

Hardness of Network Satisfaction for Relation Algebras with Normal Representations

Manuel Bodirsky* and Simon Knäuer **

Institut für Algebra, TU Dresden, 01062 Dresden, Germany

Abstract. We study the computational complexity of the general network satisfaction problem for a finite relation algebra A with a normal representation B . If B contains a non-trivial equivalence relation with a finite number of equivalence classes, then the network satisfaction problem for A is NP-hard. As a second result, we prove hardness if B has domain size at least three and contains no non-trivial equivalence relations but a symmetric atom a with a forbidden triple (a, a, a) , that is, $a \not\leq a \circ a$. We illustrate how to apply our conditions on two small relation algebras.

1 Introduction

Many computational problems in temporal and spatial reasoning can be formulated as network satisfaction problems for a fixed finite relation algebra [Dün05, RN07, BJ17]. Famous examples of finite relation algebras that have been studied in this context are the Point Algebra, the Left Linear Point Algebra, Allen's Interval Algebra, RCC5, and RCC8, just to name a few; much more material about relation algebras can be found in [HH02]. Robin Hirsch [Hir96] asked in 1996 the *Really Big Complexity Problem (RBCP)*: can we classify the computational complexity of the network satisfaction problem for every finite relation algebra? For example, the network satisfaction problem for the Point Algebra and the Left Linear Point Algebra are polynomial-time tractable [VKvB89, BK07], while it is NP-complete for the other relation algebras mentioned above [All83, RN99]. A finite relation algebra with an undecidable network satisfaction problem has been found by Hirsch [Hir99].

An important notion in the theory of representability of finite relation algebras are *normal representations*, i.e., representations that are fully universal, square, and homogeneous [Hir96]. The network satisfaction problem for a relation algebra with a normal representation can be seen as the constraint satisfaction problem for an infinite structure \mathfrak{B} that is homogeneous and finitely bounded (these concepts from model theory will be introduced in Section 3). The network satisfaction problem is in this case in NP and a complexity dichotomy has been conjectured [BPP14]. There is even a promising candidate condition for the boundary between NP-completeness and containment in P; the condition can be phrased in several equivalent ways [BKO⁺17, Bod18]. However, this conjecture has not yet been verified for the homogeneous finitely bounded structures that arise as the normal representation of a finite relation algebra.

We present some first steps towards a solution to the RBCP for relation algebras \mathbf{A} with a normal representation \mathfrak{B} . Our approach is to study the automorphism group $\text{Aut}(\mathfrak{B})$ of \mathfrak{B} and to identify properties that imply hardness. Because of the homogeneity of \mathfrak{B} , one can translate back and forth between properties of \mathbf{A} and properties of $\text{Aut}(\mathfrak{B})$. For example, $\text{Aut}(\mathfrak{B})$ is primitive if and only if \mathbf{A} contains no equivalence relation which is different from the *trivial* equivalence relations Id and 1 . Specifically, we show that the network satisfaction problem for \mathbf{A} is NP-complete if

- $\text{Aut}(\mathfrak{B})$ is primitive, $|B| > 2$ and \mathbf{A} has a symmetric atom a with a forbidden triple (a, a, a) , that is, $a \not\leq a \circ a$ (Section 5);

* The author has received funding from the European Research Council (Grant Agreement no. 681988, CSP-Infinity)

** The author is supported by DFG Graduiertenkolleg 1763 (QuantLA).

- $\text{Aut}(\mathfrak{B})$ has a congruence with at least two but finitely many equivalence classes (Section 6).

In our proof we use the so-called *universal-algebraic approach* which has recently led to a full classification of the computational complexity of constraint satisfaction problems for \mathfrak{B} if the domain of \mathfrak{B} is *finite* [Bul17, Zhu17]. The central insight is that the complexity of the CSP is for finite \mathfrak{B} fully determined by the *polymorphism clone* $\text{Pol}(\mathfrak{B})$ of \mathfrak{B} . This result extends to homogeneous structures with finite relational signature (more generally, to ω -categorical structures [BN06]). Both of our hardness proofs come from the technique of factoring $\text{Pol}(\mathfrak{B})$ with respect to a congruence with finitely many classes, and using known hardness conditions from corresponding finite-domain constraint satisfaction problems. The article is fully self-contained: we introduce the network satisfaction problem (Section 2), normal representations (Section 3), and the universal algebraic approach (Section 4).

2 The (General) Network Satisfaction Problem

Network satisfaction problems have been introduced in [LM94], capturing well-known computational problems, e.g., for Allen’s Interval Algebra [All83]; see [Dün05] for a survey. An *algebra* in the sense of universal algebra is a set together with operations on this set, each equipped with an arity $n \in \mathbb{N}$. In this context, operations of arity zero are viewed as constants. The *type* of an algebra is a tuple that represents the arities of the operations. For the definitions concerning relation algebras, we basically follow [Mad06].

Definition 1. Let D be a set and $E \subseteq D^2$ an equivalence relation. Let $(\mathcal{P}(E); \cup, \bar{\cdot}, 0, 1, \text{Id}, \bar{\cdot}, \circ)$ be an algebra of type $(2, 1, 0, 0, 0, 1, 2)$ with the following operations:

1. $A \cup B := \{(x, y) \mid (x, y) \in A \text{ or } (x, y) \in B\}$,
2. $\bar{A} := E \setminus A$,
3. $0 := \emptyset$,
4. $1 := E$,
5. $\text{Id} := \{(x, x) \mid x \in D\}$,
6. $A^\sim := \{(x, y) \mid (y, x) \in A\}$,
7. $A \circ B := \{(x, z) \mid \exists y \in D : (x, y) \in A \text{ and } (y, z) \in B\}$.

A subalgebra of $(\mathcal{P}(E); \cup, \bar{\cdot}, 0, 1, \text{Id}, \bar{\cdot}, \circ)$ is called a proper relation algebra.

A *representable relation algebra* is an algebra of type $(2, 1, 0, 0, 0, 1, 2)$ that is isomorphic (as an algebra) to a proper relation algebra. We denote algebras by bold letters, like \mathbf{A} ; the underlying domain of an algebra \mathbf{A} is denoted with the regular letter A . An algebra \mathbf{A} is finite if A is finite. We do not need the more general definition of an *(abstract) relation algebra* (for a definition see for example [Mad06]) because the network satisfaction problem for relation algebras that are not representable is trivial. We use the language of model theory to define *representations* of relation algebras; the definition is essentially the same as the one given in [Mad06].

