

The Complexity of Network Satisfaction Problems for Symmetric Relation Algebras with a Flexible Atom

Manuel Bodirsky

Institut für Algebra, TU Dresden, 01062 Dresden, Germany

MANUEL.BODIRSKY@TU-DRESDEN.DE

Simon Knäuer

Institut für Algebra, TU Dresden, 01062 Dresden, Germany

SIMON.KNAEUER@TU-DRESDEN.DE

Abstract

Robin Hirsch posed in 1996 the *Really Big Complexity Problem*: classify the computational complexity of the network satisfaction problem for all finite relation algebras \mathbf{A} . We provide a complete classification for the case that \mathbf{A} is symmetric and has a flexible atom; in this case, the problem is NP-complete or in P. The classification task can be reduced to the case where \mathbf{A} is integral. If a finite integral relation algebra has a flexible atom, then it has a normal representation \mathfrak{B} . We can then study the computational complexity of the network satisfaction problem of \mathbf{A} using the universal-algebraic approach, via an analysis of the polymorphisms of \mathfrak{B} . We also use a Ramsey-type result of Nešetřil and Rödl and a complexity dichotomy result of Bulatov for conservative finite-domain constraint satisfaction problems.

1. Introduction

One of the earliest approaches to formalise constraint satisfaction problems over infinite domains is based on *relation algebras* (Ladkin & Maddux, 1994; Hirsch, 1997). We think about the elements of a relation algebra as binary relations; the algebra has operations for intersection, union, complement, converse, and composition of relations, and constants for the empty relation, the full relation, and equality, and is required to satisfy certain axioms. Important examples of relation algebras are the Point Algebra, the Left Linear Point Algebra, Allen’s Interval Algebra, RCC5, and RCC8, just to name a few.

The most important computational problems related to relation algebras that have been studied are called *network satisfaction problems (NSPs)*. A finite network is given in which the nodes are linked via the elements of a fixed relation algebra. Then the computational task of the NSP for this relation algebra is to decide whether the network satisfies a strong notion of consistency with respect to the laws of the fixed relation algebra. Such NSPs for a finite relation algebra can be used to model many computational problems in temporal and spatial reasoning (Düntsch, 2005; Renz & Nebel, 2007; Bodirsky & Jonsson, 2017).

More than two decades ago, Hirsch (1996) asked 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 complexity of the network satisfaction problem for the Point Algebra and the Left Linear Point Algebra is in P (Vilain, Kautz, & van Beek, 1990; Bodirsky & Kutz, 2007), while it is NP-complete for all of the other examples mentioned above (Allen, 1983; Renz & Nebel, 1999). There also exist relation algebras where the complexity of the network satisfaction problem is not in NP: Hirsch gave an example of a finite relation algebra with an undecidable network satisfaction prob-

lem (Hirsch, 1999). This result might be surprising at first sight; it is related to the fact that the representation of a finite relation algebra by concrete binary relations over some set can be quite complicated. We also mention that not every finite relation algebra has a representation (Lyndon, 1950). There are even non-representable relation algebras that are symmetric (Maddux, 2006a); a relation algebra is *symmetric* if every element is its own converse.

A simple condition that implies that a finite relation algebra \mathbf{A} has a representation is the existence of a so-called *flexible atom* (Comer, 1984; Maddux, 1982). A flexible atom is an element of \mathbf{A} that is maximally unconstrained in its interaction with other elements of the relation algebra; the formal definitions can be found in Section 2.

Such relation algebras have been studied intensively, for example in the context of the so-called *flexible atoms conjecture* (Maddux, 1994; Alm, Maddux, & Manske, 2008). We will see that integral relation algebras with a flexible atom even have a *normal representation*, i.e., a representation which is fully universal, square, and homogeneous (Hirsch, 1996). The network satisfaction problem for a relation algebra with a normal representation can be seen as a *constraint satisfaction problem* for an infinite structure \mathfrak{B} that is well-behaved from a model-theoretic point of view; in particular, we may choose \mathfrak{B} to be *homogeneous* and *finitely bounded*.

Constraint satisfaction problems over finite domains have been studied intensively in the past two decades, and tremendous progress has been made concerning systematic findings about their computational complexity. As a highlighting result, Bulatov (2017) and Zhuk (2017, 2020) proved the famous Feder-Vardi dichotomy conjecture which states that every finite-domain CSP is in P or NP-complete. Both proofs build on an important connection between the computational complexity of constraint satisfaction problems and universal algebra.

