \inf_{s \in C^t} \theta(t_1, \ldots, t_n, s)$$

and hence $\inf_{a \in C^n} \phi = \inf_{c \in C^{n+t}} \theta$; here we used that the infima are realized. Since we replace each summand by an expression whose size is constant (since Γ is fixed and finite) the expression θ can be computed in polynomial time, which shows the statement.

Suppose that $R^\Delta = aS^\Gamma + b$ where $a \in \mathbb{Q}_{\geq 0}, b \in \mathbb{Q}$. Let $c \in \mathbb{Z}_{\geq 0}$ and $d \in \mathbb{Z}_{>0}$ be coprime integers such that $c/d = a$. Let (ϕ, u) be an instance of VCSP(Δ) where $\phi(x_1, \ldots, x_n) = \sum_{i=1}^{\ell} \phi_i + \sum_{j=1}^{k} \psi_j$, the summands ϕ_i contain only symbols from τ, and each ψ_j involves the symbol R. Let ψ'_j be the expression obtained from ψ_j by replacing R with S. We replace in ϕ for every $i \in \{1, \ldots, \ell\}$ ϕ_i with d copies of itself and for every $j \in \{1,\dots,k\}$, replace ψ_j with c copies of ψ'_j; let $\phi'(x_1,\dots,x_n)$ be the resulting τ-expression. Define $u' := d(u-kb)$. Then for every $t \in C^n$ the following are equivalent:

$$\phi(t) = \sum_{i=1}^{\ell} \phi_i + \sum_{j=1}^{k} \Big(\frac{c}{d}\psi'_j + b\Big) \leq u$$
$$\phi'(t) = d\sum_{i=1}^{\ell} \phi_i + c\sum_{j=1}^{k} \psi'_j \leq du - dkb = u'$$

Since (ϕ', u') can be computed from (ϕ, u) in polynomial time, this provides the desired reduction.

Now suppose that $R^\Delta = \mathrm{Feas}(S^\Gamma)$ for some $S \in \tau$. Let (ϕ, u) be an instance of VCSP(Δ), i.e., $\phi(x_1,\dots,x_n) = \sum_{i=1}^{\ell} \phi_i + \sum_{j=1}^{k} \psi_j$ where ψ_j, $j \in \{1,\dots,k\}$ are all the atomic expressions in ϕ that involve R. If $R^\Delta = (\emptyset)_0^\infty$, then the statement follows from the reduction for $(\emptyset)_0^\infty$. Therefore, suppose that this not the case. Since τ is finite and $\mathrm{Aut}(\Gamma)$ is oligomorphic, we may assume without loss of generality that all valued relations attain only non-negative values; otherwise we shift the values, which by the previous case does not affect the complexity up to polynomial-time reductions. Let w be the maximum finite weight assigned by S. Note that there are only finitely many values that the ℓ atoms ϕ_i may take and therefore only finitely many values that $\sum_{i=1}^{\ell} \phi_i$ may take. Let v be the smallest of these values such that $v > u$ and let $d = v - u$; if v does not exist, let $d = 1$. To simplify the notation, set $a = \lceil (kw)/d \rceil + 1$. Let ψ'_j be the τ-expression resulting from ψ_j by replacing the symbol R by the symbol S. Let ϕ' be the τ-expression obtained from ϕ by replacing each atom ϕ_i with a copies of it and replacing every atom ψ_j by ψ'_j. Let $(\phi', au + kw)$ be the resulting instance of VCSP(Γ); note that it can be computed in polynomial time.

We claim that for every $t \in C^n$, the following are equivalent:

$$\phi(t) = \sum_{i=1}^{\ell} \phi_i + \sum_{j=1}^{k} \psi_j \leq u \tag{1.1}$$
$$\phi'(t) = a \cdot \sum_{i=1}^{\ell} \phi_i + \sum_{j=1}^{k} \psi'_j \leq au + kw \tag{1.2}$$

If (1.1) holds, then by the definition of Feas we must have $\psi_j = 0$ for every $j \in \{1,\dots,k\}$. Thus $\sum_{i=1}^{\ell} \phi_i \leq u$ and $\sum_{j=1}^{k} \psi'_j \leq kw$, which implies (1.2). Conversely, if (1.2) holds, then ψ'_j is finite for every $j \in \{1,\dots,k\}$ and hence $\psi_j = 0$. Moreover, (1.2) implies

$$\sum_{i=1}^{\ell} \phi_i \leq u + \frac{kw}{a}.$$

Note that if v exists, then $u + (kw)/a < v$. Therefore (regardless of the existence of v), this implies $\sum_{i=1}^{\ell} \phi_i \leq u$, which together with what we have observed previously shows (1.1).

Finally, we consider the case that $R^\Delta = \mathrm{Opt}(S^\Gamma)$ for some relation symbol $S \in \tau$. Similarly to the previous case, we may assume without loss of generality that the minimum weight of all valued relations in Δ equals 0; otherwise, we subtract the smallest weight assigned to a tuple by some valued relation in Δ. This transformation does not affect the computational complexity of the VCSP (up to polynomial-time reductions). We may also assume that S^Γ takes finite positive values, because otherwise $\mathrm{Opt}(S^\Gamma) = S^\Gamma$ and the statement is trivial. Let m be the smallest positive weight assigned by S^Γ and let M be the largest finite weight assigned by any valued relation of Γ (we again use that τ is finite and that $\mathrm{Aut}(\Gamma)$ is oligomorphic). Let (ϕ, u), where $\phi(x_1, \dots, x_n) = \sum_{i=1}^k \phi_i$, be an instance of $\mathrm{VCSP}(\Delta)$. For $i \in \{1, \dots, k\}$, if ϕ_i involves the symbol R, then replace it by $k \cdot \lceil M/m \rceil + 1$ copies and replace R by S. Let ϕ' be the resulting τ-expression. We claim that $t \in C^n$ is a solution to the instance $(\phi', \min(kM, u))$ of $\mathrm{VCSP}(\Gamma)$ if and only if it is the solution to (ϕ, u).

If $t \in C^n$ is such that $\phi(t) \leq u$ then for every $i \in \{1, \dots, k\}$ such that ϕ_i involves R we have $\phi_i(t) = 0$. Recall that the minimal value attained by S^Γ equals 0 by our assumption, and hence $\phi'(t) = \phi(t) \leq u$ and therefore $\phi'(t) \leq \min(kM, u)$ by the choice of M. Now suppose that $\phi(t) > u$. Then $\phi'(t) > u \geq \min(kM, u)$ or there exists an $i \in \{1, \dots, k\}$ such that $\phi_i(t) = \infty$. If ϕ_i does not involve the symbol R, then $\phi'(t) = \infty$ as well. If ϕ_i involves the symbol R, then $\phi'(t) \geq (k \cdot \lceil M/m \rceil + 1)m > kM$. In any case, $\phi'(t) > \min(kM, u)$. Since ϕ' can be computed from ϕ in polynomial time, this concludes the proof. $\square$

Lemma 1.27 above provides tools for polynomial-time reductions between VCSPs on the same domain. We give two examples of how reductions from CSPs based on this lemma can be utilized for hardness proofs.

Example 1.28. *Recall the structure Γ_{max} from Example 1.2. We have seen in Example 1.7 that $\mathrm{VCSP}(\Gamma_{max})$ is the directed max-cut problem. Note that, for all $x, y, z \in \{0,1\}$,*

$$\mathrm{NAE}(x,y,z) = \mathrm{Opt}\left((<_2)_0^1(x,y) + (<_2)_0^1(y,z) + (<_2)_0^1(z,x)\right).$$

Since $\mathrm{CSP}(\{0,1\}, \mathrm{NAE})$ is an NP-hard variant of the 3-SAT-problem (see Example 1.20), this provides an alternative proof of the NP-hardness of the directed max-cut problem via Lemma 1.27.

Example 1.29. *We revisit the valued structure Γ_{LCC} from Example 1.11. Recall that $\mathrm{VCSP}(\Gamma_{LCC})$ is the least correlation clustering problem with partial information and that $\mathrm{Aut}(\Gamma_{LCC})$ is oligomorphic. Let $\mathfrak{B}$ be the relational structure $(\mathbb{N}; R)$, where*

$$R := \{(x,y,z) \in \mathbb{N}^3 \mid (x = y \wedge y \neq z) \vee (x \neq y \wedge y = z)\}.$$

Note that

$$R(x,y,z) = \mathrm{Opt}((\neq)_0^1(x,z) + (\neq)_0^1(x,z) + (=)_0^1(x,y) + (=)_0^1(y,z)).$$

Clearly, $\mathfrak{B}$ is the same valued structure as $(\mathbb{Q}; \mathrm{Dis})$ from Example 1.21 up to renaming elements. Since $\mathrm{CSP}(\mathbb{Q}; \mathrm{Dis})$ is NP-hard, this provides an alternative proof of NP-hardness of the least correlation clustering problem with partial information via Lemma 1.27.

Note that we can replace $(\neq)_0^1(x,z)+(\neq)_0^1(x,z)$ *in the definition of* R *by* $(\neq)_0^\infty(x,z)$. *This shows that even* $\mathrm{VCSP}(\mathbb{N};(=)_0^1,(\neq)_0^\infty)$ *is NP-hard.*

1.5.1 The underlying crisp structure

For understanding the complexity of VCSPs, it is often crucial to understand the complexity of the crisp relations that are expressible in the template. We therefore introduce the following notation.

Definition 1.30. *Let* Γ *be a valued structure on the domain* C. *Then* $\langle\Gamma\rangle_0^\infty$ *denotes the set of valued relations*

$$\{R^\Gamma \in \langle\Gamma\rangle \mid R \text{ of arity } k, \forall t \in C^k \colon R(t) \in \{0,\infty\}\}.$$

We sometimes refer to $(C;\langle\Gamma\rangle_0^\infty)$ *as the* underlying crisp structure *of* Γ.

In words, $\langle\Gamma\rangle_0^\infty$ contains all crisp relations that can be expressed in Γ. The following example shows that some information is lost in the transition from a valued structure to its underlying crisp structure; in particular, they might have different automorphisms groups.

Example 1.31. *Let* $\Gamma=(\mathbb{Q};R)$ *where* R *is a binary valued relation on* $\mathbb{Q}$ *defined by*

$$R(x,y)=\begin{cases}0 & x=y,\\ 1 & x<y,\\ 2 & x>y.\end{cases}$$

Let $\Gamma'=(\mathbb{Q};\langle\Gamma\rangle_0^\infty)$. *Clearly,* $\mathrm{Aut}(\Gamma)=\mathrm{Aut}(\mathbb{Q};<)\subseteq\mathrm{Aut}(\Gamma')$. *We claim that* $\mathrm{Aut}(\Gamma')=\mathrm{Sym}(\mathbb{Q})\supsetneq\mathrm{Aut}(\Gamma)$.

A relation $S\subseteq\mathbb{Q}^k$ *will be called* dull *if* S *is primitively positively definable in* $(\mathbb{Q};=)$ *and either* $S=\emptyset$ *or* $(a,\dots,a)\in S$ *for every* $a\in\mathbb{Q}$. *We claim that for every* $S\in\langle\Gamma\rangle$, $\mathrm{Feas}(S)$ *and* $\mathrm{Opt}(S)$ *are dull, which implies that every* $S\in\langle\Gamma\rangle_0^\infty$ *is dull (recall that for a crisp relation* $S=\mathrm{Feas}(S)=\mathrm{Opt}(S)$*) and hence* $\mathrm{Aut}(\Gamma')=\mathrm{Sym}(\mathbb{Q})$.

We prove the claim by induction on the depth of expressions that define S. *As a base step, we note that* $(\emptyset)_0^\infty$, $(\neq)_0^\infty$, $\mathrm{Feas}(R)$ *and* $\mathrm{Opt}(R)$ *are dull, therefore all relations expressed by atomic expressions over* $\{R\}$ *satisfy the claim.*

For the induction step, let $S,S'\in\langle\Gamma\rangle$ *of arity* k *and* k', *respectively, and suppose that* $\mathrm{Feas}(S)$, $\mathrm{Opt}(S)$, $\mathrm{Feas}(S')$ *and* $\mathrm{Opt}(S')$ *are dull. Then, trivially,* $\mathrm{Feas}(S)$ *and* $\mathrm{Opt}(S)$ *satisfy the claim. Let* $a\in\mathbb{Q}_{\geq 0}$ *and* $b\in\mathbb{Q}$. *Then* $\mathrm{Feas}(S)=\mathrm{Feas}(aS)=\mathrm{Feas}(S+b)$ *and* $\mathrm{Opt}(S)=\mathrm{Opt}(aS)=\mathrm{Opt}(S+b)$, *and therefore* aS *and* $S+b$ *satisfy the claim. Let* $\rho\colon[k]\to[k'']$, $\sigma\colon[k']\to[k'']$ *and let* S'' *be a valued relation of arity* k'' *defined by*

$$S''(x_1,\dots,x_{k''})=S(x_{\rho(1)},\dots,x_{\rho(k)})+S'(x_{\sigma(1)},\dots,x_{\sigma(k')}).$$

Then

$$\mathrm{Feas}(S'')(x_1,\dots,x_{k''})=\mathrm{Feas}(S)(x_{\rho(1)},\dots,x_{\rho(k)})\wedge\mathrm{Feas}(S')(x_{\sigma(1)},\dots,x_{\sigma(k')})$$

and hence $\mathrm{Feas}(S'')$ *is dull by the assumption, because either one of* $\mathrm{Feas}(S)$ *or* $\mathrm{Feas}(S')$ *is empty, or they both contain all constant tuples. The argument for* $\mathrm{Opt}(S'')$ *is analogous. Finally, let* $k'' < k$ *and*

$$S''(x_1, \ldots, x_{k''}) := \inf_{x_{k''+1},\ldots,x_k} S(x_1, \ldots, x_k).$$

Then

$$\mathrm{Feas}(S'')(x_1, \ldots, x_{k''}) := \exists x_{k''+1}, \ldots, x_k \, \mathrm{Feas}(S)(x_1, \ldots, x_k)$$

and therefore $\mathrm{Feas}(S'')$ *is dull.* $\mathrm{Opt}(S'')$ *can be expressed analogously.*

It follows that every $S \in \langle \Gamma \rangle_0^\infty$ *is dull and hence* $\mathrm{Aut}(\Gamma') = \mathrm{Sym}(\mathbb{Q})$. *Therefore,* $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\mathbb{Q}; <) \subsetneq \mathrm{Aut}(\Gamma')$.

However, under some assumptions, the situation from Example 1.31 can be avoided; this will be shown in Proposition 3.7. We remark that essentially crisp valued structures with behave like relational structures.

Remark 1.32. *Suppose that* Γ *is an essentially crisp valued* τ*-structure. For every* $R \in \tau$*, let* $a_R \in \mathbb{Q}$ *be such that* R^Γ *only attains values in* $\{a_R, \infty\}$*; such an* a_R *exists because* Γ *is essentially crisp. Then* $R^\Gamma = \mathrm{Feas}(R^\Gamma) + a_R$*. Therefore,* $\langle \Gamma \rangle = \langle \mathrm{Feas}(\Gamma) \rangle$ *and, by Remark 1.24,* $\langle \Gamma \rangle_0^\infty$ *consists of precisely those relations that are primitively positively definable in* $\mathrm{Feas}(\Gamma)$*. By Lemma 1.27 and Remark 1.8, there is a polynomial-time reduction from* $\mathrm{VCSP}(\Gamma)$ *to* $\mathrm{CSP}(\mathrm{Feas}(\Gamma))$ *and vice versa.*

Chapter 2

Pp-constructability for valued structures

A universal-algebraic theory of VCSPs for finite-domain valued structures has been developed in [58], 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 [6]. We adapt this perspective to the algebraic theory of VCSPs and develop (parts of) this approach here. As in [6], we immediately step from valued structures with a finite domain to the more general case of valued structures with an oligomorphic automorphism group. Most of the original results in this chapter were published in [30] or announced in the preprint [12].

2.1 Pp-powers

We start with defining the concept of pp-powers.

Definition 2.1 (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*

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

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

Proof. Let d be the dimension of the pp-power and let τ be the signature of Γ. By Remark 1.26, $\mathrm{Aut}(\Gamma) \subseteq \mathrm{Aut}(\Delta)$ and thus $\mathrm{Aut}(\Delta)$ is oligomorphic. By Lemma 1.27, we may suppose that for every valued relation R of arity k of Δ the valued relation $S \in \langle \Gamma \rangle$ of arity dk from the definition of a pp-power equals S^Γ for some $S \in \tau$. Let (ϕ, u) be an instance of $\mathrm{VCSP}(\Delta)$. For each variable x of ϕ we introduce d new variables $x_1, \dots, x_d$. For each summand $R(y^1, \dots, y^k)$ of ϕ we introduce a summand $S(y^1_1, \dots, y^1_d, \dots, y^k_1, \dots, y^k_d)$; let ψ be the resulting τ-expression. It is now straightforward to verify that (ϕ, u) has a solution with respect to Δ if and only if (ψ, u) has a solution with respect to Γ. □

2.2 Fractional homomorphisms

In this thesis, we work with a generalization of maps that is useful in the valued setting, so-called *fractional maps*. These are probability distributions on maps. To establish some invariance properties that are used for classifying complexity of VCSPs, we need to compute expected values of random variables with respect to these distributions. Since the probability distributions we consider are possibly on uncountable sets, we have to use Lebesgue integrals to define the expected values. The following sections introduce some topological notions and Lebesgue integrals, which allows to define expected values of random variables (Section 2.2.1–2.2.3). Finally, in Section 2.2.4 we introduce fractional maps and *fractional homomorphisms*. To keep this material more accessible, all of the notions are specialized to our setting where the topological spaces are spaces of functions on discrete sets.

2.2.1 Topology

If C and D are sets, 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

$$\mathscr{S}_{s,t} := \{f \in C^D \mid f(s) = t\}$$

for $s \in D^k$ and $t \in C^k$ for some $k \in \mathbb{N}$, and f is applied componentwise. For $S \subseteq C^D$, we denote by $\overline{S}$ the closure of S in the topology of pointwise convergence.

We write $[0,1]$ for the set $\{x \in \mathbb{R} \mid 0 \leq x \leq 1\}$. 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 $\{x \in \mathbb{R} \mid x > a\} \cup \{\infty\}$.

2.2.2 Lebesgue integral

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

It will be convenient to use an additional value $-\infty$ that has the usual properties:

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

The sum of ∞ and $-\infty$ is undefined.

Let C and D be sets. We define the Lebesgue integration over the space C^D of all functions from D to C. We often work with the special case $D = C^\ell$, i.e. the space is the space of all operations on C of arity $\ell \in \mathbb{N}$, which we denote by $\mathscr{O}_C^{(\ell)}$.

To define the Lebesgue integral, we need the definition of a *simple function*: this is a function $Y \colon C^D \to \mathbb{R}$ given by

$$\sum_{k=1}^{n} a_k 1_{S_k}$$

where $n \in \mathbb{N}$, $S_1, S_2, \ldots S_n$ are disjoint elements of $\mathcal{B}(C^D)$, $a_1, \ldots, a_n \in \mathbb{R}$, and $1_S \colon C^D \to \{0,1\}$ denotes the indicator function for $S \subseteq C^D$. If Y is a such a simple function, then the Lebesgue integral is defined as follows:

$$\int_{C^D} Y d\omega := \sum_{k=1}^{n} a_k \omega(S_k).$$

If X and Y are two random variables, then we write $X \leq Y$ if $X(f) \leq Y(f)$ for every $f \in C^D$. We say that X is *non-negative* if $0 \leq X(f)$ for every $f \in C^D$ and we write $0 \leq X$. If X is a non-negative measurable function, then the Lebesgue integral is defined as

$$\int_{C^D} X d\omega := \sup \left\{ \int_{C^D} Y d\omega \mid 0 \leq Y \leq X, Y \text{ simple} \right\}.$$

For an arbitrary measurable function X, we write $X = X^+ - X^-$, where

$$X^+(x) := \begin{cases} X(x) & \text{if } X(x) > 0, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$X^-(x) := \begin{cases} -X(x) & \text{if } X(x) < 0, \\ 0 & \text{otherwise.} \end{cases}$$

Then both X^+ and X^- are measurable, and both $\int_{C^D} X^- d\omega$ and $\int_{C^D} X^+ d\omega$ take values in $\mathbb{R}_{\geq 0} \cup \{\infty\}$. If both take value ∞, then the integral is undefined. Otherwise, define

$$\int_{C^D} X d\omega := \int_{C^D} X^+ d\omega - \int_{C^D} X^- d\omega.$$

In particular, note that for $X \geq 0$ the integral is always defined.

2.2.3 Expected values

If C, D are sets and $X : C^D \to \mathbb{R} \cup \{\infty\}$ 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

$$E_\omega[X] := \int_{C^D} X d\omega.$$

Recall that the Lebesgue integral $\int_{C^D} X d\omega$ need not exist, in which case $E_\omega[X]$ is undefined; otherwise, the integral equals a real number, ∞, or $-\infty$. Note that if C and D are infinite sets, then C^D is uncountable and there are probability distributions ω such that $\omega(\{f\}) = 0$ for every $f \in C^D$. In that case, we need Lebesgue integrals to be able to correctly define $E_\omega[X]$.

It follows easily from the definition of Lebesgue integral 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 $X \leq Y$, then $E_\omega[X] \leq E_\omega[Y]$.

2.2.4 Fractional maps

We can now define fractional maps and homomorphisms, which form the basis for the algebraic approach to classifying complexity of VCSPs.

Definition 2.3 (fractional map)**.** *Let C and D be sets. A* fractional map *from D to C is a probability distribution*

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

that is, $\omega(C^D) = 1$ and ω is countably additive: if $A_1, A_2, \dots \in \mathcal{B}(C^D)$ 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 $\{f\} \in \mathcal{B}(C^D)$ for every $f : D \to C$. If there is a single $f \in C^D$ such that $\omega(f) = 1$, we typically write f instead of ω.

Let ω be a fractional map from D to C, let $R \in \mathscr{R}_C^{(k)}$ be a valued relation, and let $s \in D^k$. Then $X\colon C^D \to \mathbb{R} \cup \{\infty\}$ given by

$$f \mapsto R(f(s))$$

is a random variable: if (a, b) is a basic open subset of $\mathbb{R} \cup \{\infty\}$, then

$$\begin{aligned} X^{-1}((a,b)) &= \{f \in C^D \mid R(f(s)) \in (a,b)\} \\ &= \bigcup_{t \in C^k, R(t) \in (a,b)} \mathscr{S}_{s,t} \end{aligned}$$

is a union of basic open sets in C^D, hence open. The argument for the other basic open sets in $\mathbb{R} \cup \{\infty\}$ is similar. Random variables of this form will play a crucial role in this thesis, starting from the definition below.

Definition 2.4 (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 ω from D to C such that for every $R \in \tau$ of arity k and every tuple $s \in D^k$ it holds for the random variable $X \colon C^D \to \mathbb{R} \cup \{\infty\}$ given by*

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

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

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

If the set C is countable and $X : f \mapsto R(f(s))$ for some $R \in \mathscr{R}_C^{(k)}$ and $s \in D^k$, we may express $E_\omega[X]$ as a sum, which is useful in proofs throughout the thesis. If $E_\omega[X]$ exists, then it is equal to

$$\begin{aligned} E_\omega[X] &= \int_{C^D} X^+ d\omega - \int_{C^D} X^- d\omega \\ &= \sup\left\{\int_{C^D} Y d\omega \mid 0 \leq Y \leq X^+, Y \text{ simple}\right\} \\ &\quad - \sup\left\{\int_{C^D} Y d\omega \mid 0 \leq Y \leq X^-, Y \text{ simple}\right\} \\ &= \sum_{t \in C^k, R(t) \geq 0} R(t)\omega(\mathscr{S}_{s,t}) + \sum_{t \in C^k, R(t) < 0} R(t)\omega(\mathscr{S}_{s,t}) \\ &= \sum_{t \in C^k} R(t)\omega(\mathscr{S}_{s,t}). \end{aligned} \tag{2.1}$$

Note that if the fractional map $\omega = f$ for some $f \in C^D$, then for every $R \in \tau$ of arity k and $s \in D^k$, the expected value in Definition 2.4 exists and is equal to $R^{\mathfrak{A}}(f(s))$. If Γ and Δ are crisp structures and $f \in C^D$ is a fractional homomorphism from Δ to Γ, then we also call f a *homomorphism*; it is straightforward to check that this definition coincides with the standard definition of a homomorphism between relational structures. For results about relational structures and studying resilience problems we will often use the following lemma that can be proved using König's tree lemma.

Lemma 2.5 ([9, Lemma 4.1.7])**.** *Let $\mathfrak{B}$ be a relational τ-structure with an oligomorphic automorphism group and $\mathfrak{A}$ be a countable relational structure over the same signature τ. If there is no homomorphism (embedding) from $\mathfrak{A}$ to $\mathfrak{B}$, then there is a finite substructure of $\mathfrak{A}$ that does not homomorphically map (embed) to $\mathfrak{B}$.*

In the example below, we illustrate that there are valued structures Γ and Δ such that there is a fractional homomorphism from Δ to Γ, but no fractional homomorphism of the form $f \in C^D$.

Example 2.6. *Let $\tau = \{R, S\}$, where R, S are binary relational symbols. We consider two valued τ-structures Γ and Δ. Let $\Gamma = (\{0,1\}; R^\Gamma, S^\Gamma)$, where*

$$R^\Gamma(x,y) = \begin{cases} 0 & x = 0, \\ 1 & x \neq 0, \end{cases}$$

$$S^\Gamma(x,y) = \begin{cases} 0 & y = 1, \\ 1 & y \neq 1. \end{cases}$$

Let $\Delta = (\{0,1\}^2; R^\Delta, S^\Delta)$ where

$$R^\Delta((x,u),(y,v)) = \frac{1}{2}(R^\Gamma(x,y) + R^\Gamma(u,v))$$

and S^Δ is defined analogously from S^Γ; note that Δ is a pp-power of Γ. Let $\min_2$ and $\max_2$ denote the minimum and maximum operation on $\{0,1\}$, respectively. Let ω be a fractional map from $\{0,1\}^2$ to $\{0,1\}$ defined by $\omega(\min_2) = \omega(\max_2) = 1/2$. Then it is easy to verify that for all $x, y, u, v \in \{0,1\}$

$$\begin{aligned} &E_\omega[f \mapsto R^\Gamma(f(x,u), f(y,v))] \\ =&\frac{1}{2}R^\Gamma\left(\min_2(x,u), \min_2(y,v)\right)) + \frac{1}{2}R^\Gamma\left(\max_2(x,u), \max_2(y,v)\right) \\ \leq&\frac{1}{2}(R^\Gamma(x,y) + R^\Gamma(u,v)) \\ =&R^\Delta((x,u),(y,v)), \end{aligned}$$

(we avoid componentwise application in the expressions above and prefer to write the expressions more precisely). Analogously we obtain that

$$E_\omega[f \mapsto S^\Gamma(f(x,u), f(y,v))] \leq S^\Delta((x,u),(y,v)).$$

It follows that ω is a fractional homomorphism from Δ to Γ.

Suppose there is $f : \{0,1\}^2 \to \{0,1\}$ such that f is a fractional homomorphism from Δ to Γ. Then we have

$$R^\Gamma(f(0,1), f(1,0)) \leq R^\Delta((0,1),(1,0)) = \frac{1}{2} \text{ and}$$
$$S^\Gamma(f(1,0), f(0,1)) \leq S^\Delta((1,0),(0,1)) = \frac{1}{2}.$$

The inequality for R implies $f(0,1) = 0$ and the inequality for S implies $f(0,1) = 1$, a contradiction. Therefore, such f does not exist.

We remark that the expected value from Definition 2.4 may be undefined and give an example below.

Example 2.7. *Consider $C = D = \mathbb{N}$ and the unary valued relation on $\mathbb{N}$ defined by*

$$R(x) = \begin{cases} -2^x & \text{if } x \text{ is even,} \\ 2^x & \text{otherwise.} \end{cases}$$

Let $s \in \mathbb{N}$ and define $X \colon \mathbb{N}^{\mathbb{N}} \to \mathbb{R} \cup \{\infty\}$ by $f \mapsto R(f(s))$. Let ω be a unary fractional operation such that for every $t \in \mathbb{N}$ we have $\omega(\mathscr{S}_{s,t}) = \frac{1}{2^{t+1}}$. Then

$$\begin{aligned} \int_{\mathbb{N}^{\mathbb{N}}} X^+ d\omega &= \sup\{\int_{\mathbb{N}^{\mathbb{N}}} Y d\omega \mid 0 \leq Y \leq X^+, Y \text{ simple}\} \\ &= \sum_{t \in C, R(t) \geq 0} R(t)\omega(\mathscr{S}_{s,t}) \\ &= \sum_{t \in C, R(t) \geq 0} \frac{1}{2} = \infty \end{aligned}$$

and, similarly, $\int_{\mathbb{N}^{\mathbb{N}}} X^- d\omega = \infty$. It follows that

$$E_\omega[X] = \int_{\mathbb{N}^{\mathbb{N}}} X^+ d\omega - \int_{\mathbb{N}^{\mathbb{N}}} X^- d\omega$$

is undefined.

In contrast to Example 2.7, the following lemma shows that the oligomorphicity assumption on $\mathrm{Aut}(\Gamma)$ is enough to guarantee the existence of the expected value from Definition 2.4.

Lemma 2.8. *Let C and D be sets, $s \in D^k$, $R \in \mathscr{R}_C^{(k)}$. Let $X \colon C^D \to \mathbb{R} \cup \{\infty\}$ be the random variable given by*

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

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

Proof. It is enough to show that $\int_{C^D} X^- d\omega \neq \infty$. Since $\mathrm{Aut}(C; R)$ is oligomorphic, there are only finitely many orbits of k-tuples in $\mathrm{Aut}(C; R)$. Let $O_1, \dots, O_m$ be all orbits of k-tuples of $\mathrm{Aut}(C; R)$ on which R is negative. For every $i \in \{1, \dots, m\}$, let $t^i \in O_i$.

Then we obtain (see (2.1) for a detailed derivation of the first equality)

$$\begin{aligned}\int_{C^D} X^- d\omega &= \sum_{t \in C^k, R(t)<0} -R(t)\omega(\mathscr{S}_{s,t}) \\ &= -\sum_{i=1}^{m} R(t^i) \sum_{t \in O_i} \omega(\mathscr{S}_{s,t}) \\ &= -\sum_{i=1}^{m} R(t^i)\, \omega\left(\bigcup_{t \in O_i} \mathscr{S}_{s,t}\right) \\ &\leq -\sum_{i=1}^{m} R(t^i) < \infty. \end{aligned}$$ □

In the next lemma we show that fractional homomorphisms compose.

Lemma 2.9. *Let Γ_1, Γ_2, Γ_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.*

Proof of Lemma 2.9. Let C_1, C_2, C_3 be the domains of Γ_1, Γ_2, and Γ_3, respectively. If $s \in C_1^k$ and $u \in C_3^k$, for some $k \in \mathbb{N}$, then define

$$\omega_3(\mathscr{S}_{s,u}) := \sum_{t \in C_2^k} \omega_1(\mathscr{S}_{s,t})\omega_2(\mathscr{S}_{t,u}).$$

Note that on sets of this form, i.e., on basic open sets in $C_3^{C_1}$, ω_3 is countably additive. Since our basis of open sets is closed under intersection, this definition extends uniquely to all of $\mathcal{B}(C_3^{C_1})$ by Dynkin's π-λ theorem. □

The following proposition shows that fractional homomorphisms improve costs of instances. The statement of the proposition was shown for valued structures over finite domains in [34, Proposition 8.4].

Proposition 2.10. *Let Γ and Δ be valued τ-structures with domains C and D and with a fractional homomorphism ω from Δ to Γ. Then the cost of every* VCSP *instance ϕ with respect to Γ is at most the cost of ϕ with respect to Δ.*

Proof. Let

$$\phi(x_1, \dots, x_n) = \sum_{i=1}^{m} R_i(x_{j_1^i}, \dots, x_{j_{k_i}^i})$$

be a τ-expression, where $R_i \in \tau$ is of arity k_i and $j_1^i, \dots, j_{k_i}^i \in \{1, \dots, n\}$ for every $i \in \{1, \dots m\}$. To simplify the notation in the proof, if $v = (v_1, \dots, v_p)$ is a p-tuple of elements of some domain and $i_1, \dots, i_q \in \{1, \dots, p\}$, we will write $v_{i_1,\dots,i_q}$ for the tuple $(v_{i_1}, \dots, v_{i_q})$.

Let $\varepsilon > 0$. From the definition of infimum, there exists $s^* \in D^n$ such that

$$\phi^{\Delta}(s^*) \leq \inf_{s \in D^n} \phi^{\Delta}(s) + \varepsilon/2 \tag{2.2}$$

and $f^* \in C^D$ such that

$$\phi^{\Gamma}(f^*(s^*)) \leq \inf_{f \in C^D} \phi^{\Gamma}(f(s^*)) + \varepsilon/2. \tag{2.3}$$

For every $i \in \{1, \ldots, m\}$, $E_\omega[f \mapsto R_i^{\Gamma}(f(s^*)_{j_1^i, \ldots, j_{k_i}^i})]$ exists by the definition of a fractional homomorphism. Suppose first that $\sum_{i=1}^m E_\omega[f \mapsto R_i^{\Gamma}(f(s^*)_{j_1^i, \ldots, j_{k_i}^i})]$ is defined. Then by the monotonicity and linearity of E_ω and since ω is a fractional homomorphism we obtain

$$\begin{aligned}
\inf_{t \in C^n} \phi^{\Gamma}(t) &\leq \phi^{\Gamma}(f^*(s^*)) \\
&\leq \inf_{f \in C^D} \phi^{\Gamma}(f(s^*)) + \varepsilon/2 && \text{(by (2.3))} \\
&\leq E_\omega[f \mapsto \phi^{\Gamma}(f(s^*))] + \varepsilon/2 && \text{(by monotonicity of } E_\omega) \\
&= \sum_{i=1}^{m} E_\omega[f \mapsto R_i^{\Gamma}(f(s^*)_{j_1^i, \ldots, j_{k_i}^i})] + \varepsilon/2 && \text{(by linearity of } E_\omega) \\
&\leq \sum_{i=1}^{m} R_i^{\Delta}(s^*_{j_1^i, \ldots, j_{k_i}^i}) + \varepsilon/2 && (\omega \text{ is a frac. homomorphism}) \\
&= \phi^{\Delta}(s^*) + \varepsilon/2 \\
&\leq \inf_{s \in D^n} \phi^{\Delta}(s) + \varepsilon && \text{(by (2.2)).}
\end{aligned}$$

Since $\varepsilon > 0$ was chosen arbitrarily, it follows that the cost of ϕ with respect to Γ is at most the cost of ϕ with respect to Δ.

Suppose now that $\sum_{i=1}^m E_\omega[f \mapsto R_i^{\Gamma}(f(s^*)_{j_1^i, \ldots, j_{k_i}^i})]$ is not defined. Then there exists $i \in \{1, \ldots, m\}$ such that

$$E_\omega[f \mapsto R_i^{\Gamma}(f(s^*)_{j_1^i, \ldots, j_{k_i}^i})] = \infty.$$

By the definition of a fractional homomorphism, this implies that $R_i^{\Delta}(s^*_{j_1^i, \ldots, j_{k_i}^i}) = \infty$ and hence $\sum_{i=1}^m R_i^{\Delta}(s^*_{j_1^i, \ldots, j_{k_i}^i}) = \infty$. Therefore, we obtain as above that

$$\inf_{t \in C^n} \phi^{\Gamma}(t) \leq \inf_{s \in D^n} \phi^{\Delta}(s),$$

which is what we wanted to prove. □

Remark 2.11. *For finite domains, the converse of Proposition 2.10 is true as well [34, Proposition 8.4].*

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. By Lemma 2.9, fractional homomorphic equivalence is transitive and hence an equivalence relation on valued structures of the same signature.

Corollary 2.12. *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 equal as computational problems.*

Proof. By Proposition 2.10, for every instance ϕ, the values of ϕ with respect to Γ and Δ are equal. By Lemma 1.12, the cost is realized by some tuple in both structures and hence every instance ϕ with a threshold u has a solution with respect to Γ if and only if it has a solution with respect to Δ. $\square$

For fractionally homomorphically equivalent valued structures with an oligomorphic automorphism group, we can prove the following result of similar flavor as Proposition 2.10.

Proposition 2.13. *Let Γ and Δ be valued τ-structures with oligomorphic automorphism groups that are fractionally homomorphically equivalent. Let ω be a fractional homomorphism from Δ to Γ. Let $R \in \langle\Delta\rangle$ be of arity k and $R' \in \langle\Gamma\rangle$ be the valued relation obtained when the expression for R in Δ is interpreted over Γ. Then for every $s \in D^k$,*

$$E_\omega[f \mapsto R'(f(s))] \leq R(s). \tag{2.4}$$

Proof. Recall that by Lemma 2.8, the expected values from (2.4) exist. By the definition of a fractional homomorphism, (2.4) holds for every pair $(R, R') = (S^\Delta, S^\Gamma)$ where $S \in \tau$. Clearly, the same is true for $R = R_\emptyset$. To see that (2.4) holds for $R = R_=$, let $s \in D^2$. Note that either $s_1 = s_2$ in which case $f(s_1) = f(s_2)$ for every $f \in C^D$, and hence both sides of (2.4) are equal to 0, or $s_1 \neq s_2$, in which case $R(s) = \infty$ and (2.4) is again satisfied.

We will show that every valued relation R obtained from a valued relation satisfying (2.4) by a single operator from Definition 1.23 satisfies (2.4); the general statement then follows by induction. This is clear for valued relations R obtained by non-negative scaling and addition of constants, since these operations preserve (2.4) by the linearity of expectation. The assumption that Γ and Δ are fractionally homomorphically equivalent (rather than the existence of ω) is needed only for the operator Opt.