Definition 2. A relational structure \mathfrak{B} is called a representation of a relation algebra \mathbf{A} if

- \mathfrak{B} is an A -structure with domain B (i.e., each element $a \in A$ is used as a relation symbol denoting a binary relation $a^{\mathfrak{B}}$ on B);
- there exists an equivalence relation $E \subseteq B^2$ such that the set of relations of \mathfrak{B} is the domain of a subalgebra of $(\mathcal{P}(E); \cup, \bar{\cdot}, 0, 1, \text{Id}, \bar{\cdot}, \circ)$;
- the map that sends $a \in A$ to $a^{\mathfrak{B}}$ is an isomorphism between \mathbf{A} and this subalgebra.

Remark 3. For a relation algebra $\mathbf{A} = (A; \cup, \bar{\cdot}, 0, 1, \text{Id}, \bar{\cdot}, \circ)$ the algebra $(A; \cup, \bar{\cdot}, 0, 1)$ is a Boolean algebra. With respect to this algebra there is a partial ordering on the elements of a relation algebra. We denote this with \subseteq since in proper relation algebras this ordering is with respect to set inclusion. The minimal non-empty relations with respect to \subseteq are called the *atomic relations* or *atoms*; we denote the set of atoms of \mathbf{A} by A_0 .

Definition 4. Let \mathbf{A} be a relation algebra. An \mathbf{A} -network $(V; f)$ is a finite set of nodes V together with a function $f: V \times V \rightarrow A$.

Let \mathfrak{B} be a representation of \mathbf{A} . An \mathbf{A} -network $(V; f)$ is satisfiable in \mathfrak{B} if there exists an assignment $s: V \rightarrow B$ such that for all $x, y \in V$

$$(s(x), s(y)) \in f(x, y)^{\mathfrak{B}}.$$

An \mathbf{A} -network $(V; f)$ is satisfiable if there exists some representation \mathfrak{B} of \mathbf{A} such that $(V; f)$ is satisfiable in \mathfrak{B} .

Definition 5. The (general) network satisfaction problem for a finite relation algebra \mathbf{A} , denoted by $\text{NSP}(\mathbf{A})$, is the problem of deciding whether a given \mathbf{A} -network is satisfiable.

3 Normal Representations and CSPs

We recall a connection between network satisfaction problems and constraint satisfaction problems that is presented in more detail in [BJ17, Bod18].

Definition 6 (from [Hir96]). Let \mathbf{A} be a relation algebra. An \mathbf{A} -network $(V; f)$ is called atomic if the image of f only contains atoms and if

$$f(a, c) \subseteq f(a, b) \circ f(b, c).$$

The last line ensures a “local consistency” of the atomic \mathbf{A} -network with respect to the multiplication rules in the relation algebra \mathbf{A} . This property is in the literature sometimes called “closedness” of an \mathbf{A} -network [Hir97].

Definition 7 (from [Hir96]). A representation \mathfrak{B} of a relation algebra \mathbf{A} is called

- fully universal if every atomic \mathbf{A} -network is satisfiable in \mathfrak{B} ;
- square if $1^{\mathfrak{B}} = B^2$;
- homogeneous if every isomorphism of finite substructures of \mathfrak{B} can be extended to an automorphism;
- normal if it is fully universal, square and homogeneous.

If a relation algebra \mathbf{A} has a normal representation \mathfrak{B} then the problem of deciding whether an \mathbf{A} -network is satisfiable in *some* representation reduces to a question whether it is satisfiable in the concrete representation \mathfrak{B} . Such decision problems are known as constraint satisfaction problems, which are formally defined in the following.

Definition 8. Let \mathfrak{B} be a τ -structure for a finite relational signature τ . The constraint satisfaction problem of \mathfrak{B} is the problem of deciding for a given finite τ -structure \mathfrak{C} whether there exists a homomorphism from \mathfrak{C} to \mathfrak{B} .

To formulate the connection between NSPs and CSPs, we have to give a translation between networks and structures. On the one hand we may view an \mathbf{A} -network $(V; f)$ as an A -structure \mathfrak{C} with domain $C := V$ where $(a, b) \in f(a, b)^{\mathfrak{C}}$. On the other hand we can transform an A -structure \mathfrak{C} into an \mathbf{A} -network $(V; f)$ with $V = C$ and by defining the network function $f(x, y)$ for $x, y \in C$ as follows: let X be the set of all relations that hold on (x, y) in \mathfrak{C} . If X is non-empty we define $f(x, y) := \bigcup X$; otherwise $f(x, y) := 1$.

Proposition 9 (see [Bod18]). Let \mathfrak{B} be a normal representation of a finite relation algebra \mathbf{A} . Then $\text{NSP}(\mathbf{A})$ and $\text{CSP}(\mathfrak{B})$ are the same problem (up to the translation showed above).

The following is an important notion in model theory and the study of infinite-domain CSPs. Let \mathcal{F} be a finite set of finite τ -structures. Then $\text{Forb}(\mathcal{F})$ is the class of all finite τ -structures that embed no $\mathfrak{C} \in \mathcal{F}$. A class \mathcal{C} of finite τ -structures is called *finitely bounded* if $\mathcal{C} = \text{Forb}(\mathcal{F})$ for a finite set \mathcal{F} . A structure \mathfrak{B} is called *finitely bounded* if the class of finite structures that embed into \mathfrak{B} is finitely bounded.

Proposition 10 (see [Bod18]). Let \mathbf{A} be a finite relation algebra with a normal representation \mathfrak{B} . Then \mathfrak{B} is finitely bounded and $\text{CSP}(\mathfrak{B})$ and $\text{NSP}(\mathbf{A})$ are in NP.

4 The Universal Algebraic Approach

This section gives a short overview of the important notions and concepts for the universal-algebraic approach to the computational complexity of CSPs.

4.1 Clones

We start with the definition of an operation clone.