The universal-algebraic approach can also be applied to study the computational complexity of countably infinite homogeneous structures \mathfrak{B} with finite relational signature (Bodirsky & Nešetřil, 2006). For an introduction to the field, we refer to (Bodirsky, 2021). If \mathfrak{B} is finitely bounded, then $\text{CSP}(\mathfrak{B})$ is contained in NP (see, e.g. Bodirsky, 2012). If \mathfrak{B} is homogeneous and finitely bounded then a complexity dichotomy has been conjectured, along with algebraic criteria that distinguish NP-complete from polynomial-time solvable problems (Bodirsky, Pinsker, & Pongrácz, 2019). The exact formulation of the conjecture from Bodirsky et al. (2019) in full generality requires concepts that we do not need to prove our results. In Theorem 9.1 we verify these conjectures for all normal representations of finite integral symmetric relation algebras with a flexible atom, and thereby also solve Hirsch’s RBCP for symmetric relation algebras with a flexible atom. Phrased in the terminology of relation algebras, our result is the following.

Theorem 1.1. *Let \mathbf{A} be a finite symmetric representable relation algebra with a flexible atom, and let A_0 be the set of atoms of \mathbf{A} . Then*

- *there exists an operation $f: A_0^6 \rightarrow A_0$ that preserves the allowed triples of \mathbf{A} , satisfies*

$$\forall x_1, \dots, x_6 \in A_0. f(x_1, \dots, x_6) \in \{x_1, \dots, x_6\}$$

and satisfies the Siggers identity

$$\forall x, y, z \in A_0. f(x, x, y, y, z, z) = f(y, z, x, z, x, y);$$

in this case the network satisfaction problem for \mathbf{A} is in P, or

- *the network satisfaction problem for \mathbf{A} is NP-complete.*

Moreover, the satisfiability of the Siggers identity in Theorem 1.1 is a decidable criterion for \mathbf{A} that is a sufficient condition for the polynomial-time tractability of the network satisfaction problem of \mathbf{A} . We want to mention that there are several other equivalent criteria that could be used instead of the first item in the theorem, namely all characterizations of Taylor Algebras for finite conservative algebras (see e.g., Bulatov, 2003).

This article is the full version of results described in a conference article (Bodirsky & Knäuer, 2021).

1.1 Proof Strategy

Every finite integral representable relation algebra \mathbf{A} with a flexible atom has a normal representation \mathfrak{B} ; for completeness, and since we are not aware of a reference for this fact, we include a proof in Section 3. It follows that the classification question about the complexity of the network satisfaction problem of \mathbf{A} can be translated into a question about the complexity of the constraint satisfaction problem for the relational structure \mathfrak{B} .

We then associate a finite relational structure \mathfrak{D} to \mathfrak{B} and show that $\text{CSP}(\mathfrak{B})$ can be reduced to $\text{CSP}(\mathfrak{D})$ in polynomial-time (Section 4). If the structure \mathfrak{D} satisfies the condition of the first case in Theorem 1.1, then known results about finite-domain CSPs imply that $\text{CSP}(\mathfrak{D})$ is in P (Bulatov, 2003, 2016; Barto, 2011), and hence $\text{CSP}(\mathfrak{B})$ is in P, too. If the first case in Theorem 1.1 does not apply, then known results about finite-domain algebras imply that there are $a, b \in A_0$ such that the *canonical polymorphisms* of \mathfrak{B} act as a projection on $\{a, b\}$ (Bulatov, 2003, 2016; Barto, 2011). We first observe NP-hardness of $\text{CSP}(\mathfrak{B})$ if \mathfrak{B} does not have a binary injective polymorphism (Section 6). If \mathfrak{B} has a binary injective polymorphism, we use results from structural Ramsey theory to show that \mathfrak{B} must even have a binary injective polymorphism which is canonical (Section 7). This implies that none of a, b equals $\text{Id} \in A$. We then prove that \mathfrak{B} does not have a binary $\{a, b\}$ -symmetric polymorphism; also in this step, we apply Ramsey theory. In Section 8 we show that this in turn implies that all polymorphisms of \mathfrak{B} must be canonical on $\{a, b\}$. Finally, we show that \mathfrak{B} cannot have a polymorphism which acts as a majority or as a minority on $\{a, b\}$, and thus by Schaefer’s theorem all polymorphisms of \mathfrak{B} act as a projection on $\{a, b\}$. This is again implied by results from Section 7. Finally it follows that $\text{CSP}(\mathfrak{B})$ is NP-hard. This concludes the proof of Theorem 1.1.