Let $\phi(x_1, \dots, x_k, y_1, \dots, y_n)$ be a τ-expression. Let R be the k-ary valued relation defined by $R(x) = \inf_{y \in D^n} \phi^\Delta(x, y)$ for every $x \in D^k$. Since ϕ is a τ-expression, there are $R_i \in \tau$ such that

$$\phi(x_1, \dots, x_k, y_1, \dots, y_n) = \sum_{i=1}^{m} R_i(x_{p^i_1}, \dots, x_{p^i_{k_i}}, y_{q^i_1}, \dots, y_{q^i_{n_i}})$$

for some $k_i, n_i \in \mathbb{N}$, $p^i_1, \dots, p^i_{k_i} \in \{1, \dots, k\}$ and $q^i_1, \dots, q^i_{n_i} \in \{1, \dots, n\}$. In this proof, if $v = (v_1, \dots, v_N)$ is a tuple and $i_1, \dots, i_\ell \in \{1, \dots, N\}$, we will write $v_{i_1,\dots,i_\ell}$ for the tuple $(v_{i_1}, \dots, v_{i_\ell})$ for short.

Let $s \in D^k$. By the oligomorphicity of $\mathrm{Aut}(\Delta)$, there is $t \in D^n$ such that $R(s) = \phi^\Delta(s,t)$. Moreover, for every $f \in C^D$,

$$R'(f(s)) \leq \phi^\Gamma(f(s), f(t)).$$

By the linearity and monotonicity of expectation, we obtain

$$\begin{aligned} E_\omega[f \mapsto R'(f(s))] &\leq E_\omega[f \mapsto \phi^\Gamma(f(s), f(t))] \\ &= E_\omega[f \mapsto \sum_{i=1}^m R_i^\Gamma((f(s))_{p_1^i,\dots,p_{k_i}^i}, (f(t))_{q_1^i,\dots,q_{n_i}^i})] \\ &= \sum_{i=1}^m E_\omega[f \mapsto R_i^\Gamma((f(s))_{p_1^i,\dots,p_{k_i}^i}, (f(t))_{q_1^i,\dots,q_{n_i}^i})]. \end{aligned}$$

Since ω is a fractional homomorphism, the last row of the inequality above is at most

$$\sum_{i=1}^m R_i^\Delta(s_{p_1^i,\dots,p_{k_i}^i}, t_{q_1^i,\dots,q_{n_i}^i}) = \phi^\Delta(s,t) = R(s).$$

It follows that (2.4) holds for R.

Next, we prove the statement for $R = \mathrm{Feas}(S^\Delta)$ for some $S \in \tau$ of arity k. Let $s \in D^k$. If $R(s) = \infty$, then (2.4) is trivially true. So suppose that $R(s) = 0$, i.e., $S^\Delta(s) < \infty$. Since ω is a fractional homomorphism, we have

$$E_\omega[f \mapsto S^\Gamma(f(s))] \leq S^\Delta(s) \tag{2.5}$$

and hence the expected value on the left-hand side is finite as well. By (2.1),

$$E_\omega[f \mapsto S^\Gamma(f(s))] = \sum_{t \in C^k} S^\Gamma(t)\omega(\mathscr{S}_{s,t}), \tag{2.6}$$

which implies that $S^\Gamma(t)$ is finite unless $\omega(\mathscr{S}_{s,t}) = 0$, and hence $R'(t) = 0$. Consequently (again by (2.1)),

$$E_\omega[f \mapsto R'(f(s))] = \sum_{t \in C^k} R'(t)\omega(\mathscr{S}_{s,t}) = 0 = R(s).$$

It follows that (2.4) holds for R.

Finally, suppose that $R = \mathrm{Opt}(S^\Delta)$. Let $s \in D^k$; note that we may again assume that $R(s) = 0$ as we did in the previous case. This means that $S^\Delta(s) \leq S^\Delta(s')$ for every $s' \in D^k$. Let $u \in C^k$ be such that $S^\Gamma(u)$ is minimal; such a u exists by the oligomorphicity of $\mathrm{Aut}(\Gamma)$. By the monotonicity of expected value and since ω is a fractional homomorphism, we have

$$S^\Gamma(u) \leq E_\omega[f \mapsto S^\Gamma(f(s))] = \sum_{t \in C^k} S^\Gamma(t)\omega(\mathscr{S}_{s,t}) \leq S^\Delta(s). \tag{2.7}$$

If ω' is a fractional homomorphism from Γ to Δ (which exists by the assumption) and since $s \in \mathrm{Opt}(S^\Delta)$, we also have

$$S^\Delta(s) \leq E_{\omega'}[g \mapsto S^\Delta(g(u))] \leq S^\Gamma(u).$$

This implies that $S^\Delta(s) = S^\Gamma(u)$. Since ω is a probability distribution, we obtain from (2.7) that $S^\Gamma(t) = S^\Gamma(u)$ unless $\omega(\mathscr{S}_{s,t}) = 0$, and hence $R'(t) = \mathrm{Opt}(S^\Gamma)(t) = 0$. Therefore,

$$E_\omega[f \mapsto R'(f(s))] = \sum_{t \in C^k} R'(t)\omega(\mathscr{S}_{s,t}) = 0 = R(s).$$

This concludes the proof. □

2.3 Pp-constructions

Definition 2.14 (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.

Instead of 'Δ has a pp-construction in Γ', we often say 'Γ pp-constructs Δ'. Combining Proposition 2.2 and Corollary 2.12 yields the following.

Corollary 2.15. *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 1.28 and 1.29 are special cases of Corollary 2.15. We give a more involved example in Example 2.18. We first show that the relation of pp-constructibility on the class of countable valued structures is transitive.

Lemma 2.16. *Let* Γ_1*,* Γ_2*, and* Γ_3 *be valued structures, each with a countable domain. Suppose that* Γ_1 *pp-constructs* Γ_2 *and* Γ_2 *pp-constructs* Γ_3*. Then* Γ_1 *pp-constructs* Γ_3*.*

Proof. Clearly, a pp-power of a pp-power is again a pp-power, and fractional homomorphic equivalence is transitive by Lemma 2.9. We are therefore left to prove that if Γ and Δ are valued structures such that Δ is a d-dimensional pp-power of Γ, and if Γ' is fractionally homomorphically equivalent to Γ via fractional homomorphisms $\omega_1 \colon \Gamma \to \Gamma'$ and $\omega_2 \colon \Gamma' \to \Gamma$, then Δ also has a pp-construction in Γ'.

Let C and C' be the domains of Γ and Γ', respectively. Take the τ-expressions that define the valued relations of Δ over Γ, and interpret them over Γ' instead of Γ; let Δ' be the resulting valued structure. Note that Δ' is a d-dimensional pp-power of Γ'. For a map $f \colon \Gamma \to \Gamma'$, let $\tilde{f} \colon \Delta \to \Delta'$ be given by $(x_1, \dots, x_d) \mapsto (f(x_1), \dots, f(x_d))$. Then for all $S \in \mathcal{B}((C')^C)$ we define

$$\tilde{\omega}_1(\{\tilde{f} \mid f \in S\}) := \omega_1(S)$$

and
$$\tilde{\omega}_1(\tilde{S}) := \tilde{\omega}_1(\tilde{S} \cap \{\tilde{f} \mid f \in (C')^C\})$$
for all $\tilde{S} \in B\big(((C')^d)^{C^d}\big)$. We argue that $\tilde{\omega}_1$ is a fractional homomorphism from Δ to Δ'. To see this, let R be a valued relation of Δ and R' be the corresponding valued relation of Δ'. If we view R as an element of $\langle\Gamma\rangle$, then Proposition 2.13 applied to ω_1 implies precisely that $\tilde{\omega}_1$ is a fractional homomorphism from Δ to Δ'. Analogously we obtain from ω_2 a fractional homomorphism $\tilde{\omega}_2$ from Δ' to Δ. Therefore, Δ is fractionally homomorphically equivalent to Δ', which is a pp-power of Γ'. In other words, Δ has a pp-construction in Γ'. ☐

Combining Corollary 2.15, Lemma 2.16 together with the NP-hardness of $\mathrm{CSP}(K_3)$ and of $\mathrm{CSP}(\{0,1\}; \mathrm{OIT})$ yields the following corollary.

Corollary 2.17. *Let Γ be a valued structure with a finite signature and oligomorphic automorphism group such that K_3 or $(\{0,1\}; \mathrm{OIT})$ has a pp-construction in Γ. Then Γ pp-constructs all relational structures on a finite domain and $\mathrm{VCSP}(\Gamma)$ is NP-hard.*

Proof. K_3 and $(\{0,1\}; \mathrm{OIT})$ are known to pp-construct all relational structures on a finite domain, see, e.g., [9, Corollary 6.4.4]. Therefore Γ pp-constructs all finite relational structures by Lemma 2.16. Recall that $\mathrm{CSP}(K_3)$ and $\mathrm{CSP}(\{0,1\}; \mathrm{OIT})$ are NP-complete (Example 1.9 and 1.19). Therefore, $\mathrm{VCSP}(\Gamma)$ is NP-hard by Corollary 2.15. ☐

Example 2.18. *Recall the valued structure $\Gamma_{max} = (\{0,1\}; (<_2)^1_0)$ from Example 1.2. We give a pp-construction of $(\{0,1\}; \mathrm{OIT})$ in Γ_{max}, which gives an alternative hardness proof of $\mathrm{VCSP}(\Gamma_{max})$ by Corollary 2.17.*

Let $R(x,y,z) := \mathrm{Opt}((<_2)^1_0(x,y) + (<_2)^1_0(y,z))$ be a valued relation over Γ_{max} and observe that $(x,y,z) \in R$ if and only if

- *$x <_2 y$ and $y \geq_2 z$, or*
- *$x \geq_2 y$ and $y <_2 z$.*

Let $\Delta = (\{0,1\}^2, \mathrm{OIT}^\Delta)$ be a pp-power of Γ_{max} where

$$\begin{aligned}\mathrm{OIT}^\Delta((u,v),(u',v'),(u'',v'')) := \mathrm{Opt}\big(&(<_2)^1_0(v,v') + (<_2)^1_0(v',v'') + (<_2)^1_0(v'',v)\\ &R(u,v,v') + R(u',v',v'') + R(u'',v'',v)\big). \qquad (2.8)\end{aligned}$$

A visualisation of the expression in (2.8) is in Figure 2.1. Note that Γ_{max} can be visualised as a single arrow from 0 to 1. The idea behind the definition of OIT^Δ is that the shape in the figure has to be folded into a single arrow to be satisfiable in Γ_{max} with removing as little arrows as possible. The optimal ways to do so then encode the tuples that lie in OIT^Δ.

Note that $(<_2)^1_0(v,v') + (<_2)^1_0(v',v'') + (<_2)^1_0(v'',v)$ is always at least 2 and therefore the optimal value is when it is exactly 2 and precisely one of (v,v'), (v',v'') and (v'',v) lies in $<_2$. Therefore, $((u,v),(u',v'),(u'',v'')) \in \mathrm{OIT}^\Delta$ if and only if exactly one of the three pairs does not lie in $<_2$.

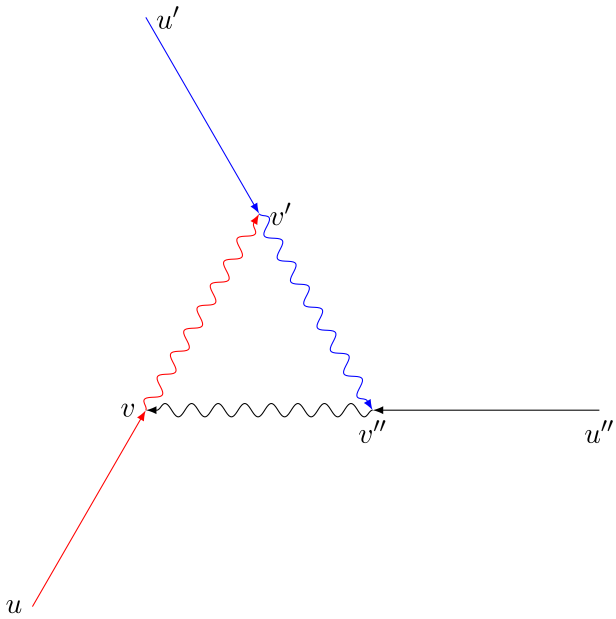

Figure 2.1: Visualisation of the definition of OIT^{Δ} from Example 2.18.

We show that Δ is homomorphically equivalent to $(\{0,1\};\mathrm{OIT})$; note that they are both relational structures and thus we may talk about homomorphisms. Let $f:\{0,1\}^2 \to \{0,1\}$ be such that $f(x,y)=0$ if $x <_2 y$ and $f(x,y)=1$ otherwise. Let $g:\{0,1\} \to \{0,1\}^2$ be such that $g(0)=(0,1)$ and $g(1)=(0,0)$. Then f is a homomorphism from Δ to $(\{0,1\};\mathrm{OIT})$ and g is a homomorphism from $(\{0,1\};\mathrm{OIT})$ to Δ. It follows that Γ_{max} pp-constructs $(\{0,1\};\mathrm{OIT})$.

Using Corollary 2.17, we can give a stronger formulation of Conjecture 0.3.

Conjecture 2.19. *Let $\mathfrak{A}$ be a relational structure with a countable domain. Suppose that $\mathfrak{A}$ is a reduct of a finitely bounded homogeneous relational structure $\mathfrak{B}$. If $\mathfrak{A}$ does not pp-construct K_3, then $\mathrm{CSP}(\mathfrak{A})$ is in P (and otherwise $\mathrm{CSP}(\mathfrak{A})$ is NP-complete by Corollary 2.17 and Theorem 1.14).*

2.3.1 Pp-constructing relational structures

In this section we study the properties of fractional homomorphisms between relational structures and pp-constructions of relational structures. This is relevant in typical NP-hardness proofs for VCSPs, as they are mostly based on Corollary 2.17, that is, pp-constructing the relational structure K_3 or $(\{0,1\};\mathrm{OIT})$.

Lemma 2.20. *Let Γ and Δ be valued τ-structures on countable domains such that $\mathrm{Aut}(\Gamma)$ is oligomorphic. If there exists a fractional homomorphism ω from Δ to Γ, then there also exists a homomorphism f from $\mathrm{Feas}(\Delta)$ to $\mathrm{Feas}(\Gamma)$. In particular, if Γ and Δ*

are crisp, then there is a fractional homomorphism from Δ *to* Γ *if and only if there is a homomorphism.*

Proof. Suppose that there exists a fractional homomorphism ω from Δ to Γ. Since D is countable and $\mathrm{Aut}(\Gamma) \subseteq \mathrm{Aut}(\mathrm{Feas}(\Gamma))$ is oligomorphic, by Lemma 2.5 it suffices to show that every finite substructure $\mathfrak{F}$ of $\mathrm{Feas}(\Delta)$ has a homomorphism to $\mathrm{Feas}(\Gamma)$. Let $s_1, \ldots, s_n$ be the elements of $\mathfrak{F}$ and $s := (s_1, \ldots, s_n)$. By the countable additivity of probability distributions, there exists $t \in C^n$ such that $\omega(\mathscr{S}_{s,t}) > 0$. Let $f : F \to C$ be the map such that $f(s) = t$. Suppose that there exists $R \in \tau$ of arity k and $s' \in F^k$ such that $R^\Gamma(f(s')) = \infty$. Since $\omega(\mathscr{S}_{s,t}) > 0$, we have $\omega(\mathscr{S}_{s',f(s')}) > 0$, and thus $E_\omega[g \mapsto R^\Gamma(g(s'))] = \infty$ by (2.1). Then $R^\Delta(s') = \infty$, because ω is a fractional polymorphism. Hence, for every $R \in \tau$ of arity k and $s' \in F^k$ we have $R^\Gamma(f(s')) < \infty$ whenever $R^\Delta(s') < \infty$. Therefore, f is a homomorphism from $\mathfrak{F}$ to $\mathrm{Feas}(\Gamma)$.

The final statement follows from the fact that every homomorphism is a fractional homomorphism and that $\mathfrak{A} = \mathrm{Feas}(\mathfrak{A})$ for every crisp structure $\mathfrak{A}$. □

An analogue of the last statement of Lemma 2.20 for general valued structures is false: as shown in Example 2.6, there are valued structures Γ and Δ such that there is a fractional homomorphism from Δ to Γ, but no fractional homomorphism of the form $f \in C^D$.

Remark 2.21. *Let* $\mathfrak{A}$ *and* $\mathfrak{B}$ *be relational structures with oligomorphic automorphism groups, and suppose that* $\mathfrak{A}$ *pp-constructs* $\mathfrak{B}$*, that is, there is a pp-power* Γ *of* $\mathfrak{A}$ *which is fractionally homomorphically equivalent to* $\mathfrak{B}$*. It follows from Remark 1.24 that* Γ *is essentially crisp and* $\mathrm{Feas}(\Gamma)$ *is a pp-power of* $\mathfrak{A}$ *with all relations primitively positively definable in* $\mathfrak{A}$ *when viewed over the domain* A*. By Lemma 2.20,* $\mathrm{Feas}(\Gamma)$ *is homomorphically equivalent to* $\mathfrak{B}$*. Hence, our definition of pp-constructability between two relational structures with oligomorphic automorphism groups coincides with the definition for relational structures from [6].*

The following proposition relates pp-constructability in a valued structure Γ with pp-constructability in the relational structure $(C; \langle\Gamma\rangle_0^\infty)$.

Proposition 2.22. *Let* Γ *be a valued structure and let* $\mathfrak{B}$ *be a relational* τ*-structure on countable domains* C*,* B*, respectively. Then* Γ *pp-constructs* $\mathfrak{B}$ *if and only if* $(C; \langle\Gamma\rangle_0^\infty)$ *pp-constructs* $\mathfrak{B}$*.*

Proof. Clearly, whenever $(C; \langle\Gamma\rangle_0^\infty)$ pp-constructs $\mathfrak{B}$, then Γ pp-constructs $\mathfrak{B}$. Suppose that Γ pp-constructs $\mathfrak{B}$. Then there exists $d \in \mathbb{N}$ and a pp-power Δ on the domain $D = C^d$ of Γ which is fractionally homomorphically equivalent to $\mathfrak{B}$. We claim that $\mathrm{Feas}(\Delta)$ is fractionally homomorphically equivalent to $\mathfrak{B}$ as well, witnessed by the same fractional homomorphisms.

Let ω_1 be a fractional homomorphism from Δ to $\mathfrak{B}$ and ω_2 be a fractional homomorphism from $\mathfrak{B}$ to Δ. Let $R \in \tau$ be of arity k and $s \in D^k$. By the definition of a fractional homomorphism,

$$E_{\omega_1}[f \mapsto R^{\mathfrak{B}}(f(s))] \leq R^\Delta(s).$$

We claim that

$$E_{\omega_1}[f \mapsto R^{\mathfrak{B}}(f(s))] \leq \mathrm{Feas}(R^{\Delta})(s). \tag{2.9}$$

This is clear if $\mathrm{Feas}(R^{\Delta})(s) = \infty$. Otherwise, $\mathrm{Feas}(R^{\Delta})(s) = 0$, and therefore $R^{\Delta}(s)$ is finite. Hence,

$$E_{\omega_1}[f \mapsto R^{\mathfrak{B}}(f(s))] = \sum_{t \in B^k} \omega_1(\mathscr{S}_{s,t}) R^{\mathfrak{B}}(t)$$

is finite. Since $R^{\mathfrak{B}}$ attains only values 0 and ∞, it follows that $E_{\omega_1}[f \mapsto R^{\mathfrak{B}}(f(t))] = 0$ and therefore (2.9) holds. Since R and s were arbitrary, it follows that ω_1 is a fractional homomorphism from $\mathrm{Feas}(\Delta)$ to $\mathfrak{B}$. The proof that ω_2 is a fractional homomorphism from $\mathfrak{B}$ to $\mathrm{Feas}(\Delta)$ is similar.

Note that $\mathrm{Feas}(\Delta)$ is a pp-power of $(C; \langle\Gamma\rangle_0^{\infty})$: every relation R of Δ of arity k lies in $\langle\Gamma\rangle$ when viewed as a relation of arity dk and therefore $\mathrm{Feas}(R)$ lies in $\langle\Gamma\rangle_0^{\infty}$ in this sense. Hence, $(C; \langle\Gamma\rangle_0^{\infty})$ pp-constructs $\mathfrak{B}$ as we wanted to prove. □

2.4 Open questions

Note that in all examples that we considered so far (in fact, all concrete examples that appear in the thesis), it was sufficient to work with fractional homomorphisms ω that are *finitary*, i.e., there are finitely many operations $f_1, \ldots, f_n \in \mathscr{O}_C$ such that $\sum_{i=1}^{n} \omega(f_i) = 1$. It is therefore possible that all fractional homomorphisms relevant to the complexity classification of VCSPs of valued structures Γ on a countable domain such that $\mathrm{Aut}(\Gamma)$ contains an automorphism group of a finitely bounded homogeneous structure are finitary. This motivates the following question.

Question 2.23. *Does our notion of pp-constructability change if we restrict ourselves to finitary fractional homomorphisms ω?*

Using pp-constructions, we can give a more concrete version of Conjecture 0.3: if $\mathfrak{A}$ is a reduct of a finitely bounded homogeneous structure on a countable domain, then $\mathrm{CSP}(\mathfrak{A})$ is NP-complete if $\mathfrak{A}$ pp-constructs K_3 and in P otherwise. We formulate a generalization of this formulation of the conjecture to valued structures.

Conjecture 2.24. *Let Γ be a valued structure with finite signature and a countable domain such that* $\mathrm{Aut}(\Gamma)$ *contains* $\mathrm{Aut}(\mathfrak{B})$ *for some finitely bounded homogeneous structure $\mathfrak{B}$. If K_3 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 1.14 and Corollary 2.17).*

Note that if the conjecture is true, then the complexity of $\mathrm{VCSP}(\Gamma)$ is completely determined by the underlying crisp structure of Γ by Proposition 2.22; this is true for finite-domain valued structures, see Corollary 3.40. A more concrete version of the conjecture, specialized to valued structures arising from resilience problems, can be found in Section 5.7.

One might hope to prove Conjecture 2.24 under the assumption of Conjecture 2.19. Recall that also the finite-domain VCSP classification was first proven conditionally on the finite-domain tractability conjecture [57,58], which was only confirmed later [32,83].

Finally, we remark that unlike in the CSP setting, proofs for several statements about pp-constructions exploited the oligomorphicity assumption on the automorphism group of some of the valued structures in question, for example, the proofs of Proposition 2.13, Corollary 2.15 and Lemma 2.20. In some cases, this assumption can be replaced by a less elegant assumption that all valued relations attain finitely many values, which is always the case in the CSP setting. It would be of interest to investigate further minimal assumptions for some of the results in this chapter.

Question 2.25. *Is Corollary 2.15 or Lemma 2.20 true without the oligomorphicity assumption (possibly replaced by a weaker assumption)?*

Chapter 3

Fractional polymorphisms

In this chapter, we address one of the crucial concepts of the theory of VCSPs, namely, *fractional polymorphisms*. These generalize polymorphisms, which proved extremely useful in understanding the theory of CSPs. Similarly to polymorphisms, even though in a more intricate manner, fractional polymorphisms can provide polynomial-time algorithms for VCSPs. We start this chapter by introducing polymorphisms and then proceed to generalize them to fractional polymorphisms. After discussing some of their important properties and special cases, we focus on the tractability proofs for VCSPs that can be obtained from fractional polymorphisms under various conditions on them and the VCSP templates in question. We finish the chapter with open questions related to this topic. Most of the original results in this chapter were published in [30] or announced in the preprint [12].

3.1 Polymorphisms and endomorphisms

Let A be a set and $R \subseteq A^k$. An operation $f : A^\ell \to A$ on the set A *preserves* R, if for every $t^1, \dots, t^\ell \in R$, $f(t^1, \dots, t^\ell) \in R$, where f is applied componentwise. 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 $\mathrm{Pol}(\mathfrak{A})$ and is closed under arbitrary composition of operations. We write $\mathrm{Pol}^{(\ell)}(\mathfrak{A})$ for the set $\mathrm{Pol}(\mathfrak{A}) \cap \mathscr{O}_A^{(\ell)}$, $\ell \in \mathbb{N}$. Unary polymorphisms are called *endomorphisms* and $\mathrm{Pol}^{(1)}(\mathfrak{A})$ is also denoted by $\mathrm{End}(\mathfrak{A})$.

Example 3.1. *Let $\ell \in \mathbb{N}$ and $i \in [\ell]$. Let $\pi_i^\ell \in \mathscr{O}_A^{(\ell)}$ be the i-th projection of arity ℓ, which is given by $\pi_i^\ell(x_1, \dots, x_\ell) = x_i$ (the domain is always clear from the context). Note that for every relational structure $\mathfrak{A}$, $\pi_i^\ell \in \mathrm{Pol}(\mathfrak{A})$ for all i and ℓ.*

Recall that primitively positively definable relations can be added to a relational structure without changing the complexity of its CSP. Polymorphisms play an important role in the algebraic approach to CSPs, since they capture computational complexity, as follows from the following theorem from [24].

Theorem 3.2 ([24, Theorem 5.1]). *Let $\mathfrak{A}$ be a relational structure with an oligomorphic automorphism group and let R be a relation on A. Then R is primitively positively definable in $\mathfrak{A}$ if and only if it preserved by* $\mathrm{Pol}(\mathfrak{A})$.

The theorem above (and its less general version for finite-domain structures) has been crucial in classifying the complexity CSPs, and in particular in the proof of the complexity dichotomy for finite-domain CSPs. To rephrase the finite-domain dichotomy theorem, we need a notion of a cyclic polymorphism.

Definition 3.3. *An operation $f\colon C^\ell \to C$ 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 $\mathrm{Cyc}_C^{(\ell)} \subseteq \mathscr{O}_C^{(\ell)}$ be the set of all operations on C of arity ℓ that are cyclic.

Using polymorphisms, we can rephrase Theorem 0.2 with algebraic conditions characterizing the respective complexities.

Theorem 3.4. *Let $\mathfrak{A}$ be a relational structure with a finite domain. Then exactly one of the following applies.*

1. *$\mathfrak{A}$ has a cyclic polymorphism. In this case,* $\mathrm{CSP}(\mathfrak{A})$ *is in P.*
2. *$\mathfrak{A}$ pp-constructs K_3. In this case,* $\mathrm{CSP}(\mathfrak{A})$ *is NP-complete.*

We finish this section with a definition and some useful properties of an important model-theoretic notion based on endomorphisms and automorphisms.

Definition 3.5. *A structure $\mathfrak{B}$ with an oligomorphic automorphism group is a* model-complete core *if $\overline{\mathrm{Aut}(\mathfrak{B})} = \mathrm{End}(\mathfrak{B})$.*

The definition of a model-complete core that we use is not the standard one, see, e.g., [9, Section 2.6.1-2.6.2]; several equivalent characterizations for structures with an oligomorphic automorphism group can be found in [9, Theorem 4.5.1]. Note that the definition implies that every endomorphism of $\mathfrak{B}$ is an embedding. As follows from the following proposition, one might often restrict their attention to model-complete cores when studying CSPs.

Proposition 3.6 ([8, Theorem 16], [9, Proposition 4.7.7]). *Let $\mathfrak{B}$ be a relational structure with an oligomorphic automorphism group. Then there exists a model-complete core $\mathfrak{C}$ homomorphically equivalent to $\mathfrak{B}$, which is unique up to isomomorphism. If $\mathfrak{B}$ is homogeneous, then $\mathfrak{C}$ is homogeneous as well.*

We refer to the structure $\mathfrak{C}$ from Proposition 3.6 structure as *the model-complete core of* $\mathfrak{B}$. We finish this section with a proposition, which gives a sufficient condition for the automorphism group of a valued structure and the automorphism group of its underlying crisp structure to be equal (recall from Example 1.31 that this might not be the case).

Proposition 3.7. *Let Γ be a valued structure with a domain C and an oligomorphic automorphism group. If $(C; \langle\Gamma\rangle_0^\infty)$ is a model-complete core, then* $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(C; \langle\Gamma\rangle_0^\infty)$.

Proof. Let $\Gamma' := (C; \langle\Gamma\rangle_0^\infty)$. By Remark 1.26, $\mathrm{Aut}(\Gamma) \subseteq \mathrm{Aut}(\Gamma')$. We prove the reverse inclusion. Suppose for contradiction that $\mathrm{Aut}(\Gamma) \subsetneq \mathrm{Aut}(\Gamma')$. Then there is a valued relation R of Γ of arity k that is not preserved by $\mathrm{Aut}(\Gamma')$. Let $O \subseteq C^k$ be an orbit of $\mathrm{Aut}(\Gamma')$ such that there exist $\alpha \in \mathrm{Aut}(\Gamma')$ and $s \in O$ such that $R(s) < R(\alpha(s))$ and choose $s \in O$ with the property that $R(s)$ is minimal. Note that O is preserved by $\mathrm{Aut}(\Gamma')$ and therefore by the assumption also by $\mathrm{End}(\Gamma')$. Since O consists of a single orbit of k-tuples, this implies that O is preserved by $\mathrm{Pol}(\Gamma')$, because every application of a polymorphism on tuples from O can be viewed as a composition of an endomorphism with automorphisms. By Theorem 3.2, O is primitively positively definable in Γ', and hence $O_0^\infty \in \langle\mathfrak{A}\rangle$.

Let S be a relation defined by

$$S(x_1, \dots, x_k) := \mathrm{Opt}\left(R(x_1, \dots, x_k) + O_0^\infty(x_1, \dots, x_k)\right);$$

then $S \in \langle\Gamma\rangle_0^\infty$. By the choice of s, $s \in S$, because $\mathrm{Aut}(\Gamma')$ is transitive on O. Since $R(\alpha(s)) > R(s)$, $\alpha(s) \notin S$. Therefore, α does not preserve S, a contradiction. It follows that $\mathrm{Aut}(\Gamma') = \mathrm{Aut}(\Gamma)$. □

3.2 Fractional operations

We now introduce *fractional operations* and *fractional polymorphisms* of valued structures; they are an important tool for formulating tractability results and complexity classifications for 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. [79]). Our approach is different from the one of Schneider and Viola [77, 80] and Viola and Živný [81] in that we work with arbitrary probability spaces instead of distributions with finite support or countable additivity property. As we will see in Section 3.7, 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.

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

$$\mathscr{S}_{s^1,\dots,s^\ell,t} := \{f \in \mathscr{O}_C^{(\ell)} \mid f(s^1, \dots, s^\ell) = t\} \tag{3.1}$$

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

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

Remark 3.8. *Let A be a countable set and $\mathscr{S} \subseteq \mathscr{O}_A$. Then $\mathscr{S}$ is equal to* $\mathrm{Pol}(\mathfrak{A})$ *for a relational structure $\mathfrak{A}$ on the domain A if and only if $\mathscr{S}$ contains π_i^n for every $n \in \mathbb{N}$ and $i \in \{1, \dots, n\}$, is closed under arbitrary composition of operations and is closed in the topology of pointwise convergence (note that the last condition is void if A is finite).*

We can now define the main object of study in this chapter, a fractional operation.

Definition 3.9 (fractional operation)**.** *Let $\ell \in \mathbb{N}$. A* fractional operation *on a set C of arity ℓ is a probability distribution*

$$\big(\mathscr{O}_C^{(\ell)}, \mathcal{B}(\mathscr{O}_C^{(\ell)}), \omega \colon \mathcal{B}(\mathscr{O}_C^{(\ell)}) \to [0,1]\big).$$

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

Definition 3.10. *A fractional operation $\omega \in \mathscr{F}_C^{(\ell)}$* improves *a k-ary valued relation $R \in \mathscr{R}_C^{(k)}$ if for all $t^1, \dots, t^\ell \in C^k$*

$$E_\omega[f \mapsto R(f(t^1, \dots, t^\ell))]$$

exists and

$$E_\omega[f \mapsto R(f(t^1, \dots, t^\ell))] \leq \frac{1}{\ell} \sum_{j=1}^{\ell} R(t^j). \tag{3.2}$$

Note that (3.2) has the interpretation that the expected value of $R(f(t^1, \dots, t^\ell))$ is at most the average of the values $R(t^1), \dots, R(t^\ell)$. Also note that if R is a crisp relation improved by a fractional operation ω and $\omega(f) > 0$ for $f \in \mathscr{O}^{(\ell)}$, then f preserves R as defined in Section 3.1. It follows from Lemma 2.8 that if $\mathrm{Aut}(C; R)$ is oligomorphic, then $E_\omega[f \mapsto R(f(t^1, \dots, t^\ell))]$ always exists and is greater than $-\infty$.

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

Definition 3.11 (fractional polymorphism)**.** *Let Γ be a valued structure. If ω improves every valued relation of 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 3.12. *Our notion of fractional polymorphism coincides with the previously used notions of fractional polymorphisms with finite support [77, 80] or the countable additivity property [81], since in this case the expected value on the left-hand side of* (3.2) *is equal to the weighted arithmetic mean.*

Remark 3.13. *Let Γ be a valued τ-structure with the domain C. A fractional polymorphism of arity ℓ of 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: Γ^ℓ is a valued τ-structure with the domain C^ℓ such that for every $R \in \tau$ of arity k we have*

$$R^{\Gamma^\ell}((t_1^1, \dots, t_\ell^1), \dots, (t_1^k, \dots, t_\ell^k)) := \frac{1}{\ell} \sum_{i=1}^{\ell} R^\Gamma(t_i^1, \dots, t_i^k).$$

We give a few examples of fractional polymorphisms.

Example 3.14. *The fractional operation Id_ℓ of arity ℓ such that $\mathrm{Id}_\ell(\pi_i^\ell) = \frac{1}{\ell}$ for every $i \in \{1, \dots, \ell\}$ is a fractional polymorphism of every valued structure.*

Example 3.15. *Recall the valued structures Γ and Δ and the fractional homomorphism ω from Δ to Γ from Example 2.6. Note that $\Delta = \Gamma^2$ and therefore ω is in fact a binary fractional polymorphism of Γ. Valued relations on the domain $\{0,1\}$ improved by ω are called* submodular. *The example also shows that there is no binary fractional polymorphism of Γ where a single operation has probability* 1.

Example 3.16. *Let Γ be a valued structure and $\alpha \in \mathrm{Aut}(\Gamma)$. The fractional operation $\omega \in \mathscr{F}_C^{(1)}$ defined by $\omega(\alpha) = 1$ is a fractional polymorphism of Γ.*

In fact, automorphisms need to preserve relations that they improve, as follows from the following lemma.

Lemma 3.17. *Let $R \in \mathscr{R}_C^{(k)}$ and let Γ be a valued structure. Suppose there exists $\alpha \in \mathrm{Aut}(\Gamma)$ which does not preserve R. Then $R \notin \mathrm{Imp}(\mathrm{fPol}(\Gamma)^{(1)})$.*

Proof. Since α does not preserve R, there exists $t \in C^k$ such that $R(t) \neq R(\alpha(t))$. If $R(\alpha(t)) > R(t)$, then let $\omega \in \mathscr{F}_C^{(1)}$ be the fractional operation defined by $\omega(\alpha) = 1$. Then ω improves every valued relation of Γ and does not improve R. If $R(\alpha(t)) < R(t)$, then the fractional polymorphism ω of Γ given by $\omega(\alpha^{-1}) = 1$ does not improve R. $\square$

3.3 Expressibility and fractional polymorphisms

Recall from Theorem 3.2 that a relation is primitively positively definable in a relational structure $\mathfrak{A}$ if and only if it is preserved by $\mathrm{Pol}(\mathfrak{A})$. In Lemma 3.18, we prove a generalization of the simpler of the two inclusions for valued structures. In Section 3.8 we discuss what is known about the other inclusion. Parts of the arguments in the proof of Lemma 3.18 can be found in the proof of [80, 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. The proof is similar to the proof of Proposition 2.13, where we also needed to assume that the automorphism groups are oligomorphic.

Lemma 3.18. *For every valued τ-structure Γ over a countable domain C we have*

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

Proof. Let $\omega \in \mathrm{fPol}(\Gamma)^{(\ell)}$. By definition, ω improves every valued relation R of Γ. It is clear that ω also preserves $(\emptyset)_0^\infty$. To see that ω preserves $(=)_0^\infty$, let $t^1, \dots, t^\ell \in C^2$. Note that either $t_1^i = t_2^i$ for every $i \in \{1, \dots, \ell\}$, in which case $f(t_1^1, \dots, t_1^\ell) = f(t_2^1, \dots, t_2^\ell)$ for every $f \in \mathscr{O}_C^{(\ell)}$, and hence

$$E_\omega[f \mapsto (=)_0^\infty(f(t^1, \dots, t^\ell))] = 0 = \frac{1}{\ell} \sum_{j=1}^{\ell} (=)_0^\infty(t^j),$$

or $t_1^i \neq t_2^i$ for some $i \in \{1, \dots, \ell\}$, in which case $\frac{1}{\ell} \sum_{j=1}^{\ell} (=)_0^\infty(t^j) = \infty$ and the inequality (3.2) is again satisfied.

The statement is also clear for valued relations obtained from valued relations in Γ by non-negative scaling and addition of constants, since these operations preserve the inequality in (3.2) by the linearity of expectation.

Let $\phi(x_1, \dots, x_k, y_1, \dots, y_n)$ be a τ-expression. We need to show that the fractional operation ω improves the k-ary valued relation R defined for every $t \in C^k$ by $R(t) = \inf_{s \in C^n} \phi^\Gamma(t, s)$. Since ϕ is a τ-expression, there are $R_i \in \tau$ such that

$$\phi(x_1, \dots, x_k, y_1, \dots, y_n) = \sum_{i=1}^{m} R_i(x_{p_1^i}, \dots, x_{p_{k_i}^i}, y_{q_1^i}, \dots, y_{q_{n_i}^i})$$

for some $k_i, n_i \in \mathbb{N}$, $p_1^i, \dots, p_{k_i}^i \in \{1, \dots, k\}$ and $q_1^i, \dots, q_{n_i}^i \in \{1, \dots, n\}$.