Definition 11. Let B be some set. Then $\mathcal{O}_B^{(n)}$ denotes the set of n -ary operations on B and $\mathcal{O}_B := \bigcup_{n \in \mathbb{N}} \mathcal{O}_B^{(n)}$. A set $\mathcal{C} \subseteq \mathcal{O}_B$ is called a operation clone (on B) if it contains all projections and is closed under composition, that is, for every $f \in \mathcal{C} \cap \mathcal{O}_B^{(k)}$ and all $g_1, \dots, g_k \in \mathcal{C} \cap \mathcal{O}_B^{(n)}$ the n -ary operation $f(g_1, \dots, g_k)$ with

$$f(g_1, \dots, g_k)(x_1, \dots, x_n) := f(g_1(x_1, \dots, x_n), \dots, g_k(x_1, \dots, x_n))$$

is also in \mathcal{C} . We denote the k -ary operations of \mathcal{C} by $\mathcal{C}^{(k)}$.

Definition 12. Let \mathfrak{B} be a relational structure. Then f preserves a relation R of \mathfrak{B} if the component-wise application of f on tuples $r_1, \dots, r_k \in R$ results in a tuple of the relation. If f preserves all relations of \mathfrak{B} then f is called a polymorphism of \mathfrak{B} . The set of all polymorphisms of arity $k \in \mathbb{N}$ is denoted by $\text{Pol}^{(k)}(\mathfrak{B})$ and $\text{Pol}(\mathfrak{B}) := \bigcup_{k \in \mathbb{N}} \text{Pol}^{(k)}(\mathfrak{B})$ is called the polymorphism clone of \mathfrak{B} .

Polymorphisms are closed under the composition and a projection is always a polymorphism, therefore a polymorphism clone is indeed an operation clone.

Definition 13. Let \mathcal{C} and \mathcal{D} be operation clones. A function $\mu: \mathcal{C} \rightarrow \mathcal{D}$ is called minor-preserving if it maps every operation to an operation of the same arity and satisfies for every $f \in \mathcal{C}^{(k)}$ and all projections p_1, \dots, p_k the following identity:

$$\mu(f(p_1, \dots, p_k)) = \mu(f)(p_1, \dots, p_k).$$

Operation clones \mathcal{C} on countable sets B can be equipped with the following complete ultrametric d . Assume that $B = \mathbb{N}$. For two polymorphisms f and g of different arity we define $d(f, g) = 1$. If f and g are both of arity k we have

$$d(f, g) := 2^{-\min\{n \in \mathbb{N} \mid \exists s \in \{1, \dots, n\}^k : f(s) \neq g(s)\}}.$$

The following is a straightforward consequence of the definition.

Lemma 14. Let \mathcal{D} be an operation clone on B and \mathcal{C} an operation clone on C and let $\nu: \mathcal{D} \rightarrow \mathcal{C}$ a map. Then ν is uniformly continuous (u.c.) if and only if

$$\forall n \geq 1 \exists \text{ finite } F \subset D \forall f, g \in \mathcal{D}^{(n)} : f|_F = g|_F \Rightarrow \nu(f) = \nu(g).$$

In order to demonstrate the use of polymorphisms in the study of CSPs we have to define primitive positive formulas. Let τ be a relational signature. A first-order formula $\varphi(x_1, \dots, x_n)$ is called *primitive positive* if it has the form

$$\exists x_{n+1}, \dots, x_m (\varphi_1 \wedge \dots \wedge \varphi_s)$$

where $\varphi_1, \dots, \varphi_s$ are atomic formulas, i.e., formulas of the form $R(y_1, \dots, y_l)$ for $R \in \tau$ and $y_i \in \{x_1, \dots, x_m\}$, of the form $y = y'$ for $y, y' \in \{x_1, \dots, x_m\}$, or of the form *false* and *true*. We have the following correspondence between polymorphisms and primitive positive formulas (or relations that are defined by them). Note that all of the statements in the following hold in a more general setting, but we only state them here for normal representations of finite relation algebras.

Theorem 15 (follows from [BN06]). *Let \mathfrak{B} be a normal representation of a finite relation algebra \mathbf{A} . Then the set of primitive positive definable relations in \mathfrak{B} is exactly the set of relations that are preserved by $\text{Pol}(\mathfrak{B})$.*

A special type of polymorphism plays an important role in our analysis.

Definition 16. *Let f be an n -ary operation on a countable set X . Then f is called cyclic if*

$$\forall x_1, \dots, x_n \in X : f(x_1, \dots, x_n) = f(x_n, x_1, \dots, x_{n-1}).$$

We write Proj for the operation clone on a two-element set that consists of only the projections.

Theorem 17 (from [BK12, BOP18]). *Let \mathcal{C} be an operation clone on a finite set C . If there exists no minor-preserving map $\mathcal{C} \rightarrow \text{Proj}$ then \mathcal{C} contains for every prime $p > |C|$ a p -ary cyclic operation.*

Note that every map between operation clones on finite domains is uniformly continuous.

Theorem 18 (from [BOP18]). *Let \mathfrak{B} be normal representation of a finite relation algebra. If there is a uniformly continuous minor-preserving map $\text{Pol}(\mathfrak{B}) \rightarrow \text{Proj}$, then $\text{CSP}(\mathfrak{B})$ is NP-complete.*

4.2 Canonical Functions

Let \mathfrak{B} be a normal representation of a finite relation algebra \mathbf{A} .

Definition 19. *Let $a_1, \dots, a_k \in A$. Then $(a_1, \dots, a_k)^{\mathfrak{B}}$ denotes a binary relation on B^k such that for $x, y \in B^k$*

$$(a_1, \dots, a_k)^{\mathfrak{B}}(x, y) : \Leftrightarrow \bigwedge_{i \in \{1, \dots, k\}} a_i^{\mathfrak{B}}(x_i, y_i).$$

Recall that A_0 denotes the set of atoms of a representable relation algebra \mathbf{A} .

Definition 20. *Let $x, y \in B^k$. Since \mathfrak{B} is square there are unique $a_1, \dots, a_k \in A_0$ such that $(a_1, \dots, a_k)^{\mathfrak{B}}(x, y)$. Then we call $(a_1, \dots, a_k)^{\mathfrak{B}}$ the configuration of (x, y) . If $a_1, \dots, a_k \in X \subseteq A_0$ then (a_1, \dots, a_k) is called an X -configuration.*

We specialise the concept of *canonical functions* (see, e.g., [BP16]) to our setting.