Our proof follows a strategy that was applied several times in the study of infinite-domain constraint satisfaction problems and recently described and generalized by Mottet and Pinsker (2022). We give some details about this in Section 10.

1.2 Organisation of the Article

Section 2 introduces all the basic concepts and tools that are used in this article. In Section 3 we define flexible atoms and obtain first results about representable relation algebras with a flexible atom. Section 4 is dedicated to the atom structure and the polynomial-time tractability results. In Section 5 we provide an additional perspective on the class of computational problems under consideration; in this section we define those problems completely

without the use of the relation algebra framework. The reader can get better intuition of the class of problems studied in this article, however our results and proofs do not rely on that section. The Sections 6-8 contain the main parts of the proof as outlined in the previous paragraph. In Section 9 we put everything together and prove the main theorem. We end with a conclusion and a small discussion of our result.

2. Preliminaries

We recall some basic definitions and results about universal algebra, representable relation algebras, constraint satisfaction, model theory, and structural Ramsey theory.

2.1 Algebras and Structures

Let \mathbb{N} denote the natural numbers starting with 0 and define $\mathbb{N}_+ := \mathbb{N} \setminus \{0\}$. A *signature* τ is a set of relation symbols and function symbols. Each symbol is associated with a natural number, called the *arity* of the symbol. Function symbols of arity 0 are called *constant symbols*. A τ -structure is a tuple $\mathfrak{A} = (A; (Q^{\mathfrak{A}})_{Q \in \tau})$ where A is a set, called the *domain* of \mathfrak{A} , such that for every $Q \in \tau$:

- if Q is a relation symbol of arity $n \in \mathbb{N}$, then $Q^{\mathfrak{A}}$ is a subset of A^n ,
- if Q is a function symbol of arity $n \in \mathbb{N}$, then $Q^{\mathfrak{A}}$ is an operation $A^n \rightarrow A$.

Note that by $A^0 = \{\emptyset\}$ a subset of A^0 can be seen as a Boolean value and the operation $f: A^0 \rightarrow A$ can be interpreted as a constant. As long as there is no risk of confusion we will often use the function symbols for the corresponding functions, and the relation symbols for the corresponding relations, i.e., we use Q instead of $Q^{\mathfrak{A}}$.

Let \mathfrak{A} and \mathfrak{B} be two τ -structures. A *homomorphism* h from \mathfrak{A} to \mathfrak{B} is a function $h: A \rightarrow B$ such that

- for every relation symbol Q of arity $n \in \mathbb{N}$ and every tuple $(a_1, \dots, a_n) \in A^n$, we have that $(a_1, \dots, a_n) \in Q^{\mathfrak{A}} \Rightarrow (h(a_1), \dots, h(a_n)) \in Q^{\mathfrak{B}}$;
- for every function symbol Q of arity $n \in \mathbb{N}$ and every tuple $(a_1, \dots, a_n) \in A^n$, we have that $h(Q^{\mathfrak{A}}(a_1, \dots, a_n)) = Q^{\mathfrak{B}}(h(a_1), \dots, h(a_n))$.

In the case that τ contains only function symbols and h is surjective, then \mathfrak{B} is called *homomorphic image* of \mathfrak{A} . In general, the homomorphism h is called an *embedding* if h is injective and satisfies

- for every relation symbol Q of arity $n \in \mathbb{N}$ and every tuple $(a_1, \dots, a_n) \in A^n$, we have that $(a_1, \dots, a_n) \in Q^{\mathfrak{A}} \Leftrightarrow (h(a_1), \dots, h(a_n)) \in Q^{\mathfrak{B}}$.

A surjective embedding is called an *isomorphism*. An *endomorphism* of a τ -structure \mathfrak{A} is a homomorphism from \mathfrak{A} to \mathfrak{A} and an *automorphism* of \mathfrak{A} is an isomorphism from \mathfrak{A} to \mathfrak{A} . We denote by $\text{Aut}(\mathfrak{A})$ the group of all *automorphisms* of \mathfrak{A} .