In this paragraph, if $v = (v_1, \dots, v_t) \in C^t$ and $i_1, \dots, i_s \in \{1, \dots, t\}$, we will write $v_{i_1, \dots, i_s}$ for the tuple $(v_{i_1}, \dots, v_{i_s})$ for short. Let $t^1, \dots, t^\ell \in C^k$. Let $\varepsilon > 0$ be a rational number. From the definition of an infimum, for every $j \in \{1, \dots, \ell\}$, there is $s^j \in C^n$ such that

$$R(t^j) \leq \phi(t^j, s^j) < R(t^j) + \varepsilon.$$

Moreover, for every $f \in \mathscr{O}_C^{(\ell)}$,

$$R(f(t^1, \dots, t^\ell)) \leq \phi(f(t^1, \dots, t^\ell), f(s^1, \dots, s^\ell)).$$

By linearity and monotonicity of expectation, we obtain

$$\begin{aligned}
E_\omega[f \mapsto R(f(t^1, \dots, t^\ell))] &\leq E_\omega[f \mapsto \phi(f(t^1, \dots, t^\ell), f(s^1, \dots, s^\ell))] \\
&= E_\omega[f \mapsto \sum_{i=1}^{m} R_i((f(t^1, \dots, t^\ell))_{p_1^i, \dots, p_{k_i}^i}, (f(s^1, \dots, s^\ell))_{q_1^i, \dots, q_{n_i}^i})] \\
&= \sum_{i=1}^{m} E_\omega[f \mapsto R_i((f(t^1, \dots, t^\ell))_{p_1^i, \dots, p_{k_i}^i}, (f(s^1, \dots, s^\ell))_{q_1^i, \dots, q_{n_i}^i})].
\end{aligned}$$

Since ω improves R_i for every $i \in \{1, \dots, m\}$, the last row of the inequality above is at most

$$\begin{aligned}\sum_{i=1}^{m} \frac{1}{\ell} \sum_{j=1}^{\ell} R_i(t^j_{p^i_1,\dots,p^i_{k_i}}, s^j_{q^i_1,\dots,q^i_{n_i}}) &= \frac{1}{\ell} \sum_{j=1}^{\ell} \sum_{i=1}^{m} R_i(t^j_{p^i_1,\dots,p^i_{k_i}}, s^j_{q^i_1,\dots,q^i_{n_i}}) \\ &= \frac{1}{\ell} \sum_{j=1}^{\ell} \phi(t^j, s^j) \\ &< \frac{1}{\ell} \sum_{j=1}^{\ell} R(t^j) + \varepsilon.\end{aligned}$$

Since ε was arbitrary, it follows that ω improves R.

Finally, we prove that Imp(fPol(Γ)) is closed under Feas and Opt. Let $R \in \tau$ be of arity k and define $S = \mathrm{Feas}(R)$ and $T = \mathrm{Opt}(R)$. We aim to show that $S, T \in$ Imp(fPol(Γ)). Let $s^1, \dots, s^\ell \in C^k$. If $S(s^i) = \infty$ for some $i \in \{1, \dots, \ell\}$, then $\frac{1}{\ell} \sum_{j=1}^{\ell} S(s^j) = \infty$ and hence ω satisfies (3.2) (with R replaced by S) for the tuples $s^1, \dots, s^\ell$. So suppose that $S(s^i) = 0$ for all $i \in \{1, \dots, \ell\}$, i.e., $R(s^i)$ is finite for all i. Since ω improves R it holds that

$$E_\omega[f \mapsto R(f(s^1, \dots, s^\ell))] \leq \frac{1}{\ell} \sum_{j=1}^{\ell} R(s^j) \tag{3.3}$$

and hence the expected value on the left-hand side is finite as well. By (2.1),

$$E_\omega[f \mapsto R(f(s^1, \dots, s^\ell))] = \sum_{t \in C^k} R(t)\omega(\mathscr{S}_{s^1,\dots,s^\ell,t}), \tag{3.4}$$

which implies that $R(t)$ is finite and $S(t) = 0$ unless $\omega(\mathscr{S}_{s^1,\dots,s^\ell,t}) = 0$. Consequently (again by (2.1)),

$$E_\omega[f \mapsto S(f(s^1, \dots, s^\ell))] = \sum_{t \in C^k} S(t)\omega(\mathscr{S}_{s^1,\dots,s^\ell,t}) = 0 = \frac{1}{\ell} \sum_{j=1}^{\ell} S(s^j).$$

It follows that ω improves S.

Moving to the valued relation T, we may again assume without loss of generality that $T(s^i) = 0$ for every $i \in \{1, \dots, \ell\}$ as we did for S. This means that $c := R(s^1) = \cdots = R(s^\ell) \leq R(b)$ for every $b \in C^k$. Therefore, the right-hand side in (3.3) is equal to c and by combining it with (3.4) we get

$$\sum_{t \in C^k} R(t)\omega(\mathscr{S}_{s^1,\dots,s^\ell,t}) \leq c.$$

Together with the assumption that $R(t) \geq c$ for all $t \in C^k$ and ω being a probability distribution we obtain that $R(t) = c$ and $T(t) = 0$ unless $\omega(\mathscr{S}_{s^1,\dots,s^\ell,t}) = 0$, and hence

$$E_\omega[f \mapsto T(f(s^1,\dots,s^\ell))] = \sum_{t \in C^k} T(t)\omega(\mathscr{S}_{s^1,\dots,s^\ell,t}) = 0 = \frac{1}{\ell}\sum_{j=1}^{\ell} T(s^j).$$

This concludes the proof that ω improves T. $\square$

The following example shows an application of Lemma 3.18.

Example 3.19. *Recall the valued structure* Γ_{max} *from Example 1.2. By definition,* $\langle \Gamma_{max} \rangle$ *contains* $\mathrm{Opt}((<_2)_0^1)$*, which is equal to* $<_2$*. Let* $\min_2$ *denote the minimum operation on* $\{0,1\}$*. Note that* $\min_2 \in \mathrm{fPol}(\{0,1\}; <_2)$*. However,*

$$(<_2)_0^1 \left(\min_2 \left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \end{pmatrix} \right) \right) = (<_2)_0^1(0,0) = 1,$$

while

$$\frac{1}{2}\big((<_2)_0^1(0,1) + (<_2)_0^1(0,0)\big) = \frac{1}{2}.$$

This shows that $\min_2$ *does not improve* $(<_2)_0^1$ *and hence* $(<_2)_0^1 \notin \langle(\{0,1\}; <_2)\rangle$ *by Lemma 3.18.*

Remark 3.20. *Recall that* $\langle \Gamma \rangle = \langle \mathrm{Feas}(\Gamma) \rangle$ *for essentially crisp structures* Γ*. Therefore, Lemma 3.18 implies that* $\mathrm{fPol}(\Gamma) = \mathrm{fPol}(\mathrm{Feas}(\Gamma))$*.*

Remarkably, for finite-domain valued structures, the converse of Lemma 3.18 is also true.

Theorem 3.21 ([47, Theorem 3.3])**.** *Let* Γ *be a valued structure with a finite domain. Then* $\langle \Gamma \rangle = \mathrm{Imp}(\mathrm{fPol}(\Gamma))$*.*

3.4 Crisp and essentially crisp structures

Let $\mathfrak{A}$ be a *relational* τ-structure. Then it is easy to see that $\mathrm{Pol}(\mathfrak{A}) \subseteq \mathrm{fPol}(\mathfrak{A})$ (using the convention that an operation on A can be viewed as a fractional operation). We describe $\mathrm{fPol}(\mathfrak{A})$ more concretely.

Proposition 3.22. *Let* $\mathfrak{A}$ *be a relational* τ*-structure. Then, for every* $\ell \in \mathbb{N}$, $\mathrm{fPol}^{(\ell)}(\mathfrak{A})$ *consists of precisely the fractional operations* ω *of arity* ℓ *such that* $\omega(\mathrm{Pol}^{(\ell)}(\mathfrak{A})) = 1$*.*

Proof. To see this, note that

$$\mathscr{O}_A^{(\ell)} \setminus \mathrm{Pol}^{(\ell)}(\mathfrak{A}) = \bigcup_{t^1,\dots,t^\ell \in R, s \in A^k \setminus R} \mathscr{S}_{t^1,\dots,t^\ell,s} \tag{3.5}$$

and therefore a measurable set. Hence, $\mathrm{Pol}^{(\ell)}(\mathfrak{A})$ is also measurable. Let $R \in \tau$ be of arity k and let $t^1, \dots, t^\ell \in A^k$. We want to prove

$$E_\omega[f \mapsto R(f(t^1, \dots, t^\ell))] \leq \frac{1}{\ell} \sum_{j=1}^{\ell} R(t^j). \tag{3.6}$$

Recall from (2.1) that

$$E_\omega[f \mapsto R(f(t^1, \dots, t^\ell))] = \sum_{s \in A^k} \omega(\mathscr{S}_{t^1,\dots,t^\ell,s}) R(s).$$

Therefore, if $\omega(\mathrm{Pol}^{(\ell)}(\mathfrak{A})) = 1$, either $R(t^j) = \infty$ for some j and (3.6) holds trivially, or $t^1, \dots, t^\ell \in R$ and $\omega(\mathscr{S}_{t^1,\dots,t^\ell,s}) = 0$ whenever $s \notin R$, in which case (3.6) holds because both sides of the inequality are equal to 0. On the other hand, if $\omega(\mathscr{O}_A^{(\ell)} \setminus \mathrm{Pol}^{(\ell)}(\mathfrak{A})) > 0$, then by (3.5) there exist $t^1, \dots, t^\ell \in R$ and $s \in A^k \setminus R$ such that $\omega(\mathscr{S}_{t^1,\dots,t^\ell,s}) > 0$. Then $E_\omega[f \mapsto R(f(t^1, \dots, t^\ell))] = \infty$, which contradicts (3.2) since $R(t^j) = 0$ for all j. □

We observe that essentially crisp structures are characterizated by having projections as fractional polymorphisms.

Lemma 3.23. *Let $\ell \in \mathbb{N}$, $\ell \geq 2$ and $i \in [\ell]$. Let Γ be a valued structure. Then* $\mathrm{fPol}(\Gamma)$ *contains π_i^ℓ if and only if Γ is essentially crisp.*

Proof. Suppose that Γ contains a valued relation R of arity k which takes two finite values a and b with $a < b$. For the sake of notation, assume $i = 1$. Let $s \in C^k$ be such that $R(s) = a$ and $t \in C^k$ be such that $R(t) = b$. Then

$$R(\pi_1^\ell(t, s, \dots, s)) = R(t) = b > \frac{1}{\ell}(b + (\ell - 1) \cdot a) = \frac{1}{\ell}(R(t) + (\ell - 1)R(s)),$$

and hence $\pi_1^\ell \notin \mathrm{fPol}(\mathfrak{A})$. The reverse implication follows from Remark 3.20, since $\mathrm{Feas}(\Gamma)$ is a relational structure and therefore $\pi_i^\ell \in \mathrm{Pol}(\mathrm{Feas}(\Gamma)) \subseteq \mathrm{fPol}(\mathrm{Feas}(\Gamma))$. □

3.5 Tractability from a constant fractional polymorphism

Lemma 3.24. *Let Γ be a valued structure with a finite relational signature τ such that there exists a unary constant operation $c \in \mathrm{fPol}(\Gamma)$. Then* $\mathrm{VCSP}(\Gamma)$ *is in P.*

Proof. Suppose that there exists $b \in C$ such that the unary operation c defined by $c(a) = b$ for all $a \in C$ is a fractional polymorphism of Γ. Then for every $R \in \tau$ of arity k and $t \in C^k$, we have $R(b, \dots, b) = R(c(t)) \leq R(t)$. Let $\phi(x_1, \dots, x_n) = \sum_i \phi_i$ be an instance of $\mathrm{VCSP}(\Gamma)$, where each ϕ_i is an atomic τ-expression. Then the cost of ϕ equals $\sum_i \phi_i(b, \dots, b)$ and hence $\mathrm{VCSP}(\Gamma)$ is in P. □

Example 3.25. *Recall the structure* Γ *from Example 1.31:* $\Gamma = (\mathbb{Q}; R)$ *where* R *is a valued relation on* $\mathbb{Q}$ *defined by*

$$R(x,y) = \begin{cases} 0 & x = y, \\ 1 & x < y, \\ 2 & x > y. \end{cases}$$

Let $c : \mathbb{Q} \to \mathbb{Q}$ *be defined by* $c(a) = 0$ *for every* $a \in \mathbb{Q}$*, then* c *is a unary constant operation. Clearly,* $c \in \mathrm{fPol}(\Gamma)$*, hence,* $\mathrm{VCSP}(\Gamma)$ *is in P.*

3.6 Fractional operations on finite domains

In this section we focus on fractional operations on finite domains and some of their specific properties. Most importantly, we derive a reformulation of the complexity dichotomy for finite-domain VCSPs, using the notion of pp-constructions (Theorem 3.36).

Definition 3.26. *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. An operation $f : C^\ell \to C$ is called *idempotent* if $f(x, \dots, x) = x$ for all $x \in C$; this notion plays an important role in [58] whose results on finite-domain VCSPs we build on. We will now focus on cyclic fractional operations and implications of having a cyclic fractional polymorphism for a VCSP of a valued structure on a finite domain.

Definition 3.27. *A fractional operation* ω *on a finite set* C *is called* cyclic *or* idempotent *if all operations in its support are cyclic or idempotent, respectively.*

The definition of a cyclic fractional operation above is in accordance with the definition from [58]. Cyclic operations of arity 2 are sometimes called *symmetric.*

Example 3.28. *Recall the binary minimum and maximum operations* $\min_2$ *and* $\max_2$ *on the set* $\{0,1\}$*. The fractional operation* ω *on* $\{0,1\}$ *defined by* $\omega(\min_2) = \omega(\max_2) = 1/2$ *is an idempotent and cyclic fractional operation.*

Lemma 3.29. *Let* Γ *and* Δ *be valued structures that are fractionally homomorphically equivalent.*

- *If Γ has a cyclic fractional polymorphism, then Δ has a cyclic fractional polymorphism of the same arity.*
- *Suppose that the domains of Γ and Δ are finite. 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.*

Proof. Let C be the domain of Γ and let D be the domain of Δ. Let χ_1 be a fractional homomorphism from Γ to Δ, and let χ_2 be a fractional homomorphism from Δ to Γ. Define χ_2' as the fractional homomorphism from Δ^ℓ to Γ^ℓ as follows. If $f\colon D \to C$, then f' denotes the map from D^ℓ to C^ℓ given by $(c_1, \ldots, c_\ell) \mapsto (f(c_1), \ldots, f(c_\ell))$. Define $\chi_2'(f') := \chi_2(f)$ and $\chi_2'(h) = 0$ for all other $h\colon D^\ell \to C^\ell$; since C and D are finite, this defines a fractional operation. It is straightforward to verify that χ_2' is a fractional homomorphism from Δ^ℓ to Γ^ℓ.

Suppose that ω is a fractional polymorphism of Γ of arity ℓ. Then $\omega' := \chi_1 \circ \omega \circ \chi_2'$ is a fractional homomorphism from Δ^ℓ to Δ (see Lemma 2.9), and hence a fractional polymorphism of Δ (see Remark 3.13). Note that if ω is cyclic, then ω' is cyclic; this shows that first statement of the lemma.

Next, suppose that C and D are finite and that there exists $\omega \in \mathrm{fPol}^{(\ell)}(\Gamma)$ such that $\mathrm{Supp}(\omega)$ contains a cyclic operation g of arity ℓ. Since the domain C of Γ is finite, there exists a function $f_1\colon C \to D$ such that $\chi_1(f_1) > 0$ and a function $f_2\colon D \to C$ such that $\chi_2(f_2) > 0$. Note that $f_1 \circ g \circ f_2'\colon D^\ell \to D$ is cyclic since g is cyclic, and that $\omega'(f_1 \circ g \circ f_2') > 0$ where ω' is defined as above. □

The following definition is taken from [58].

Definition 3.30 (core). *A valued structure Γ over a finite domain is called a* core *if all operations in $\mathrm{Supp}(\mathrm{fPol}(\Gamma))^{(1)}$ are injective.*

Note that the definition above and the definition of model-complete cores (Definition 3.5) specialize to the same concept for relational structures over finite domains. We have been unable to find an explicit reference for Proposition 3.32, but it should be considered to be known; we present a proof based on the following lemma from [59].

Lemma 3.31 ([59, Lemma 15]). *Let Γ be a valued structure with a finite domain. Let $f \in \mathrm{Supp}(\mathrm{fPol}(\Gamma))$ be unary and $\phi(x_1, \ldots, x_n)$ be an instance of $\mathrm{VCSP}(\Gamma)$. If $t \in C^n$ is a solution for ϕ, then $f(t)$ is a solution for ϕ as well.*

Proposition 3.32. *Let Γ be a valued structure with a finite domain. Then there exists a core valued structure Δ over a finite domain which is fractionally homomorphically equivalent to Γ.*

Proof. Let C be the domain of Γ. If Γ itself is a core then there is nothing to be shown, so we may assume that there exists a non-injective $f \in \mathrm{Supp}(\mathrm{fPol}^{(1)}(\Gamma))$. Since C is finite, we have that $D := f(C) \neq C$; let Δ be the substructure of Γ induced on $D \subseteq C$. It then follows from Lemma 3.31 that for every VCSP instance ϕ, the cost over Γ is the same as the cost over Δ. By Remark 2.11, Γ and Δ are fractionally homomorphically equivalent. After applying this construction finitely many times, we obtain a core valued structure that is fractionally homomorphically equivalent to Γ. □

The following Lemma 3.34 is a variation of Proposition 17 from [58], which is phrased only for valued structures Γ that are cores and for idempotent cyclic operations. We spell the proposition out for the convenience of the reader.

Proposition 3.33 ([58, Proposition 17]). *Let Γ be a core valued structure with a finite domain. Then Γ has an idempotent cyclic fractional polymorphism if and only if* $\mathrm{Supp}(\mathrm{fPol}(\Gamma))$ *contains an idempotent cyclic operation.*

Lemma 3.34. *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.*

Proof. The forward implication is trivial. For the reverse implication, let Δ be a core valued structure over a finite domain that is homomorphically equivalent to Γ, which exists by Proposition 3.32. By Lemma 3.29, $\mathrm{Supp}(\mathrm{fPol}(\Delta))$ contains a cyclic operation. Then $\mathrm{Supp}(\mathrm{fPol}(\Delta))$ contains even an idempotent cyclic operation: If $f \in \mathrm{Supp}(\mathrm{fPol}(\Delta))$ is cyclic, then the operation $f_0 \colon x \mapsto f(x, \ldots, x)$ is in $\mathrm{Supp}(\mathrm{fPol}(\Delta))$ as well. Since Δ is a finite core, f_0 is bijective and therefore f_0^{-1} (which is just a finite power of f_0) and the idempotent cyclic operation $f_0^{-1} \circ f$ lie in $\mathrm{Supp}(\mathrm{fPol}(\Delta))$. By Proposition 3.33, Δ has a cyclic fractional polymorphism and by Lemma 3.29, Γ also has one. □

We proceed to prove Theorem 3.36, which classifies the complexity of all finite-domain VCSPs. We rely on a crucial algorithmic result for VCSPs from [57], originally conditional on the later proven Theorem 3.4.

Theorem 3.35 ([57]). *Let Γ be a valued structure with a finite domain and a finite signature. If Γ has a cyclic fractional polymorphism, then* $\mathrm{VCSP}(\Gamma)$ *is in P.*

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 we explain how to derive it from results in [32, 57, 58, 82, 83].

Theorem 3.36. *Let Γ be a valued structure with a finite signature and a finite domain C. Then exactly one of the following applies:*

- *Γ pp-constructs K_3. In this case,* $\mathrm{VCSP}(\Gamma)$ *is NP-complete.*
- *Γ has a cyclic fractional polymorphism. In this case,* $\mathrm{VCSP}(\Gamma)$ *is in P.*

Proof. If K_3 has a pp-construction in Γ, then the NP-hardness of $\mathrm{VCSP}(\Gamma)$ follows from Corollary 2.17. Since every finite-domain VCSP is clearly in NP, $\mathrm{VCSP}(\Gamma)$ is NP-complete. Assume therefore that K_3 does not have a pp-construction in Γ.

Note that $\mathrm{Supp}(\mathrm{fPol}(\Gamma))$ contains the projections by Remark 3.14 and is closed under composition by Lemma 2.9 and Remark 3.13. By Remark 3.8, there exists a relational structure $\mathfrak{C}$ on the domain C such that $\mathrm{Pol}(\mathfrak{C}) = \mathrm{Supp}(\mathrm{fPol}(\Gamma))$. By Proposition 3.22, $\mathrm{fPol}(\Gamma) \subseteq \mathrm{fPol}(\mathfrak{C})$ and hence Theorem 3.21 implies that every relation of $\mathfrak{C}$ lies in $\langle \Gamma \rangle$. Since Γ does not pp-construct K_3, neither does $\mathfrak{C}$, and in particular, $\mathfrak{C}$ does not pp-construct K_3 in the relational setting (see Remark 2.21). By Theorem 3.4, $\mathrm{Pol}(\mathfrak{C})$ contains a cyclic operation. Since $\mathrm{Supp}(\mathrm{fPol}(\Gamma)) = \mathrm{Pol}(\mathfrak{C})$ contains a cyclic operation, by Lemma 3.34, Γ has a cyclic fractional polymorphism. Then $\mathrm{VCSP}(\Gamma)$ is in P by Theorem 3.35. □

Remark 3.37. *The problem of deciding for a given valued structure Γ with a finite domain and a finite signature whether Γ satisfies the tractability condition given in Theorem 3.36 can be solved in exponential time [56].*

Example 3.38. *Recall the valued structure Γ from Example 2.6. In Example 3.15 we observed that Γ has a binary fractional polymorphism ω. Note that ω cyclic (Example 3.28). By Theorem 3.36,* $\mathrm{VCSP}(\Gamma)$ *is in P.*

Remark 3.39. *In some articles on finite-domain VCSPs, such as [58] on whose results we build, a different definition is used instead of the concept of a fractional operation and fractional polymorphism. Let C be a finite set. A* weighting *on C of arity ℓ is $\rho \colon \mathscr{O}_C^{(\ell)} \to \mathbb{R}$ such that*

$$\sum_{f \in \mathscr{O}_C^{(\ell)}} \rho(f) = 0$$

and if $\rho(f) < 0$ then $f = \pi_i^\ell$ for some $i \in [\ell]$. A weighting ρ on C of arity ℓ is a weighted polymorphism *of a valued structure Γ if for every valued relation R of Γ of arity k and $t^1, \dots, t^\ell \in C^k$ we have*

$$\sum_{f \in \mathscr{O}_C^{(\ell)}} \rho(f) R(f(t^1, \dots, t^\ell)) \leq 0.$$

A weighted polymorphism is called cyclic *if $\rho(f) > 0$ implies that f is cyclic.*

We argue that Γ has a cyclic weighted polymorphism if and only if it has a cyclic fractional polymorphism; this justifies our use of results from [58]. The reverse direction is simple: if ω is a cyclic fractional polymorphism of Γ of arity ℓ, then define $\rho(f) := \omega(f)$ if f is not a projection and $\rho(\pi_i^\ell) = -1/\ell$ for every $i \in [\ell]$; it straightforward to verify that ρ is a weighted polymorphism.

Suppose that ρ is a weighted polymorphism of ω. Without loss of generality,

$$\rho(\pi_1^\ell) + \dots + \rho(\pi_\ell^\ell) = -1,$$

otherwise we can scale the weights. For every $f \in \mathscr{O}_C^{(\ell)}$, let $f_0 := f$, and

$$f_{i+1}(x_1, \dots, x_\ell) := f_i(x_2, \dots, x_\ell, x_1), i \in \{0, \dots, \ell - 2\}.$$

For every $f \in \mathscr{O}_C^{(\ell)}$, let ρ_i be the weighting defined by $\rho_i(f_i) = \rho(f)$, $i \in \{0, \dots, \ell - 1\}$. Note that ρ_i is a cyclic weighted polymorphism of Γ for every i. Finally, let ω be *a fractional operation defined by*

$$\omega(f) = \frac{1}{\ell}(\rho_0(f) + \cdots + \rho_{\ell-1}(f))$$

for every f that is not a projection and $\omega(\pi^\ell_i) = 0$ for every $i \in [\ell]$; note that this indeed defines a fractional operation. It is straightforward to verify that ω is a cyclic fractional polymorphism of Γ.

Some insight into the benefits of using weightings instead of fractional operations is given in [60, Remark on p. 242].

We also give a reformulation of Theorem 3.36 in terms of the underlying crisp structure of Γ. This is considered to be known, but it does not appear in this form in literature.

Corollary 3.40. *Let Γ be a valued structure with a finite signature and a finite domain. Let $\mathfrak{C}$ be the underlying crisp structure of Γ. Then exactly one of the following holds:*

- *$\mathfrak{C}$ pp-constructs K_3. In this case,* VCSP(Γ) *is NP-complete.*
- *$\mathfrak{C}$ has a cyclic polymorphism. In this case,* VCSP(Γ) *is in P.*

Proof. By Proposition 2.22, $\mathfrak{C}$ pp-constructs K_3 if and only if Γ does. Hence, if $\mathfrak{C}$ pp-constructs K_3, then VCSP(Γ) is NP-complete by Theorem 3.36. Otherwise, Γ does not pp-construct K_3 and again by Theorem 3.36, Γ has a cyclic fractional polymorphism and VCSP(Γ) is in P. In this case, $\mathfrak{C}$ has a cyclic fractional polymorphism by Lemma 3.18 and, by Lemma 3.34, Supp(fPol($\mathfrak{C}$)) contains a cyclic operation. By Proposition 3.22, Pol($\mathfrak{C}$) = Supp(fPol($\mathfrak{C}$)), because the domain C is finite. □

The importance of Corollary 3.40 is that it shows that for a valued structure Γ on a finite domain, the complexity of VCSP(Γ) is completely captured in the CSPs that can be encoded in VCSP(Γ) via expressibility (up to polynomial-time equivalence).

3.7 Tractability from canonical fractional polymorphisms

Building on the complexity dichotomy for finite-domain VCSPs (Theorem 3.36), we derive a sufficient condition for VCSP(Γ) to be in P, where Γ is a valued structure with a countable domain such that Aut(Γ) contains an automorphism group of a finitely bounded homogeneous relational structure. Recall that this assumption implies that Aut(Γ) is oligomorphic and that VCSP(Γ) is in NP (Theorem 1.14). To exploit the finite-domain VCSP dichotomy, the key ingredient is a polynomial-time reduction from VCSPs of such valued structures to VCSPs of finite-domain structures. This reduction is inspired by a similar reduction in the classical relational setting [21].

Definition 3.41 ($\Gamma^*_{\mathfrak{B},m}$). *Let Γ be a valued τ-structure with a countable domain such that* Aut(Γ) *contains the automorphism group of a finitely bounded homogeneous structure $\mathfrak{B}$. Let m be at least as large as the maximal arity of the relations of Γ. Let $\Gamma^*_{\mathfrak{B},m}$ be the following valued structure.*

- *The domain of* $\Gamma^*_{\mathfrak{B},m}$ *is the set of orbits of* m*-tuples of* $\mathrm{Aut}(\mathfrak{B})$.
- *For every* $R \in \tau$ *of arity* $k \leq m$ *the signature of* $\Gamma^*_{\mathfrak{B},m}$ *contains a unary relation symbol* R^**, which denotes in* $\Gamma^*_{\mathfrak{B},m}$ *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 $\mathrm{Aut}(\mathfrak{B})$ has finitely many orbits of k-tuples for every $k \in \mathbb{N}$ and therefore $\Gamma^*_{\mathfrak{B},m}$ has a finite domain. The following reduction is inspired by a known reduction for CSPs from [21].

Theorem 3.42. *Let* Γ *be a valued structure with a finite signature and a countable domain such that* $\mathrm{Aut}(\Gamma)$ *contains 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 of* 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* $\mathrm{VCSP}(\Gamma^*_{\mathfrak{B},m})$.

Proof. Let τ, τ^*, σ be the signatures of Γ, $\Gamma^*_{\mathfrak{B},m}$, and $\mathfrak{B}$, respectively. Let ϕ be an instance of $\mathrm{VCSP}(\Gamma)$ with threshold u and let V be the variables of ϕ. Create a variable $y(\bar{x})$ for every $\bar{x} = (x_1, \dots, x_m) \in V^m$. For every summand $R(x_1, \dots, x_k)$ of ϕ and we create a summand $R^*(y(x_1, \dots, x_k, \dots, x_k))$; this makes sense since $m \geq r$. For every $\bar{x}, \bar{x}' \in V^m$, $p \in \{1, \dots, m\}$, and $i, j \colon \{1, \dots, p\} \to \{1, \dots, m\}$, add the summand $C_{i,j}(y(\bar{x}), y(\bar{x}'))$ if $(x_{i(1)}, \dots, x_{i(p)}) = (x'_{j(1)}, \dots, x'_{j(p)})$; we will refer to these as *compatibility constraints*. Let ϕ^* be the resulting τ^*-expression. Clearly, ϕ^* can be computed from ϕ in polynomial time.

Suppose first that (ϕ, u) has a solution; it will be notationally convenient to view the solution as a function f from the variables of ϕ to the elements of Γ (rather than a tuple). We claim that the map f^* which maps $y(\bar{x})$ to the orbit of $f(\bar{x})$ in $\mathrm{Aut}(\mathfrak{B})$ is a solution for (ϕ^*, u). And indeed, each of the summands involving a symbol $C_{i,j}$ evaluates to 0, and $(\phi^*)^{\Gamma^*_{\mathfrak{B},m}}$ equals ϕ^Γ.

Now suppose that (ϕ^*, u) has a solution f^*. To construct a solution f to (ϕ, u), we first define an equivalence relation $\sim$ on V. For $x_1, x_2 \in V$, define $x_1 \sim x_2$ if a (equivalently: every) tuple t in $f^*(y(x_1, x_2, \dots, x_2))$ satisfies $t_1 = t_2$. To see that $\sim$ is reflexive, we consider i to be the identity map on $\{1,2\}$ and $j : \{1,2\} \to \{1\}$ to be the constant map, then the reflexivity follows from the constraint $C_{i,j}(y(x_1, \dots, x_1), y(x_1, \dots, x_1))$. To see that $\sim$ is symmetric, suppose that $x_1 \sim x_2$. Then the compatibility constraint

$$C_{i,j}(y(x_1, x_2, \dots, x_2), y(x_2, x_1, \dots, x_1)),$$

where i is the identity map on $\{1,2\}$ and $j : \{1,2\} \to \{1,2\}$ is defined by $j(a) = 3 - a$ implies $x_2 \sim x_1$. Finally we verify that $\sim$ is transitive. Suppose that $x_1 \sim x_2$ and

$x_2 \sim x_3$. In the following we use that $m \geq 3$. Let i be the identity map on $\{1, 2\}$, let $j\colon \{1, 2\} \to \{2, 3\}$ be given by $x \mapsto x+1$, and let $j'\colon \{1, 2\} \to \{1, 3\}$ be given by $j'(1) = 1$ and $j'(2) = 3$. Then ϕ^* contains the conjuncts

$$C_{i,i}(y(x_1, x_2, x_2, \dots, x_2), y(x_1, x_2, x_3, \dots, x_3)),$$
$$C_{i,j}(y(x_2, x_3, x_3, \dots, x_3), y(x_1, x_2, x_3, \dots, x_3)),$$
$$C_{i,j'}(y(x_1, x_3, x_3, \dots, x_3), y(x_1, x_2, x_3, \dots, x_3)).$$

Let t be a tuple from $f^*(y(x_1, x_2, x_3, \dots, x_3))$. Then it follows from the conjuncts with the relation symbols $C_{i,i}$ and $C_{i,j}$ that $t_1 = t_2$ and $t_2 = t_3$, and therefore $t_1 = t_3$. Thus we obtain from the conjunct with $C_{i,j'}$ that $x_1 \sim x_3$.

Claim 0. For all equivalence classes $[x_1]_\sim, \dots, [x_m]_\sim$, tuple $t \in f^*(y(x_1, \dots, x_m))$, $S \in \sigma$ of arity k, and a map $j\colon \{1, \dots, k\} \to \{1, \dots, m\}$, whether $\mathfrak{B} \models S(t_{j(1)}, \dots, t_{j(k)})$ does not depend on the choice of the representatives $x_1, \dots, x_m$. It suffices to show this statement if we choose another representative x'_i for $[x_i]_\sim$ for some $i \in \{1, \dots, m\}$, because the general case then follows by induction.

Suppose that for every $t \in f^*(y(x_1, \dots, x_m))$ we have $\mathfrak{B} \models S(t_{j(1)}, \dots, t_{j(k)})$ and we have to show that for every tuple $t' \in f^*(y(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_m))$ holds $\mathfrak{B} \models S(t'_{j(1)}, \dots, t'_{j(k)})$. If $i \notin \{j(1), \dots, j(k)\}$, then ϕ^* contains

$$C_{j,j}(y(x_1, \dots, x_m), y(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_m))$$

and hence $\mathfrak{B} \models S(t'_{j(1)}, \dots, t'_{j(k)})$. Suppose $i \in \{j(1), \dots, j(k)\}$; for the sake of notation we suppose that $i = j(1)$. By the definition of $\sim$, and since $x_{j(1)} \sim x'_{j(1)}$, every $t'' \in f^*(y(x_{j(1)}, x'_{j(1)}, \dots, x'_{j(1)}))$ satisfies $t''_1 = t''_2$. Let $\tilde{t}$ be a tuple from

$$f^*(y(x_{j(1)}, \dots, x_{j(k)}, x'_{j(1)}, \dots, x'_{j(1)})).$$

(Here we use that $m \geq r + 1$.)

- $\mathfrak{B} \models S(\tilde{t}_1, \dots, \tilde{t}_k)$, because we have a compatibility constraint in ϕ^* between $y(x_1, \dots, x_m)$ and $y(x_{j(1)}, \dots, x_{j(k)}, x'_{j(1)}, \dots, x'_{j(1)})$;
- $\tilde{t}_1 = \tilde{t}_{k+1}$ because of $x_{j(1)} \sim x'_{j(1)}$ and a compatibility constraint between the variables $y(x_{j(1)}, \dots, x_{j(k)}, x'_{j(1)}, \dots, x'_{j(1)})$ and $y(x_{j(1)}, x'_{j(1)}, \dots, x'_{j(1)})$ in ϕ^*;
- hence, $\mathfrak{B} \models S(\tilde{t}_{k+1}, \tilde{t}_2, \dots, \tilde{t}_k)$;
- $\mathfrak{B} \models S(t'_{j(1)}, t'_{j(2)}, \dots, t'_{j(k)})$ due to a compatibility constraint between the variables $y(x_{j(1)}, \dots, x_{j(k)}, x'_{j(1)}, \dots, x'_{j(1)})$ and $y(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_m)$ in ϕ^*: namely, consider the map $j'\colon \{1, \dots, k\} \to \{1, \dots, m\}$ that coincides with the identity map except that $j'(1) := k + 1$, then ϕ^* contains

$$C_{j,j'}\big(y(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_m), y(x_{j(1)}, \dots, x_{j(k)}, x'_{j(1)}, \dots, x'_{j(1)})\big).$$

This concludes the proof of Claim 0.

Now we can define a structure $\mathfrak{C}$ in the signature σ on the equivalence classes of $\sim$. If $S \in \sigma$ has arity k, $j_1, \dots, j_k \in \{1, \dots, m\}$, and $[x_1]_\sim, \dots, [x_m]_\sim$ are equivalence classes of $\sim$ such that the tuples t in $f^*(y(x_1, \dots, x_m))$ satisfy $S^{\mathfrak{B}}(t_{j_1}, \dots, t_{j_k})$ for some representatives $x_1, \dots x_m$ (equivalently, for all representatives, by Claim 0), then add $([x_{j_1}]_\sim, \dots, [x_{j_k}]_\sim)$ to $S^{\mathfrak{C}}$. No other tuples are contained in the relations of $\mathfrak{C}$.

Claim 1. If $[x_1]_\sim, \dots, [x_m]_\sim$ are equivalence classes of $\sim$, and $t \in f^*(y(x_1, \dots, x_m))$, then $[x_i]_\sim \mapsto t_i$, for $i \in \{1, \dots, m\}$, is an isomorphism between a substructure of $\mathfrak{C}$ and a substructure of $\mathfrak{B}$ for any choice of representatives $x_1, \dots, x_m$. First note that $[x_i]_\sim = [x_j]_\sim$ if and only if $t_i = t_j$, so the map is well-defined and bijective. Let $S \in \sigma$ be of arity k and $j \colon \{1, \dots, k\} \to \{1, \dots, m\}$. If $\mathfrak{B} \models S(t_{j(1)}, \dots, t_{j(k)})$, then $\mathfrak{C} \models S([x_{j(1)}]_\sim, \dots, [x_{j(k)}]_\sim)$ by the definition of $\mathfrak{C}$. Conversely, suppose that $\mathfrak{C} \models S([x_{j(1)}]_\sim, \dots [x_{j(k)}]_\sim)$. By Claim 0 and the definition of $\mathfrak{C}$, there is $t' \in f^*(y(x_1, \dots, x_m))$ such that $\mathfrak{B} \models S(t'_{j(1)}, \dots, t'_{j(k)})$. Since $f^*(y(x_1, \dots, x_m))$ is an orbit of $\mathrm{Aut}(\mathfrak{B})$, we have $\mathfrak{B} \models S(t_{j(1)}, \dots, t_{j(k)})$ as well.