Definition 21. *Let f be a k -ary operation on B . Let $X \subseteq A_0$ and let T be the set of all X -configurations. Then f is called X -canonical if there exists a map $\bar{f}: T \rightarrow A_0$ such that for every $(a_1, \dots, a_k) \in T$ and $(x, y) \in (a_1, \dots, a_k)^{\mathfrak{B}}$ we have $(f(x), f(y)) \in (\bar{f}(a_1, \dots, a_k))^{\mathfrak{B}}$. If $X = A_0$ then f is called canonical.*

An operation $f: B^n \rightarrow B$ is called *conservative* if for all $x_1, \dots, x_n \in B$

$$f(x_1, \dots, x_n) \in \{x_1, \dots, x_n\}.$$

If \mathfrak{B} is a finite structure such that every polymorphism of \mathfrak{B} is conservative, then $\text{CSP}(\mathfrak{B})$ has been classified already before the proof of the Feder-Vardi conjecture, and there are several proofs [Bul03, Bul14, Bar11]. The polymorphisms of normal representations of finite relation algebras satisfy a strong property that resembles conservativity.

Proposition 22. *Let \mathfrak{B} be a normal representation. Then every $f \in \text{Pol}^{(n)}$ is edge-conservative, that is, for all $x, y \in B^n$ with configuration $(a_1, \dots, a_n)^{\mathfrak{B}}$ it holds that*

$$(f(x), f(y)) \in \left(\bigcup_{i \in \{1, \dots, n\}} a_i \right)^{\mathfrak{B}}.$$

Proof. By definition, $b := \bigcup_{i \in \{1, \dots, n\}} a_i$ is part of the signature of \mathfrak{B} . Moreover, for every $i \in \{1, \dots, n\}$ we have that $(x_i, y_i) \in b^{\mathfrak{B}}$ by the assumption on the configuration of x and y . Then $(f(x), f(y)) \in b^{\mathfrak{B}}$ because f preserves $b^{\mathfrak{B}}$. \square

5 Finitely Many Equivalence Classes

In the following, \mathbf{A} denotes a finite relation algebra with a normal representation \mathfrak{B} .

Theorem 23. *Suppose that $e \in A$ is such that $e^{\mathfrak{B}}$ is a non-trivial equivalence relation with finitely many classes. Then $\text{CSP}(\mathfrak{B})$ is NP-complete.*

Proof. We use the notation $n := 1 \setminus e$. Let $\{c_1, \dots, c_m\}$ be a set of representatives of the equivalence classes of $e^{\mathfrak{B}}$. We denote the equivalence class of c_i by $\overline{c_i}$. A k -ary polymorphism $f \in \text{Pol}(\mathfrak{B})$ induces an operation \overline{f} of arity k on $C = \{\overline{c_1}, \dots, \overline{c_m}\}$ in the following way:

$$\overline{f}(\overline{d_1}, \dots, \overline{d_k}) := \overline{f(d_1, \dots, d_k)}$$

for all $\overline{d_1}, \dots, \overline{d_k} \in \{\overline{c_1}, \dots, \overline{c_m}\}$. This definition is independent from the choice of the representatives since the polymorphisms preserve the relation $e^{\mathfrak{B}}$. We denote the set of all operations that are induced in this way by operations from $\text{Pol}(\mathfrak{B})$ by \mathcal{C} . It is easy to see that \mathcal{C} is an operation clone on a finite set. Moreover, the mapping $\mu: \text{Pol}(\mathfrak{B}) \rightarrow \mathcal{C}$ defined by $\mu(f) := \overline{f}$ is a minor-preserving map. To show that μ is uniformly continuous, we use Lemma 14; it suffices to observe that if two k -ary operations $f, g \in \text{Pol}(\mathfrak{B})$ are equal on $F := \{c_1, \dots, c_m\}$, then they induce the same operation on the equivalence classes.

Suppose for contradiction that \mathcal{C} contains a p -ary cyclic operation for every prime $p > m$.

Case 1: $m = 2$. By assumption there exists a ternary cyclic operation $\overline{f} \in \mathcal{C}$. Since $e^{\mathfrak{B}}$ is non-trivial, one of the equivalence classes of $e^{\mathfrak{B}}$ must have size at least two. So we may without loss of generality assume that $\overline{c_1}$ contains at least two elements. Let $c'_1 \in \overline{c_1}$ with $c_1 \neq c'_1$. We have that $\overline{f}(c_1, c_1, c_2) = \overline{f}(c_2, c_1, c_1)$ which means that

$$(f(c_1, c_1, c_2), f(c_2, c_1, c_1)) \in e^{\mathfrak{B}}. \quad (1)$$

On the other hand $(n, \text{Id}, n)^{\mathfrak{B}}((c_1, c_1, c_2), (c_2, c_1, c_1))$. Since f is an edge conservative polymorphism we have that

$$(f(c_1, c_1, c_2), f(c_2, c_1, c_1)) \in (n \cup \text{Id})^{\mathfrak{B}}. \quad (2)$$

Combining (1) and (2) we obtain that

$$f(c_1, c_1, c_2) = f(c_2, c_1, c_1). \quad (3)$$

Similarly, $\overline{f}(c_2, c_1, c_1) = \overline{f}(c_1, c_2, c_1)$. Since f preserves the equivalence relation $e^{\mathfrak{B}}$ we also have $(f(c_1, c_2, c_1), f(c'_1, c_2, c_1)) \in e^{\mathfrak{B}}$. But then $(f(c_2, c_1, c_1), f(c'_1, c_2, c_1)) \in e^{\mathfrak{B}}$ holds. Also note that $(n, n, \text{Id})^{\mathfrak{B}}((c_2, c_1, c_1), (c'_1, c_2, c_1))$ implies that $(f(c_2, c_1, c_1), f(c'_1, c_2, c_1)) \in (n \cup \text{Id})^{\mathfrak{B}}$. These two facts together imply $f(c_2, c_1, c_1) = f(c'_1, c_2, c_1)$. By (3) and the transitivity of equality we get $f(c_1, c_1, c_2) = f(c'_1, c_2, c_1)$. But this is impossible because $(e, n, n)^{\mathfrak{B}}((c_1, c_1, c_2), (c'_1, c_2, c_1))$ implies that $f(c_1, c_1, c_2) \neq f(c'_1, c_2, c_1)$.