A τ -structure \mathfrak{A} is a *substructure* of a τ -structure \mathfrak{B} if

- $A \subseteq B$;

- for every relation symbol Q of arity $n \in \mathbb{N}$ and every tuple $(a_1, \dots, a_n) \in A^n$, we have that $(a_1, \dots, a_n) \in Q^{\mathfrak{A}}$ if and only if $(a_1, \dots, a_n) \in Q^{\mathfrak{B}}$;
- for every function symbol Q of arity $n \in \mathbb{N}$ and every tuple $(a_1, \dots, a_n) \in A^n$, we have that $Q^{\mathfrak{A}}(a_1, \dots, a_n) = Q^{\mathfrak{B}}(a_1, \dots, a_n)$.

Note that for every subset $A' \subseteq B$ of the domain of a τ -structure \mathfrak{B} there exists a unique substructure \mathfrak{A} of \mathfrak{B} with smallest domain A and $A' \subseteq A$. We call this the *substructure of \mathfrak{B} induced by A* .

Let \mathfrak{A} and \mathfrak{B} be τ structures. The *direct product* $\mathfrak{C} = \mathfrak{A} \times \mathfrak{B}$ is the τ -structure where

- $A \times B$ is the domain of \mathfrak{C} ;
- for every relation symbol Q of arity $n \in \mathbb{N}$ and every tuple $((a_1, b_1), \dots, (a_n, b_n)) \in (A \times B)^n$, we have that $((a_1, b_1), \dots, (a_n, b_n)) \in Q^{\mathfrak{C}}$ if and only if $(a_1, \dots, a_n) \in Q^{\mathfrak{A}}$ and $(b_1, \dots, b_n) \in Q^{\mathfrak{B}}$;
- for every function symbol Q of arity $n \in \mathbb{N}$ and every tuple $((a_1, b_1), \dots, (a_n, b_n)) \in (A \times B)^n$, we have that $Q^{\mathfrak{C}}((a_1, b_1), \dots, (a_n, b_n)) = (Q^{\mathfrak{A}}(a_1, \dots, a_n), Q^{\mathfrak{B}}(b_1, \dots, b_n))$.

We denote the direct product $\mathfrak{A} \times \mathfrak{A}$ by \mathfrak{A}^2 . The k -fold product $\mathfrak{A} \times \dots \times \mathfrak{A}$ is defined analogously and denoted by \mathfrak{A}^k .

Structures with a signature that only contains function symbols are called *algebras* and structure with purely relational signature are called *relational structures*. Since we do not deal with signatures of mixed type in this article, we will from now on use the term structure for relational structures only.

2.2 Representable Relation Algebras and Network Satisfaction Problems

Relation algebras that are not representable have a trivial network satisfaction problem, namely the class of yes-instances of their network satisfaction problem is empty (see Definition 2.5). We thus omit the definition of relation algebras and start immediately with the simpler definition of *representable relation algebras*; here we basically follow the textbook by Maddux (2006b).

Definition 2.1. *Let D be a set and $E \subseteq D^2$ an equivalence relation on D . Let $(\mathcal{P}(E); \cup, \bar{}, 0, 1, \text{Id}, \smile, \circ)$ be an algebra with the following operations:*

1. $a \cup b := \{(x, y) \mid (x, y) \in a \vee (x, y) \in b\}$,
2. $\bar{a} := E \setminus a$,
3. $0 := \emptyset$,
4. $1 := E$,
5. $\text{Id} := \{(x, x) \mid x \in D\}$,
6. $a \smile := \{(x, y) \mid (y, x) \in a\}$,
7. $a \circ b := \{(x, z) \mid \exists y \in D : (x, y) \in a \text{ and } (y, z) \in b\}$.

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

The class of *representable relation algebras*, denoted by RRA, consists of all algebras with signature $\tau = \{\cup, \bar{}, 0, 1, \text{Id}, \smile, \circ\}$ with corresponding arities 2, 1, 0, 0, 0, 1 and 2 that are isomorphic to some proper relation algebra. We use bold letters (such as \mathbf{A}) to denote algebras from RRA and the corresponding roman letter (such as A) to denote the domain of the algebra. An algebra is called *finite* if its domain is finite. We call $\mathbf{A} \in \text{RRA}$ *symmetric* if all its elements are symmetric, i.e., $a^\smile = a$ for every $a \in A$.

According to the previous definition, an algebra \mathbf{A} is in RRA if it has an isomorphism to a proper relation algebra. Such an isomorphism is usually called the representation of \mathbf{A} . To link the theory of relation algebras with model theory it will be convenient to view representations of algebras in RRA as relational structures and to not use the classical notation here. However, it is easy to see that the existence of a representation is equivalent under both definitions.