Claim 2. $\mathfrak{C}$ embeds into $\mathfrak{B}$. It suffices to verify that $\mathfrak{C}$ satisfies each conjunct of the universal sentence ψ. Let $\psi'(x_1, \dots, x_q)$ be such a conjunct, and let $[c_1]_\sim, \dots, [c_q]_\sim$ be elements of $\mathfrak{C}$. Consider the orbit $f^*(y(c_1, \dots, c_q, \dots, c_q))$ of $\mathrm{Aut}(\Gamma)$; this makes sense since $m \geq v$. Let $t \in f^*(y(c_1, \dots, c_q, \dots, c_q))$. Since $t_1, \dots, t_q$ are elements of $\mathfrak{B}$, the tuple $(t_1, \dots, t_q)$ satisfies ψ'. Claim 1 then implies that $([c_1]_\sim, \dots, [c_q]_\sim)$ satisfies ψ'.

Let e be an embedding of $\mathfrak{C}$ to $\mathfrak{B}$. For every $x \in V$, define $f(x) = e([x]_\sim)$. Note that for every summand $R(x_1, \dots, x_k)$ in ϕ and $t \in f^*(y(x_1, \dots, x_k, \dots, x_k))$, we have

$$\begin{aligned} R^*(f^*(y(x_1, \dots, x_k, \dots, x_k))) &= R(t_1, \dots, t_k) \\ &= R(e([x_1]_\sim), \dots, e([x_k]_\sim)) \\ &= R(f(x_1), \dots, f(x_k)), \end{aligned}$$

where the middle equality follows from $t_i \mapsto e([x_i]_\sim)$ being a partial isomorphism of $\mathfrak{B}$ by Claim 1 and 2, which by the homogeneity of $\mathfrak{B}$ extends to an automorphism of $\mathfrak{B}$ and therefore also an automorphism of Γ. Since f^* is a solution to (ϕ^*, u), it follows from the construction of ϕ^* that f is a solution to (ϕ, u). □

To use the reduction from Theorem 3.42 together with Theorem 3.36 to prove tractability of VCSP(Γ), we need a condition on fractional polymorphisms of Γ that would translate to $\Gamma^*_{\mathfrak{B},m}$ having a cyclic fractional polymorphism. To this end, we introduce canonical and pseudo cyclic fractional operations.

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

Definition 3.43. *Let G be a permutation group on the set C. An operation $f\colon C^\ell \to C$ 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 $\mathrm{PC}_G^{(\ell)} \subseteq \mathscr{O}_C^{(\ell)}$ be the set of all operations on C of arity ℓ that are pseudo cyclic with respect to G.

Note that $\mathrm{PC}_G^{(\ell)} \in \mathcal{B}(\mathscr{O}_C^{(\ell)})$: its complement is a countable union of sets of the form $\mathscr{S}_{s^1,\dots,s^\ell,t}$ where for all $f \in \mathscr{O}_C^{(\ell)}$ the tuples $f(s^1, \dots, s^\ell)$ and $f(s^2, \dots, s^\ell, s^1)$ lie in different orbits with respect to G.

Definition 3.44. *Let G be a permutation group with domain C. An operation $f\colon C^\ell \to C$ for $\ell \geq 2$ is called* canonical with respect to G *if for all $k \in \mathbb{N}$ and $s^1, \dots, s^\ell \in C^k$ the orbit of the k-tuple $f(s^1, \dots, s^\ell)$ only depends on the orbits of $s^1, \dots, s^\ell$ with respect to G. Let $\mathrm{Can}_G^{(\ell)} \subseteq \mathscr{O}_C^{(\ell)}$ be the set of all operations on C of arity ℓ that are canonical with respect to G.*

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

Remark 3.45. *Note that if h is an operation on 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. Moreover, if h is pseudo cyclic with respect to G, then h^* is cyclic.*

Definition 3.46. *Let C be a set and G a permutation group on C. A fractional operation ω on C is called* pseudo cyclic with respect to G *if for every $A \in \mathcal{B}(\mathscr{O}_C^{(\ell)})$ we have $\omega(A) = \omega(A \cap \mathrm{PC}_G^{(\ell)})$.* Cyclicity*,* symmetry*, and* canonicity with respect to G *for fractional operations are defined analogously.*

Equivalently, a fractional operation is pseudo cyclic (with respect to G) if and only if it is pseudo cyclic almost everywhere (and analogously for cyclicity, symmetry and canonicity). Note that on finite domains, the definition of cyclicity coincides with Definition 3.27 from Section 3.6. We may omit the specification 'with respect to G' when G is clear from the context. Clearly, every cyclic fractional operation is pseudo cyclic. By definition, every operation and hence every fractional operation is canonical with respect to the trivial permutation group on the same set.

Before giving an example of a pseudo cyclic canonical operation (Example 3.48) we prove Lemma 3.47, which is a practical tool for verifying pseudo cyclicity. Let G be a permutation group that contains the automorphism group of a finitely bounded homogeneous structure $\mathfrak{B}$ of maximal arity at most m. A fractional operation ω over the domain C of Γ of arity ℓ which is canonical with respect to G induces a fractional operation ω^* on the orbits of m-tuples of G, given by

$$\omega^*(A) := \omega\big(\{f \in \mathrm{Can}_G^{(\ell)} \mid f^* \in A\}\big),$$

for every subset A of the set of operations of arity ℓ on the set of orbits of m-tuples of G (all such subsets are measurable). Note that $\{f \in \mathrm{Can}_G^{(\ell)} \mid f^* \in A\}$ is a measurable subset of $\mathscr{O}_C^{(\ell)}$. Also note that if ω is pseudo cyclic, then ω^* is cyclic.

Lemma 3.47. *Let G be the automorphism group of a homogeneous structure $\mathfrak{B}$ with a relational signature of maximal arity at most m. If $\omega \in \mathscr{F}_C^{(\ell)}$ is canonical with respect to G such that ω^* (defined on the orbits of m-tuples of G) is cyclic, then ω is pseudo cyclic with respect to G.*

Proof. It follows from Lemma 10.1.5 in [9] (also see the proof of Proposition 6.6 in [27]) that if $f \in \mathscr{O}^{(\ell)}$ is canonical with respect to G such that f^* (defined on the orbits of m-tuples) is cyclic, then f is pseudo cyclic; we use this fact below. Let C be the domain of Γ and let $s^1, \dots, s^\ell, t \in C^m$. It suffices to show that $\omega(\mathscr{S}_{s^1,\dots,s^\ell,t} \cap \mathrm{PC}_G^{(\ell)}) = \omega(\mathscr{S}_{s^1,\dots,s^\ell,t})$. Indeed,

$$\begin{aligned}
\omega(\mathscr{S}_{s^1,\dots,s^\ell,t}) &= \omega(\mathscr{S}_{s^1,\dots,s^\ell,t} \cap \mathrm{Can}_G^{(\ell)}) && \text{(canonicity of } \omega) \\
&= \omega^*\big(\{f^* \mid f \in \mathscr{S}_{s^1,\dots,s^\ell,t} \cap \mathrm{Can}_G^{(\ell)}\}\big) && \text{(definition of } \omega^*) \\
&= \omega^*\big(\{f^* \mid f \in \mathscr{S}_{s^1,\dots,s^\ell,t} \cap \mathrm{Can}_G^{(\ell)}\} \cap \mathrm{Cyc}_C^{(\ell)}\big) && \text{(by assumption)} \\
&= \omega^*\big(\{f^* \mid f \in \mathscr{S}_{s^1,\dots,s^\ell,t} \cap \mathrm{Can}_G^{(\ell)} \cap \mathrm{PC}_G^{(\ell)}\}\big) && \text{(fact above \& Rem. 3.45)} \\
&= \omega(\mathscr{S}_{s^1,\dots,s^\ell,t} \cap \mathrm{Can}_G^{(\ell)} \cap \mathrm{PC}_G^{(\ell)}) \\
&= \omega(\mathscr{S}_{s^1,\dots,s^\ell,t} \cap \mathrm{PC}_G^{(\ell)}).
\end{aligned}$$

□

We can now give an example of a pseudo cyclic canonical operation.

Example 3.48. *Recall the countable random graph $(\mathbb{V}; E)$ from Example 1.17. Consider a binary injective operation f on $\mathbb{V}$ with following property: if $(x, y) \in E$ or $(u, v) \in E$, then $(f(x, u), f(y, v)) \in E$; such an operation can be constructed inductively using the homogeneity and universality of $(\mathbb{V}; E)$. We claim that f is canonical and pseudo cyclic with respect to $\mathrm{Aut}(\mathbb{V}; E)$. To see that f is pseudo cyclic, we verify that f^* defined on orbits of pairs is cyclic. This is clear from the definition of f, since the orbits of pairs of $\mathrm{Aut}(\mathbb{V}; E)$ are $\{(x, x) | x \in \mathbb{V}\}$, E, and $\{(x, y) \in \mathbb{V}^2 \mid x \neq y, \neg E(x, y)\}$. Hence, f is pseudo cyclic by Lemma 3.47. To see that f is canonical, recall that the orbit of a tuple $t \in \mathbb{V}^k$ under the action of $\mathrm{Aut}(\mathbb{V}; E)$ only depends on the atomic formulas satisfied by the entries of t (because $(\mathbb{V}; E)$ is homogeneous). By the definition of f, the atomic formulas satisfied by the entries $f(s, t)$ where $s, t \in \mathbb{V}^k$ only depend on the atomic formulas satisfied by s and t. In other words, f is canonical. It follows that f is canonical and pseudo cyclic; this remains true when we view f as a fractional operation.*

Statements about the fractional polymorphisms of $\Gamma^*_{\mathfrak{B},m}$ lift back to statements about fractional polymorphisms of Γ via Lemma 3.50, which is a generalization of the following lemma for classical operations.

Lemma 3.49 ([21, Lemma 4.9]; also see [9, Lemma 10.5.12])**.** *Let $\mathfrak{B}$ be a finitely bounded homogeneous relational structure on a countable domain. Let $\mathfrak{A}$ be a relational structure with a finite signature such that* $\mathrm{Aut}(\mathfrak{B}) \subseteq \mathrm{Aut}(\mathfrak{A})$ *and let $m \in \mathbb{N}$ be as in Theorem 3.42 for $\Gamma = \mathfrak{A}$. Then for every $f \in \mathrm{Pol}(\mathfrak{A}^*_{\mathfrak{B},m})$ there exists a canonical operation $g \in \mathrm{Pol}(\mathfrak{A})$ such that $g^* = f$.*

Lemma 3.50. *Let Γ be a valued structure with a finite signature and a countable domain such that* $\mathrm{Aut}(\Gamma)$ *contains the automorphism group G of a finitely bounded homogeneous structure $\mathfrak{B}$ and let m be as in Theorem 3.42. Let $\chi \in \mathrm{fPol}(\Gamma^*_{\mathfrak{B},m})$. Then there exists $\omega \in \mathrm{fPol}(\Gamma)$ which is canonical with respect to G such that $\omega^* = \chi$.*

Proof. Let C be the domain of Γ, let D be the domain of $\Gamma^*_{\mathfrak{B},m}$, and let ℓ be the arity of χ. Suppose that $\chi(f) > 0$ for some operation f and note that $f \in \mathrm{Pol}(\mathrm{Feas}(\Gamma)^*_{\mathfrak{B},m})$ (see Proposition 3.22). Since $\mathrm{Aut}(\mathfrak{B}) \subseteq \mathrm{Aut}(\Gamma) \subseteq \mathrm{Aut}(\mathrm{Feas}(\Gamma))$, by Lemma 3.49, there exists a function $g\colon C^\ell \to C$ which is canonical with respect to G such that $g^* = f$. For every such f, choose g such that $g^* = f$ and define $\omega(g) := \chi(f)$ and $\omega(h) := 0$ for all other $h \in \mathscr{O}_C^{(\ell)}$. Since the domain of $\Gamma^*_{\mathfrak{B},m}$ is finite, this correctly defines a fractional operation ω of the same arity ℓ as χ. Then ω improves every valued relation R of Γ: if R has arity k and $t^1, \dots, t^\ell \in C^k$, then $E_\omega[h \mapsto R(h(t^1, \dots, t^\ell))]$ can be expressed as

$$\begin{aligned}
&\sum_{h \in \mathscr{O}_C^{(\ell)}} \omega(h) R(h(t^1, \dots, t^\ell)) \\
&= \sum_{f \in \mathscr{O}_D^{(\ell)}} \chi(f) R^*(f(t^1, \dots, t^\ell)_1, \dots, f(t^1, \dots, t^\ell)_k, \dots, f(t^1, \dots, t^\ell)_k) \\
&\leq \frac{1}{\ell} \sum_{j=1}^{\ell} R^*(t^j_1, \dots, t^j_k, \dots, t^j_k) \\
&= \frac{1}{\ell} \sum_{j=1}^{\ell} R(t^j_1, \dots, t^j_k).
\end{aligned}$$ □

We combine the results of the previous lemmas and give equivalent conditions for $\mathrm{fPol}(\Gamma)$, $\mathrm{fPol}(\Gamma^*_{\mathfrak{B},m})$ and $\mathrm{Supp}(\mathrm{fPol}(\Gamma^*_{\mathfrak{B},m}))$. To be able to discuss decidability of the conditions, we need the following notion. An operation $f\colon C^4 \to C$ is called *Siggers* if $f(a, r, e, a) = f(r, a, r, e)$ for all $a, r, e \in C$.

Proposition 3.51. *Let $\mathfrak{B}$ be a finitely bounded homogeneous structure with a countable domain and let Γ be a valued structure with a finite signature such that* $\mathrm{Aut}(\mathfrak{B}) \subseteq \mathrm{Aut}(\Gamma)$. *Let m be as in Theorem 3.42. 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. $\mathrm{fPol}(\Gamma^*_{\mathfrak{B},m})$ *contains a cyclic fractional operation;*
3. $\mathrm{Supp}(\mathrm{fPol}(\Gamma^*_{\mathfrak{B},m}))$ *contains a cyclic operation.*

4. $\mathrm{Supp}(\mathrm{fPol}(\Gamma^*_{\mathfrak{B},m}))$ *contains a Siggers operation.*

Proof. First, we prove that (1) implies (2). If ω is a fractional polymorphism of Γ, then ω^* is a fractional polymorphism of $\Gamma^*_{\mathfrak{B},m}$: the fractional operation ω^* improves R^* because ω improves R, and ω^* improves $C_{i,j}$ for all i, j because ω is canonical with respect to $\mathrm{Aut}(\mathfrak{B})$. Finally, if ω is pseudo cyclic with respect to $\mathrm{Aut}(\mathfrak{B})$, then ω^* is cyclic.

The implication from (2) to (1) is a consequence of Lemma 3.50 and Lemma 3.47. The equivalence of (2) and (3) follows from Lemma 3.34. The equivalence of (3) and (4) is proved in [9, Theorem 6.9.2]; the proof is based on [5, Theorem 4.1]. □

Remark 3.52. *Note that item (4) in the previous proposition can be decided algorithmically for a given valued structure $\Gamma^*_{\mathfrak{B},m}$ (which has a finite domain and finite signature) by testing all 4-ary operations on $\Gamma^*_{\mathfrak{B},m}$ (see [56] for a more efficient algorithm).*

Combining Proposition 3.51 with Theorem 3.36, we can now prove a new sufficient condition for tractability of VCSPs.

Theorem 3.53. *If the conditions from Proposition 3.51 hold, then* $\mathrm{VCSP}(\Gamma)$ *is in P.*

Proof. If $\Gamma^*_{\mathfrak{B},m}$ has a cyclic fractional polymorphism of arity $\ell \geq 2$, then the polynomial-time tractability of $\mathrm{VCSP}(\Gamma^*_{\mathfrak{B},m})$ follows from Theorem 3.36. By Theorem 3.42, there is a polynomial-time reduction from $\mathrm{VCSP}(\Gamma)$ to $\mathrm{VCSP}(\Gamma^*_{\mathfrak{B},m})$, which concludes the proof. □

Example 3.54. *Recall the structure Γ_{graph} from Example 1.17.* $\mathrm{Aut}(\Gamma_{graph}) = \mathrm{Aut}(\mathbb{V}; E)$ *and $(\mathbb{V}; E)$ is a homogeneous and finitely bounded relational structure. Also recall the binary operation f from Example 3.48, which is pseudo cyclic and canonical with respect to* $\mathrm{Aut}(\mathbb{V}; E)$*. We show that $f \in \mathrm{fPol}(\Gamma_{graph})$. Let $(x, y), (u, v) \in \mathbb{V}^2$. Then*

$$E\left(f\left(\begin{pmatrix}x\\y\end{pmatrix}, \begin{pmatrix}u\\v\end{pmatrix}\right)\right) \leq \frac{1}{2}(E(x,y) + E(u,v)),$$

because if $x \neq y$ or $u \neq v$, then $f(x, u) \neq f(y, v)$, and if $(x, y) \in E$ or $(u, v) \in E$, then $(f(x, u), f(y, v)) \in E$. It follows that Γ_{graph} has a canonical and pseudo cyclic fractional polymorphism and therefore, by Theorem 3.53, $\mathrm{VCSP}(\Gamma_{graph})$ *is in P.*

3.8 Open questions

In Section 3.3 we proved that if Γ is a valued structure with an oligomorphic automorphism group and $R \in \langle \Gamma \rangle$, then R is preserved by all fractional polymorphisms of Γ (Lemma 3.18). We do not know whether the converse is true. It is known to hold for the special cases of finite-domain valued structures [41, 47] and for relational structures with 0-∞ valued relations (CSP setting) having an oligomorphic automorphism group (Theorem 3.2).

Question 3.55. *Let* Γ *be a valued structure with an oligomorphic automorphism group. Is it true that* $R \in \langle \Gamma \rangle$ *if and only if* $R \in \operatorname{Imp}(\operatorname{fPol}(\Gamma))$*?*

A natural attempt to positively answer Question 3.55 would be to combine the proof strategy for finite-domain valued structures from [41,47] with the one for relational structures with oligomorphic automorphism group from [24]. However, since non-improving of R is not a closed condition, the compactness argument from [24] cannot be used to construct an operation from $\operatorname{fPol}(\Gamma)$ that does not improve R. A positive answer to Question 3.55 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 (see Chapter 5), is fully determined by the fractional polymorphisms of Γ.

Another aspect of the algebraic approach to VCSPs that attracts attention is the concept of cores (see Definition 3.30). By Proposition 3.32, for every finite-domain valued structure Γ, there exists a core valued structure Γ' fractionally homomorphically equivalent to Γ. If Γ is a relational structure, then it is known that Γ' is unique up to isomorphism, see, e.g., [9, Proposition 1.1.11]. For relational structures with an oligomorphic automorphism group, the concept of cores is generalized to model-complete cores, see Definition 3.5. Let us call a valued structure Γ on a countable domain a *core*, if for every $\omega \in \operatorname{fPol}^{(1)}(\Gamma)$ and every $S \in \mathcal{B}(\mathscr{O}_C^{(1)})$

$$\omega(S) = \omega(S \cap \overline{\operatorname{Aut}(\Gamma)});$$

one can show that this definition generalizes both Definition 3.5 and Definition 3.30. Since cores played an essential role in the classification of finite-domain VCSPs (for example, in Proposition 3.33), they appear to be a natural object of study also on infinite domains, motivating the following question.

Question 3.56. *Is every valued structure* Γ *with an oligomorphic automorphism group fractionally homomorphically equivalent to a core* Γ'*? Is* Γ' *unique up to isomorphism?*

Finally, similarly as for pp-constructability (Question 2.23), we may ask whether it is enough to consider finitary fractional polymorphisms to distinguish relations that are not expressible in a valued structures.

Question 3.57. *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*?*

Chapter 4

Temporal VCSPs

This chapter is devoted to the complexity classification of temporal VCSPs, that is, VCSPs of valued structures Γ whose automorphism group contains $\mathrm{Aut}(\mathbb{Q}; <)$. Recall that the structure $(\mathbb{Q}; <)$ is finitely bounded and homogeneous (Example 1.13). Therefore, by Theorem 1.14, $\mathrm{VCSP}(\Gamma)$ is in NP for every temporal valued structure Γ. Note that our classification result is incomparable with the result of [20] which shows that that submodular PLH functions form a maximally tractable class of PLH cost functions, because the class of temporal valued structure is a proper subclass of PLH and contains structures with tractable VCSPs that are not submodular.

We first introduce some notation in Section 4.1. In Section 4.2, we focus on a special case of valued structures preserved by $\mathrm{Sym}(\mathbb{Q})$, which serves both as a warm-up and a building block for the general case. In Sections 4.3–4.5 we treat temporal VCSPs in full generality and we finish with a discussion and outlook to the future in Section 4.6. The results presented in this chapter were announced in [12].

4.1 Order types

Let Γ be a valued structure such that $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$. Since the group $\mathrm{Aut}(\mathbb{Q}; <)$ has only finitely many orbits of k-tuples for every k, so does $\mathrm{Aut}(\Gamma)$. In particular, if $k = 2$ and $a, b \in \mathbb{Q}$, $a < b$, then the values $\phi(a, a), \phi(a, b)$ and $\phi(b, a)$ do not depend on the choice of a and b and if $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\mathbb{Q}; =)$, then $\phi(a, b) = \phi(b, a)$.

We will often use the following notation in the proofs in this chapter.

Definition 4.1 (E_t, N_t, O_t)**.** *If $t \in \mathbb{Q}^k$ for some $k \in \mathbb{N}$, we define*

$$\begin{aligned} E_t &:= \{(p, q) \in [k]^2 \mid t_p = t_q\}, \\ N_t &:= \{(p, q) \in [k]^2 \mid t_p \neq t_q\}, \text{ and} \\ O_t &:= \{(p, q) \in [k]^2 \mid t_p < t_q\}, \end{aligned}$$

where $<$ is the standard order on $\mathbb{Q}$.

As we already alluded to, we will repeatedly use the fact that in temporal VCSPs, any valued relation R is such that $R(t)$ only depends on the *order type* of the tuple t.

Observation 4.2. *Let Γ be a valued structure with a finite signature τ such that* $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\Gamma)$ *and let $R \in \tau$ be any valued relation of arity k. Then for every $t, t' \in \mathbb{Q}^k$ such that $E_t = E_{t'}$ and $O_t = O_{t'}$ it holds that $R(t) = R(t')$.*

4.2 Equality VCSPs

An *equality relational structure* is a relational structure whose automorphism group is the group of all permutations of its domain [17]; we define an *equality valued structure* analogously. In this section, we prove that for every equality valued structure Γ, $\mathrm{VCSP}(\Gamma)$ is in P or NP-complete. This generalises the P versus NP-complete dichotomy for equality min-CSPs from [73]. If the domain of Γ is finite, then this is already known (see Theorem 0.5). It is easy to see that classifying the general infinite case reduces to the countably infinite case. For notationally convenient use in the later sections, we work with the domain $\mathbb{Q}$, but we could have used any other countably infinite set instead.

Recall the relation Dis from Example 1.21. It follows from Theorem 4.3 below that the structure $(\mathbb{Q}; \mathrm{Dis})$ pp-constructs K_3. Let $\mathrm{const}\colon \mathbb{Q} \to \mathbb{Q}$ be the constant zero operation, given by $\mathrm{const}(x) := 0$ for all $x \in \mathbb{Q}$. Let $\mathrm{inj}\colon \mathbb{Q}^2 \to \mathbb{Q}$ be a fixed injective operation.

Theorem 4.3 ([9,17])**.** *If $\mathfrak{A}$ is a relational structure such that* $\mathrm{Aut}(\mathfrak{A}) = \mathrm{Sym}(\mathbb{Q})$*, then exactly one of the following cases applies.*

- $\mathrm{const} \in \mathrm{Pol}(\mathfrak{A})$ *or* $\mathrm{inj} \in \mathrm{Pol}(\mathfrak{A})$*. In this case, for every reduct $\mathfrak{A}'$ of $\mathfrak{A}$ with a finite signature,* $\mathrm{CSP}(\mathfrak{A}')$ *is in P.*
- *The relation* Dis *has a primitive positive definition in $\mathfrak{A}$. In this case, $\mathfrak{A}$ pp-constructs K_3 and $\mathfrak{A}$ has a reduct $\mathfrak{A}'$ with a finite signature such that* $\mathrm{CSP}(\mathfrak{A}')$ *is NP-complete.*

We prove the following general lemma that assumes only $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\Gamma)$ to avoid repeating the proof in Section 4.4. The case $\mathrm{Aut}(\Gamma) = \mathrm{Sym}(\mathbb{Q})$ is a special case.

Lemma 4.4. *Let Γ be a valued structure such that* $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\Gamma)$*. If* $\mathrm{const} \notin \mathrm{fPol}(\Gamma)$*, then $(\neq)_0^\infty \in \langle \Gamma \rangle$ or $(<)_0^\infty \in \langle \Gamma \rangle$. In particular, if* $\mathrm{Aut}(\Gamma) = \mathrm{Sym}(\mathbb{Q})$*, then $(\neq)_0^\infty \in \langle \Gamma \rangle$.*

Proof. By assumption, there exists $R \in \tau$ of arity k and $t \in A^k$ such that $m := R(t) < R(0, \dots, 0)$. For $i \in \{1, \dots, k\}$, define $\psi_i(x_1, \dots, x_i)$ to be $R(x_1, \dots, x_i, x_i, \dots, x_i)$. Choose t and i such that i is minimal with the property that $\psi_i(t_1, \dots, t_i) < R(0, \dots, 0)$. Note that such an i exists, because for $i = k$ we have $\psi_i(t_1, \dots, t_i) = R(t_1, \dots, t_k) < R(0, \dots, 0)$. Moreover, $i > 1$, since for every $a \in A$ there exists $\alpha \in \mathrm{Aut}(\Gamma)$ such that $\alpha(a) = 0$ and hence $R(a, \dots, a) = R(\alpha(a), \dots, \alpha(a)) = R(0, \dots, 0)$. Also note that $t_{i-1} \neq t_i$, by the minimality of i.

From all the such pairs (t, i) that minimise i, choose a pair (t, i) where $\psi_i(t_1, \dots, t_i)$ is minimal. Such a t exists because R attains only finitely many values. Define

$$\psi(x_{i-1}, x_i) := \min_{x_1, \dots, x_{i-2}} \psi_i(x_1, \dots, x_{i-2}, x_{i-1}, x_i).$$

Let $a \in \mathbb{Q}$ and note that, by observation 4.2, value $\psi(a,a)$ does not depend on the choice of a. By our choice of i, $\psi(a,a) > \psi(t_{i-1},t_i)$; otherwise, there are $a_1,\dots,a_{i-2} \in \mathbb{Q}$ such that $\psi_i(a_1,\dots,a_{i-2},a,a) \leq \psi_i(t_1,\dots,t_i)$, in contradiction to the choice of (t,i) such that i is minimal. We distinguish three cases (recall that $t_{i-1} \neq t_i$):

(1) $\psi(t_{i-1},t_i) = \psi(t_i,t_{i-1})$,

(2) $\psi(t_{i-1},t_i) < \psi(t_i,t_{i-1})$ and $t_{i-1} < t_i$,

(3) $\psi(t_{i-1},t_i) < \psi(t_i,t_{i-1})$ and $t_i < t_{i-1}$.

Note that for $a,b \in A$ such that $a < b$, the values $\psi(a,b)$ and $\psi(b,a)$ do not depend on the choice of a, b. In case (1), $\mathrm{Opt}(\psi)$ expresses $(\neq)_0^\infty$. In case (2), $\mathrm{Opt}(\psi)$ expresses $(<)_0^\infty$. Finally, in case (3) $\mathrm{Opt}(\psi)$ expresses $(>)_0^\infty$, which expresses $(<)_0^\infty$ by exchanging the input variables.

The last statement follows from the fact that $\mathrm{Sym}(\mathbb{Q})$ does not preserve $(<)_0^\infty$. $\square$

The following is a reformulation of what was observed in Example 1.29 and of Theorem 4.3.

Lemma 4.5. *Let Γ be a valued structure such that $\mathrm{Aut}(\Gamma) = \mathrm{Sym}(\mathbb{Q})$. If $\langle\Gamma\rangle$ contains $(=)_0^1$ and $(\neq)_0^\infty$, then $\mathrm{Dis} \in \langle\Gamma\rangle$. In particular, Γ pp-constructs K_3, and, if the signature of Γ is finite, $\mathrm{VCSP}(\Gamma)$ is NP-complete.*

As a next step, we prove two lemmas that will be used in the classification of equality VCSPs.

Lemma 4.6. *Let Γ be a valued structure such that $\mathrm{Aut}(\Gamma) = \mathrm{Sym}(\mathbb{Q})$. Suppose that $\mathrm{fPol}(\Gamma)$ contains inj. Then $\mathrm{VCSP}(\Gamma)$ is in P.*

Proof. Let (ϕ,u) be an instance of $\mathrm{VCSP}(\Gamma)$ with variable set V. We first check whether ϕ contains summands with at most one variable that evaluate to ∞ for some (equivalently, for all) assignment; in this case, the minimum of ϕ is above every rational threshold and the algorithm rejects. Otherwise, we propagate (crisp) forced equalities: if ϕ contains a summand $R(x_1,\dots,x_k)$ and for all $f\colon V \to A$ we have that if $R(f(x_1),\dots,f(x_k))$ is finite, then $f(x_i) = f(x_j)$ for some $i < j$, then we say that $x_i = x_j$ is *forced*. In this case, we replace all occurrences of x_j in ϕ by x_i and repeat this process (including the check for unary summands that evaluate to ∞); clearly, this procedure must terminate after finitely many steps. Let V' be the resulting set of variables, and let ϕ' be the resulting instance of $\mathrm{VCSP}(\Gamma)$. Clearly, the minimum for ϕ' equals the minimum for ϕ. Fix any injective $g\colon V' \to \mathbb{Q}$; we claim that g minimises ϕ'. To see this, let $f\colon V' \to \mathbb{Q}$ be any assignment and let $\psi(x) = \psi(x_1,\dots,x_k)$ be a summand of ϕ'. We show that $\psi(g(x)) \leq \psi(f(x))$. The statement is trivially true if $k = 1$ by the transitivity of $\mathrm{Aut}(\Gamma)$. Assume therefore that $k \geq 2$.

We first prove that $\psi(g(x))$ is finite. Let $t^1,\dots,t^n$ be an enumeration of representatives of all orbits of k-tuples such that $\psi(t^i) < \infty$ and note that $n \geq 1$, because otherwise the algorithm would have rejected the instance. If for some distinct $p,q \in \{1,\dots,k\}$

we have $(t^i)_p = (t^i)_q$ for all $i \in \{1,\dots,n\}$, then the algorithm would have replaced all occurrences of x_p by x_q or vice versa. So for all distinct $p,q \in \{1,\dots,k\}$ there exists $i \in \{1,\dots,n\}$ such that $(t^i)_p \neq (t^i)_q$. Therefore, since inj is injective, the tuple $\mathrm{inj}(t^1, \mathrm{inj}(t^2,\dots,\mathrm{inj}(t^{n-1},t^n)\dots))$ lies in the same orbit as $g(x)$. Since $\mathrm{inj} \in \mathrm{fPol}(\Gamma)$, we have $\psi(g(x)) < \infty$.

Note that

- $2\psi(\mathrm{inj}(g(x), f(x)) \leq \psi(g(x)) + \psi(f(x))$, because $\mathrm{inj} \in \mathrm{fPol}(\Gamma)$, and
- $\mathrm{inj}(g(x), f(x))$ lies in the same orbit of $\mathrm{Aut}(\Gamma)$ as $g(x)$, and thus $\psi(\mathrm{inj}(g(x), f(x)) = \psi(g(x))$.

Combining, we obtain that $\psi(g(x)) \leq \psi(f(x))$. It follows that g minimises ϕ'. Recall that the minimum of ϕ and ϕ' are equal. Therefore, the algorithm accepts if the evaluation of ϕ' under g is at most u and rejects otherwise. Since checking whether a summand forces an equality can be done in constant time, and there is a linear number of variables, the propagation of forced equalities can be done in polynomial time. It follows that VCSP(Γ) is in P. $\square$

Lemma 4.7. *Let Γ be a valued structure such that $\mathrm{Aut}(\Gamma) = \mathrm{Sym}(\mathbb{Q})$. Suppose that $\mathrm{inj} \notin \mathrm{fPol}(\Gamma)$. Then $(=)_0^1 \in \langle(\Gamma, (\neq)_0^\infty)\rangle$ or $\mathrm{Dis} \in \langle(\Gamma, (\neq)_0^\infty)\rangle$.*

Proof. By assumption, there exists $R \in \tau$ with arity k which is not improved by inj, that is, there exist $s,t \in \mathbb{Q}^k$ such that

$$R(s) + R(t) < 2R(\mathrm{inj}(s,t)). \tag{4.1}$$

Note that, in particular, $R(s), R(t) < \infty$. Suppose first that $\mathrm{inj}(s,t) = \infty$. In this case, the inequality above implies that $\mathrm{Feas}(R)$ is not improved by inj. It follows that $\mathrm{Pol}(\mathbb{Q}; \mathrm{Feas}(R), (\neq)_0^\infty)$ contains neither const nor inj. Hence, by Theorem 4.3, $(\mathbb{Q}; \mathrm{Feas}(R), (\neq)_0^\infty)$ primitively positively defines Dis and thus $\mathrm{Dis} \in \langle(\Gamma, (\neq)_0^\infty\rangle$. We may therefore assume in the rest of the proof that $R(\mathrm{inj}(s,t)) < \infty$.

Inequality (4.1) implies that $R(s) < R(\mathrm{inj}(s,t))$ or $R(t) < R(\mathrm{inj}(s,t))$. Without loss of generality, assume $R(t) < R(\mathrm{inj}(s,t))$. Since $R(\mathrm{inj}(s,t))$ is finite, this implies

$$R(\mathrm{inj}(s,t)) + R(t) < 2R(\mathrm{inj}(s,t)) = 2R(\mathrm{inj}(\mathrm{inj}(s,t),t)),$$

where the last equality follows from the fact that $\mathrm{inj}(s,t)$ and $\mathrm{inj}(\mathrm{inj}(s,t),t)$ lie in the same orbit of $\mathrm{Aut}(\Gamma) = \mathrm{Sym}(\mathbb{Q})$. This is an inequality of the same form as (4.1) (with $\mathrm{inj}(s,t)$ in the role of s), which implies that we can assume without loss of generality that s and $\mathrm{inj}(s,t)$ lie in the same orbit of $\mathrm{Aut}(\Gamma)$. Then (4.1) implies that $R(t) < R(\mathrm{inj}(s,t)) = R(s)$. We show that in this case $(=)_0^1 \in \langle(\Gamma, (\neq)_0^\infty)\rangle$.

Out of all pairs $(s,t) \in (A^k)^2$ such that s and $\mathrm{inj}(s,t)$ lie in the same orbit and $R(t) < R(s)$, we choose (s,t) such that $t_p \neq t_q$ holds for as many pairs (p,q) as possible. Note that s and t cannot lie in the same orbit, and by the injectivity of inj there exist $i,j \in \{1,\dots,k\}$ such that $t_i = t_j$ and $s_i \neq s_j$. Note that since s and $\mathrm{inj}(s,t)$ lie in the same orbit, we have $E_s \subseteq E_t$ and $N_t \subseteq N_s$ (recall Definition 4.1). For the sake of notation, assume that $i = 1$ and $j = 2$. Consider the expression

$$\phi(x_1, x_2) := \min_{x_3,\dots,x_k} R(x_1, \dots, x_k) + \sum_{(p,q)\in E_s} (=)_0^\infty(x_p, x_q) + \sum_{(p,q)\in N_t} (\neq)_0^\infty(x_p, x_q).$$

Then $\phi(x, y)$ attains at most two values by the 2-transitivity of $\mathrm{Aut}(\Gamma)$. For every $x \in \mathbb{Q}$, we have $\phi(x, x) =: m \leq R(t)$. Let $\ell := \phi(x, y)$ for some distinct $x, y \in \mathbb{Q}$; this value does not depend on the choice of x and y by the 2-transitivity of $\mathrm{Aut}(\Gamma)$. Suppose for contradiction that $\ell \leq m$. Then there exists a tuple $u \in \mathbb{Q}^k$ such that

(i) $R(u) \leq R(t)$,

(ii) $u_i \neq u_j$,

(iii) $E_s \subseteq E_u$, and

(iv) $N_t \subseteq N_u$.

By (iii), $\mathrm{inj}(s, u)$ lies in the same orbit as s. By (i), we get that $R(u) \leq R(t) < R(s)$. By (ii) and (iv), u satisfies $u_p \neq u_q$ for more pairs (p, q) than t, which contradicts our choice of t. Therefore, $m < \ell$. It follows that $\phi(x_1, x_2)$ is equivalent to $(=)_m^\ell$ with $m < \ell$. Recall that $(=)_0^\infty \in \langle\Gamma\rangle$ by definition and therefore $(=)_m^\ell \in \langle\Gamma, (\neq)_0^\infty\rangle$. Note that $\ell \leq R(s) < \infty$. Hence, shifting ϕ by $-m$ and scaling it by $1/(\ell - m)$ shows that $(=)_0^1 \in \langle\Gamma, (\neq)_0^\infty\rangle$, as we wanted to prove. □

We are now ready to prove the classification theorem for equality VCSPs.

Theorem 4.8. *Let Γ be a valued structure with a countably infinite domain $\mathbb{Q}$ over a finite relational signature such that $\mathrm{Aut}(\Gamma) = \mathrm{Sym}(\mathbb{Q})$. Then exactly one of the following two cases applies.*

1. $(\neq)_0^\infty \in \langle\Gamma\rangle$ *and* $\mathrm{Dis} \in \langle\Gamma\rangle$. *In this case, Γ pp-constructs K_3, and* $\mathrm{VCSP}(\Gamma)$ *is NP-complete.*

2. $\mathrm{const} \in \mathrm{fPol}(\Gamma)$ *or* $\mathrm{inj} \in \mathrm{fPol}(\Gamma)$. *In both of these cases,* $\mathrm{VCSP}(\Gamma)$ *is in P.*

Proof. If $\mathrm{const} \in \mathrm{fPol}(\Gamma)$, then Lemma 3.24 implies that $\mathrm{VCSP}(\Gamma)$ is in P. If $\mathrm{const} \notin \mathrm{fPol}(\Gamma)$, then Γ can express $(\neq)_0^\infty$ by Lemma 4.4.