Case 2: $m > 2$. Let f be a p -ary cyclic operation for some prime $p > m$. Consider the representatives c_1, c_2 and c_3 . By the cyclicity of \overline{f} we have

$$\overline{f}(c_1, c_2, \dots, c_1, c_2, c_3) = \overline{f}(c_3, c_1, c_2, \dots, c_1, c_2)$$

and therefore

$$(f(c_1, c_2, \dots, c_1, c_2, c_3), f(c_3, c_1, c_2, \dots, c_1, c_2)) \in e^{\mathfrak{B}}. \quad (4)$$

On the other hand,

$$(n, n, n, \dots, n, n)^{\mathfrak{B}}((c_1, c_2, \dots, c_1, c_2, c_3), (c_3, c_1, c_2, \dots, c_1, c_2))$$

and since f preserves $n^{\mathfrak{B}}$ we get that

$$(f(c_1, c_2, \dots, c_1, c_2, c_3), f(c_3, c_1, c_2, \dots, c_1, c_2)) \in n^{\mathfrak{B}},$$

contradicting (4).

We showed that there exists a prime $p > m$ such that \mathcal{C} does not contain a p -ary cyclic polymorphism and therefore Theorem 17 implies the existence of a (uniformly continuous) minor-preserving map $\nu: \mathcal{C} \rightarrow \text{Proj}$. Since the composition of uniformly continuous minor-preserving maps is again uniformly continuous and minor-preserving, there exists a uniformly continuous minor-preserving map $\nu \circ \mu: \text{Pol}(\mathfrak{B}) \rightarrow \text{Proj}$. This map implies the NP-hardness of $\text{CSP}(\mathfrak{B})$ by Theorem 18. \square

6 No Non-Trivial Equivalence Relations

In this section \mathbf{A} denotes a finite relation algebra with a normal representation \mathfrak{B} with $|B| > 2$.

Definition 24. *The automorphism group $\text{Aut}(\mathfrak{C})$ of a relational structure \mathfrak{C} is called primitive if $\text{Aut}(\mathfrak{C})$ does not preserve a non-trivial equivalence relation, i.e., the only equivalence relations that are preserved by $\text{Aut}(\mathfrak{C})$ are Id and C^2 .*

Proposition 25. *Let a be an atom of \mathbf{A} . If $\text{Aut}(\mathfrak{B})$ is primitive then $a \subseteq \text{Id}$ implies $a = \text{Id}$.*

Proof. If $a \subsetneq \text{Id}$ then

$$c := \text{Id} \cup (a \circ 1 \circ a)$$

would be such that $c^{\mathfrak{B}}$ is a non-trivial equivalence relation. \square

Proposition 26. *Let a be a symmetric atom of \mathbf{A} with $a \cap \text{Id} = 0$. If $\text{Aut}(\mathfrak{B})$ is primitive then $a^{\mathfrak{B}} \circ a^{\mathfrak{B}} \neq \text{Id}$.*

Proof. Assume for contradiction $a^{\mathfrak{B}} \circ a^{\mathfrak{B}} = \text{Id}^{\mathfrak{B}}$. This implies $(\text{Id} \cup a)^{\mathfrak{B}} \circ (\text{Id} \cup a)^{\mathfrak{B}} \subset (\text{Id} \cup a)^{\mathfrak{B}}$ and therefore $(\text{Id} \cup a)^{\mathfrak{B}}$ is an equivalence relation. Since \mathfrak{B} is primitive $(\text{Id} \cup a)^{\mathfrak{B}} = B^2$. By assumption B contains at least 3 elements. These elements are now all connected by the atomic relation $a^{\mathfrak{B}}$. This is a contradiction to our assumption $a^{\mathfrak{B}} \circ a^{\mathfrak{B}} = \text{Id}^{\mathfrak{B}}$. \square

Higman's lemma states that a permutation group G on a set B is primitive if and only if for every two distinct elements $x, y \in B$ the undirected graph with vertex set B and edge set $\{\{\alpha(x), \alpha(y)\} \mid \alpha \in G\}$ is connected (see, e.g., [Cam99]). We need the following variant of this result for $\text{Aut}(\mathfrak{B})$; we also present its proof since we are unaware of any reference in the literature. If $a \in A$ then a sequence $(b_0, \dots, b_n) \in B^{n+1}$ is called an a -walk (of length n) if $(b_i, b_{i+1}) \in a^{\mathfrak{B}}$ for every $i \in \{0, \dots, n-1\}$ (we count the number of traversed edges rather than the number of vertices when defining the length).

Lemma 27. *Let $a \in A$ be a symmetric atom of \mathbf{A} with $a \cap \text{Id} = 0$ and suppose that $\text{Aut}(\mathfrak{B})$ is primitive. Then there exists an $a^{\mathfrak{B}}$ -walk of even length between any $x, y \in B$. Moreover, there exists $k \in \mathbb{N}$ such that for all $x, y \in B$ there exists an $a^{\mathfrak{B}}$ -walk of length $2k$ between x and y .*

Proof. If R is a binary relation then $R^k = R \circ R \circ \dots \circ R$ denotes the k -th relational power of R . The sequence of binary relations $L_n := \text{Id}^{\mathfrak{B}} \cup \bigcup_{k=1}^n (a^{\mathfrak{B}})^{2k}$ is non-decreasing by definition and terminates because all binary relations are unions of at most finitely many atoms. Therefore, there exists $k \in \mathbb{N}$ such for all $n \geq k$ we have $L_n = L_k$. Note that L_k is an equivalence relation, namely the relation "there exists an $a^{\mathfrak{B}}$ -walk of even length between x and y ". Since \mathfrak{B} is primitive L_k must be trivial. If $L_k = B^2$ then there exists an $a^{\mathfrak{B}}$ -walk of length $2k$ between any two $x, y \in B$ and we are done. Otherwise,

$$L_k = \{(x, x) \mid x \in B\} = \text{Id}^{\mathfrak{B}}.$$

Since a is symmetric $a^{\mathfrak{B}} \circ a^{\mathfrak{B}} \neq 0$ and $a^{\mathfrak{B}} \circ a^{\mathfrak{B}}$ contains therefore an atom. But then $a^{\mathfrak{B}} \circ a^{\mathfrak{B}} \subseteq L_k$ implies by Proposition 25 $a^{\mathfrak{B}} \circ a^{\mathfrak{B}} = L_k$. This is a contradiction to Proposition 26. \square