Definition 2.2. Let $\mathbf{A} \in \text{RRA}$. Then a representation of \mathbf{A} is a relational structure \mathfrak{B} such that

- \mathfrak{B} is an A -structure, i.e., the elements of A are binary relation symbols of \mathfrak{B} ;
- The map $a \mapsto a^\mathfrak{B}$ is an isomorphism between \mathbf{A} and the proper relation algebra induced by the relations of \mathfrak{B} in $(\mathcal{P}(1^\mathfrak{B}); \cup, \bar{}, 0, 1, \text{Id}, \smile, \circ)$.

Let $\mathbf{A} = (A; \cup, \bar{}, 0, 1, \text{Id}, \smile, \circ)$ be a representable relation algebra and let us define a new operation $x \cap y := \overline{\overline{x} \cup \overline{y}}$ on the set A . Then the algebra $(A; \cup, \cap, \bar{}, 0, 1)$ is by definition a Boolean algebra and induces therefore a partial order \leq on A , which is defined by $x \leq y \Leftrightarrow x \cup y = y$. Note that for proper relation algebras, this ordering coincides with the set-inclusion order. The minimal elements of this order in $A \setminus \{0\}$ are called *atoms*. The set of atoms of \mathbf{A} is denoted by A_0 . A tuple $(x, y, z) \in (A_0)^3$ is called an *allowed triple* if $z \leq x \circ y$. Otherwise, (x, y, z) is called a *forbidden triple* and in this case $\overline{z} \cup \overline{x \circ y} = 1$. We say that a relational A -structure \mathfrak{B} *induces a forbidden triple (from \mathbf{A})* if there exists $b_1, b_2, b_3 \in B$ and $(x, y, z) \in (A_0)^3$ such that $x(b_1, b_2), y(b_2, b_3)$ and $z(b_1, b_3)$ hold and (x, y, z) is a forbidden triple. Note that a representation of \mathbf{A} does not induce a forbidden triple.

Definition 2.3. Let $\mathbf{A} \in \text{RRA}$. An \mathbf{A} -network $(V; f)$ is a finite set V together with a partial function $f: E \subseteq V^2 \rightarrow A$, where E is the domain of f . An \mathbf{A} -network $(V; f)$ is satisfiable in a representation \mathfrak{B} of \mathbf{A} if there exists an assignment $s: V \rightarrow B$ such that for all $(x, y) \in E$ the following holds:

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

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

We give an example of an instance of the NSP for a representable relation algebra. The numbering of the representable relation algebra is by Andr eka and Maddux (1994).

Example 2.4 (An instance of NSP of representable relation algebra #17). Let \mathbf{A} be the symmetric representable relation algebra with the set of atoms $\{\text{Id}, a, b\}$ and the values for

\circ	Id	a	b
Id	Id	a	b
a	a	Id \cup b	$a \cup b$
b	b	$a \cup b$	Id \cup $a \cup b$

Figure 1: Multiplication table of the representable relation algebra #17.

the composition operation \circ on these atoms be given by Table 1. Note that this determines the composition operation on the whole domain of \mathbf{A} , which is the following set:

$$A = \{\emptyset, \text{Id}, a, b, \text{Id} \cup a, \text{Id} \cup b, a \cup b, \text{Id} \cup a \cup b\}.$$

Let $V := \{x_1, x_2, x_3\}$ be a set. Consider the map $f: V^2 \rightarrow A$ given by

$$\begin{aligned} f(x_i, x_i) &= \text{Id} \quad \text{for all } i \in \{1, 2, 3\} \\ f(x_1, x_2) &= f(x_2, x_1) = a \\ f(x_1, x_3) &= f(x_3, x_1) = \text{Id} \cup a \\ f(x_2, x_3) &= f(x_3, x_2) = b \cup a. \end{aligned}$$

The tuple $(V; f)$ is an example of an instance of NSP of $\mathbf{A} \in \text{RRA}$.

We will in the following assume that for an \mathbf{A} -network $(V; f)$ it holds that $f(V^2) \subseteq A \setminus \{0\}$. Otherwise, $(V; f)$ is not satisfiable. Note that every \mathbf{A} -network $(V; f)$ can be viewed as an A -structure \mathfrak{C} on the domain V : for all $x, y \in V$ and $a \in A$ the relation $a^{\mathfrak{C}}(x, y)$ holds if and only if $f(x, y) = a$.