If $\mathrm{inj} \in \mathrm{fPol}(\Gamma)$, then Lemma 4.6 implies that $\mathrm{VCSP}(\Gamma)$ is in P. If $\mathrm{inj} \notin \mathrm{fPol}(\Gamma)$, then Γ can express Dis or $(=)_0^1$ by Lemma 4.7. If $\mathrm{Dis} \in \langle\Gamma\rangle$, then the statement follows from Theorem 4.3. If $(=)_0^1 \in \langle\Gamma\rangle$, then we obtain that $\mathrm{Dis} \in \langle\Gamma\rangle$ by Lemma 4.5 and again the statement follows. Note that neither const nor inj improves Dis. Therefore, the two cases in the statement of the theorem are disjoint. □

Remark 4.9. *If Γ is as in Theorem 4.8, then Γ has a quantifier-free first-order definition in $(\mathbb{Q}; =)$ as introduced in Section 1.3. If Γ is given by such a first-order definition over $(\mathbb{Q}; =)$, then it is decidable which of the conditions (1) and (2) in Theorem 4.8 applies: we can just test whether all valued relations of Γ are improved by* const, *and we can test whether all of them are improved by* inj.

Remark 4.10. *The complexity classification of equality minCSPs from [74], which can be viewed as VCSPs of valued structures where each relation attains only values* 0 *and* 1*, can be obtained as a special case of Theorem 4.8. Suppose that* Γ *is such a valued structure. If* $\mathrm{const} \in \mathrm{fPol}(\Gamma)$*, then* Γ *is constant (in the terminology of [74]) and* $\mathrm{VCSP}(\Gamma)$ *is in P. If* $\mathrm{inj} \in \mathrm{fPol}(\Gamma)$*, then it is immediate that* Γ *is Horn (in the terminology of [74]) and even strictly negative: otherwise, as observed in [74], we have* $(=)_0^1 \in \langle \Gamma \rangle$*. But this is in contradiction to the assumption that* $\mathrm{inj} \in \mathrm{fPol}(\Gamma)$*, since* inj *does not improve* $(=)_0^1$*. Otherwise, it follows from Theorem 4.8 that* $\mathrm{VCSP}(\Gamma)$ *is NP-hard.*

4.3 Preliminaries on temporal CSPs

We first define several important relations on $\mathbb{Q}$ that already played a role in the classification of temporal CSPs [18].

Definition 4.11. *Let*

$$\begin{aligned}
\mathrm{Betw} &:= \{(x,y,z) \in \mathbb{Q}^3 \mid (x<y<z) \vee (z<y<x)\},\\
\mathrm{Cycl} &:= \{(x,y,z) \in \mathbb{Q}^3 \mid (x<y<z) \vee (y<z<x) \vee (z<x<y)\},\\
\mathrm{Sep} &:= \{(x_1,y_1,x_2,y_2) \in \mathbb{Q}^4 \mid \quad (x_1<x_2<y_1<y_2) \vee (x_1<y_2<y_1<x_2)\\
&\qquad \vee (y_1<x_2<x_1<y_2) \vee (y_1<y_2<x_1<x_2)\\
&\qquad \vee (x_2<x_1<y_2<y_1) \vee (x_2<y_1<y_2<x_1)\\
&\qquad \vee (y_2<x_1<x_2<y_1) \vee (y_2<y_1<x_2<x_1)\},\\
T_3 &:= \{(x,y,z) \in \mathbb{Q}^3 \mid (x=y<z) \vee (x=z<y)\}.
\end{aligned}$$

The following theorem provides an important case distinction for temporal relational structures.

Theorem 4.12 (Theorem 20 in [18])**.** *Let* $\mathfrak{A}$ *be a relational structure with a finite signature such that* $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\mathfrak{A})$*. Then it satisfies at least one of the following:*

- $\mathfrak{A}$ *primitively positively defines* Betw*,* Cycl*, or* Sep*.*
- $\mathrm{const} \in \mathrm{Pol}(\mathfrak{A})$*.*
- $\mathrm{Aut}(\mathfrak{A}) = \mathrm{Sym}(\mathbb{Q})$*.*
- *There is a primitive positive definition of* $<$ *in* $\mathfrak{A}$*.*

We need the following operations on $\mathbb{Q}$. By min and max we refer to the binary minimum and maximum operation on the set $\mathbb{Q}$, respectively.

Definition 4.13. *Let* $e_{<0}, e_{>0}$ *be any endomorphisms of* $(\mathbb{Q};<)$ *satisfying* $e_{<0}(x) < 0$ *and* $e_{>0}(x) > 0$ *for every* $x \in \mathbb{Q}$*. We denote by* $\pi\pi$ *the binary operation on* $\mathbb{Q}$ *defined by*

$$\pi\pi(x,y) = \begin{cases} e_{<0}(x) & x \le 0,\\ e_{>0}(y) & x > 0. \end{cases}$$

The operation lex *is any binary operation on* $\mathbb{Q}$ *satisfying* $\mathrm{lex}(x,y) < \mathrm{lex}(x',y')$ *iff* $x < x'$, *or* $x = x'$ *and* $y < y'$ *for all* $x, x', y, y' \in \mathbb{Q}$. *We denote by* ll *the binary operation on* $\mathbb{Q}$ *defined by*

$$\mathrm{ll}(x,y) = \begin{cases} \mathrm{lex}(e_{<0}(x), e_{<0}(y)) & x \leq 0, \\ \mathrm{lex}(e_{>0}(y), e_{>0}(x)) & x > 0. \end{cases}$$

Definition 4.14. *Let* $e_<$, $e_=$ *and* $e_>$ *be any endomorphisms of* $(\mathbb{Q};<)$ *satisfying for all* $x, \varepsilon \in \mathbb{Q}$, $\varepsilon > 0$,

$$e_=(x) < e_>(x) < e_<(x) < e_=(x+\varepsilon).$$

We denote by mi *the binary operation on* $\mathbb{Q}$ *defined by*

$$\mathrm{mi}(x,y) = \begin{cases} e_<(x) & x < y, \\ e_=(x) & x = y, \\ e_>(y) & x > y. \end{cases}$$

Definition 4.15. *Let* $e_=$ *and* $e_{\neq}$ *be any endomorphisms of* $(\mathbb{Q};<)$ *satisfying for all* $x, \varepsilon \in \mathbb{Q}$, $\varepsilon > 0$,

$$e_{\neq}(x) < e_=(x) < e_{\neq}(x+\varepsilon).$$

We denote by mx *the binary operation on* $\mathbb{Q}$ *defined by*

$$\mathrm{mx}(x,y) = \begin{cases} e_{\neq}(\min(x,y)) & x \neq y, \\ e_=(x) & x = y. \end{cases}$$

The construction of endomorphisms that appear in Definitions 4.14 and 4.15 can be found for example in [9, Section 12.5]. The following was observed and used in [18].

Lemma 4.16. *If* $\mathfrak{A}$ *is a relational structure such that* $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\mathfrak{A})$ *and* $\mathfrak{A}$ *is preserved by a binary injective operation* f*, then it is also preserved by the operation defined as one of* $\mathrm{lex}(x,y)$, $\mathrm{lex}(-x,y)$, $\mathrm{lex}(x,-y)$*, or* $\mathrm{lex}(-x,-y)$. *In particular, if* f *preserves* $\leq$ *(for example,* ll*), then* $\mathfrak{A}$ *is preserved by* lex.

Definition 4.17. *The* dual *of an operation* $g\colon \mathbb{Q}^k \to \mathbb{Q}$ *is the operation*

$$g^*\colon (x_1, \ldots, x_k) \mapsto -g(-x_1, \ldots, -x_k).$$

The dual *of a relation* $R \subseteq \mathbb{Q}^\ell$ *is the relation*

$$-R = \{(-a_1, \ldots, -a_\ell) \mid (a_1, \ldots, a_\ell) \in R\}.$$

Note that $\min^* = \max$ and the relation $-(>)$ is equal to $<$. Statements about operations and relations on $\mathbb{Q}$ can be naturally dualized and we will often use dual versions of results on temporal (V)CSPs; we do not state it explicitly if it is clear from the context.

By combining Theorem 50, Corollary 51, Corollary 52 and the accompanying remarks in [18], we obtain the following; see also Theorem 12.10.1 in [9].

Theorem 4.18. *Let $\mathfrak{A}$ be a relational structure such that* $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\mathfrak{A})$. *Then exactly one of the following is true.*

1. *At least one of the operations* const, min, mx, mi, ll, *or one of their duals lies in* $\mathrm{Pol}(\mathfrak{A})$. *In this case, for every reduct $\mathfrak{A}'$ of $\mathfrak{A}$ with a finite signature,* $\mathrm{CSP}(\mathfrak{A}')$ *is P.*

2. *$\mathfrak{A}$ primitively positively defines one of the relations* Betw, Cycl, Sep, T_3, $-T_3$, *or* Dis. *In this case, $\mathfrak{A}$ has a reduct $\mathfrak{A}'$ with a finite signature such that* $\mathrm{CSP}(\mathfrak{A}')$ *is NP-complete.*

Moreover, it is decidable whether (1) or (2) holds.

Let G be a permutation group on a set A and $\ell \in \mathbb{N}$. An operation $f : A^\ell \to A$ is called a *pseudo weak near unanimity (pwnu)* operation (with respect to G) if there exist $e_1, \ldots, e_\ell \in \overline{G}$ such that for every $x, y \in A$,

$$e_1 f(y, x, \ldots, x) = e_2 f(x, y, x, \ldots, x) = \cdots = e_\ell f(x, \ldots, x, y).$$

We say that a relational structure $\mathfrak{A}$ has a *pwnu polymorphism*, if there exists $f \in \mathrm{Pol}(\mathfrak{A})$ which is a pwnu operation with respect to $\mathrm{Aut}(\mathfrak{A})$. Using this notion, we can formulate an alternative version of Theorem 4.18.

Theorem 4.19 (Theorem 12.0.1 in [9]; see also Theorem 7.24 in [29])**.** *Let $\mathfrak{A}$ be a relational structure such that* $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\mathfrak{A})$. *Then exactly one of the following is true:*

1. *$\mathfrak{A}$ has a pwnu polymorphism. In this case, for every reduct $\mathfrak{A}'$ of $\mathfrak{A}$ with a finite signature,* $\mathrm{CSP}(\mathfrak{A}')$ *is in P.*

2. *$\mathfrak{A}$ pp-constructs K_3. In this case, there exists a reduct $\mathfrak{A}'$ of $\mathfrak{A}$ with a finite signature such that* $\mathrm{CSP}(\mathfrak{A}')$ *is NP-complete.*

The following proposition provides a connection between the hardness conditions from Theorem 4.18 and 4.19.

Proposition 4.20. *Each of the relational structures* $(\mathbb{Q}; \mathrm{Betw})$, $(\mathbb{Q}; \mathrm{Cycl})$, $(\mathbb{Q}; \mathrm{Sep})$, $(\mathbb{Q}; T_3)$, *and* $(\mathbb{Q}; -T_3)$ *pp-constructs K_3.*

Proof. This is proved in the proof of [9, Theorem 12.0.1]. In fact, the proof shows that each of these structures pp-interprets all finite structures. Since K_3 is finite and a pp-interpretation is a special case of a pp-construction, the statement follows. □

We finish this section by stating several results about polymorphisms of temporal relational structures.

Proposition 4.21 (Proposition 25, 27, and 29 in [18])**.** *Let $\mathfrak{A}$ be a relational structure such that* $\mathrm{Pol}(\mathfrak{A})$ *contains* min, mi, *or* mx. *Then* $\mathrm{Pol}(\mathfrak{A})$ *contains* $\pi\pi$.

Proposition 4.22 (Lemma 12.4.4 in [9]). *Let $\mathfrak{A}$ be a relational structure such that* $\mathrm{Pol}(\mathfrak{A})$ *contains* lex *and* $\pi\pi$. *Then* $\mathrm{Pol}(\mathfrak{A})$ *contains* ll.

Note that there exists $\alpha \in \mathrm{Aut}(\mathbb{Q}; <)$ such that, for all $x, y \in \mathbb{Q}$, $\mathrm{lex}^*(x, y) = \alpha(\mathrm{lex}(x, y))$. Hence, whenever $\mathfrak{A}$ is a relational structure such that $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\mathfrak{A})$, then $\mathrm{Pol}(\mathfrak{A})$ contains lex if and only if it contains lex*. This is relevant for dualising statements like Proposition 4.22. We also need the following result from [13].

Theorem 4.23 (Theorem 5.1 in [13]). *Let $\mathfrak{A}$ be an expansion of $(\mathbb{Q}; <)$ by first-order definable relations. If $\pi\pi \in \mathrm{Pol}(\mathfrak{A})$ and $\mathrm{ll} \notin \mathrm{Pol}(\mathfrak{A})$, then the relation*

$$R^{\mathrm{mix}} = \{(x, y, z) \in \mathbb{Q}^3 \mid (x = y) \vee (z < x \wedge z < y)\}$$

has a primitive positive definition in $\mathfrak{A}$.

4.4 Expressibility of temporal valued relations

In this section we consider *valued* structures Γ such that $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$, and study expressibility of valued relations in Γ. For $\alpha, \beta, \gamma \in \mathbb{Q} \cup \{\infty\}$, define the binary valued relation $R_{\alpha,\beta,\gamma}$ on $\mathbb{Q}$:

$$R_{\alpha,\beta,\gamma}(x, y) := \begin{cases} \alpha & x = y \\ \beta & x < y \\ \gamma & x > y \end{cases}$$

Note that $R_{0,1,1}$ is equal to $(=)_0^1$, $R_{1,0,0}$ is equal to $(\neq)_0^1$, $R_{1,0,1}$ is equal to $(<)_0^1$, and $R_{0,0,1}$ is equal to $(\leq)_0^1$.

Lemma 4.24. *Let Γ be a valued structure such that $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$ and $\alpha > \frac{1}{3}$. If $\langle \Gamma \rangle$ contains $R_{\alpha,0,1}$, then $\mathrm{Cycl} \in \langle \Gamma \rangle$.*

Proof. Note that $\mathrm{Cycl}(x, y, z) = \mathrm{Opt}(R_{\alpha,0,1}(x, y) + R_{\alpha,0,1}(y, z) + R_{\alpha,0,1}(z, x))$. Therefore, $\mathrm{Cycl} \in \langle \Gamma \rangle$. □

Assuming $(<)_0^\infty \in \langle \Gamma \rangle$, we can relax the assumptions on the parameters α, β, γ and still get $\mathrm{Cycl} \in \langle \Gamma \rangle$.

Lemma 4.25. *Let Γ be a valued structure such that $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$. Let $(<)_0^\infty \in \langle \Gamma \rangle$. Let $\alpha, \beta, \gamma \in \mathbb{Q} \cup \{\infty\}$ be such that*

- *$\alpha < \min(\beta, \gamma) < \infty$, or*
- *$\beta \neq \gamma$ and $\beta, \gamma < \infty$.*

If $R_{\alpha,\beta,\gamma} \in \langle \Gamma \rangle$, then $\mathrm{Cycl} \in \langle \Gamma \rangle$.

Proof. Without loss of generality, we may assume that $\beta \leq \gamma$, because otherwise we may consider $R_{\alpha,\gamma,\beta}(x,y) = R_{\alpha,\beta,\gamma}(y,x)$. In the first case of the statement we have that $\alpha < \min(\beta,\gamma) = \beta < \infty$. Then

$$(<)_0^1(x,y) = \frac{1}{\beta-\alpha} \min_z \left(R_{\alpha,\beta,\gamma}(z,x) + (<)_0^\infty(z,y) - \alpha\right).$$

Suppose now that we are in the second case, i.e., $\beta < \gamma < \infty$, and additionally not in the first case, i.e., $\alpha \geq \min(\beta,\gamma) = \beta$. Then for every $x,y \in \mathbb{Q}$

$$(<)_0^1(x,y) = \frac{1}{\gamma-\beta} \min_z \left(R_{\alpha,\beta,\gamma}(x,z) + (<)_0^\infty(z,y) - \beta\right).$$

Therefore, in both cases, $(<)_0^1 \in \langle\Gamma\rangle$. Since $(<)_0^1$ is equal to $R_{1,0,1}$, the statement follows from Lemma 4.24. $\square$

The following lemma implies that, when analyzing complexity of VCSPs of temporal structures, we can assume that $(<)_0^\infty \in \langle\Gamma\rangle$ without loss of generality (taking into account the results we already presented in this chapter and Chapter 3).

Lemma 4.26. *Let* Γ *be a valued structure such that* $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\Gamma) \neq \mathrm{Sym}(\mathbb{Q})$*. If* $\mathrm{const} \notin \mathrm{fPol}(\Gamma)$*, then* $\langle\Gamma\rangle$ *contains* Betw*,* Cycl*,* Sep*, or* $(<)_0^\infty$*.*

Proof. By Lemma 4.4, $\langle\Gamma\rangle$ contains $(<)_0^\infty$ or $(\neq)_0^\infty$. If $(<)_0^\infty \in \langle\Gamma\rangle$, then we are done. Assume therefore $(\neq)_0^\infty \in \langle\Gamma\rangle$. Let R be a valued relation of Γ of arity k such that there exists an orbit O of the action of $\mathrm{Sym}(\mathbb{Q})$ on $\mathbb{Q}^k$ and $s,t \in O$ with $R(s) < R(t)$. Let $s \in O$ be such that $R(s)$ is minimal. Note that O is not the orbit of constant tuples, because $\mathrm{Aut}(\mathbb{Q};<)$ is transitive.

Consider the crisp relation $S \in \langle\Gamma\rangle_0^\infty$ defined by

$$S(x_1,\dots,x_k) := \mathrm{Opt}\left(R(x_1,\dots,x_k) + \sum_{(p,q)\in E_s} (=)_0^\infty(x_p,x_q) + \sum_{(p,q)\in N_s} (\neq)_0^\infty(x_p,x_q)\right).$$

Clearly, $s \in S$. Note that a tuple $u \in \mathbb{Q}^k$ lies in O if and only if $E_s \subseteq E_u$ and $N_s \subseteq N_u$. In particular, $S \subseteq O$. Since $R(s) < R(t)$, we have $t \notin S$. It follows that S is not preserved by $\mathrm{Sym}(\mathbb{Q})$. Moreover, S is not preserved by const, because O is not the orbit of constant tuples. By Theorem 4.12, the relational structure $(\mathbb{Q};S)$ admits a primitive positive definition of Betw, Cycl or Sep, or a primitive positive definition of $<$. Since $S \in \langle\Gamma\rangle$, the statement of the lemma follows. $\square$

In Lemma 4.29 below we present a polynomial-time algorithm for VCSPs of valued structures Γ improved by lex provided that Γ cannot express any crisp relation that prevents tractability. In fact, to check whether the algorithm can be applied, it suffices to check whether a certain structure $\hat{\Gamma}$ with a finite signature has a tractable CSP, instead of considering all relations in $\langle\Gamma\rangle_0^\infty$. We define $\hat{\Gamma}$ below.

Definition 4.27. *Let A be a set and let R be a valued relation on A of arity k. Let $\ell \in \mathbb{N}$, $\ell \leq k$, and let $\sigma\colon [k] \to [\ell]$ be a map. Then R_σ is the valued relation on A of arity ℓ defined by $R_\sigma(x_1, \dots, x_\ell) = R(x_{\sigma(1)}, \dots, x_{\sigma(k)})$ for all $x_1, \dots, x_\ell \in A$. If S is a valued relation of some arity $\ell \leq k$ such that there exists $\sigma\colon [k] \to [\ell]$ and $S = R_\sigma$, we call S a* minor *of R.*

Let Γ be a valued τ-structure. Then $\hat{\Gamma}$ denotes the relational structure on the same domain which contains the relations $\operatorname{Feas}(R^\Gamma)$ and $\operatorname{Opt}((R^\Gamma)_\sigma)$ for every $R \in \tau$ of arity k, $\ell \leq k$, and $\sigma\colon [k] \to [\ell]$.

Note that $R_\sigma \in \langle (A; R) \rangle$ for every valued relation R of arity k and every $\sigma\colon [k] \to [\ell]$.

Remark 4.28. *Note that we do not need to include relations of the form $\operatorname{Feas}((R^\Gamma)_\sigma)$ in $\hat{\Gamma}$, because for every valued relation R on a set C of arity k and $\sigma\colon [k] \to [\ell]$, we have*

$$\operatorname{Feas}(R_\sigma) = \operatorname{Feas}(R)_\sigma$$

and therefore $\operatorname{Feas}(R_\sigma) \in \langle (C; \operatorname{Feas}(R)) \rangle$. The same is not true for the operator Opt.

Lemma 4.29. *Let Γ be a valued structure with a finite signature such that $\operatorname{Aut}(\mathbb{Q}; <) \subseteq \operatorname{Aut}(\Gamma)$. Suppose that $\operatorname{lex} \in \operatorname{fPol}(\Gamma)$ and that $\hat{\Gamma}$ is preserved by one of the operations* const*,* min*,* mx*,* mi*,* ll *or one of their duals. Then $\operatorname{VCSP}(\Gamma)$ is in P.*

Proof. If const $\in \operatorname{fPol}(\Gamma)$, then $\operatorname{VCSP}(\Gamma)$ is in P by Lemma 3.24. We may therefore assume that const $\notin \operatorname{fPol}(\Gamma)$. Let $R \in \langle \Gamma \rangle$ be of arity k. Since lex $\in \operatorname{fPol}(\Gamma)$, it improves R. Therefore, for every injective tuple $s \in \mathbb{Q}^k$ and any $t \in \mathbb{Q}^k$, it holds that

$$R(s) = R(\operatorname{lex}(s,t)) \leq 1/2 \cdot (R(s) + R(t)),$$

where the first equality follows from s and $\operatorname{lex}(s,t)$ being in the same orbit of $\operatorname{Aut}(\Gamma)$. Therefore, if $R(s) < \infty$, then $R(s) \leq R(t)$. In particular, there is $m_R \in \mathbb{Q}$ such that for every injective tuple $s \in \mathbb{Q}^k$, we have $R(s) = m_R$ or $R(s) = \infty$. Note that if there is at least one injective tuple s with $R(s) = m_R$, then $\operatorname{Opt}(R)$ is the crisp relation that consists of all the tuples t such that $R(t) = m_R$.

Let (ϕ, u) be an instance of $\operatorname{VCSP}(\Gamma)$ with variable set $V = \{v_1, \dots, v_N\}$. Note that ϕ interpreted over $\operatorname{Feas}(\Gamma)$ can be seen as an instance of $\operatorname{CSP}(\operatorname{Feas}(\Gamma))$ where each summand $R(x_1, \dots, x_k)$ of ϕ is interpreted as $\operatorname{Feas}(R^\Gamma)(x_1, \dots, x_k)$. By the assumption on $\hat{\Gamma}$, $\operatorname{Feas}(\Gamma)$ is preserved by one of the operations const, min, mx, mi, ll, or one of their duals. Since lex $\in \operatorname{fPol}(\Gamma)$, by Lemma 3.18 and Proposition 3.22, lex $\in \operatorname{Pol}(\operatorname{Feas}(\Gamma))$. Since const $\notin \operatorname{fPol}(\Gamma)$, by Lemma 4.4, const $\notin \operatorname{Pol}(\operatorname{Feas}(\Gamma))$. Then min, mx, mi, or one of their duals preserves $\operatorname{Feas}(\Gamma)$, and, by Proposition 4.21, $\pi\pi \in \operatorname{Pol}(\operatorname{Feas}(\Gamma))$ or $\pi\pi^* \in \operatorname{Pol}(\operatorname{Feas}(\Gamma))$. Therefore, by Proposition 4.22, we always have that ll or ll* preserves $\operatorname{Feas}(\Gamma)$. Hence, by Theorem 4.18, $\operatorname{CSP}(\operatorname{Feas}(\Gamma))$ is solvable in polynomial time and we can use the polynomial-time algorithm from [18] based on the operation ll or ll* to solve $\operatorname{CSP}(\operatorname{Feas}(\Gamma))$. If ϕ, viewed as a primitive positive formula, is not satisfiable over $\operatorname{Feas}(\Gamma)$, then the minimum of ϕ is above every rational threshold and (ϕ, u) is rejected. Otherwise, we may compute the set $E \subseteq V^2$ of all pairs (x, y) such that $f(x) = f(y)$ in every solution of $f\colon V \to \mathbb{Q}$ of ϕ over Feas(Γ) (we may assume without loss of generality that Feas(Γ) contains the relation $(\neq)_0^\infty$; since Feas(Γ) is preserved by lex it suffices to test the unsatisfiability of $\phi \wedge x \neq y$ for each of these pairs). It follows from the definition of Feas that for every $g\colon V \to \mathbb{Q}$, if ϕ evaluates to a finite value in Γ under the assignment g, then $g(x) = g(y)$ for every $(x, y) \in E$. Moreover, for every $(x, y) \in V^2 \setminus E$, there exists $g\colon V \to \mathbb{Q}$ such that ϕ evaluates to a finite value under g and $g(x) \neq g(y)$.

We create a new τ-expression ϕ' from ϕ by replacing each occurrence of v_j by v_i for every $(v_i, v_j) \in E$ such that $i < j$. Let V' be the set of variables of ϕ'. By the discussion above, the minimum for ϕ' over Γ equals the minimum for ϕ. Moreover, for every $(x, y) \in (V')^2$, there exists $g'\colon V' \to \mathbb{Q}$ such that ϕ' evaluates to a finite value under g' and $g'(x) \neq g'(y)$. Let $\phi' := \phi'_1 + \cdots + \phi'_n$ where for every $j \in [n]$ the summand ϕ'_j is an atomic τ-expression. We execute the following procedure for each $j \in [n]$. Let $\phi'_j = R(x_1, \dots, x_k)$. Let $y_1^j, \dots, y_{\ell_j}^j$ be an enumeration of all distinct variables that appear in $\{x_1, \dots, x_k\}$ and let S_j be a valued relation of arity ℓ_j defined by $S_j(y_1^j, \dots, y_{\ell_j}^j) = R(x_1, \dots, x_k)$. Clearly, S_j is a minor of R. Note that the relation S_j might be different for every summand, even if they contain the same relation symbol R, due to possibly different variable identifications. Observe that, by the properties of ϕ', there exists an injective tuple $s^j \in \mathbb{Q}^{\ell_j}$ such that $S_j(s^j)$ is finite. Note that $S_j \in \langle\Gamma\rangle$, and let $m_j := m_{S_j}$. By the discussion in the beginning of the proof, $S_j(s^j) = m_j$ and $\mathrm{Opt}(S_j) \in \langle\Gamma\rangle_0^\infty$ consists of all tuples that evaluate to m_j in S_j. Since S_j attains only finitely many values, we can identify m_j in polynomial time for every j.

Let $\mathfrak{B}$ be the relational structure with domain $\mathbb{Q}$ and relations $\mathrm{Opt}(S_1), \dots, \mathrm{Opt}(S_n)$. Let ψ be the instance of CSP($\mathfrak{B}$) obtained from ϕ' by replacing the summand ϕ'_j by $\mathrm{Opt}(S_j)(y_1^j, \dots, y_{\ell_j}^j)$ for all $j \in [n]$; all relations in ψ are crisp and hence it can be seen as a primitive positive formula. Note that the variable set of ψ is equal to V'. By assumption, $\hat{\Gamma}$ is preserved by one of the operations const, min, mx, mi, ll, or one of their duals and, in particular, $\mathfrak{B}$ is preserved by one of them. Hence, CSP($\mathfrak{B}$) is in P by Theorem 4.18. Therefore, the satisfiability of ψ over $\mathfrak{B}$ can be tested in polynomial time. We claim that if ψ is unsatisfiable, then the minimum of ϕ is above every rational threshold and the algorithm rejects.

We prove the claim by contraposition. Suppose that the minimum of ϕ over Γ is finite. Then the minimum of ϕ' over Γ is finite and hence there exists $f'\colon V' \to \mathbb{Q}$ such that ϕ' evaluates to a finite value under f'. From all f' with this property, choose f' with the property that $f'(x) \neq f'(y)$ holds for as many pairs $(x, y) \in (V')^2$ as possible. We first show that f' is in fact injective. Suppose that there are $v, w \in V'$ such that $f'(v) = f'(w)$. Let $g'\colon V' \to \mathbb{Q}$ be such that ϕ' evaluates to a finite value under g' and $g'(v) \neq g'(w)$; recall that such g' must exist by the construction of ϕ'. Consider the assignment $\mathrm{lex}(f', g')\colon V' \to \mathbb{Q}$ and note that $\mathrm{lex}(f', g')(x) \neq \mathrm{lex}(f', g')(y)$ holds for all pairs (x, y) such that $f'(x) \neq f'(y)$ and also $\mathrm{lex}(f', g')(v) \neq \mathrm{lex}(f', g')(w)$. Moreover, ϕ' evaluates to a finite value under $\mathrm{lex}(f', g')$: for every $j \in [n]$, if ϕ'_j is of the form $R(x_1, \dots, x_k)$, then, since $\mathrm{lex} \in \mathrm{fPol}(\Gamma)$,

$$R(\mathrm{lex}(f', g')(x_1, \dots, x_k)) \leq 1/2 \cdot (R(f'(x_1, \dots x_k)) + R(g'(x_1, \dots, x_k))) < \infty.$$

This contradicts our choice of f'. Therefore, f' is injective.

Note that for every $j \in [n]$ we have $S_j(f'(y_1^j), \dots f'(y_{\ell_j}^j)) < \infty$, because ϕ'_j evaluates to a finite value under f'. Since $(f'(y_1^j), \dots f'(y_{\ell_j}^j))$ is an injective tuple, this implies $S_j(f'(y_1^j), \dots f'(y_{\ell_j}^j)) = m_j$ and $(f'(y_1^j), \dots f'(y_{\ell_j}^j)) \in \mathrm{Opt}(S_j)$ for every $j \in [n]$. It follows that f' is a satisfying assignment to ψ. Therefore, we proved that whenever ψ unsatisfiable, the algorithm correctly rejects, because there is no assignment to ϕ of finite cost.

Finally, suppose that there exists a solution $h' \colon V' \to \mathbb{Q}$ to the instance ψ of $\mathrm{CSP}(\mathfrak{B})$. Then, for every $j \in [n]$, ϕ'_j takes under h' the value $S_j(h'(y_1^j), \dots, h'(y_{\ell_j}^j))$. By the definition of Opt, $(h'(y_1^j), \dots, h'(y_{\ell_j}^j))$ minimizes S_j and therefore h' minimizes ϕ'_j. It follows that h' minimizes ϕ' and that the cost of ϕ' under h' is equal to $m_1 + \dots + m_n$. Since the cost of ϕ' is equal to the cost of ϕ, the algorithm accepts if $m_1 + \dots + m_n \leq u$ and rejects otherwise. This completes the algorithm and its correctness proof. It follows that $\mathrm{VCSP}(\Gamma)$ is in P. $\square$

The following lemma provides a useful case distinction concerning expressible relations in temporal structures.

Lemma 4.30. *Let Γ be a valued τ-structure such that $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\mathbb{Q}; <)$ and $(<)_0^\infty \in \langle \Gamma \rangle$. Suppose that Γ is not essentially crisp. Then one of the following holds:*

- $\mathrm{Cycl} \in \langle \Gamma \rangle$, *or*
- $(\neq)_0^1 \in \langle \Gamma \rangle$, *or*
- $R_{1,0,\infty} \in \langle \Gamma \rangle$.

Proof. Let R be a valued relation of Γ of arity k that attains at least two finite values. Let $m, \ell \in \mathbb{Q}$, with $m < \ell$, be the two smallest finite values attained by R. Let $t \in \mathbb{Q}^k$ be such that $R(t) = \ell$. Choose $s \in \mathrm{Opt}(R)$ so that $|(E_s \cap E_t) \cup (O_s \cap O_t)|$ is maximal (recall definition 4.1). By the definition of Opt, $R(s) = m$.

Let $\sim \subseteq (\mathbb{Q}^2)^2$ be the equivalence relation with the classes $=$, $<$, and $>$. Since $R(s) \neq R(t)$, there exist distinct i, j such that $(s_i, s_j) \not\sim (t_i, t_j)$. For the sake of notation, assume that $(i, j) = (1, 2)$. Let $\phi(x_1, x_2)$ be defined by

$$\min_{x_3, \dots, x_k} \left(R(x_1, \dots, x_k) + \sum_{(p,q) \in E_s \cap E_t} (=)_0^\infty(x_p, x_q) + \sum_{(p,q) \in O_s \cap O_t} (<)_0^\infty(x_p, x_q) \right).$$

Observe that $\phi(x, y) \geq m$ for all $x, y \in \mathbb{Q}$ and hence whenever $(x, y) \sim (s_1, s_2)$ we have $\phi(x, y) = m$. Let $(x, y) \sim (t_1, t_2)$. Then $\phi(x, y) \leq \ell$. By the choice of s, there is no $s' \in \mathrm{Opt}(R)$ that satisfies $(s'_1, s'_2) \sim (t_1, t_2)$, $(E_s \cap E_t) \subseteq E_{s'}$ and $(O_s \cap O_t) \subseteq O_{s'}$. Therefore, $\phi(x, y) > m$. It follows that $\phi(x, y) = \ell$.

Let

$$S(x, y) := \frac{1}{\ell - m}(\phi(x, y) - m).$$

By the construction, $S \in \langle\Gamma\rangle$, $S(x,y) = 0$ for $(x,y) \sim (s_1,s_2)$, and $S(x,y) = 1$ for $(x,y) \sim (t_1,t_2)$. Note that $\mathrm{Aut}(\mathbb{Q};<)$ has three orbits of pairs, two of which are represented by (s_1,s_2) and (t_1,t_2). Let $(u_1,u_2) \in \mathbb{Q}^2$ be a representative of the third orbit and let $\alpha = S(u_1,u_2)$. It follows that S is equal to one of the relations $R_{0,1,\alpha}$, $R_{0,\alpha,1}$, $R_{1,0,\alpha}$, $R_{1,\alpha,0}$, $R_{\alpha,0,1}$ or $R_{\alpha,1,0}$. By the choice of m and ℓ, we have that $\alpha = 0$ or $\alpha \geq 1$. By Lemma 4.25, this implies that $\mathrm{Cycl} \in \langle\Gamma\rangle$ unless $S = R_{1,0,0}$, $S = R_{1,0,\infty}$, or $S = R_{1,\infty,0}$. Since $R_{1,0,0}$ is equal to $(\neq)_0^1$ and $R_{1,0,\infty}(x,y) = R_{1,\infty,0}(y,x)$, the statement follows. □

We finish this section with two lemmas that will provide hardness criterions for the classification in Section 4.5.

Lemma 4.31. *Let Γ be a valued structure with $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\mathbb{Q};<)$ and $(<)_0^\infty \in \langle\Gamma\rangle$. Suppose that $\mathrm{lex} \notin \mathrm{Pol}(\mathbb{Q};\langle\Gamma\rangle_0^\infty)$ and that Γ is not essentially crisp. Then $\langle\Gamma\rangle$ contains one of the relations Betw, Cycl, Sep, T_3, $-T_3$, or Dis.*

Proof. Let $\Gamma' := (\mathbb{Q};\langle\Gamma\rangle_0^\infty)$. Note that $\mathrm{const} \notin \mathrm{Pol}(\Gamma')$, because const does not preserve $(<)_0^\infty$. If $\langle\Gamma\rangle_0^\infty$ contains one of the relations Betw, Cycl, Sep, T_3, $-T_3$, or Dis, then we are done. Suppose that this is not the case. Then $\mathrm{Pol}(\Gamma')$ contains min, mx, mi, ll, or one of their duals by Theorem 4.18. Suppose first that $\mathrm{Pol}(\Gamma')$ contains min, mx, mi, or ll. Since $\mathrm{lex} \notin \mathrm{Pol}(\Gamma')$, we have $\mathrm{ll} \notin \mathrm{Pol}(\Gamma')$ (Lemma 4.16). Then, by Proposition 4.21, $\mathrm{Pol}(\Gamma')$ contains $\pi\pi$. Note that Γ' is preserved by $\mathrm{Aut}(\mathbb{Q};<)$ and contains $(<)_0^\infty$, and therefore is a first-order expansion of $(\mathbb{Q};<)$. By Theorem 4.23, Γ' primitively positively defines, equivalently, contains the relation R^{mix}. By Lemma 4.30, we have that $\langle\Gamma\rangle$ contains Cycl, $(\neq)_0^1$, or $R_{1,0,\infty}$. If $\mathrm{Cycl} \in \langle\Gamma\rangle$, then we are done. Suppose therefore that $(\neq)_0^1 \in \langle\Gamma\rangle$ or $R_{1,0,\infty} \in \langle\Gamma\rangle$. Note that for every $x,y \in \mathbb{Q}$, we have

$$(<)_0^1(x,y) = \min_z \big(R^{\mathrm{mix}}(y,z,x) + (\neq)_0^1(y,z)\big) = \min_z \big(R^{\mathrm{mix}}(y,z,x) + R_{1,0,\infty}(y,z)\big).$$

Indeed, if $x < y$, then by choosing $z > y$ we get $R^{\mathrm{mix}}(y,z,x) + (\neq)_0^1(y,z) = R^{\mathrm{mix}}(y,z,x) + R_{1,0,\infty}(y,z) = 0$, which is clearly the minimal value that can be obtained. If $x \geq y$, then by choosing $z = y$ we get $R^{\mathrm{mix}}(y,z,x) + (\neq)_0^1(y,z) = R^{\mathrm{mix}}(y,z,x) + R_{1,0,\infty}(y,z) = 1$, which is clearly the minimal value, because if $z \neq y$ we obtain $R^{\mathrm{mix}}(y,z,x) = \infty$.