Lemma 28. *Let $a \in A$ be a symmetric atom of \mathbf{A} such that $\text{Aut}(\mathfrak{B})$ is primitive and (a, a, a) is forbidden. Then all polymorphisms of \mathfrak{B} are $\{\text{Id}, a\}$ -canonical.*

In the proof, we need the following notation. Let $a_1, \dots, a_k \in A$ be such that $a_1 = \dots = a_j$ and $a_{j+1} = \dots = a_k$. Instead of writing $(a_1, \dots, a_n)^{\mathfrak{B}}$ we use the shortcut $(a_1|_j a_{j+1})^{\mathfrak{B}}$.

Proof (of Lemma 28). The following ternary relation R on B is primitive positive definable in \mathfrak{B} .

$$R := \{(x_1, x_2, x_3) \in B^3 \mid (a \cup \text{Id})^{\mathfrak{B}}(x_1, x_2) \wedge (a \cup \text{Id})^{\mathfrak{B}}(x_2, x_3) \wedge a^{\mathfrak{B}}(x_1, x_3)\}$$

Observe that $c \in R$ if and only if $a^{\mathfrak{B}}(c_1, c_2) \wedge \text{Id}^{\mathfrak{B}}(c_2, c_3)$ or $\text{Id}^{\mathfrak{B}}(c_1, c_2) \wedge a^{\mathfrak{B}}(c_2, c_3)$.

Let f be a polymorphism of \mathfrak{B} of arity n . Let $x, y, u, v \in B^n$ be arbitrary such that (x, y) and (u, v) have the same $\{\text{Id}, a\}$ -configuration. Without loss of generality we may assume that $(a|_j \text{Id})^{\mathfrak{B}}(x, y)$ and $(a|_j \text{Id})^{\mathfrak{B}}(u, v)$. Now consider $p, q \in B^n$ such that $(\text{Id}|_j a)^{\mathfrak{B}}(p, q)$ holds.

Note that by the edge-conservativeness of f the following holds:

$$(f(x), f(y)) \in (a \cup \text{Id})^{\mathfrak{B}}, (f(u), f(v)) \in (a \cup \text{Id})^{\mathfrak{B}} \text{ and } (f(p), f(q)) \in (\text{Id}|_j a)^{\mathfrak{B}}.$$

By Lemma 27 there exists a $k \in \mathbb{N}$ such that for every $i \in \{1, \dots, n\}$ there exists an $a^{\mathfrak{B}}$ -walk (s_i^0, \dots, s_i^k) with $s_i^0 = y_i$ and $s_i^k = p_i$. Now consider the following walk in B^n :

$$\begin{aligned} & (a|_j \text{Id})^{\mathfrak{B}}(x, y) \\ & (\text{Id}|_j a)^{\mathfrak{B}}(y, (s_1^0, \dots, s_j^0, s_{j+1}^1, \dots, s_n^1)) \\ & (a|_j \text{Id})^{\mathfrak{B}}((s_1^0, \dots, s_j^0, s_{j+1}^1, \dots, s_n^1), (s_1^1, \dots, s_j^1, s_{j+1}^1, \dots, s_n^1)) \\ & \quad \vdots \\ & (a|_j \text{Id})^{\mathfrak{B}}((s_1^i, \dots, s_j^i, s_{j+1}^{i+1}, \dots, s_n^{i+1}), (s_1^{i+1}, \dots, s_j^{i+1}, s_{j+1}^{i+1}, \dots, s_n^{i+1})) \\ & (\text{Id}|_j a)^{\mathfrak{B}}((s_1^{i+1}, \dots, s_j^{i+1}, s_{j+1}^{i+1}, \dots, s_n^{i+1}), (s_1^{i+1}, \dots, s_j^{i+1}, s_{j+1}^{i+2}, \dots, s_n^{i+2})) \\ & \quad \vdots \\ & (a|_j \text{Id})^{\mathfrak{B}}((s_1^{k-1}, \dots, s_j^{k-1}, s_{j+1}^k, \dots, s_n^k), p) \\ & (\text{Id}|_j a)^{\mathfrak{B}}(p, q) \end{aligned}$$

Every three consecutive elements on this walk are component wise in the relation R . Since R is primitive positive definable the polymorphism f preserves R by Theorem 15. This means that f maps this walk on a walk where the atomic relations are an alternating sequence of $a^{\mathfrak{B}}$ and $\text{Id}^{\mathfrak{B}}$, which implies

$$(f(x), f(y)) \in a^{\mathfrak{B}} \Leftrightarrow (f(p), f(q)) \in \text{Id}^{\mathfrak{B}}.$$

If we repeat the same argument with a walk from q to v we get:

$$(f(p), f(q)) \in a^{\mathfrak{B}} \Leftrightarrow (f(u), f(v)) \in \text{Id}^{\mathfrak{B}}.$$

Combining these two equivalences gives us

$$(f(x), f(y)) \in a^{\mathfrak{B}} \Leftrightarrow (f(u), f(v)) \in a^{\mathfrak{B}}.$$

Since the tuples $x, y, u, v \in B^n$ were arbitrary this shows that f is $\{\text{Id}, a\}$ -canonical. \square

Theorem 29. *Let $\text{Aut}(\mathfrak{B})$ be primitive and let a be a symmetric atom of \mathbf{A} such that (a, a, a) is forbidden. Then $\text{CSP}(\mathfrak{B})$ is NP-hard.*

Proof. By Lemma 28 we know that all polymorphisms of \mathfrak{B} are $\{a, \text{Id}\}$ -canonical. This means that every $f \in \text{Pol}(\mathfrak{B})$ induces an operation \bar{f} of the same arity on the set $\{a, \text{Id}\}$. Let \mathcal{C}_2 be the set of induced operations. Note that \mathcal{C}_2 is an operation clone on a Boolean domain. The mapping $\mu: \text{Pol}(\mathfrak{B}) \rightarrow \mathcal{C}_2$ defined by $\mu(f) := \bar{f}$ is a uniformly continuous minor-preserving map.