Definition 2.5. The (general) network satisfaction problem for a finite representable relation algebra \mathbf{A} , denoted by $\text{NSP}(\mathbf{A})$, is the problem of deciding whether a given \mathbf{A} -network is satisfiable.

2.3 Normal Representations and CSPs

In this section we consider a subclass of RRA introduced by Hirsch in 1996. For representable relation algebras \mathbf{A} from this class, $\text{NSP}(\mathbf{A})$ corresponds naturally to a constraint satisfaction problem (CSP). In the last two decades a rich and fruitful theory emerged to analyse the computational complexity of CSPs. We use this theory to obtain results about the computational complexity of NSPs.

In the following let \mathbf{A} be in RRA. An \mathbf{A} -network $(V; f)$ is called *closed* (transitively closed in the work by Hirsch (1997)) if it is total and for all $x, y, z \in V$ it holds that $f(x, x) \leq \text{Id}$, $f(x, y) = f(y, x)^\smile$, and $f(x, z) \leq f(x, y) \circ f(y, z)$. It is called *atomic* if the range of f only contains atoms from \mathbf{A} .

Definition 2.6 (from Hirsch (1996)). Let \mathfrak{B} be a representation of \mathbf{A} . Then \mathfrak{B} is called

- fully universal, if every atomic closed \mathbf{A} -network is satisfiable in \mathfrak{B} ;
- square, if $1^{\mathfrak{B}} = B^2$;

- homogeneous, if for every isomorphism between finite substructures of \mathfrak{B} there exists an automorphism of \mathfrak{B} that extends this isomorphism;
- normal, if it is fully universal, square and homogeneous.

Definition 2.7. Let τ be a relational signature. A first-order formula $\varphi(x_1, \dots, x_n)$ is called primitive positive (pp) 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.

As usual, formulas without free variables are called *sentences*. A relational structure is called *connected* if it is not the disjoint union of two structures. A *connected component* of a relational structure \mathfrak{C} is a substructure \mathfrak{B} of \mathfrak{C} that is maximal with respect to domain inclusion and is connected. For every primitive positive τ -sentence φ with variable set X the *canonical database* $D(\varphi)$ is defined as the τ -structure on X where $x \in X^m$ is in the relation $R^{D(\varphi)}$ for $R \in \tau$ if and only if $R(x)$ is a conjunct from φ . Conversely every relational τ -structure \mathfrak{A} induces a primitive positive τ -sentence on the variable set A , the so-called *conjunctive query*, simply by taking the conjunction all atomic formulas that hold in \mathfrak{A} . We say that the primitive positive τ -sentences $\varphi_1, \dots, \varphi_n$ are the *connected components* of a primitive positive τ -sentence φ if they are the conjunctive queries of the connected components of the canonical database $D(\varphi)$.

Definition 2.8. Let τ be a finite relational signature and let \mathfrak{B} be a τ -structure. Then the *constraint satisfaction problem of \mathfrak{B}* ($\text{CSP}(\mathfrak{B})$) is the computational problem of deciding whether a given primitive positive τ -sentence holds in \mathfrak{B} .

Consider the following translation which associates to each \mathbf{A} -network $(V; f)$ a primitive positive A -sentences φ as follows: the variables of φ are the elements of V and φ contains for every (x, y) in the domain of f the conjunct $a(x, y)$ if and only if $f(x, y) = a$ holds. For the other direction let φ be an A -sentence with variable set X and consider the \mathbf{A} -network $(X; f)$ with the following definition: for every $x, y \in X$, if (x, y) does not appear in any conjunct from φ we leave $f(x, y)$ undefined, otherwise let $a_1(x, y), \dots, a_n(x, y)$ all the conjuncts from φ that contain (x, y) . We compute in \mathbf{A} the element $a := a_1 \cap \dots \cap a_n$ and define $f(x, y) := a$.

The following theorem, which subsumes the connection between network satisfaction problems and constraint satisfaction problems is based on this natural 1-to-1 correspondence between \mathbf{A} -networks and A -sentences.