It follows that $(<)_0^1 \in \langle\Gamma\rangle$. Observe that $(<)_0^1$ equals $R_{1,0,1}$. Therefore, $\mathrm{Cycl} \in \langle\Gamma\rangle$ by Lemma 4.24, as we wanted to prove. If $\mathrm{Pol}(\Gamma')$ contains min^*, mx^*, mi^*, or ll^*, we use the dual versions of Proposition 4.21 and Theorem 4.23 to analogously prove that $(>)_0^1 \in \langle\Gamma\rangle$. Since $(<)_0^1(x,y) = (>)_0^1(y,x)$ we obtain $\mathrm{Cycl} \in \langle\Gamma\rangle$ by Lemma 4.24. □

Lemma 4.32. *Let Γ be a valued structure with $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(\mathbb{Q};<)$ and $(<)_0^\infty \in \langle\Gamma\rangle$. Suppose that $\mathrm{lex} \notin \mathrm{fPol}(\Gamma)$ and $\mathrm{lex} \in \mathrm{Pol}(\mathbb{Q};\langle\Gamma\rangle_0^\infty)$. Then $\mathrm{Cycl} \in \langle\Gamma\rangle$.*

Proof. Let R be a valued relation of Γ of arity k that is not improved by lex. Then there exist $s,t \in \mathbb{Q}^k$ such that

$$R(s) + R(t) < 2R(\mathrm{lex}(s,t)).$$

In particular, $R(s), R(t) < \infty$. Since $\mathrm{Feas}(R) \in \langle\Gamma\rangle_0^\infty$ is improved by lex, we have $R(\mathrm{lex}(s,t)) < \infty$. Let $u := \mathrm{lex}(s,t)$. Note that we must have $R(s) < R(u)$ or $R(t) <$

$R(u)$. Moreover, $E_u = E_s \cap E_t$. Let $v \in \{s,t\}$ be such that $R(v) < R(u) < \infty$. Note that we have $E_u \subseteq E_v$. Let O be a maximal subset of O_u such that there exists $w \in \mathbb{Q}^k$ satisfying

- $R(w) \leq R(v)$,
- $E_u \subseteq E_w$, and
- $O \subseteq O_w$,

and let w be any such witness for O. Such a maximal set O must exist, because v satisfies these conditions for $O = \emptyset$.

Since $R(w) \neq R(u)$ and $E_u \subseteq E_w$, there exist $i, j \in [k]$ such that $w_i \leq w_j$ and $u_i > u_j$. Without loss of generality me may assume $(i,j) = (1,2)$, because otherwise we permute the entries of R. Let $\phi(x_1, x_2)$ be defined by

$$\min_{x_3,\ldots,x_k} \left(R(x_1,\ldots,x_k) + \sum_{(p,q)\in E_u} (=)_0^\infty(x_p,x_q) + \sum_{(p,q)\in O} (<)_0^\infty(x_p,x_q) \right). \tag{4.2}$$

Let $a, b \in \mathbb{Q}$ such that $a < b$. Then $\phi(b,a) = \phi(u_1,u_2) \leq R(u)$, because $O \subseteq O_u$. Suppose that $\phi(b,a) \leq R(w)$. Then there exists $w' \in \mathbb{Q}^k$ such that $w'_1 > w'_2$ and $w'_3, \ldots, w'_k$ realize the minimum for $\phi(b,a)$ in (4.2) and hence $\phi(b,a) = R(w') \leq R(w) \leq R(v)$. In particular, the sums in (4.2) are finite. Therefore, $O \cup \{(2,1)\} \subseteq O_{w'}$ and $E_u \subseteq E_{w'}$. Since $(2,1) \in O_u \setminus O$, this contradicts the choice of O and w. Therefore, $\phi(b,a) > R(w)$. Note that $\phi(w_1,w_2) \leq R(w)$. If $w_1 < w_2$, then $\phi(a,b) \leq R(w)$ and ϕ expresses $R_{\alpha,\beta,\gamma}$ where $\beta = \phi(a,b)$ and $\gamma = \phi(b,a)$. In particular, $\beta < \gamma < \infty$. Therefore, by Lemma 4.25, Cycl $\in \langle\Gamma\rangle$. Otherwise we have $w_1 = w_2$. Then $\phi(a,a) \leq R(w)$ and ϕ expresses $R_{\alpha,\beta,\gamma}$ where $\alpha = \phi(a,a) \leq R(w) < \phi(b,a) = \gamma$. If $\beta \geq \gamma$, then $\alpha < \min(\beta,\gamma)$, and otherwise $\beta < \gamma < \infty$. In both cases, Cycl $\in \langle\Gamma\rangle$ by Lemma 4.25. □

4.5 Classification

In this section we generalize the complexity dichotomy from Theorem 4.8 to temporal VCSPs, which is the main result of this chapter. We first phrase the classification with 4 cases, where we distinguish between the tractable cases that are based on different algorithms. As a next step, we formulate two corollaries each of which provides two concise mutually disjoint conditions that correspond to NP-completeness and polynomial-time tractability, respectively.

Theorem 4.33. *Let Γ be a valued structure such that* $\mathrm{Aut}(\mathbb{Q};<) \subseteq \mathrm{Aut}(\Gamma)$. *Then at least one of the following holds:*

1. $\langle\Gamma\rangle$ *contains one of the relations* Betw, Cycl, Sep, T_3 *(see Definition 4.11),* $-T_3$, *or* Dis *(see Example 1.21). In this case,* Γ *has a reduct* Γ' *over a finite signature such that* VCSP(Γ') *is NP-complete.*

2. const $\in$ fPol(Γ). *In this case, for every reduct* Γ' *of* Γ *over a finite signature,* VCSP(Γ') *is in P.*

3. lex $\in$ fPol(Γ) *and* Pol($\hat{\Gamma}$) *contains* min*,* mx*,* mi*,* ll*, or one of their duals. In this case, for every reduct* Γ' *of* Γ *over a finite signature,* VCSP(Γ') *is in P.*

4. $\pi_1^2 \in$ fPol(Γ) *and* fPol(Γ) *contains* min*,* mx*,* mi*,* ll*, or one of their duals. In this case, for every reduct* Γ' *of* Γ *over a finite signature,* VCSP(Γ') *is in P.*

Proof. Note that for every reduct Γ' of Γ, the automorphism group Aut(Γ') contains Aut(Γ) and hence is oligomorphic. If $\langle\Gamma\rangle$ contains one of the relations Betw, Cycl, Sep, T_3, $-T_3$, or Dis, then there is a reduct Γ' of Γ over a finite signature such that VCSP(Γ') is NP-hard by Lemma 1.27 and Theorem 4.18. By Theorem 1.14, VCSP(Γ') is in NP, therefore it is NP-complete. If const $\in$ fPol(Γ), then const $\in$ fPol(Γ') for every reduct Γ' of Γ over a finite signature, and VCSP(Γ') is in P by Lemma 3.24. Suppose therefore that const $\notin$ fPol(Γ) and that $\langle\Gamma\rangle$ does not contain any of the relations Betw, Cycl, Sep, T_3, $-T_3$, or Dis. By Lemma 4.4, $\langle\Gamma\rangle$ contains $(\neq)_0^\infty$ or $(<)_0^\infty$, and hence const $\notin$ Pol($\mathbb{Q}; \langle\Gamma\rangle_0^\infty$). Recall that $\langle\Gamma\rangle_0^\infty$ contains all relations primitively positively definable in ($\mathbb{Q}; \langle\Gamma\rangle_0^\infty$) (Remark 1.32). By Theorem 4.18, Pol($\mathbb{Q}; \langle\Gamma\rangle_0^\infty$) (and thus Pol($\hat{\Gamma}$)) contains min, mx, mi, ll, or one of their duals.

Let Γ' be a reduct of Γ over a finite signature. If Aut(Γ) = Sym($\mathbb{Q}$), then by Theorem 4.8, inj $\in$ fPol(Γ) $\subseteq$ fPol(Γ') and VCSP(Γ') is in P. By Lemma 4.16, and since Aut(Γ) contains $x \mapsto -x$, we have lex $\in$ fPol(Γ) and therefore satisfy (3).

Finally, suppose that Aut(Γ) $\neq$ Sym($\mathbb{Q}$). By Lemma 4.26 we have $(<)_0^\infty \in \langle\Gamma\rangle$, and hence Aut($\Gamma$) = Aut($\mathbb{Q}; <$). By Lemma 4.31 we have that Γ is essentially crisp or lex $\in$ Pol($\mathbb{Q}; \langle\Gamma\rangle_0^\infty$). If Γ is not essentially crisp, we have lex $\in$ Pol($\mathbb{Q}; \langle\Gamma\rangle_0^\infty$), and lex $\in$ fPol(Γ) $\subseteq$ fPol(Γ') by Lemma 4.32. Then VCSP(Γ') is in P by Lemma 4.29 and Condition (3) holds. Suppose that Γ is essentially crisp. Then by Lemma 3.23 we have $\pi_1^2 \in$ fPol(Γ). Since const $\notin$ fPol(Γ), we have const $\notin$ Pol(Feas(Γ)) (see Remark 3.20). Since $\langle\Gamma\rangle$ does not contain any of the relations Betw, Cycl, Sep, T_3, $-T_3$, or Dis, none of these relations are primitively positively definable in Feas(Γ). By Theorem 4.18, Pol(Feas(Γ)) $\subseteq$ Pol(Feas(Γ')) contains min, mx, mi, ll, or one of their duals and CSP(Feas(Γ')) is in P. By Remark 1.32, VCSP(Γ') is in P. By Remark 3.20, fPol(Γ) contains min, mx, mi, ll, or one of their duals. Therefore, (4) holds. □

Recall from Section 1.3 that a valued structure Γ with Aut($\mathbb{Q}; <$) $\subseteq$ Aut(Γ) has a quantifier-free first-order definition in Aut($\mathbb{Q}; <$) with the defining formulas being disjunctions of conjunctions of atomic formulas over ($\mathbb{Q}; <$). We continue by proving that the complexity dichotomy we gave in Theorem 4.33 is decidable, using the representation of Γ by a first-order definition in ($\mathbb{Q}; <$) of the form described above.

Remark 4.34. *We also obtain decidability if arbitrary first-order formulas may be used for defining the valued relations, because every first-order formula can be effectively transformed into a quantifier-free formula. This holds more generally over so-called* finitely bounded homogeneous *structures; see, e.g., [75, Proposition 7]. Without the finite boundedness assumption, the problem can become undecidable [28].*

Proposition 4.35. *Given a first-order definition of a valued structure* Γ *with a finite signature in* $(\mathbb{Q}; <)$*, it is decidable whether* VCSP(Γ) *is in P or NP-complete.*

Proof. Recall that if Γ has a first-order definition in $(\mathbb{Q}; <)$, then $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$ and, in particular, VCSP(Γ) is in NP by Theorem 1.14. By Theorem 4.33, VCSP(Γ) is in P or NP-complete. If P = NP, then the decision problem is trivial. Suppose that P $\neq$ NP. Then in the statement of Theorem 4.33, item (1) and the union of (2)–(4) is disjoint. Since Γ has a finite signature, we can decide whether const improves Γ, i.e., whether (2) holds. Similarly, we can decide whether lex improves Γ. By the last sentence of Theorem 4.18 applied to $\hat{\Gamma}$ we can decide whether one of the operations min, mx, mi, ll, or one of their duals preserves $\hat{\Gamma}$. Therefore, we can decide whether (3) holds. Finally, we can decide whether π_1^2 improves Γ. If yes, Γ is essentially crisp by Lemma 3.23. In this case fPol(Γ) contains min, mx, mi, ll, or one of their duals if and only if Pol(Feas(Γ)) does (Proposition 3.22 and Remark 3.20), which can be decided by Theorem 4.18. It follows that we can decide whether union of (2)–(4) holds, which implies the statement. □

We reformulate Theorem 4.33 with two mutually exclusive cases that capture the respective complexities of the VCSPs.

Corollary 4.36. *Let* Γ *be a valued structure such that* $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$ *and let the signature of* Γ *be finite. Then exactly one of the following holds.*

1. $\langle \Gamma \rangle$ *contains one of the relations* Betw, Cycl, Sep, T_3, $-T_3$, *or* Dis. *In this case,* VCSP(Γ) *is NP-complete.*

2. $(\mathbb{Q}; \langle \Gamma \rangle_0^\infty)$ *is preserved by one of the operations* const, min, mx, mi, ll, *or one of their duals. In this case,* VCSP(Γ) *is in P.*

Proof. Let $\Gamma' := (\mathbb{Q}; \langle \Gamma \rangle_0^\infty)$. Theorem 4.18 states that either Γ' primitively positively defines one of the relations Betw, Cycl, Sep, T_3, $-T_3$, Dis, or Pol(Γ') contains const, min, mx, mi, ll, or one of their duals. Clearly, Γ' primitively positively defines a relation R is and only if $R \in \langle \Gamma \rangle_0^\infty$, which is the case if and only if $R \in \langle \Gamma \rangle$.

It remains to discuss the implications for the complexity of VCSP(Γ). If (1) holds, then VCSP(Γ) is NP-complete by Theorem 4.33. On the other hand, if (1) does not hold, one of the cases (2)–(4) in Theorem 4.33 applies and VCSP(Γ) is in P. □

Note that the corollary above implies that if $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$, then the complexity of VCSP(Γ) is up to polynomial-time reductions determined by the complexity of the crisp relations Γ can express. Loosely speaking, the complexity of such a VCSP is determined solely by the CSPs that can be encoded in this VCSP via expressibility. We formulate an alternative and more concise variant of the previous result.

Corollary 4.37. *Let* Γ *be a valued structure such that* $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$ *and let the signature of* Γ *be finite. Then exactly one of the following holds.*

1. Γ *pp-constructs* K_3. *In this case,* VCSP(Γ) *is NP-complete.*

2. $(\mathbb{Q}; \langle\Gamma\rangle_0^\infty)$ *has a pwnu polymorphism. In this case,* $\mathrm{VCSP}(\Gamma)$ *is in P.*

Proof. Let $\Gamma' := (\mathbb{Q}; \langle\Gamma\rangle_0^\infty)$. By Theorem 1.14, $\mathrm{VCSP}(\Gamma)$ is in NP. By Proposition 2.22, Γ pp-constructs K_3 if and only if Γ' pp-constructs K_3 and in this case, $\mathrm{VCSP}(\Gamma)$ is NP-complete by Corollary 2.17. Hence, it follows from Theorem 4.19 applied on Γ' that either (1) holds or $(\mathbb{Q}; \langle\Gamma\rangle_0^\infty)$ has a pwnu polymorphism. Hence, if $(\mathbb{Q}; \langle\Gamma\rangle_0^\infty)$ has a pwnu polymorphism, then Γ does not pp-construct K_3. By Proposition 4.20 and Theorem 4.3, Betw, Cycl, Sep, T_3, $-T_3$, Dis $\notin \langle\Gamma\rangle$ and therefore, item (2) from Corollary 4.36 applies and $\mathrm{VCSP}(\Gamma)$ is in P. □

Conjecture 2.24 states that, under some structural assumptions on Γ, $\mathrm{VCSP}(\Gamma)$ is in P whenever Γ does not pp-construct K_3 (and is NP-hard otherwise). All temporal structures satisfy the assumptions of the conjecture and hence Corollary 4.37 confirms the conjecture for the class of temporal VCSPs.

4.6 Discussion and open questions

We proved a complexity dichotomy for temporal VCSPs: $\mathrm{VCSP}(\Gamma)$ is in P or NP-complete for every valued structure Γ such that $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$. Moreover, we showed that the meta-problem of deciding whether $\mathrm{VCSP}(\Gamma)$ is in P or NP-complete for a given Γ is decidable. As a side product of our proof, we obtain that the complexity of every such VCSP is captured by the classical relations that it can express, in other words, by the CSPs that are encoded in this VCSP. Our results confirm Conjecture 2.24 for all temporal valued structures.

The proof of our decidability result (Proposition 4.35) is based on the distinction of two cases depending on whether P=NP. Typical results on decidability of such meta-problems in the theory of (V)CSPs are rather formulated by deciding the algebraic conditions that imply the respective complexities, more concretely, deciding the presence of certain (fractional) polymorphisms. This can often be checked by the naive approach, as long as the signature of the structure is finite. However, if we wanted to do so in our case, we would have to check for polymorphisms of the structure $(\mathbb{Q}; \langle\Gamma\rangle_0^\infty)$, which has an infinite signature by definition. This motivates the following question.

Question 4.38. *Let* Γ *be a valued structure such that* $\mathrm{Aut}(\mathbb{Q}; <) \subseteq \mathrm{Aut}(\Gamma)$ *and let the signature of* Γ *be finite. Given a first-order definition of* Γ *in* $(\mathbb{Q}; <)$*, is it decidable whether* Γ *pp-constructs* K_3*, equivalently, whether item (2) in Corollary 4.36 holds?*

In analogy to the development of the results on infinite-domain CSPs, we propose the class of valued structures that are preserved by all automorphisms of the countable random graph (see Example 1.17) as a natural next step in the complexity classification of VCSPs on infinite domains.

Question 4.39. *Does the class of VCSPs of all valued structures* Γ *over a finite signature such that* $\mathrm{Aut}(\Gamma)$ *contains the automorphism group of the countable random graph exhibit a P vs. NP-complete dichotomy? In particular, is* $\mathrm{VCSP}(\Gamma)$ *in P whenever* Γ *does not pp-construct* K_3*?*

A positive answer to the second question in Question 4.39 would confirm Conjecture 2.24 for valued structures preserved by all automorphisms of the countable random graph.

Chapter 5

Resilience

In this chapter, we study 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 resilience problem is parameterized by a query μ. The input to the problem is a finite database $\mathfrak{A}$ and the question is how many tuples need to be removed from the 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). We introduce the problem formally in the next section.

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 [45, 46, 65]. 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. In this chapter, we present a surprising link between the resilience problem for (unions of) conjunctive queries under bag semantics and VCSPs and show how the theory of VCSPs can be applied to classify complexity of resilience problems. The original results in this chapter have been published in [30].

5.1 Conjunctive queries and bag databases

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. 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 relational τ-structure $\mathfrak{A}$ with a finite domain, 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 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

[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 thesis.

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 5.1: The resilience problem considered in this paper.

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*; this variant of the problem appears first in [65]. Formally, 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 this thesis, we focus on bag databases, which are of importance because they represent SQL databases more faithfully than set databases [37]. 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 on set databases as well. Regarding the converse, Makhija and Gatterbauer [65] 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.

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 5.9). The resilience problem that we study is summarized in Figure 5.1. Note that this problem is always in NP independently from μ.

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 $\bar{x} \in R^{\mathfrak{A}}$ for $R \in \tau$ if and only if μ contains the conjunct $R(\bar{x})$. 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 $t \in R^{\mathfrak{A}}$ the conjunct $R(t)$.

Remark 5.1. *All terminology introduced for τ-structures also applies to conjunctive queries with signature τ: by definition, the query has the property if the canonical*

Figure 5.2: The query μ from Example 5.2 (on the left) and the corresponding structure $\mathfrak{B}$ (on the right).

database has the property.

We now give an example of how a resilience problem can be represented as a VCSP using an appropriately chosen valued structure.

Example 5.2. *The following query is taken from [66]; the authors 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 5.2). We turn $\mathfrak{B}$ into the valued structure Γ with domain $\{0, 1\}$ where $R^{\Gamma} = (R^{\mathfrak{B}})^1_0$ and $S^{\Gamma} = (S^{\mathfrak{B}})^1_0$; note that Γ is the valued structure Γ from Example 2.6. Then VCSP(Γ) *is precisely the resilience problem for μ (with multiplicities). We will reprove the result from [65] that even with multiplicities, the problem can be solved in polynomial time (see Example 5.14).*

Example 5.3. *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 [66]; also see Fact 3.18 in [44]). Our results reprove the result from [65] that this problem remains polynomial-time solvable with multiplicities (see Example 5.20).*

Remark 5.4. *Consider the computational problem of* finding *tuples to be removed from the input database $\mathfrak{A}$ so that $\mathfrak{A} \not\models \mu$. We observe that if the resilience problem (with or without multiplicities) for a union μ of conjunctive queries is in P, then this problem also is in P. To see this, let $u \in \mathbb{N}$ be threshold. If $u = 0$, then no tuple needs to be found and we are done. Otherwise, for every tuple t in a relation $R^{\mathfrak{A}}$, we remove all copies of t from $R^{\mathfrak{A}}$ and test the resulting database with the threshold $u - m$, where m is the multiplicity of t. If the modified instance is accepted, then t is a correct tuple to be removed and we may proceed to find a solution of this modified instance. Otherwise we return a step back and try to remove a different tuple.*

5.2 Connectivity

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

Lemma 5.5. *Let $\nu_1, \dots, \nu_k$ be conjunctive queries such that ν_i does not imply ν_j if $i \neq j$. Then the resilience problem for $\nu := \nu_1 \wedge \cdots \wedge \nu_k$ is NP-hard if the resilience problem for one of the ν_i is NP-hard. Conversely, if the resilience problem is in P for each ν_i, then the resilience problem for ν is in P as well. The same is true in the setting without multiplicities and/or exogeneous relations.*

Proof. We first present a polynomial-time reduction from the resilience problem of ν_i, for some $i \in \{1, \dots, k\}$, to the resilience problem of ν. Given an instance $\mathfrak{A}$ of the resilience problem for ν_i, let m be the number of tuples in relations of $\mathfrak{A}$. Let $\mathfrak{A}'$ be the disjoint union of $\mathfrak{A}$ with m copies of the canonical database of ν_j for every $j \in \{1, \dots, k\} \setminus \{i\}$. Observe that $\mathfrak{A}'$ can be computed in polynomial time in the size of $\mathfrak{A}$ and that the resilience of $\mathfrak{A}$ with respect to ν_i equals the resilience of $\mathfrak{A}'$ with respect to ν.

Conversely, if the resilience problem is in P for each ν_i, then also the resilience problem for ν is in P: given an instance $\mathfrak{A}$ of the resilience problem for ν, we compute the resilience of $\mathfrak{A}_j$ with respect to ν_i for every $i \in \{1, \dots, k\}$, and the minimum of all the resulting values.

The same proof works in the setting without multiplicities. $\square$

Corollary 5.6. *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 for each μ_i, then the resilience problem for μ is in P as well. The same is true in the setting without multiplicities and/or exogeneous relations.*

Proof. Follows immediately from Lemma 5.5. $\square$

By applying Corollary 5.6 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.

5.3 Translating resilience problems to VCSPs

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. The conjunctive query in Example 5.2, for instance, even has a *finite* dual, namely, the structure $\mathfrak{B}$ from the same example. As will follow from Theorem 5.17, every connected μ has a countable dual; will discuss the existence of duals in detail in Section 5.4 and 5.5. To construct valued structures from duals, we introduce the following notation.

Definition 5.7. *Let $\mathfrak{B}$ be a relational 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 $R \in \tau \setminus \sigma$, $R^{\Gamma(\mathfrak{B},\sigma)} := (R^{\mathfrak{B}})_0^1$, and*
- *for each $R \in \sigma$, $R^{\Gamma(\mathfrak{B},\sigma)} := (R^{\mathfrak{B}})_0^\infty$.*

Note that $\text{Aut}(\mathfrak{B}) = \text{Aut}(\Gamma(\mathfrak{B}, \sigma))$ for any relational τ-structure $\mathfrak{B}$ and any σ. The following proposition shows how a dual of μ provides a valued constraint satisfaction problem that is polynomial-time equivalent to the resilience problem for μ.

Proposition 5.8. *Let μ be a union of connected conjunctive queries with a signature τ and $\sigma \subseteq \tau$. Then the resilience problem for μ where the relations from σ are exogenous is polynomial-time equivalent to* $\text{VCSP}(\Gamma(\mathfrak{B}, \sigma))$ *for any dual $\mathfrak{B}$ of μ.*

Proof. Let $\mathfrak{B}$ be a dual of μ. For every bag database $\mathfrak{A}$ with a signature τ and with exogenous relations from σ, let ϕ be the τ-expression obtained by adding atomic τ-expressions $S(x_1, \ldots, x_n)$ according to the multiplicity of the tuples $(x_1, \ldots, x_n)$ in $S^{\mathfrak{A}}$ for all $S \in \tau$. Note that ϕ can be computed in polynomial time. Then the resilience of $\mathfrak{A}$ with respect to μ is at most u if and only if (ϕ, u) has a solution over $\Gamma(\mathfrak{B}, \sigma)$.

To prove a polynomial-time reduction in the other direction, let ϕ be a τ-expression. We construct a bag database $\mathfrak{A}$ with a signature τ. The domain of $\mathfrak{A}$ are the variables that appear in ϕ and for every $S \in \tau$, we put a tuple $(x_1, \ldots, x_n)$ in $S^{\mathfrak{A}}$ with a multiplicity equal to the occurences of $S(x_1, \ldots x_n)$ as a summand of ϕ. The relations $S^{\mathfrak{A}}$ with $S \in \sigma$ are exogenous in $\mathfrak{A}$, the remaining ones are endogenous. Again, $\mathfrak{A}$ can be computed in polynomial time and the resilience of $\mathfrak{A}$ with respect to μ is at most u if and only if (ϕ, u) has a solution over $\Gamma(\mathfrak{B}, \sigma)$. ☐

In [65] one may find a seemingly more general notion of exogenous tuples in resilience problems, where in a single relation there might be both endogenous and exogenous tuples. However, using the operator Opt and a similar reduction as in Proposition 5.8, one can show that classifying the complexity of resilience problems according to our original definition also entails a classification of this variant.

Remark 5.9. *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 $(R^x)^\Gamma = (R^{\mathfrak{B}})_0^\infty$, i.e., a crisp relation that takes values 0 and ∞. The resilience problem for μ with exogenous tuples specified as above is polynomial-time equivalent to* $\text{VCSP}(\Gamma)$ *by analogous reductions as in the proof of Proposition 5.8. Note that $(R^x)^\Gamma = \text{Opt}\left(R^{\Gamma(\mathfrak{B},\sigma)}\right)$ for every $R \in \rho$, and therefore by Lemma 1.27,* $\text{VCSP}(\Gamma)$ *is polynomial-time equivalent to* $\text{VCSP}(\Gamma(\mathfrak{B}, \sigma))$ *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 (Proposition 5.8). 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.*

5.4 Finite duals

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 5.10. *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 $R(t)$ where $R \in \tau$ has arity k, $t \in A^k$, and $\mathfrak{A} \models R(t)$. An edge $e_{a,i,R(t)}$ joins $a \in A$ with $R(t)$ if $t_i = a$. A relational 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 5.2.*

The following theorem characterizes the existence of finite duals.

Theorem 5.11 ([72]; see also [43, 62])**.** *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 the query μ from Example 5.3 does not have a finite dual, since the query given there is not incidence-acyclic and hence cannot be homomorphically equivalent to a tree.

In the following result we combine the correspondence between resilience problems and VCSPs (Proposition 5.8) with the finite-domain VCSP dichotomy theorem (Theorem 3.36), obtaining a complexity dichotomy for resilience problems of unions of incidence-acyclic conjunctive queries.

Theorem 5.12. *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 μ is in P.*

If μ is a union of conjunctive queries, each of which is homomorphically equivalent to a tree and $\mathfrak{B}$ is a finite dual of μ (which exists by Theorem 5.11), then exactly one of the following applies:

- $\Gamma(\mathfrak{B}, \sigma)$ *has a cyclic fractional polymorphism. In this case, the resilience problem for μ with exogeneous relations from σ is in P.*
- $\Gamma(\mathfrak{B}, \sigma)$ *pp-constructs K_3. In this case, the resilience problem for μ with exogeneous relations from σ is NP-complete.*

Proof. By virtue of Corollary 5.6, we may assume for the P versus NP-complete dichotomy that each of the conjunctive queries in μ is connected and thus a tree. The same is true also for the algebraic dichotomy since replacing a conjunctive query in a union with a homomorphically equivalent one does not affect the complexity of resilience. Let $\Gamma := \Gamma(\mathfrak{B}, \sigma)$. By Proposition 5.8, VCSP($\Gamma(\mathfrak{B}, \sigma)$) is polynomial-time equivalent to the resilience problem for μ with exogeneous relations from σ. The algebraic and complexity dichotomy therefore follows from Theorem 3.36.

Concerning the decidability of the tractability condition, it is known that the finite dual of μ, and hence also Γ, can be effectively computed from μ (e.g., the construction of the dual in [72] is effective). The existence of a fractional cyclic polymorphism for a given valued structure Γ with finite domain and finite signature can be decided (see Remark 3.37). □

Remark 5.13. *We mention that Theorem 5.12 also applies to (2-way) regular path queries, which can be shown to always have a finite dual, more details can be found in [31, Appendix B].*

5.4.1 Examples

We give a few examples to illustrate the use of Theorem 5.12. We first revisit a known tractable resilience problem from [44–46, 66] and show that the corresponding valued structure has a fractional cyclic polymorphism.

Example 5.14. *We revisit Example 5.2. Consider again the conjunctive query*

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

There is a finite dual $\mathfrak{B}$ of μ with domain $\{0, 1\}$, as described in Example 5.2 and the example also describes a valued structure Γ, which is in fact $\Gamma(\mathfrak{B}, \emptyset)$.

Let ω be the fractional operation given by $\omega(\min_2) = \omega(\max_2) = \frac{1}{2}$. The operation ω is cyclic, see Example 3.28. Finally observe that in Example 2.6 and 3.15 we proved that ω is a fractional polymorphism of Γ. Hence, ω is a cyclic fractional polymorphism of Γ and by Theorem 5.12 the resilience problem for μ is in P, which reproves the results from [45] (without multiplicities) and [65] (with multiplicities).

In the following example we generalize the approach from Example 5.14.

Example 5.15. *Consider for $n \geq 1$ the conjunctive query*

$$\mu := \exists x_0, \ldots, x_n(R_1(x_0, x_1) \wedge \cdots \wedge R_n(x_{n-1}, x_n)).$$

Let $\tau = \{R_1, \ldots, R_n\}$ be the signature of μ. We describe a finite dual $\mathfrak{B}$ of μ. The domain B consists of 2^{n-1} elements v_S, which we index by a subset S of $[n-1]$. Then the dual satisfies $R_i(v_S, v_T)$ for $S, T \subseteq [n-1]$ and $i \in \{1, \ldots, n\}$ if and only if

- *$i < n$ and $i \notin T$, or*

- $i > 1$ *and* $(i-1) \in S$.

To see that $\mathfrak{B}$ *is indeed a dual of* μ, *let* $\mathfrak{A}$ *be a finite relational* τ*-structure that does not satisfy* μ, *and let* $a \in A$. *Let* S_a *be the set of all elements* $i \in [n-1]$ *such that there exist elements* $a_{i+1}, \dots, a_n \in A$ *such that*

$$\mathfrak{A} \models R_{i+1}(a, a_{i+1}) \wedge R_{i+2}(a_{i+1}, a_{i+2}) \wedge \cdots \wedge R_n(a_{n-1}, a_n).$$

We claim that $a \mapsto v_{S_a}$, $a \in A$, *defines a homomorphism* h *from* $\mathfrak{A}$ *to* $\mathfrak{B}$.

Suppose that $a, b \in A$ *are such that* $R_j(h(a), h(b))$ *does not hold in* $\mathfrak{B}$ *for some* $j \in [n]$. *This implies that*

- $j = n$ *and* $(n-1) \notin S_a$, *or*
- $j = 1$ *and* $1 \in S_b$, *or*
- $j \in \{2, \dots, n-1\}$, $(j-1) \notin S_a$, *and* $j \in S_b$.

In all three cases, the definition of S_a *and* S_b *and* $\mathfrak{A} \not\models \mu$ *imply that* $R_j(a, b)$ *does not hold in* $\mathfrak{A}$. *This shows that* h *is a homomorphism.*

Now suppose that $\mathfrak{A}$ *has a homomorphism to* $\mathfrak{B}$. *We have to show that* $\mathfrak{A}$ *does not satisfy* μ. *It suffices to show that* $\mathfrak{B}$ *does not satisfy* μ. *Suppose for contradiction that there are elements* $v_{S_0}, \dots, v_{S_n}$ *such that* $\mathfrak{B}$ *satisfies* $R_1(v_{S_0}, v_{S_1}) \wedge \cdots \wedge R_n(v_{S_{n-1}}, v_{S_n})$. *Then* $R_n(v_{S_{n-1}}, v_{S_n})$ *implies that* $(n-1) \in S_{n-1}$. *By induction, we obtain from* $R_i(v_{S_{i-1}}, v_{S_i})$ *that* $(i-1) \in S_{i-1}$ *for every* $i \in \{2, \dots, n\}$. *In particular,* $1 \in S_1$. *However, note that* $R_1(v_{S_0}, v_{S_1})$ *by definition implies that* $1 \notin S_1$. *We reached a contradiction, hence,* $\mathfrak{B}$ *does not satisfy* μ.

Let f *be the symmetric binary operation that maps* (v_S, v_T) *to* $v_{S \cup T}$, *and let* g *be the symmetric binary operation that maps* (v_S, v_T) *to* $v_{S \cap T}$. *Define the binary symmetric fractional operation* ω *by setting* $\omega(f) = \omega(g) = \frac{1}{2}$. *We claim that* ω *is a fractional polymorphism of* $\Gamma(\mathfrak{B}, \emptyset)$.

Let $i \in \{1, \dots, n\}$. *We have to show that for all* $v_S, v_T, v_P, v_Q \in B$, *it holds in* Γ *that*

$$\begin{aligned} R_i(v_S, v_T) + R_i(v_P, v_Q) &\geq R_i(f(v_S, v_P), f(v_T, v_Q)) + R_i(g(v_S, v_P), g(v_T, v_Q)) \\ &= R_i(v_{S \cup P}, v_{T \cup Q}) + R_i(v_{S \cap P}, v_{T \cap Q}). \end{aligned} \tag{5.1}$$

First suppose that $1 < i < n$. *If the left-hand side in* (5.1) *is equal to* 2, *then the disequality holds trivially. Whenever the left-hand side in* (5.1) *is equal to* 1, *we must have* $(i-1) \in S$, $(i-1) \in P$, $i \notin T$, *or* $i \notin Q$. *Hence* $(i-1) \in S \cup P$ *or* $i \notin T \cap Q$ *and* (5.1) *holds. Finally, if the left-hand side in* (5.1) *is* 0, *then* $R_i(v_S, v_T) = R_i(v_P, v_Q) = 0$. *Thus we have* $(i-1) \in S$ *or* $i \notin T$ *and at the same time* $(i-1) \in P$ *or* $i \notin Q$. *Therefore, we have* $(i-1) \in S \cap P$, $i \notin T \cup Q$, *or both* $(i-1) \in S \cup P$ *and* $i \notin T \cap Q$. *It follows that* $R_i(v_{S \cup P}, v_{T \cup Q}) = R_i(v_{S \cap P}, v_{T \cap Q}) = 0$. *The case that* $i = 1$ *and the case that* $i = n$ *can be treated similarly. Therefore,* ω *is a cyclic fractional polymorphism of* $\Gamma(\mathfrak{B}, \emptyset)$, *which by Theorem 5.12 implies that the resilience problem for* μ *is in P.*

We finish this section with an example of an NP-complete resilience problem, where the dual of the query is finite.

Example 5.16. *Let*

$$\mu := \exists x, y, z(R(x, y) \wedge R(y, z)).$$

There is a finite dual $\mathfrak{B}$ *of* μ*:* $\mathfrak{B} = (\{0,1\}; R^{\mathfrak{B}})$ *where* $R^{\mathfrak{B}} = \{(0,1)\}$*. Note that* $\Gamma(\mathfrak{B}, \emptyset) = \Gamma_{max}$ *from Example 1.2. We have seen in Example 2.18 that* Γ_{max} *pp-constructs* $(\{0,1\}; \mathrm{OIT})$ *and hence by Theorem 5.12, the resilience of* μ *is NP-complete.*

5.5 Infinite duals

As shown in Section 5.4, conjunctive queries might not have a finite dual (see Theorem 5.11 and Example 5.3), but unions of connected conjunctive queries always have a countably infinite dual. Cherlin, Shelah and Shi [38] showed that in this case we may even find a dual with an oligomorphic automorphism group (see Theorem 5.17 below). This is the key insight to utilize Proposition 5.8 to phrase resilience problems as VCSPs of valued structures with oligomorphic automorphism groups. The not necessarily connected case again reduces to the connected case by Corollary 5.6.

In Theorem 5.17 below we state a variant of a theorem of Cherlin, Shelah, and Shi [38] (also see [9, 53, 54]). If $\mathfrak{B}$ is a relational 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 5.1) 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 5.17. *For every union* μ *of connected conjunctive queries over a finite relational signature* τ *there exists a relational* 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.*

Proof. The construction of a structure $\mathfrak{B}_\mu$ with the given properties follows from a proof of Hubička and Nešetřil [53, 54] of the theorem of Cherlin, Shelah, and Shi [38], and can be found in [9, Theorem 4.3.8]. Properties (1), (2) and property (3) restricted to finite structures $\mathfrak{A}$ are explicitly stated in [9, Theorem 4.3.8]. Property (3) restricted to finite structures clearly implies property (4). Property (5) holds because reducts of homogeneous structures with a finite relational signature have an oligomorphic automorphism group. Property (3) for countable structures now follows from Lemma 2.5.

Since we are not aware of a reference for (6) in the literature, we present a proof here. Let σ be the signature of $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)}$. We claim that the following universal σ-sentence ψ describes the structures in the age of $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)}$. If ϕ is a σ-sentence, then ϕ' denotes the τ-sentence obtained from ϕ by replacing every occurrence of $R(\bar{x})$, for $R \in \sigma \setminus \tau$, by the primitive positive τ-formula $\eta(\bar{x})$ for which R was introduced in $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)}$. Then ψ is a conjunction of all σ-sentences $\neg\phi$ such that ϕ is primitive positive, ϕ' has at most $|\mu|$ variables, and ϕ' implies μ. Clearly, there are finitely many conjuncts of this form.