\circ	Id	a	b
Id	Id	a	b
a	a	$\neg b$	b
b	b	b	$\neg b$

\circ	Id	a	b
Id	Id	a	b
a	a	$\neg a$	$0'$
b	b	$0'$	1

Fig. 1. Multiplication tables of relation algebras #13 (left) and #17 (right).

Assume for contradiction that there exists a ternary cyclic polymorphism \bar{s} in \mathcal{C}_2 . Let $x, y, z \in B^3$ be such that

$$\begin{aligned} & (a, a, \text{Id})^{\mathfrak{B}}(x, y), \\ & (\text{Id}, a, a)^{\mathfrak{B}}(y, z), \\ & \text{and } (a, \text{Id}, a)^{\mathfrak{B}}(x, z). \end{aligned}$$

By the cyclicity of the operation \bar{s} and the edge-conservativeness of s we have that either

$$(s(x), s(y)) \in a^{\mathfrak{B}}, (s(y), s(z)) \in a^{\mathfrak{B}} \text{ and } (s(x), s(z)) \in a^{\mathfrak{B}}$$

or

$$(s(x), s(y)) \in \text{Id}^{\mathfrak{B}}, (s(y), s(z)) \in \text{Id}^{\mathfrak{B}} \text{ and } (s(x), s(z)) \in \text{Id}^{\mathfrak{B}}.$$

Since (a, a, a) is forbidden, the second case holds. Note that \mathbf{A} must have an atom $b \neq \text{Id}$ such that the triple (a, a, b) is allowed, because otherwise a would be an equivalence relation. Now consider $u, v, w \in B^3$ such that

$$\begin{aligned} & (a, a, \text{Id})^{\mathfrak{B}}(u, v), \\ & (\text{Id}, a, a)^{\mathfrak{B}}(v, w), \\ & \text{and } (a, b, a)^{\mathfrak{B}}(u, w). \end{aligned}$$

Since s is $\{a, \text{Id}\}$ -canonical and with the observation from before we have

$$(s(u), s(v)) \in \text{Id}^{\mathfrak{B}} \text{ and } (s(v), s(w)) \in \text{Id}^{\mathfrak{B}}.$$

Now the transitivity of equality contradicts $(s(u), s(w)) \in (a \cup b)^{\mathfrak{B}}$.

We conclude that \mathcal{C}_2 does not contain a ternary cyclic operation. Since the domain of \mathcal{C}_2 has size two, Theorem 17 implies the existence of a u.c. minor-preserving map $\nu: \mathcal{C}_2 \rightarrow \text{Proj}$. The composition $\nu \circ \mu: \text{Pol}(\mathfrak{B}) \rightarrow \text{Proj}$ is also a u.c. minor-preserving map and therefore by Theorem 18 the CSP(\mathfrak{B}) is NP-hard. \square

7 Examples

Andréka and Maddux classified *small relation algebras*, i.e., finite relation algebras with at most 3 atoms [AM94]. We consider the complexity of the network satisfaction problem of two of them, namely the relation algebras #13 and #17 (we use the enumeration from [AM94]). Both relation algebras have normal representations (see below) and fall into the scope of our hardness criteria. Cristani and Hirsch [CH04] classified the complexities of the network satisfaction problems for small relation algebras, but due to a mistake the algebras #13 and #17 were left open.

Example 30 (Relation Algebra #13). The relation algebra #13 is given by the multiplication table in Fig. 1. This finite relation algebra has a normal representation \mathfrak{B} defined as follows. Let V_1 and V_2 be countable, disjoint sets. We set $B := V_1 \cup V_2$ and define the following atomic relations:

$$\begin{aligned} \text{Id}^{\mathfrak{B}} &:= \{(x, x) \in B^2\}, \\ a^{\mathfrak{B}} &:= \{(x, y) \in B^2 \setminus \text{Id}^{\mathfrak{B}} \mid (x \in V_1 \wedge y \in V_1) \vee (x \in V_2 \wedge y \in V_2)\}, \\ b^{\mathfrak{B}} &:= \{(x, y) \in B^2 \setminus \text{Id}^{\mathfrak{B}} \mid (x \in V_1 \wedge y \in V_2) \vee (x \in V_2 \wedge y \in V_1)\}. \end{aligned}$$

It is easy to check that this structure is a square representation for #13. Moreover, this structure is fully universal for #13 and homogeneous, and therefore a normal representation.

Note that the relation $(\text{Id} \cup a)^{\mathfrak{B}}$ is an equivalence relation where V_1 and V_2 are the two equivalence classes. Therefore we get by Theorem 23 that the (general) network satisfaction problem for the relation algebra #13 is NP-hard. We mention that this result can also be deduced from the results in [BMPP19].

Example 31 (Relation Algebra #17). The relation algebra #17 is given by the multiplication table in Fig. 1. Let $\mathfrak{N} = (V; E^{\mathfrak{N}})$ be the countable, homogeneous, universal triangle-free, undirected graph (see [Hod97]), also called called a *Henson graph*. We use this Henson graph to obtain a square representation \mathfrak{B} with domain V for the relation algebra #17 as follows:

$$\begin{aligned}\text{Id}^{\mathfrak{B}} &:= \{(x, x) \in V^2\}, \\ a^{\mathfrak{B}} &:= \{(x, y) \in V^2 \mid (x, y) \in E^{\mathfrak{N}}\}, \\ b^{\mathfrak{B}} &:= \{(x, y) \in V^2 \setminus \text{Id}^{\mathfrak{B}} \mid (x, y) \notin E^{\mathfrak{N}}\}.\end{aligned}$$

This structure is homogeneous and fully universal since \mathfrak{N} is homogeneous and embeds every triangle free graph. It is easy to see that there exists no non-trivial equivalence relation in this relation algebra. For the atom a the triangle (a, a, a) is forbidden, which means we can apply Theorem 29 and get NP-hardness for the (general) network satisfaction problem for the relation algebra #17. Also in this case, the hardness result can also be deduced from the results in [BMPP19].