Theorem 2.9 (Bodirsky, 2012, see also Bodirsky & Jonsson, 2017; Bodirsky, 2018). *Let $\mathbf{A} \in \text{RRA}$ be finite. Then the following holds:*

1. \mathbf{A} has a representation \mathfrak{B} such that $\text{NSP}(\mathbf{A})$ and $\text{CSP}(\mathfrak{B})$ are the same problem up to the translation between \mathbf{A} -networks and A -sentences.
2. If \mathbf{A} has a normal representation \mathfrak{B} the problems $\text{NSP}(\mathbf{A})$ and $\text{CSP}(\mathfrak{B})$ are the same up to the translation between \mathbf{A} -networks and A -sentences.

2.4 Model Theory

Let τ be a finite relational signature. The class of finite τ -structures that embed into a τ -structure \mathfrak{B} is called the *age* of \mathfrak{B} , denoted by $\text{Age}(\mathfrak{B})$. If \mathcal{F} is a class of finite τ -structures, then $\text{Forb}(\mathcal{F})$ is the class of all finite τ -structures \mathfrak{A} such that no structure from \mathcal{F} embeds into \mathfrak{A} . A class \mathcal{C} of finite τ structures is called *finitely bounded* if there exists a finite set of finite τ -structures \mathcal{F} such that $\mathcal{C} = \text{Forb}(\mathcal{F})$ (see, e.g., Macpherson, 2011). The structures from \mathcal{F} are called *bounds* or *forbidden substructures*. It is easy to see that a class \mathcal{C} of τ -structures is finitely bounded if and only if it is axiomatisable by a universal τ -sentence. A structure \mathfrak{B} is called finitely bounded if $\text{Age}(\mathfrak{B})$ is finitely bounded.

Proposition 2.10 (Bodirsky, 2018). *Let \mathfrak{B} be a normal representation of a finite $\mathbf{A} \in \text{RRA}$. Then the following holds:*

- \mathfrak{B} is finitely bounded by bounds of size at most three.
- The $A_0 \setminus \{\text{Id}\}$ -reduct of \mathfrak{B} is finitely bounded by bounds of size at most three.

Definition 2.11. *A class of finite τ -structures has the amalgamation property if for all structures $\mathfrak{A}, \mathfrak{B}_1, \mathfrak{B}_2 \in \mathcal{C}$ with embeddings $e_1: \mathfrak{A} \rightarrow \mathfrak{B}_1$ and $e_2: \mathfrak{A} \rightarrow \mathfrak{B}_2$ there exist a structure $\mathfrak{C} \in \mathcal{C}$ and embeddings $f_1: \mathfrak{B}_1 \rightarrow \mathfrak{C}$ and $f_2: \mathfrak{B}_2 \rightarrow \mathfrak{C}$ such that $f_1 \circ e_1 = f_2 \circ e_2$. If additionally $f_1(B_1) \cap f_2(B_2) = f_1(e_1(A)) = f_2(e_2(A))$, then we say that \mathcal{C} has the strong amalgamation property.*

Let $\mathfrak{B}_1, \mathfrak{B}_2$ be τ -structures. Then $\mathfrak{B}_1 \cup \mathfrak{B}_2$ is the τ -structure on the domain $B_1 \cup B_2$ such that $R^{\mathfrak{B}_1 \cup \mathfrak{B}_2} := R_1^{\mathfrak{B}_1} \cup R_2^{\mathfrak{B}_2}$ for every $R \in \tau$. If Definition 2.11 holds with $\mathfrak{C} := \mathfrak{B}_1 \cup \mathfrak{B}_2$ then we say that \mathcal{C} has the *free amalgamation property*; note that the free amalgamation property implies the strong amalgamation property.

Theorem 2.12 (Fraïssé; see, e.g., Hodges, 1997). *Let τ be a finite relational signature and let \mathcal{C} be a class of finite τ -structures that is closed under taking induced substructures and isomorphisms and has the amalgamation property. Then there exists an up to isomorphism unique countable homogeneous structure \mathfrak{B} such that $\mathcal{C} = \text{Age}(\mathfrak{B})$.*

Definition 2.13. *Let \mathfrak{B} be a relational structure. A set $O \subseteq B^n$ is called an n -orbit of $\text{Aut}(\mathfrak{B})$ if O is preserved by all $\alpha \in \text{Aut}(\mathfrak{B})$ and for all $x, y \in O$ there exists $\alpha \in \text{Aut}(\mathfrak{B})$ such that $\alpha(x) = y$.*

For a structure \mathfrak{B} the automorphism group $\text{Aut}(\mathfrak{B})$ is called *transitive* if $\text{Aut}(\mathfrak{B})$ has only one orbit. We want to remark that if \mathfrak{B} is a normal representation of a finite $\mathbf{A} \in \text{RRA}$ then the 2-orbits of $\text{Aut}(\mathfrak{B})$ are exactly the relations induced by the atoms A_0 of \mathbf{A} .