Suppose that $\mathfrak{A} \in \mathrm{Age}(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)}$. Then $\mathfrak{A}$ satisfies each conjunct $\neg\phi$ of ψ, because otherwise $\mathfrak{B}_\mu$ satisfies ϕ', and thus satisfies μ, contrary to our assumptions.

The interesting direction is that if a finite σ-structure $\mathfrak{A}$ satisfies ψ, then $\mathfrak{A}$ embeds into $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)}$. Let ϕ be the canonical query of $\mathfrak{A}$. Let $\mathfrak{A}'$ be the canonical database of the τ-sentence ϕ'. Suppose for contradiction that $\mathfrak{A}' \models \mu$. Let χ be a minimal subformula of ϕ' such that the canonical database of χ models μ. Then χ has at most $|\mu|$ variables and implies μ, and hence $\neg\chi$ is a conjunct of of ψ which is not satisfied by $\mathfrak{A}$, a contradiction to our assumptions. Therefore, $\mathfrak{A}' \models \neg\mu$ and by Property (2), we have that $\mathfrak{A}'_{\mathrm{pp}(|\mu|)}$ has an embedding f into $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)}$.

We claim that the restriction of f to the elements of $\mathfrak{A}$ is an embedding of $\mathfrak{A}$ into $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)}$. Clearly, if $\mathfrak{A} \models R(\bar{x})$ for some relation R that has been introduced for a primitive positive formula η, then $\mathfrak{A}'$ satisfies $\eta(\bar{x})$, and hence $\mathfrak{B}_\mu \models \eta(f(\bar{x}))$, which in turn implies that $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)} \models R(f(\bar{x}))$ as desired. Conversely, if $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)} \models R(f(\bar{x}))$, then $\mathfrak{A}'_{\mathrm{pp}(|\mu|)} \models R(\bar{x})$, and hence $\mathfrak{A}' \models \eta(\bar{x})$. This in turn implies that $\mathfrak{A} \models R(\bar{x})$. Since the restriction of f and its inverse preserve the relations from τ trivially, we conclude that $\mathfrak{A}$ embeds into $(\mathfrak{B}_\mu)_{\mathrm{pp}(|\mu|)}$. □

By Properties (1) and (6) of Theorem 5.17, $\mathfrak{B}_\mu$ is always a reduct of a finitely bounded homogeneous structure and hence $\mathrm{Aut}(\mathfrak{B}_\mu)$ contains the automorphism group of such structure. By (3), $\mathfrak{B}_\mu$ is a dual of μ. For short, we write Γ_μ for $\Gamma(\mathfrak{B}_\mu, \emptyset)$ and $\Gamma_{\mu,\sigma}$ for $\Gamma(\mathfrak{B}_\mu, \sigma)$ (see Definition 5.7). The existence of the structure $\mathfrak{B}_\mu$ motivates the following corollary of Proposition 5.8 and Theorem 3.53.

Corollary 5.18. *Let μ be a union of connected conjunctive queries over a finite signature τ and $\sigma \subseteq \tau$. Let $\mathfrak{B}$ be a reduct of a finitely bounded homogeneous structure $\mathfrak{A}$ with a countable domain. If $\mathfrak{B}$ is a dual of μ and $\Gamma(\mathfrak{B}, \sigma)$ has a canonical and pseudo cyclic fractional polymorphism with respect to* $\mathrm{Aut}(\mathfrak{A})$*, then the resilience problem for μ with exogeneous relations from σ is in P.*

Proof. The resilience problem for μ with exogeneous relations from σ is polynomial-time equivalent to $\mathrm{VCSP}(\Gamma(\mathfrak{B}, \sigma))$ by Proposition 5.8. By the assumption, $\mathrm{Aut}(\Gamma(\mathfrak{B}, \sigma)) = \mathrm{Aut}(\mathfrak{B})$ contains $\mathrm{Aut}(\mathfrak{A})$. By Theorem 3.53, $\mathrm{VCSP}(\Gamma(\mathfrak{B}, \sigma))$ is in P. □

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 homogeneous itself, which often simplifies the arguments. 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 5.19. *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 relational τ-structure $\mathfrak{C}_\mu$ such that the following statements hold:*

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

Moreover, a countable τ-structure satisfies $\neg\mu$ if and only if it embeds into $\mathfrak{C}_\mu$.

Proof. Let $\mathfrak{A}_1$ and $\mathfrak{A}_2$ be finite τ-structures that satisfy $\neg\mu$ such that the substructure induced by $A_1 \cap A_2$ in $\mathfrak{A}_1$ and $\mathfrak{A}_2$ is the same. Since the Gaifman graph of each of the conjunctive queries in μ is complete, the union of the structures $\mathfrak{A}_1$ and $\mathfrak{A}_2$ satisfies $\neg\mu$ as well. By Fraïssé's Theorem (see, e.g., [51]) there is a countable homogeneous τ-structure $\mathfrak{C}_\mu$ such that $\mathrm{Age}(\mathfrak{C}_\mu)$ is the class of all finite structures that satisfy $\neg\mu$; this shows that $\mathfrak{C}_\mu$ is finitely bounded. The final statement follows from Lemma 2.5. $\square$

Note that $\mathrm{Aut}(\mathfrak{C}_\mu)$ is oligomorphic and $\mathfrak{C}_\mu$ is a dual of μ. By Lemma 2.5, $\mathfrak{C}_\mu$ is homomorphically equivalent to $\mathfrak{B}_\mu$. Therefore, $\Gamma(\mathfrak{C}_\mu, \sigma)$ is homomorphically equivalent to $\Gamma_{\mu,\sigma}$ for any $\sigma \subseteq \tau$. We remark that for every μ satisfying the assumptions of Theorem 5.19, there is a unique dual with these properties, up to isomorphism, see, e.g., [9, Section 2.3].

5.5.1 Examples

We demonstrate how to use Corollary 5.18 to show tractability for a resilience problem.

Example 5.20. *We revisit Example 5.3. Consider the conjunctive query*

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

over the signature $\tau = \{R, S\}$. Note that the Gaifman graph of μ is complete; let $\mathfrak{C}_\mu$ be the structure from Theorem 5.19 and recall that it is finitely bounded and homogeneous. We construct a binary fractional polymorphism of $\Gamma(\mathfrak{C}_\mu, \emptyset)$ *which is canonical and pseudo cyclic with respect to* $\mathrm{Aut}(\Gamma(\mathfrak{C}_\mu, \emptyset)) = \mathrm{Aut}(\mathfrak{C}_\mu)$. *Let $\mathfrak{M}$ be the τ-structure with domain $(C_\mu)^2$ and where*

- $((t_1^1, t_1^2), (t_2^1, t_2^2)) \in R^{\mathfrak{M}}$ *if* $(t_1^1, t_2^1) \in R^{\mathfrak{C}_\mu}$ *and* $(t_1^2, t_2^2) \in R^{\mathfrak{C}_\mu}$
- $((t_1^1, t_1^2), (t_2^1, t_2^2), (t_3^1, t_3^2)) \in S^{\mathfrak{M}}$ *if* $(t_1^1, t_2^1, t_3^1) \in S^{\mathfrak{C}_\mu}$ *or* $(t_1^2, t_2^2, t_3^2) \in S^{\mathfrak{C}_\mu}$.

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

- $((t_1^1, t_1^2), (t_2^1, t_2^2)) \in R^{\mathfrak{N}}$ *if* $(t_1^1, t_2^1) \in R^{\mathfrak{C}_\mu}$ *or* $(t_1^2, t_2^2) \in R^{\mathfrak{C}_\mu}$,
- $((t_1^1, t_1^2), (t_2^1, t_2^2), (t_3^1, t_3^2)) \in S^{\mathfrak{N}}$ *if* $(t_1^1, t_2^1, t_3^1) \in S^{\mathfrak{C}_\mu}$ *and* $(t_1^2, t_2^2, t_3^2) \in S^{\mathfrak{C}_\mu}$.

Note that $\mathfrak{M} \not\models \mu$ *and* $\mathfrak{N} \not\models \mu$ *and hence there are embeddings* $f\colon \mathfrak{M} \to \mathfrak{C}_\mu$ *and* $g\colon \mathfrak{N} \to \mathfrak{C}_\mu$*. 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* $\mathrm{Aut}(\mathfrak{C}_\mu)$*; this can be proved similarly as in Example 3.48. 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* Γ*: for* $t^1, t^2 \in (C_\mu)^2$ *we have*

$$E_\omega[h \mapsto R(h(t^1, t^2))] = \frac{1}{2} R^\Gamma(f(t^1, t^2)) + \frac{1}{2} R^\Gamma(g(t^1, t^2))$$
$$= \frac{1}{2} \sum_{j=1}^{2} R^\Gamma(t^j). \tag{5.2}$$

Therefore, ω *improves* R*, and analogously,* ω *improves* S*.*

We proved that Γ *has a binary canonical pseudo cyclic fractional polymorphism. By Corollary 5.18, the resilience problem for* μ *is in P, which reproves the results from [45] (without multiplicities) and [65] (with multiplicities).*

For the following conjunctive query μ, the NP-hardness of the resilience problem without multiplicities was shown in [45]; to illustrate our method, 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 for μ (with multiplicities) is NP-hard.

Example 5.21 (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*

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

Since the Gaifman graph of μ *is NP-complete, the structure* $\mathfrak{C}_\mu$ *from Theorem 5.19 exists. Let* $\Gamma := \Gamma(\mathfrak{C}_\mu, \emptyset)$*. We provide a pp-construction of* $(\{0,1\}; \mathrm{OIT})$ *in* Γ*, which proves NP-hardness of* $\mathrm{VCSP}(\Gamma)$ *by Corollary 2.17. By Proposition 5.8, this implies that the resilience problem for* μ *is NP-complete. Since* Γ *is homomorphically equivalent to* Γ_μ*, we also obtain a pp-construction of* $(\{0,1\}; \mathrm{OIT})$ *in* Γ_μ *(see Lemma 2.16).*

Let C *be the domain of* Γ*. 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 expression*

$$\begin{aligned} &R(a,b) + 2S(b,c) + 2T(c,d) + 2R(d,e) \\ &+ 2S(e,f) + 2T(f,g) + 2R(g,h) + S(h,i) \end{aligned} \tag{5.3}$$
$$\begin{aligned} &+\ \mathrm{Opt}(T)(i,g) + \mathrm{Opt}(S)(h,f) + \mathrm{Opt}(R)(g,e) + \mathrm{Opt}(T)(f,d) \\ &+\ \mathrm{Opt}(S)(e,c) + \mathrm{Opt}(R)(d,b) + \mathrm{Opt}(T)(c,a). \end{aligned} \tag{5.4}$$

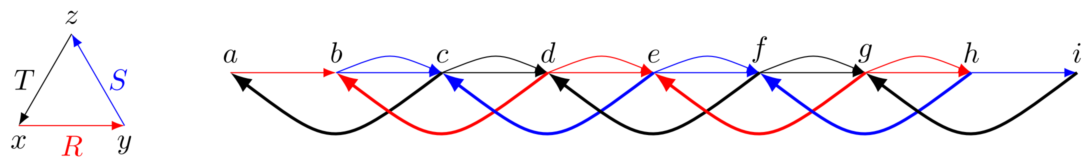

Figure 5.3: Example 5.21, visualisation of μ and ϕ. The thick edges correspond to crisp constraints.

For an illustration of μ and ϕ, see Figure 5.3. 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 being crisp.

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 (5.4) hold, and*
- *either every atom at even position or every atom at odd position in (5.3) holds,*

evaluates ϕ to 7 and hence is a solution to ϕ. Note that such an assignment exists because $\mathfrak{C}_\mu$ is homogeneous and a dual of μ.

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*

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

Note that $R(x,y)$, $S(y,z)$ *and* $T(z,x)$ *cannot hold at the same time and therefore we have* $(u,v,u',v',u'',v'') \in M$ *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* Γ *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 from* $(\{0,1\};\mathrm{OIT})$ *to* Δ *that maps* 1 *to any pair of vertices* $(u,v) \in R^{\mathfrak{C}_\mu}$ *and* 0 *to any pair of vertices* $(u,v) \notin R^{\mathfrak{C}_\mu}$. *Therefore,* Γ *pp-constructs* $(\{0,1\};\mathrm{OIT})$.

As the next example shows, by declaring a relation exogenous we may obtain a computationally easier resilience problem.

Example 5.22. *Consider the conjunctive query*

$$\mu := \exists x,y,z(R(x,y) \wedge S(y,z) \wedge T^{ex}(z,x))$$

over the signature $\tau = \{R,S,T^{ex}\}$, *where* R,S *are endogeneous and* T^{ex} *is exogeneous. Note that the Gaifman graph of* μ *is complete; let* $\mathfrak{C}_\mu$ *be the homogeneous dual of* μ *that embeds every countable structure that does not satisfy* μ. *We construct a binary fractional polymorphism of* $\Gamma := \Gamma(\mathfrak{C}_\mu, \{T^{ex}\})$, *which is canonical and pseudo cyclic with respect to* $\mathrm{Aut}(\mathfrak{C}_\mu)$. *Let* $\mathfrak{M}$ *be the* τ*-structure with the domain* C^2 *and where*

- $((t_1^1,t_1^2),(t_2^1,t_2^2)) \in R^{\mathfrak{M}}$ *if and only if* $(t_1^1,t_2^1) \in R^{\mathfrak{C}_\mu}$ *and* $(t_1^2,t_2^2) \in R^{\mathfrak{C}_\mu}$,
- $((t_1^1,t_1^2),(t_2^1,t_2^2)) \in S^{\mathfrak{M}}$ *if and only if* $(t_1^1,t_2^1) \in S^{\mathfrak{C}_\mu}$ *or* $(t_1^2,t_2^2) \in S^{\mathfrak{C}_\mu}$, *and*
- $((t_1^1,t_1^2),(t_2^1,t_2^2)) \in (T^{ex})^{\mathfrak{M}}$ *if and only if* $(t_1^1,t_2^1) \in (T^{ex})^{\mathfrak{C}_\mu}$ *and* $(t_1^2,t_2^2) \in (T^{ex})^{\mathfrak{C}_\mu}$.

Similarly, let $\mathfrak{N}$ *be the* τ*-structure with the domain* C^2 *and where*

- $((t_1^1,t_1^2),(t_2^1,t_2^2)) \in R^{\mathfrak{N}}$ *if and only if* $(t_1^1,t_2^1) \in R^{\mathfrak{C}_\mu}$ *or* $(t_1^2,t_2^2) \in R^{\mathfrak{C}_\mu}$,
- $((t_1^1,t_1^2),(t_2^1,t_2^2)) \in S^{\mathfrak{N}}$ *if and only if* $(t_1^1,t_2^1) \in S^{\mathfrak{C}_\mu}$ *and* $(t_1^2,t_2^2) \in S^{\mathfrak{C}_\mu}$, *and*
- $((t_1^1,t_1^2),(t_2^1,t_2^2)) \in (T^{ex})^{\mathfrak{N}}$ *if and only if* $(t_1^1,t_2^1) \in (T^{ex})^{\mathfrak{C}_\mu}$ *and* $(t_1^2,t_2^2) \in (T^{ex})^{\mathfrak{C}_\mu}$.

Note that $\mathfrak{M} \not\models \mu$ *and hence there exists an embedding* $f\colon \mathfrak{M} \to \mathfrak{C}_\mu$. *Similarly, there exists an embedding* $g\colon \mathfrak{N} \to \mathfrak{C}_\mu$. *Clearly, both* f *and* g *regarded as operations on the set* C *are pseudo cyclic and canonical with respect to* $\mathrm{Aut}(\mathfrak{C}_\mu)$. *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* Γ*: for* $t^1, t^2 \in (C_\mu)^2$ *we have*

$$\begin{aligned}\sum_{h \in \mathscr{O}^{(2)}} \omega(h) R^\Gamma(h(t^1,t^2)) &= \frac{1}{2} R^\Gamma(f(t^1,t^2)) + \frac{1}{2} R^\Gamma(g(t^1,t^2)) \\ &= \frac{1}{2}\sum_{j=1}^{2} R^\Gamma(t^j). \end{aligned} \tag{5.5}$$

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

Finally, ω *improves* T^{ex}, *since the right-hand side of the inequality*

$$\sum_{h \in \mathscr{O}^{(2)}} \omega(h)(T^{ex})^\Gamma(h(t^1,t^2)) = \frac{1}{2}(T^{ex})^\Gamma(f(t^1,t^2)) + \frac{1}{2}(T^{ex})^\Gamma(g(t^1,t^2)) \leq \frac{1}{2}\sum_{j=1}^{2}(T^{ex})^\Gamma(t^j),$$

is equal to ∞ *whenever* $t^j \notin (T^{ex})^{\mathfrak{C}_\mu}$ *for some* j.

It follows that Γ *has a fractional polymorphism that is canonical and pseudo cyclic with respect to* $\mathrm{Aut}(\mathfrak{C}_\mu)$, *which by Corollary 5.18 implies that the resilience problem for* μ *where* T^{ex} *is exogeneous is in P.*

We finish this section with an example that shows that for tractability proofs via Corollary 5.18 we cannot limit ourselves to binary fractional polymorphisms.

Example 5.23. *Consider the conjunctive query*

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

and observe that its resilience problem is in P: if $\mathfrak{A}$ *is the input database, and both* (a,b) *and* (b,a) *lie in* $R^{\mathfrak{A}}$, *we remove all copies of the pair with the smaller multiplicity.*

Let $\mathfrak{C}_\mu$ *be the homogeneous dual of* μ *that embeds every countable structure* $\mathfrak{A}$ *that does not satisfy* μ. *Let* C *be the domain of* $\mathfrak{C}_\mu$ *and let* $\Gamma := \Gamma(\mathfrak{C}_\mu, \emptyset)$. *We show that* Γ *has a ternary canonical pseudo cyclic fractional polymorphism (with respect to* $\mathrm{Aut}(\mathfrak{C}_\mu)$*), which implies tractability of* $\mathrm{VCSP}(\Gamma)$ *(Theorem 3.53) and of the resilience problem for* μ *(Corollary 5.18). To increase readability, we introduce a relation* $\widetilde{R}^{\mathfrak{C}_\mu} = \{(a,b) \mid (b,a) \in R^{\mathfrak{C}_\mu}\}$*; this relation is not in the signature of* $\mathfrak{C}_\mu$. *Note that since* $\mathfrak{C}_\mu \models \neg\mu$, *for* *every $a, b \in C$, we have precisely one of the following: $(a,b) \in R^{\mathfrak{C}_\mu}$, $(a,b) \in \widetilde{R}^{\mathfrak{C}_\mu}$, or $(a,b) \notin R^{\mathfrak{C}_\mu} \cup \widetilde{R}^{\mathfrak{C}_\mu}$.*

Let $\mathfrak{M}$ be an $\{R\}$-structure with the domain C^3 such that

$$\mathfrak{M} \models R((x,y,z),(u,v,w))$$

if and only if at least one of the following is true:

(1) at least two of $(x,u),(y,v),(z,w)$ lie in $R^{\mathfrak{C}_\mu}$;

(2) $(x,u) \in R^{\mathfrak{C}_\mu}$, $(y,v) \in \widetilde{R}^{\mathfrak{C}_\mu}$ and $(z,w) \notin R^{\mathfrak{C}_\mu} \cup \widetilde{R}^{\mathfrak{C}_\mu}$;

(3) $(x,u) \notin R^{\mathfrak{C}_\mu} \cup \widetilde{R}^{\mathfrak{C}_\mu}$, $(y,v) \in R^{\mathfrak{C}_\mu}$ and $(z,w) \in \widetilde{R}^{\mathfrak{C}_\mu}$;

(4) $(x,u) \in \widetilde{R}^{\mathfrak{C}_\mu}$, $(y,v) \notin R^{\mathfrak{C}_\mu} \cup \widetilde{R}^{\mathfrak{C}_\mu}$ and $(z,w) \in R^{\mathfrak{C}_\mu}$.

Note that items (2)-(4) are just cyclic shifts of the same condition. It is straightforward to verify that $\mathfrak{M} \models \neg\mu$: for example, if $\mathfrak{M} \models R((x,y,z),(u,v,w))$ because of item (1), then at least two of $(u,x),(v,y),(w,z)$ lie in $\widetilde{R}^{\mathfrak{C}_\mu}$ and hence $\mathfrak{M} \models \neg R((u,v,w),(x,y,z))$ by definition. Therefore, there is an embedding f of $\mathfrak{M}$ into $\mathfrak{C}_\mu$. By the definition of $\mathfrak{M}$, the operation f is pseudo cyclic and canonical with respect to $\mathrm{Aut}(\mathfrak{C}_\mu)$ *(this is easy to see because $\mathfrak{C}_\mu$ is homogeneous, see a similar argument in Example 3.48). The idea is that f has the behavior of a majority operation on orbits of pairs.*

Let $\mathfrak{N}$ be an $\{R\}$-structure with the domain C^3 such that

$$\mathfrak{N} \models R((x,y,z),(u,v,w))$$

if and only if at least one of the following is true:

(5) $(x,u),(y,v),(z,w) \in R^{\mathfrak{C}_\mu}$;

(6) one of $(x,u),(y,v),(z,w)$ lies in $R^{\mathfrak{C}_\mu}$, and the remaining two lie in $\widetilde{R}^{\mathfrak{C}_\mu}$;

(7) one of $(x,u),(y,v),(z,w)$ lies in $R^{\mathfrak{C}_\mu}$, and the remaining two do not lie in $R^{\mathfrak{C}_\mu} \cup \widetilde{R}^{\mathfrak{C}_\mu}$, or

(8) $(x,u) \in \widetilde{R}^{\mathfrak{C}_\mu}$, $(y,v) \in R^{\mathfrak{C}_\mu}$ and $(z,w) \notin R^{\mathfrak{C}_\mu} \cup \widetilde{R}^{\mathfrak{C}_\mu}$;

(9) $(x,u) \notin R^{\mathfrak{C}_\mu} \cup \widetilde{R}^{\mathfrak{C}_\mu}$, $(y,v) \in \widetilde{R}^{\mathfrak{C}_\mu}$ and $(z,w) \in R^{\mathfrak{C}_\mu}$, or

(10) $(x,u) \in R^{\mathfrak{C}_\mu}$, $(y,v) \notin R^{\mathfrak{C}_\mu} \cup \widetilde{R}^{\mathfrak{C}_\mu}$ and $(z,w) \in \widetilde{R}^{\mathfrak{C}_\mu}$.

Note that items (4)-(6) are just cyclic shifts of the same condition. Again one can verify that $\mathfrak{N} \models \neg\mu$, because $\mathfrak{C}_\mu \models \neg\mu$. Therefore, there is an embedding g of $\mathfrak{N}$ into $\mathfrak{C}_\mu$. By the definition of $\mathfrak{N}$, the operation g is pseudo cyclic and canonical with respect to $\mathrm{Aut}(\mathfrak{C}_\mu)$*; it has the behavior of a minority operation on orbits of pairs.*

Let ω be the ternary fractional operation defined by $\omega(f) = 2/3$ and $\omega(g) = 1/3$. Note that ω is pseudo cyclic and canonical ternary fractional operation on C. We show that $\omega \in \mathrm{fPol}(\Gamma)$. Let $(x,u),(y,v),(z,w) \in C^2$. We want to verify that

$$E_\omega\left[h \mapsto R\left(h\left(\begin{pmatrix}x\\u\end{pmatrix},\begin{pmatrix}y\\v\end{pmatrix},\begin{pmatrix}z\\w\end{pmatrix}\right)\right)\right] \leq \frac{1}{3}(R(x,u)+R(y,v)+R(z,w)),$$

equivalently

$$\begin{aligned}2R\left(f\left(\begin{pmatrix}x\\u\end{pmatrix},\begin{pmatrix}y\\v\end{pmatrix},\begin{pmatrix}z\\w\end{pmatrix}\right)\right)+R\left(g\left(\begin{pmatrix}x\\u\end{pmatrix},\begin{pmatrix}y\\v\end{pmatrix},\begin{pmatrix}z\\w\end{pmatrix}\right)\right)&\\ \leq R(x,u)+R(y,v)+R(z,w).&\end{aligned} \tag{5.6}$$

We break into cases:

- *If $(x,u),(y,v),(z,w) \in R^{\mathfrak{C}_\mu}$, then by item (1) and (5) the left-hand side of* (5.6) *evaluates to 0 and hence* (5.6) *holds.*
- *If exactly two of $(x,y),(y,v),(z,w)$ lie in $R^{\mathfrak{C}_\mu}$, then by item (1) the left-hand side of* (5.6) *is at most 1 and hence* (5.6) *holds.*
- *If exactly one of $(x,y),(y,v),(z,w)$ lies in $R^{\mathfrak{C}_\mu}$, then precisely one of the conditions (2)-(4), (6)-(10) applies and therefore the left-hand side of* (5.6) *is at most 2. Therefore,* (5.6) *holds.*
- *If $(x,u),(y,v),(z,w) \notin R^{\mathfrak{C}_\mu}$, then* (5.6) *holds trivially since the left-hand side is always at most 3.*

We conclude that $\omega \in \mathrm{fPol}(\Gamma)$.

We show that there is no binary fractional polymorphism of Γ that is canonical and pseudo cyclic with respect to $\mathrm{Aut}(\mathfrak{C}_\mu)$. Suppose for contradiction that $\chi \in \mathrm{fPol}(\Gamma)$ is binary, canonical, and pseudo cyclic. Let $(a,b) \in R^{\mathfrak{C}_\mu}$. Since $\chi \in \mathrm{fPol}(\Gamma)$, we have

$$E_\chi\left[h \mapsto R\left(h\left(\begin{pmatrix}a\\b\end{pmatrix},\begin{pmatrix}b\\a\end{pmatrix}\right)\right)\right] \leq \frac{1}{2}(R(a,b)+R(b,a)) = \frac{1}{2}.$$

This implies that

$$\chi\left(\left\{h \in \mathscr{O}^{(2)} \;\middle|\; R\left(h\left(\begin{pmatrix}a\\b\end{pmatrix},\begin{pmatrix}b\\a\end{pmatrix}\right)\right)=1\right\}\right) \leq \frac{1}{2}.$$

and hence

$$\chi\left(\left\{h \in \mathscr{O}^{(2)} \;\middle|\; R\left(h\left(\begin{pmatrix}a\\b\end{pmatrix},\begin{pmatrix}b\\a\end{pmatrix}\right)\right)=0\right\}\right) \geq \frac{1}{2}. \tag{5.7}$$

Let $h \in \mathscr{O}^{(2)}$ be binary, canonical and pseudo cyclic and such that

$$R\left(h\left(\begin{pmatrix}a\\b\end{pmatrix},\begin{pmatrix}b\\a\end{pmatrix}\right)\right)=0,$$

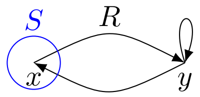

Figure 5.4: Visualisation of the query μ from (5.8).

in other words, $h((a,b),(b,a)) \in R^{\mathfrak{C}_\mu}$. By flipping the entries, we get $h((b,a),(a,b)) \in \widetilde{R}^{\mathfrak{C}_\mu}$. On the other hand, by pseudo cyclicity, $h((b,a),(a,b))$ lies in the same orbit of $\mathrm{Aut}(\mathfrak{C}_\mu)$ as $h((a,b),(b,a)) \in R^{\mathfrak{C}_\mu}$. However, $R^{\mathfrak{C}_\mu}$ and $\widetilde{R}^{\mathfrak{C}_\mu}$ are two disjoint orbits of pairs of $\mathrm{Aut}(\mathfrak{C}_\mu)$ and therefore such h cannot exist. Therefore, the set in (5.7) does not contain any operations that are both canonical and pseudo cyclic, which contradicts the assumption that χ is canonical and pseudo cyclic.

5.6 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 [46, Section 8.5]; this result was originally published in [30]. Let

$$\mu := \exists x, y (S(x) \wedge R(x,y) \wedge R(y,x) \wedge R(y,y)). \tag{5.8}$$

Let $\tau = \{R, S\}$ be the signature of μ. To study the complexity of resilience of μ, it will be convenient to work with a dual which has different model-theoretic properties than the duals $\mathfrak{B}_\mu$ from Theorem 5.17 and $\mathfrak{C}_\mu$ from Theorem 5.19, namely a dual that is a model-complete core. The advantage of working with model-complete cores is that the structure is in a sense 'minimal' and therefore easier to work with in some concrete examples.[2]

Proposition 5.24. *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 $\mathrm{Aut}(\mathfrak{B})$. 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 $\sigma \subseteq \tau$.*

Proof. Since the Gaifman graph of μ is a complete graph, there exists the structure $\mathfrak{C}_\mu$ as in Theorem 5.19. Let $\mathfrak{B}$ be the model-complete core of $\mathfrak{C}_\mu$. Note that $\mathfrak{B}$ has the property that a countable structure $\mathfrak{A}$ maps homomorphically to $\mathfrak{B}$ if and only if $\mathfrak{A} \models \neg\mu$; in particular, $\mathfrak{B}$ is a dual of μ and $\mathfrak{B} \models \neg\mu$. The structure $\mathfrak{C}_\mu$ is homogeneous, and by Proposition 3.6, $\mathfrak{B}$ is homogeneous as well. Let $\Gamma := \Gamma(\mathfrak{B}, \emptyset)$.

[2]The model-complete core of $\mathfrak{B}_\mu$ would be a natural choice for the canonical dual of μ to work with instead of $\mathfrak{B}_\mu$. However, proving that the model-complete core has a finitely bounded homogeneous expansion (so that, for example, Theorem 3.53 applies) requires introducing further model-theoretical notions from Ramsey theory [69] which we want to avoid in this thesis.

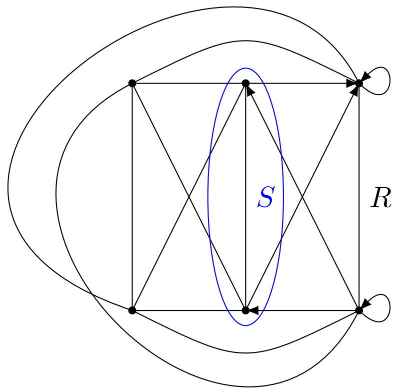

Figure 5.5: Illustration of a finite substructure of $\mathfrak{B}$ from the proof of Proposition 5.24 that contains representatives for all orbits of pairs of $\mathrm{Aut}(\mathfrak{B})$. Arrows are not drawn on undirected edges.

Note that

$$\mathfrak{B} \models \forall x\big(\neg S(x) \vee \neg R(x,x)\big) \tag{5.9}$$

$$\text{and } \mathfrak{B} \models \forall x,y\big(x = y \vee R(x,y) \vee R(y,x)\big). \tag{5.10}$$

To see (5.10), suppose for contradiction that $\mathfrak{B}$ contains distinct elements x, y such that neither (x,y) nor (y,x) is in $R^{\mathfrak{B}}$. Let $\mathfrak{B}'$ be the structure obtained from $\mathfrak{B}$ by adding (x,y) to $R^{\mathfrak{B}}$. Then $\mathfrak{B}' \models \neg\mu$ as well, and hence there is a homomorphism from $\mathfrak{B}'$ to $\mathfrak{B}$ by the properties of $\mathfrak{B}$. This homomorphism is also an endomorphism of $\mathfrak{B}$ which is not an embedding, a contradiction to the assumption that $\mathfrak{B}$ is a model-complete core.

Also observe that

$$\mathfrak{B} \models \forall x,y\big(x = y \vee (R(x,y) \wedge R(y,x)) \vee (S(x) \wedge R(y,y)) \vee (R(x,x) \wedge S(y))\big). \tag{5.11}$$

Suppose for contradiction that (5.11) does not hold for some distinct x and y. Then $\neg S(x) \vee \neg R(y,y)$ and $\neg R(x,x) \vee \neg S(y)$, i.e., $\neg S(x) \wedge \neg R(x,x)$, or $\neg S(x) \wedge \neg S(y)$, or $\neg R(y,y) \wedge \neg R(x,x)$, or $\neg R(y,y) \wedge \neg S(y)$. In each of these cases we may add both R-edges between the distinct elements x and y to $\mathfrak{B}$ and obtain a structure not satisfying μ, which leads to a contradiction as above.

For an illustration of a finite substructure of $\mathfrak{B}$ which contains a representative for every orbit of pairs under the action of $\mathrm{Aut}(\mathfrak{B})$, see Figure 5.5.

Claim 1. For every finite relational τ-structure $\mathfrak{A}$ that satisfies $\neg\mu$ and the sentences in (5.10) and (5.11), there exists a *strong* homomorphism to $\mathfrak{B}$, i.e., a homomorphism that also preserves the complements of R and S. First observe that $\mathfrak{B}$ embeds the countably infinite complete graph, where R is the edge relation and precisely one element lies in the relation S; this is because this structure maps homomorphically to $\mathfrak{B}$ and unless embedded, it contradicts $\mathfrak{B} \not\models \mu$. In particular, there are infinitely many $x \in B$ such that $\mathfrak{B} \models \neg S(x) \wedge \neg R(x,x)$ and by (5.11), for every $y \in B$, $x \neq y$, we have $\mathfrak{B} \models R(x,y) \wedge R(y,x)$.

To prove the claim, let $\mathfrak{A}$ be a finite structure that satisfies $\neg\mu$ and the sentences in (5.10) and (5.11). For a homomorphism h from $\mathfrak{A}$ to $\mathfrak{B}$, let

$$s(h) := |\{x \in A \mid \mathfrak{A} \models \neg S(x) \wedge \mathfrak{B} \models S(h(x))\}|$$

and

$$r(h) := |\{(x,y) \in A^2 \mid \mathfrak{A} \models \neg R(x,y) \wedge \mathfrak{B} \models R(h(x),h(y))\}|.$$

Let h be a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$, which exists since $\mathfrak{A} \models \neg\mu$. If $s(h)+r(h) = 0$, then h is a strong homomorphism and there is nothing to prove. Suppose therefore $s(h)+r(h) > 0$. We construct a homomorphism h' such that $r(h')+s(h') < r(h)+s(h)$. Since $r(h)+s(h)$ is finite, by applying this construction finitely many times, we obtain a strong homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.

If $s(h) > 0$, then there exists $a \in A \setminus S^{\mathfrak{A}}$ such that $h(a) \in S^{\mathfrak{B}}$. By (5.9), $\mathfrak{B} \not\models R(h(a),h(a))$ and hence $\mathfrak{A} \not\models R(a,a)$. By (5.11), $\mathfrak{A} \models R(a,a') \wedge R(a',a)$ for every $a' \in A$, $a' \neq a$. Pick $b \in B \setminus h(A)$ such that $\mathfrak{B} \models \neg S(b) \wedge \neg R(b,b)$ and define

$$h'(x) := \begin{cases} b \text{ if } x = a, \\ h(x) \text{ otherwise.} \end{cases}$$

Observe that h' is a homomorphism, $s(h') < s(h)$ and $r(h') = r(h)$. If $r(h) > 0$, then there exists $(x,y) \in A^2 \setminus R^{\mathfrak{A}}$ such that $(h(x),h(y)) \in R^{\mathfrak{B}}$. If $x = y$, the argument is similar as in the case $s(h) > 0$. Finally, if $x \neq y$, then $\mathfrak{A} \models (S(x) \wedge R(y,y)) \vee (R(x,x) \wedge S(y))$, because $\mathfrak{A}$ satisfies the sentence in (5.11). Since $\mathfrak{A}$ satisfies the sentence in (5.10), $\mathfrak{A} \models R(y,x)$. Since h is a homomorphism, we have

$$\begin{aligned}\mathfrak{B} \models\ & R(h(x),h(y)) \wedge R(h(y),h(x)) \\ & \wedge ((S(h(x)) \wedge R(h(y),h(y))) \vee (R(h(x),h(x)) \wedge S(h(y)))),\end{aligned}$$

which contradicts $\mathfrak{B} \not\models \mu$.

Claim 2. Every finite relational τ-structure $\mathfrak{A}$ that satisfies $\neg\mu$ and the universal sentences in (5.10) and (5.11) embeds into $\mathfrak{B}$, in particular, $\mathfrak{B}$ is finitely bounded. Let $\mathfrak{A}$ be such a structure. By Theorem 5.19, there is an embedding e of $\mathfrak{A}$ into $\mathfrak{C}_\mu$. Since $\mathfrak{C}_\mu$ is homogeneous and embeds every finite relational τ-structure that satisfies $\neg\mu$, there exists a finite substructure $\mathfrak{A}'$ of $\mathfrak{C}_\mu$ satisfying the sentences in (5.10) and (5.11) such that $e(\mathfrak{A})$ is a substructure of $\mathfrak{A}'$ and for all distinct $a,b \in A$ there exists $s \in S^{\mathfrak{A}'}$ such that $\mathfrak{C}_\mu \models R(e(a),s) \wedge R(s,e(b))$. By Claim 1, there is a strong homomorphism h from $\mathfrak{A}'$ to $\mathfrak{B}$.

We claim that $h \circ e$ is injective and therefore an embedding of $\mathfrak{A}$ into $\mathfrak{B}$. Suppose there exist distinct $a,b \in A$ such that $h(e(a)) = h(e(b))$. Since $e(\mathfrak{A})$ satisfies the sentence in (5.10), we have $\mathfrak{B} \models R(h(e(a)),h(e(a)))$. Let $s \in S^{\mathfrak{A}'}$ be such that $\mathfrak{C}_\mu \models R(e(a),s) \wedge R(s,e(b))$. Hence,

$$\mathfrak{B} \models S(h(s)) \wedge R(h(e(a)),h(s)) \wedge R(h(s),h(e(a))) \wedge R(h(e(a)),h(e(a))),$$

a contradiction to $\mathfrak{B} \not\models \mu$. It follows that $h \circ e$ is an embedding of $\mathfrak{A}$ into $\mathfrak{B}$.