8 Conclusion and Future Work

To the best of our knowledge the computational complexity of the (general) network satisfaction problem was previously only known for a small number of isolated finite relation algebras, for example the point algebra, Allens interval algebra, or the 18 small relation algebras from [AM94]. Both of our criteria, Theorem 23 and Theorem 29, show the NP-hardness for relatively large classes of finite relation algebras. In Section 7 we applied these results to settle the complexity status of two problems that were left open in [CH04].

To obtain our general hardness conditions we used the universal algebraic approach for studying the complexity of constraint satisfaction problems. This approach will hopefully lead to a solution of Hirsch's RBCP for all finite relation algebras \mathbf{A} with a normal representation \mathfrak{B} . It is also relatively easy to prove that the network satisfaction problem for \mathbf{A} is NP-complete if \mathfrak{B} has an equivalence relation with an equivalence class of finite size larger than two. Hence, the next steps that have to be taken with this approach are the following.

- Classify the complexity of the network satisfaction problem for finite relation algebras \mathbf{A} where the normal representation has a primitive automorphism group.
- Classify the complexity of the network satisfaction problem for relation algebras that have equivalence relations with infinitely many classes of size two.
- Classify the complexity of the network satisfaction problem for relation algebras that have equivalence relations with infinitely many infinite classes.

References

- All83. James F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.
- AM94. Hajnal Andréka and Roger D. Maddux. Representations for small relation algebras. *Notre Dame Journal of Formal Logic*, 35(4):550–562, 1994.
- Bar11. Libor Barto. The dichotomy for conservative constraint satisfaction problems revisited. In *Proceedings of the Symposium on Logic in Computer Science (LICS)*, Toronto, Canada, 2011.
- BJ17. Manuel Bodirsky and Peter Jonsson. A model-theoretic view on qualitative constraint reasoning. *Journal of Artificial Intelligence Research*, 58:339–385, 2017.

- BK07. Manuel Bodirsky and Martin Kutz. Determining the consistency of partial tree descriptions. *Artificial Intelligence*, 171:185–196, 2007.
- BK12. 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.
- BKO⁺17. Libor Barto, Michael Kompatscher, Miroslav Olšák, Trung Van Pham, and Michael Pinsker. The equivalence of two dichotomy conjectures for infinite domain constraint satisfaction problems. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science – LICS’17*, 2017. Preprint arXiv:1612.07551.
- BMPP19. Manuel Bodirsky, Barnaby Martin, Michael Pinsker, and András Pongrácz. Constraint satisfaction problems for reducts of homogeneous graphs. *SIAM Journal on Computing*, 48(4):1224–1264, 2019. A conference version appeared in the Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, pages 119:1–119:14.
- BN06. Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.
- Bod18. Manuel Bodirsky. Finite relation algebras with normal representations. In *Relational and Algebraic Methods in Computer Science - 17th International Conference, RAMiCS 2018, Groningen, The Netherlands, October 29 - November 1, 2018, Proceedings*, pages 3–17, 2018.
- BOP18. Libor Barto, Jakub Opršal, and Michael Pinsker. The wonderland of reflections. *Israel Journal of Mathematics*, 223(1):363–398, 2018.
- BP16. Manuel Bodirsky and Michael Pinsker. Canonical Functions: a Proof via Topological Dynamics. Preprint available under <http://arxiv.org/abs/1610.09660>, 2016.
- BPP14. Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Projective clone homomorphisms. Accepted for publication in the Journal of Symbolic Logic, Preprint arXiv:1409.4601, 2014.
- Bul03. Andrei A. Bulatov. Tractable conservative constraint satisfaction problems. In *Proceedings of the Symposium on Logic in Computer Science (LICS)*, pages 321–330, Ottawa, Canada, 2003.
- Bul14. Andrei A. Bulatov. Conservative constraint satisfaction revisited. Manuscript, ArXiv:1408.3690v1, 2014.
- Bul17. 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, 2017*, pages 319–330, 2017.
- Cam99. Peter J. Cameron. *Permutation Groups*. LMS Student Text 45. Cambridge University Press, Cambridge, 1999.
- CH04. Matteo Cristiani and Robin Hirsch. The complexity of the constraint satisfaction problem for small relation algebras. *Artificial Intelligence Journal*, 156:177–196, 2004.
- Dün05. Ivo Düntsch. Relation algebras and their application in temporal and spatial reasoning. *Artificial Intelligence Review*, 23:315–357, 2005.
- HH02. Robin Hirsch and Ian Hodkinson. *Relation Algebras by Games*. North Holland, 2002.
- Hir96. Robin Hirsch. Relation algebras of intervals. *Artificial Intelligence Journal*, 83:1–29, 1996.
- Hir97. Robin Hirsch. Expressive power and complexity in algebraic logic. *Journal of Logic and Computation*, 7(3):309 – 351, 1997.
- Hir99. Robin Hirsch. A finite relation algebra with undecidable network satisfaction problem. *Logic Journal of the IGPL*, 7(4):547–554, 1999.
- Hod97. Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.
- LM94. Peter B. Ladkin and Roger D. Maddux. On binary constraint problems. *Journal of the Association for Computing Machinery*, 41(3):435–469, 1994.
- Mad06. Roger Duncan Maddux. *Relation Algebras* -. Elsevier, Amsterdam, 2006.
- RN99. Jochen Renz and Bernhard Nebel. On the complexity of qualitative spatial reasoning: A maximal tractable fragment of the region connection calculus. *Artificial Intelligence*, 108(1-2):69–123, 1999.
- RN07. Jochen Renz and Bernhard Nebel. Qualitative spatial reasoning using constraint calculi. In M. Aiello, I. Pratt-Hartmann, and J. van Benthem, editors, *Handbook of Spatial Logics*, pages 161–215. Springer Verlag, Berlin, 2007.
- VKvB89. Marc Vilain, Henry Kautz, and Peter van Beek. Constraint propagation algorithms for temporal reasoning: A revised report. *Reading in Qualitative Reasoning about Physical Systems*, pages 373–381, 1989.
- Zhu17. Dmitriy 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, 2017*, pages 331–342, 2017.