We define τ-structures $\mathfrak{M}, \mathfrak{N}$ with domain B^2 as follows. For all $x_1, x_2, y_1, y_2, x, y \in \mathfrak{B}$ define

$$\begin{aligned}
\mathfrak{M}, \mathfrak{N} &\models R\big((x_1, y_1), (x_2, y_2)\big) && \text{if } \mathfrak{B} \models R(x_1, x_2) \wedge R(y_1, y_2), && (5.12)\\
\mathfrak{M}, \mathfrak{N} &\models S\big((x, y)\big) && \text{if } \mathfrak{B} \models S(x) \wedge S(y) && (5.13)\\
\mathfrak{M} &\models S\big((x, y)\big) && \text{if } \mathfrak{B} \models S(x) \vee S(y) && (5.14)\\
\mathfrak{N} &\models R\big((x, y), (x, y)\big) && \text{if } \mathfrak{B} \models R(x, x) \vee R(y, y). && (5.15)
\end{aligned}$$

Add pairs of distinct elements to $R^{\mathfrak{M}}$ and $R^{\mathfrak{N}}$ such that both $\mathfrak{M}$ and $\mathfrak{N}$ satisfy the sentence in (5.11) (note that no addition of elements to $S^{\mathfrak{M}}$ and $S^{\mathfrak{N}}$ is needed). Finally, add $((x_1, y_1), (x_2, y_2))$ to $R^{\mathfrak{M}}$ and $((x_2, y_2), (x_1, y_1))$ to $R^{\mathfrak{N}}$ if at least one of the following cases holds:

(A) $\mathfrak{B} \models S(x_1) \wedge R(x_1, x_2) \wedge R(x_2, x_2) \wedge R(y_2, y_2) \wedge R(y_2, y_1) \wedge S(y_1)$,

(B) $\mathfrak{B} \models R(x_1, x_1) \wedge R(x_1, x_2) \wedge S(x_2) \wedge y_1 = y_2 \wedge R(y_1, y_2)$,

(C) $\mathfrak{B} \models S(y_1) \wedge R(y_1, y_2) \wedge R(y_2, y_2) \wedge R(x_2, x_2) \wedge R(x_2, x_1) \wedge S(x_1)$,

(D) $\mathfrak{B} \models R(y_1, y_1) \wedge R(y_1, y_2) \wedge S(y_2) \wedge x_1 = x_2 \wedge R(x_1, x_2)$.

Conditions (A) and (B) are illustrated in Figure 5.6; conditions (C) and (D) are obtained from (A) and (B) by replacing x by y. Note that for $(x_1, y_1) = (x_2, y_2)$, none of the conditions (A)-(D) is ever satisfied. No other atomic formulas hold on $\mathfrak{M}$ and $\mathfrak{N}$. Note that both $\mathfrak{M}$ and $\mathfrak{N}$ satisfy the property stated for $\mathfrak{B}$ in (5.9).

Claim 3. $\mathfrak{M}$ and $\mathfrak{N}$ satisfy the sentence in (5.10). We prove the statement for $\mathfrak{M}$; the proof for $\mathfrak{N}$ is similar. Let $(x_1, y_1), (x_2, y_2) \in B$ be such that $(x_1, y_1) \neq (x_2, y_2)$ and $\mathfrak{M} \models \neg R((x_2, y_2), (x_1, y_1))$. Since $\mathfrak{M}$ satisfies the sentence in (5.11), we must have either that $\mathfrak{M} \models S(x_1, y_1) \wedge R((x_2, y_2), (x_2, y_2))$ or that $\mathfrak{M} \models S(x_2, y_2) \wedge R((x_1, y_1), (x_1, y_1))$. Suppose the former is true; the other case is treated analogously. Then $\mathfrak{B} \models R(x_2, x_2) \wedge R(y_2, y_2)$ and $\mathfrak{B} \models S(x_1) \vee S(y_1)$. If $\mathfrak{B} \models S(x_1)$, then $x_1 \neq x_2$ and by (5.10) we have $\mathfrak{B} \models R(x_1, x_2) \vee R(x_2, x_1)$. By (5.10) and (5.11) for (y_1, y_2), we obtain that $\mathfrak{M} \models R((x_1, y_1), (x_2, y_2))$ by (5.12) or one of the conditions (A)-(D). The argument if $\mathfrak{B} \models S(y_1)$ is similar with x and y switched.

Claim 4. $\mathfrak{M}$ and $\mathfrak{N}$ satisfy $\neg\mu$. Let $x_1, x_2, y_1, y_2 \in B$. Suppose for contradiction that

$$\mathfrak{M} \models S(x_1, y_1) \wedge R((x_1, y_1), (x_2, y_2)) \wedge R((x_2, y_2), (x_1, y_1)) \wedge R((x_2, y_2), (x_2, y_2)).$$

By the definition of $\mathfrak{M}$, we have $\mathfrak{B} \models R(x_2, x_2) \wedge R(y_2, y_2)$ and $\mathfrak{B} \models S(x_1) \vee S(y_1)$. Assume that $\mathfrak{B} \models S(x_1)$; the case $\mathfrak{B} \models S(y_1)$ is analogous.

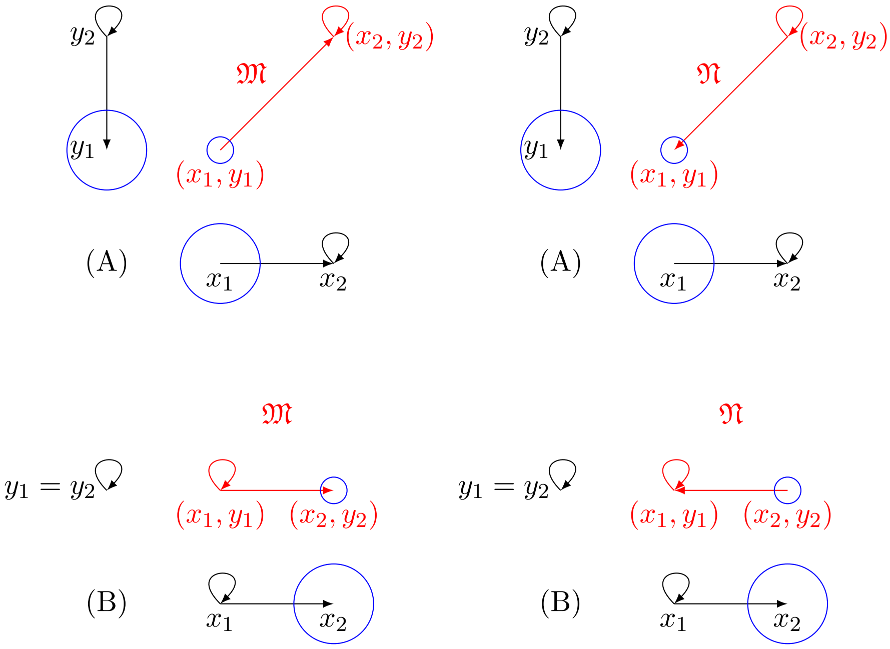

Figure 5.6: An illustration of the conditions (A) and (B) in $\mathfrak{M}$ and $\mathfrak{N}$.

By the assumption, $\mathfrak{M} \models R((x_1, y_1), (x_2, y_2))$. Then, by the definition of $\mathfrak{M}$, one of the conditions (5.12), (A)-(D) holds, or

$$\mathfrak{M} \models \neg\big(S(x_1, y_1) \wedge R((x_2, y_2), (x_2, y_2))\big)$$

(recall that $((x_1, y_1), (x_2, y_2))$ might have been added to $R^{\mathfrak{M}}$ so that $\mathfrak{M}$ satisfies the sentence in (5.11)). The last option is false by the assumption. By (5.9), $\mathfrak{B} \models \neg S(x_2) \wedge \neg S(y_2)$, and hence neither (B) nor (D) holds. Therefore, one of the conditions (5.12), (A), or (C) holds for $((x_1, y_1), (x_2, y_2))$. Similarly, we obtain that one of the conditions (5.12) or (B) holds for $((x_2, y_2), (x_1, y_1))$, since $\mathfrak{M} \models R((x_2, y_2), (x_1, y_1))$ (to exclude (D) we use the assumption that $\mathfrak{B} \models S(x_1)$ and hence $x_1 \neq x_2$). This yields six cases and in each of them we must have that $\mathfrak{B} \models R(x_1, x_2) \wedge R(x_2, x_1)$ or $\mathfrak{B} \models S(y_1) \wedge R(y_1, y_2) \wedge R(y_2, y_1)$. Since $\mathfrak{B} \models S(x_1) \wedge R(x_2, x_2) \wedge R(y_2, y_2)$, this contradicts $\mathfrak{B} \models \neg\mu$. Since $(x_1, y_1), (x_2, y_2) \in M$ were chosen arbitrarily, this shows that $\mathfrak{M} \models \neg\mu$. The argument for $\mathfrak{N}$ is similar.

Claim 5. There is an embedding f of $\mathfrak{M}$ into $\mathfrak{B}$ and an embedding g of $\mathfrak{N}$ into $\mathfrak{B}$. We show the claim for $\mathfrak{M}$; the proof for $\mathfrak{N}$ is analogous. By Lemma 2.5, it is enough to show that every finite substructure of $\mathfrak{M}$ embeds into $\mathfrak{B}$. By the definition of $\mathfrak{M}$ and Claims 3 and 4, every finite substructure $\mathfrak{M}$ satisfies (5.10), (5.11) and $\neg\mu$ and hence, by Claim 2, it embeds into $\mathfrak{B}$.

Let ω be the fractional operation over B defined by $\omega(f) = \frac{1}{2}$ and $\omega(g) = \frac{1}{2}$.

Claim 6. ω is pseudo cyclic and canonical with respect to the group $\mathrm{Aut}(\mathfrak{B}) = \mathrm{Aut}(\Gamma)$. Note that since $\mathfrak{B}$ is homogeneous in a finite relational signature, two k-tuples of elements of $\mathfrak{B}$ lie in the same orbit if and only if they satisfy the same atomic formulas. Therefore, the canonicity of f and g with respect to $\mathrm{Aut}(\mathfrak{B})$ follows from the definition of $\mathfrak{M}$ and $\mathfrak{N}$: for $(a,b) \in B^2$, whether $\mathfrak{B} \models S(f(a,b))$ only depends on whether $\mathfrak{M} \models S(a,b)$ by Claim 5, which depends only on the atomic formulas that hold on a and on b in $\mathfrak{B}$. An analogous statement is true for atomic formulas of the form $R(x,y)$ and $x = y$. Therefore, f is canonical. The argument for the canonicity of g is analogous.

To see that f and g are pseudo cyclic, we show that f^* and g^* defined on 2-orbits (using the terminology of Remark 3.45) are cyclic. By the definition of f^*, we need to show that for any $a_1, a_2, b_1, b_2 \in B$, the two pairs $(f(a_1,b_1), f(a_2,b_2))$ and $(f(b_1,a_1), f(b_2,a_2))$ satisfy the same atomic formulas. For the formulas of the form $S(x)$ and $R(x,y)$, this can be seen from Claim 5 and the definition of $\mathfrak{M}$ and $\mathfrak{N}$, since each of the conditions (5.12),(5.13),(5.14),(5.15),(5.11) and the union of (A), (B), (C), (D) is symmetric with respect to exchanging x and y. For the atomic formulas of the form $x = y$, this follows from the injectivity of f. This shows that f^* is cyclic; the argument for g^* is the same. Hence, the pseudo-cyclicity of f and g is a consequence of Lemma 3.47 for $m = 2$.

Claim 7. ω improves S. By the definition of $\mathfrak{M}$ and $\mathfrak{N}$ and Claim 5, we have for all $x, y \in B$

$$\frac{1}{2}S^\Gamma(f(x,y)) + \frac{1}{2}S^\Gamma(g(x,y)) = \frac{1}{2}(S^\Gamma(x) + S^\Gamma(y)).$$

Claim 8. ω improves R. Let $x_1, y_1, x_2, y_2 \in B$. We have to verify that

$$\begin{aligned}\frac{1}{2}R^\Gamma(f(x_1,y_1), f(x_2,y_2)) + \frac{1}{2}R^\Gamma(g(x_1,y_1), g(x_2,y_2)) \\ \leq \frac{1}{2}(R^\Gamma(x_1,x_2) + R^\Gamma(y_1,y_2)).\end{aligned} \tag{5.16}$$

We distinguish four cases.

- $\mathfrak{M}, \mathfrak{N} \models R((x_1,y_1),(x_2,y_2))$. Then Inequality (5.16) holds since the left-hand side is zero, and the right-hand side is non-negative (each valued relation in Γ is non-negative).

- $\mathfrak{M}, \mathfrak{N} \models \neg R((x_1,y_1),(x_2,y_2))$. We need to show that $\mathfrak{B} \models \neg R(x_1,x_2) \wedge \neg R(y_1,y_2)$. This is clear if $(x_1,y_1) = (x_2,y_2)$ by the definition of $\mathfrak{N}$. Suppose therefore that $(x_1,y_1) \neq (x_2,y_2)$. Since $\mathfrak{M}$ satisfies the sentence in (5.11), we have $\mathfrak{M} \models S(x_1,y_1) \wedge R((x_2,y_2),(x_2,y_2))$ or $\mathfrak{M} \models S(x_2,y_2) \wedge R((x_1,y_1),(x_1,y_1))$. Suppose that $\mathfrak{M} \models S(x_1,y_1) \wedge R((x_2,y_2),(x_2,y_2))$; the other case is analogous. Since $\mathfrak{N}$ satisfies the sentence in (5.11) as well, this implies that $\mathfrak{N} \models S(x_1,y_1) \wedge R((x_2,y_2),(x_2,y_2))$; note that if $\mathfrak{N} \models S(x_2,y_2) \wedge R((x_1,y_1),(x_1,y_1))$, we would get a contradiction with (5.9). By the definition of $\mathfrak{M}$ and $\mathfrak{N}$, we have $\mathfrak{B} \models S(x_1) \wedge S(y_1) \wedge R(x_2,x_2) \wedge R(y_2,y_2)$, in particular, by (5.9), $x_1 \neq x_2$ and $y_1 \neq y_2$.

By (5.10), there is an R-edge in $\mathfrak{B}$ between x_1 and x_2 and between y_1 and y_2. By the condition (A) and (C) for $\mathfrak{M}$ and for $\mathfrak{N}$ with 1 and 2 switched, we see that $\mathfrak{M}, \mathfrak{N} \models \neg R((x_1, y_1), (x_2, y_2))$ implies $\mathfrak{B} \models \neg R(x_1, x_2) \wedge \neg R(y_1, y_2)$. Therefore, both sides of the inequality evaluate to 1.

- $\mathfrak{M} \models \neg R((x_1, y_1), (x_2, y_2))$ and $\mathfrak{N} \models R((x_1, y_1), (x_2, y_2))$. By Claim 5, the left-hand side evaluates to $\frac{1}{2}$. By (5.12), we have $\mathfrak{B} \models \neg R(x_1, x_2)$ or $\mathfrak{B} \models \neg R(y_1, y_2)$. Therefore, the right-hand side of (5.16) is at least $\frac{1}{2}$ and the inequality holds.
- $\mathfrak{M} \models R((x_1, y_1), (x_2, y_2))$ and $\mathfrak{N} \models \neg R((x_1, y_1), (x_2, y_2))$. Similar to the previous case.

This exhausts all cases and concludes the proof of Claim 8.

It follows that ω is a binary fractional polymorphism of Γ which is canonical and pseudo cyclic with respect to $\mathrm{Aut}(\mathfrak{B})$. Polynomial-time tractability of $\mathrm{VCSP}(\Gamma)$ and of the resilience problem for μ follow from Theorem 3.53 and Corollary 5.18. The final statement follows from Remark 5.9. $\square$

5.7 Resilience tractability conjecture

In this section we present a conjecture which implies, together with Corollary 2.17 and Corollary 5.6, a P versus NP-complete dichotomy for resilience problems for finite unions of conjunctive queries.

Conjecture 5.25. *Let μ be a union of connected conjunctive queries over the signature τ, and let $\sigma \subseteq \tau$. If K_3 has no pp-construction in $\Gamma_{\mu,\sigma}$, then there exists a dual $\mathfrak{B}$ of μ such that* $\mathrm{Aut}(\mathfrak{B})$ *contains an automorphism group of a finitely bounded homogeneous structure $\mathfrak{A}$ and $\Gamma(\mathfrak{B}, \sigma)$ has a fractional polymorphism of arity $\ell \geq 2$ which is canonical and pseudo cyclic with respect to* $\mathrm{Aut}(\mathfrak{A})$ *(and in this case,* $\mathrm{VCSP}(\Gamma(\mathfrak{B}, \sigma))$ *is in P by Theorem 3.53).*

Note that all duals $\mathfrak{B}$ of μ are homomorphically equivalent to $\mathfrak{B}_\mu$ by Lemma 2.5, hence all valued structures $\Gamma(\mathfrak{B}, \sigma)$ are fractionally homomorphically equivalent to $\Gamma_{\mu,\sigma}$. Therefore, by the transitivity of pp-constructability, $\Gamma_{\mu,\sigma}$ pp-constructs K_3 if and only if $\Gamma(\mathfrak{B}, \sigma)$ does. Also note that K_3 can be equivalently replaced by $(\{0,1\}; \mathrm{OIT})$ (Corollary 2.17). If P$\neq$NP, then there cannot be a union of queries μ such that $\Gamma_{\mu,\sigma}$ that pp-constructs K_3 and there is a dual $\mathfrak{B}$ of μ satisfying the assumptions from Conjecture 5.25.

For all presented examples of resilience problems, we either provided a dual $\mathfrak{B}$ that satisfies the condition from Conjecture 5.25 and thus proves the tractability of the resilience problem or a pp-construction of $(\{0,1\}; \mathrm{OIT})$. Note that if $\mathfrak{B}$ is a relational structure on a finite domain B and $\Gamma(\mathfrak{B}, \sigma)$ has a cyclic fractional polymorphism, then this polymorphism is trivially pseudo cyclic with respect to any group and it is canonical with respect to the trivial permutation group on B. Let $\tau = \{R_b \mid b \in B\}$ be a relational signature where all symbols are unary. Note that the relational structure

$\mathfrak{A} = (B; (R_b^{\mathfrak{A}})_{b \in B})$ where $R_b^{\mathfrak{A}} = \{b\}$ for every $b \in B$ has a trivial automorphism group and is homogeneous. It is also finitely bounded: a finite relational τ-structure embeds into $\mathfrak{A}$ if and only if it satisfies

$$\forall x \left(\left(\bigvee_{b \in B} R_b(x) \right) \wedge \bigwedge_{b,c \in B, b \neq c} \big(\neg R_b(x) \vee \neg R_c(x)\big) \wedge \bigwedge_{b \in B} (R_b(x) \wedge R_b(y) \Rightarrow x = y) \right)$$

Hence, every union of conjunctive queries each of which is homomorphically equivalent to a tree satisfies the condition from Conjecture 5.25 by Theorem 5.12.

Conjecture 5.25 is intentionally formulated only for VCSPs that stem from resilience problems, because it is known to be false for the more general situation of VCSPs with a template Γ such that $\mathrm{Aut}(\Gamma)$ contains an automorphism group of a finitely bounded homogeneous structure [9, Section 12.9.1]; the counterexample is even a CSP. However, the structures $\mathfrak{B}_\mu$ from Theorem 5.17 that allow to formulate resilience problems as VCSPs are particularly well-behaved for the universal-algebraic approach and more specifically, for canonical operations (see, e.g., [11, 19, 70]), which is why we believe in the strong formulation of Conjecture 5.25.

We also believe that the 'meta-problem' of deciding whether for a given conjunctive query the resilience problem is in P is decidable. This might follow from a positive answer to Conjecture 5.25 if an appropriate $\mathfrak{B}$ can be found effectively. In that case, we can compute $\Gamma := \Gamma(\mathfrak{B}, \sigma)$ and $\Gamma^*_{\mathfrak{A},m}$ effectively and decide Item 4 of Proposition 3.51 for $\Gamma^*_{\mathfrak{A},m}$ (see Remark 3.52).

We mention that another conjecture concerning a P vs. NP-complete complexity dichotomy for resilience problems appears in [65, Conjecture 7.7]. The conjecture has a similar form as Conjecture 5.25 in the sense that it states that a sufficient hardness condition for resilience is also necessary. The relationship between our hardness condition from Corollary 2.17 and the condition from [65] remains open.

Conclusion

The focus of this thesis are VCSPs with templates with an oligomorphic automorphism group, methods to study complexity of such VCSPs and some applications of these methods to concrete classes of VCSPs. These methods are based on two influential concepts from the theory of (V)CSPs: pp-constructions introduced in [6] for the CSP setting and fractional polymorphisms, which have been essential for identifying tractable VCSPs, see, e.g., [41,57]. For tractability results, we often restrict the templates further and require that they are preserved by an automorphism group of a finitely bounded homogeneous structure. Most of the general theoretical results for VCSP templates satisfying the above mentioned restrictions were presented in a conference paper by Bodirsky, Lutz and the author of this thesis [30].

As a natural next step, we apply the presented tractability and hardness results to concrete classes of VCSPs to obtain complexity classifications. We first focus on the class of temporal VCSPs, a generalization of temporal CSPs, which are well-understood, see [13,18,25]. Our results yield a decidable complexity dichotomy for temporal VCSPs: all of them are in P or NP-complete. This part of the thesis is based on joint work of Bodirsky, Bonnet and the author [12].

The other application of our methods to a concrete computational problem has overlap with database theory: we study resilience problems in bag semantics by translating them to VCSPs and applying our tools to find their computational complexity. We obtain a sufficient hardness and tractability conditions for resilience problems and conjecture that every resilience problem satisfies one of them, which would imply a complexity dichotomy for resilience problems. These results form the core of the paper [30].

The class of infinite-domain VCSPs and, in particular, VCSPs of valued structures with an oligomorphic automorphism group, has not received much attention in the research on CSPs (see Introduction) and still lacks understanding of some algebraic properties of valued structures. In the view of the complexity classification of finite-domain VCSPs and of various classes of infinite-domain CSPs, further development of the algebraic approach for infinite-domain VCSPs seems to be necessary. A natural next step in this direction would be to explore the interplay between fractional polymorphisms and expressibility (see Question 3.55) and to search for analogues of expressibility and valued relational clones on the operational side. If $\mathcal{F}$ is a set of fractional operations on the same domain A, one can define a notion of *generation* with the property that every fractional operation generated by $\mathcal{F}$ improves all valued relations in $\mathrm{Imp}(\mathcal{F})$. If we assume that A

is finite, the reverse direction holds as well: if a fractional operation improves all valued relations in $\mathrm{Imp}(\mathcal{F})$, then it is generated by $\mathcal{F}$ [60]. An analogous notion of generation over infinite domains is not known. Another important algebraic insight could come from studying cores of infinite-domain valued structures (see Question 3.56) and their properties since the concept of cores played a significant role in the finite-domain VCSP classification and is still very prominent in the area of infinite-domain CSPs.

Also from the perspective of complexity classification, there exist several natural classes of infinite-domain VCSPs that are of particular interest. This includes VCSPs that model resilience problems (see Conjecture 5.25) and VCSPs of valued structures preserved by a fixed oligomorphic permutation group such as the automorphism group of the countable random graph (see Question 4.39). Inspired by the recent results on complexity classification transfer for CSPs [16], it appears feasible to obtain results also for VCSPs on product domains, which could lead to results for optimization variants of formalisms relevant for temporal and spatial reasoning such as Cardinal Direction Calculus [3, 63], Allen's Interval Algebra [1] or Rectangle Algebra [50, 71].

Following the development of the algebraic theory for infinite-domain VCSPs and case studies of concrete classes, one might try to investigate whether the complexity of $\mathrm{VCSP}(\Gamma)$ is completely determined by the underlying crisp structure of Γ; as discussed in Section 2.4, this is a weaker statement than Conjecture 2.24. We believe that the systematic study of VCSPs preserved by an automorphism group of a finitely bounded homogeneous structure could pave the way towards a potential proof Conjecture 2.24 conditional on Conjecture 2.19.

Bibliography

[1] James F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.

[2] Antoine Amarilli, Wolfgang Gatterbauer, Neha Makhija, and Mikaël Monet. Resilience for regular path queries: Towards a complexity classification, 2024. `arXiv:2412.09411`.

[3] Philippe Balbiani and Jean-François Condotta. Spatial reasoning about points in a multidimensional setting. *Applied Intelligence*, 17(3):221–238, 2002.

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

[5] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms and the constraint satisfaction problem. *Logical Methods in Computer Science*, 8/1(07):1–26, 2012.

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

[7] Meghyn Bienvenu, Balder ten Cate, Carsten Lutz, and Frank Wolter. Ontology-based data access: A study through disjunctive Datalog, CSP, and MMSNP. *ACM Transactions of Database Systems*, 39(4):33, 2014.

[8] Manuel Bodirsky. Cores of countably categorical structures. *Logical Methods in Computer Science (LMCS)*, 3(1):1–16, 2007.

[9] Manuel Bodirsky. *Complexity of Infinite-Domain Constraint Satisfaction*. Lecture Notes in Logic (52). Cambridge University Press, 2021. `doi:10.1017/9781107337534`.

[10] Manuel Bodirsky and Bertalan Bodor. A complexity dichotomy in spatial reasoning via ramsey theory. *ACM Trans. Comput. Theory*, 16(2), 6 2024. `doi:10.1145/3649445`.

[11] Manuel Bodirsky and Bertalan Bodor. A complexity dichotomy in spatial reasoning via Ramsey theory. *ACM Trans. Comput. Theory*, 16(2), June 2024. `doi:10.1145/3649445`.

[12] Manuel Bodirsky, Édouard Bonnet, and Žaneta Semanišinová. Temporal valued constraint satisfaction problems, 2024. URL: `https://arxiv.org/abs/2409.07285`, `arXiv:2409.07285`, `doi:10.48550/arXiv.2409.07285`.

[13] Manuel Bodirsky, Johannes Greiner, and Jakub Rydval. Tractable Combinations of Temporal CSPs. *Logical Methods in Computer Science*, Volume 18, Issue 2, May 2022. URL: `https://lmcs.episciences.org/9609`, `doi:10.46298/lmcs-18(2:11)2022`.

[14] Manuel Bodirsky and Martin Grohe. Non-dichotomies in constraint satisfaction complexity. In Luca Aceto, Ivan Damgard, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, Lecture Notes in Computer Science, pages 184 –196. Springer Verlag, July 2008.

[15] Manuel Bodirsky and Peter Jonsson. A model-theoretic view on qualitative constraint reasoning. *Journal of Artificial Intelligence Research*, 58:339–385, 2017.

[16] Manuel Bodirsky, Peter Jonsson, Barnaby Martin, Antoine Mottet, and Žaneta Semanišinová. Complexity classification transfer for CSPs via algebraic products. *SIAM Journal on Computing*, 53(5):1293–1353, 2024. `doi:10.1137/22M1534304`.

[17] Manuel Bodirsky and Jan Kára. The complexity of equality constraint languages. *Theory of Computing Systems*, 3(2):136–158, 2008. A conference version appeared in the proceedings of Computer Science Russia (CSR'06). `doi:10.1007/s00224-007-9083-9`.

[18] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2):1–41, 2009. An extended abstract appeared in the Proceedings of the Symposium on Theory of Computing (STOC). `doi:10.1145/1667053.1667058`.

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

[20] Manuel Bodirsky, Marcello Mamino, and Caterina Viola. Piecewise linear valued CSPs solvable by linear programming relaxation. *ACM Transactions of Computational Logic*, 23(1):1–35, 2022. Preprint arXiv:1912.09298. `doi:10.1145/3488721`.

[21] Manuel Bodirsky and Antoine Mottet. A dichotomy for first-order reducts of unary structures. *Logical Methods in Computer Science*, 14(2), 2018. URL: `https://doi.org/10.23638/LMCS-14(2:13)2018`.

[22] Manuel Bodirsky and Jens K. Mueller. Rooted phylogeny problems. *Logical Methods in Computer Science*, 7(4), 2011. An extended abstract appeared in the proceedings of ICDT'10.

[23] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. In *Proceedings of Computer Science Logic (CSL)*, pages 44–57, Vienna, 2003.

[24] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006. `doi:10.1093/logcom/exi083`.

[25] Manuel Bodirsky, Wied Pakusa, and Jakub Rydval. Temporal constraint satisfaction problems in fixed-point logic. In Holger Hermanns, Lijun Zhang, Naoki Kobayashi, and Dale Miller, editors, *LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Saarbrücken, Germany, July 8-11, 2020*, pages 237–251. ACM, 2020. `doi:10.1145/3373718.3394750`.

[26] Manuel Bodirsky and Michael Pinsker. Schaefer's theorem for graphs. *Journal of the ACM*, 62(3):52 pages (article number 19), 2015. A conference version appeared in the Proceedings of STOC 2011, pages 655-664. `doi:10.1145/2764899`.

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

[28] Manuel Bodirsky, Michael Pinsker, and Todor Tsankov. Decidability of definability. *Journal of Symbolic Logic*, 78(4):1036–1054, 2013. A conference version appeared in the Proceedings of the Twenty-Sixth Annual IEEE Symposium on. Logic in Computer Science (LICS 2011), pages 321-328.

[29] Manuel Bodirsky and Jakub Rydval. On the descriptive complexity of temporal constraint satisfaction problems. *J. ACM*, 70(1), dec 2022. `doi:10.1145/3566051`.

[30] Manuel Bodirsky, Žaneta Semanišinová, 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`.

[31] Manuel Bodirsky, Žaneta Semanišinová, and Carsten Lutz. The complexity of resilience problems via valued constraint satisfaction problems, 2023. `arXiv:2309.15654`.

[32] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17*, pages 319–330, 2017.

[33] Andrei A. Bulatov, Andrei A. Krokhin, and Peter G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.

[34] Silvia Butti. *Symmetries in Constraint Satisfaction: Weisfeiler-Leman Invariance and Promise Problems.* PhD thesis, Universitat Pompeu Fabra, Barcelona, 2022.

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

[36] Vaggos Chatziafratis and Konstantin Makarychev. Triplet reconstruction and all other phylogenetic CSPs are approximation resistant. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 253–284. IEEE, 2023. doi:10.1109/FOCS57990.2023.00024.

[37] Surajit Chaudhuri and Moshe Y. Vardi. Optimization of real conjunctive queries. In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, PODS '93, pages 59–70, New York, NY, USA, 1993. Association for Computing Machinery. doi:10.1145/153850.153856.

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

[39] Artem Chernikov, Daniel Palacin, and Kota Takeuchi. On n-Dependence. *Notre Dame Journal of Formal Logic*, 60(2):195 – 214, 2019. doi:10.1215/00294527-2019-0002.

[40] 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.

[41] David A. Cohen, Martin C. Cooper, and Peter Jeavons. 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*, pages 107–121, 2006. doi:10.1007/11889205_10.

[42] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM Journal on Computing*, 28:57–104, 1999.

[43] Jan Foniok. Homomorphisms and structural properties of relational systems. Doctoral Thesis, Charles University Prague, Faculty of Mathematics and Physics, 2007.

[44] Cibele Freire, Wolfgang Gatterbauer, Neil Immerman, and Alexandra Meliou. A characterization of the complexity of resilience and responsibility for self-join-free conjunctive queries. *CoRR*, abs/1507.00674, 2015. URL: http://arxiv.org/abs/1507.00674, arXiv:1507.00674.

[45] 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. URL: `http://www.vldb.org/pvldb/vol9/p180-freire.pdf`, `doi:10.14778/2850583.2850592`.

[46] 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`.

[47] Peter Fulla and Stanislav Živný. A Galois connection for weighted (relational) clones of infinite size. *TOCT*, 8(3):9:1–9:21, 2016.

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

[49] Andrew V. Goldberg and Robert E. Tarjan. A new approach to the maximum-flow problem. *J. Assoc. Comput. Mach.*, 35(4):921–940, 1988. `doi:10.1145/48014.61051`.

[50] Hans Guesgen. Spatial reasoning based on Allen's temporal logic. Report ICSI TR 89-049, International Computer Science Institute, 1989.

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

[52] J. Håstad, R. Manokaran, P. Raghavendra, and M. Charikar. Beating the random ordering is hard: Every ordering CSP is approximation resistant. *SIAM Journal on Computing*, 40(3):878–914, 2011.

[53] Jan Hubička and Jaroslav Nešetřil. Universal structures with forbidden homomorphisms. In *Logic Without Borders - Essays on Set Theory, Model Theory, Philosophical Logic and Philosophy of Mathematics*, pages 241–264. De Gruyter, 2015. arXiv:0907.4079.

[54] Jan Hubička and Jaroslav Nešetřil. Homomorphism and embedding universal structures for restricted classes. *Multiple-Valued Logic and Soft Computing*, 27(2-3):229–253, 2016. arXiv:0909.4939.

[55] Phokion G. Kolaitis and Moshe Y. Vardi. Conjunctive-query containment and constraint satisfaction. In *Proceedings of Symposium on Principles of Database Systems (PODS)*, pages 205–213, 1998.

[56] Vladimir Kolmogorov. Testing the Complexity of a Valued CSP Language. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*,

pages 77:1–77:12, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2019.77.

[57] Vladimir Kolmogorov, Andrei A. Krokhin, and Michal Rolínek. The complexity of general-valued CSPs. *SIAM J. Comput.*, 46(3):1087–1110, 2017. doi:10.1137/16M1091836.

[58] Marcin Kozik and Joanna Ochremiak. 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*, pages 846–858, 2015.

[59] Marcin Kozik and Joanna Ochremiak. Algebraic properties of valued constraint satisfaction problem, 2015. URL: https://arxiv.org/abs/1403.0476v3, arXiv:1403.0476v3.

[60] Andrei Krokhin and Stanislav Živný. The Complexity of Valued CSPs. In *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 233–266. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017. URL: http://drops.dagstuhl.de/opus/volltexte/2017/6966, doi:10.4230/DFU.Vol7.15301.233.

[61] Richard E. Ladner. On the structure of polynomial time reducibility. *Journal of the ACM*, 22(1):155–171, 1975.

[62] Benoit Larose, Cynthia Loten, and Claude Tardif. A characterisation of first-order constraint satisfaction problems. *Logical Methods in Computer Science*, 3(4:6), 2007.

[63] Gérard Ligozat. Reasoning about cardinal directions. *Journal of Visual Languages & Computing*, 9(1):23 – 44, 1998. URL: http://www.sciencedirect.com/science/article/pii/S1045926X97999997, doi:10.1006/jvlc.1997.9999.

[64] Florent Madelaine and Iain A. Stewart. Constraint satisfaction, logic and forbidden patterns. *SIAM Journal on Computing*, 37(1):132–163, 2007.

[65] 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), dec 2023. doi:10.1145/3626715.

[66] Alexandra Meliou, Wolfgang Gatterbauer, Katherine F. Moore, and Dan Suciu. The complexity of causality and responsibility for query answers and non-answers. *Proc. VLDB Endow.*, 4(1):34–45, 2010. URL: http://www.vldb.org/pvldb/vol4/p34-meliou.pdf, doi:10.14778/1880172.1880176.

[67] Alexandra Meliou, Wolfgang Gatterbauer, and Dan Suciu. Reverse data management. *Proc. VLDB Endow.*, 4(12):1490 – 1493, August 2011. doi:10.14778/3402755.3402803.

[68] Antoine Mottet, Tomáš Nagy, and Michael Pinsker. An Order out of Nowhere: A New Algorithm for Infinite-Domain CSPs. In *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 148:1–148:18, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: `https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2024.148`, `doi:10.4230/LIPIcs.ICALP.2024.148`.

[69] Antoine Mottet and Michael Pinsker. Cores over Ramsey structures. *Journal of Symbolic Logic*, 86(1):352–361, 2021.

[70] Antoine Mottet and Michael Pinsker. Smooth approximations: An algebraic approach to CSPs over finitely bounded homogeneous structures. *Journal of the ACM*, 71(5), October 2024. `doi:10.1145/3689207`.

[71] Amitabha Mukerjee and Gene Joe. A qualitative model for space. In *Proc. 8th National Conference on Artificial Intelligence (AAAI-1990)*, pages 721–727, 1990.

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

[73] George Osipov and Magnus Wahlström. Parameterized complexity of equality MinCSP. In Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman, editors, *31st Annual European Symposium on Algorithms, ESA 2023, September 4-6, 2023, Amsterdam, The Netherlands*, volume 274 of *LIPIcs*, pages 86:1–86:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. URL: `https://doi.org/10.4230/LIPIcs.ESA.2023.86`, `doi:10.4230/LIPICS.ESA.2023.86`.

[74] George Osipov and Magnus Wahlström. Parameterized complexity of equality MinCSP. In *31st Annual European Symposium on Algorithms (ESA 2023)*, volume 274 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 86:1–86:17, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: `https://drops-dev.dagstuhl.de/entities/document/10.4230/LIPIcs.ESA.2023.86`, `doi:10.4230/LIPIcs.ESA.2023.86`.

[75] Jakub Rydval. *Using Model Theory to Find Decidable and Tractable Description Logics with Concrete Domains.* PhD thesis, Dresden University of Technology, Germany, 2022. URL: `https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa2-799074`.

[76] Jakub Rydval, Žaneta Semanišinová, and Michał Wrona. Identifying Tractable Quantified Temporal Constraints Within Ord-Horn. In *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 151:1–151:20, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. `doi:10.4230/LIPIcs.ICALP.2024.151`.

[77] 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. URL: `https://www.sciencedirect.com/science/article/pii/S0022247X22006059`, `doi:10.1016/j.jmaa.2022.126591`.

[78] Pierre Simon. *A Guide to NIP Theories*. Lecture Notes in Logic. Cambridge University Press, 2015.

[79] 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.

[80] Caterina Viola. *Valued Constraint Satisfaction Problems over Infinite Domains*. PhD thesis, TU Dresden, 2020.

[81] 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`.

[82] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. `doi:10.1145/3402029`.

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

List of Figures

I herewith declare that I have produced this thesis without the prohibited assistance of third parties and without making use of aids other than those specified; notions taken over directly or indirectly from other sources have been identified as such. This thesis has not previously been presented in identical or similar form to any other German or foreign examination board.

Žaneta Semanišinová
Dresden, 24 January 2025