2.5 The Universal-Algebraic Approach

In this section we present basic notions for the so-called universal-algebraic approach to the study of CSPs.

Definition 2.14. *Let B be some set. We denote by $O_B^{(n)}$ the set of all n -ary operations on B and by $O_B := \bigcup_{n \in \mathbb{N}} O_B^{(n)}$ the set of all operations on B . A set $\mathcal{C} \subseteq O_B$ is called*

an operation clone on B if it contains all projections of all arities and if it is closed under composition, i.e., for all $f \in \mathcal{C}^{(n)} := \mathcal{C} \cap O_B^{(n)}$ and $g_1, \dots, g_n \in \mathcal{C} \cap O_B^{(s)}$ it holds that $f(g_1, \dots, g_n) \in \mathcal{C}$, where $f(g_1, \dots, g_n)$ is the s -ary function defined as follows

$$f(g_1, \dots, g_n)(x_1, \dots, x_s) := f(g_1(x_1, \dots, x_s), \dots, g_n(x_1, \dots, x_s)).$$

An operation $f: B^n \rightarrow B$ is called *conservative* if for all $x_1, \dots, x_n \in B$ it holds that $f(x_1, \dots, x_n) \in \{x_1, \dots, x_n\}$. A clone is called *conservative* if all its operations are conservative. Note that an operation clone \mathcal{C} on B can be considered as an algebra with domain B and an infinite signature that consists of symbols for all operations in \mathcal{C} (cf. Section 2.1). In this sense a conservative operation clone \mathcal{C} on B induces on every set $A \subseteq B$ a subalgebra. It is easy to see that this subalgebra corresponds to an operation clone on A . We call this the *restriction* of \mathcal{C} to A . We later need the following classical result for clones over a two-element set.

Theorem 2.15 (Post, 1941). *Let \mathcal{C} be a conservative operation clone on $\{0, 1\}$. Then either \mathcal{C} contains only projections, or at least one of the following operations:*

1. the binary function \min ,
2. the binary function \max ,
3. the minority function,
4. the majority function.

Operation clones occur naturally as polymorphism clones of relational structures. If $x_1, \dots, x_n \in B^k$ and $f: B^n \rightarrow B$, then we write $f(x_1, \dots, x_n)$ for the k -tuple obtained by applying f component-wise to the tuples x_1, \dots, x_n .

Definition 2.16. *Let \mathfrak{B} a structure with a finite relational signature τ and let $R \in \tau$. An n -ary operation preserves a relation $R^{\mathfrak{B}}$ if for all $x_1, \dots, x_n \in R^{\mathfrak{B}}$ it holds that*

$$f(x_1, \dots, x_n) \in R^{\mathfrak{B}}.$$

If f preserves all relations from \mathfrak{B} then f is called a polymorphism of \mathfrak{B} .

In order to provide an additional view on polymorphisms we give the following definition.

Definition 2.17. *Let τ be a relational signature and let \mathfrak{B} be a τ -structure. For $n \in \mathbb{N}_+$ the structure \mathfrak{B}^n is the τ -structure on the domain B^n defined as follows. Let $R \in \tau$ be of arity l . Then*

$$(x^1, \dots, x^l) \in R^{\mathfrak{B}^n} :\Leftrightarrow \forall i \in \{1, \dots, n\} : (x_i^1, \dots, x_i^l) \in R^{\mathfrak{B}}.$$

It is easy to see that the n -ary polymorphisms of \mathfrak{B} are precisely the homomorphisms from \mathfrak{B}^n to \mathfrak{B} .

The set of all polymorphisms (of all arities) of a relational structure \mathfrak{B} is an operation clone on B , which is denoted by $\text{Pol}(\mathfrak{B})$. A *Siggers operation* is an operation that satisfies the Siggers identity (see Theorem 1.1). The following result can be obtained by combining known results from the literature.

Theorem 2.18 (Siggers, 2010; Bulatov, 2003; see also Barto, 2011; Bulatov, 2016). *Let \mathfrak{B} be a finite structure with a finite relational signature such that $\text{Pol}(\mathfrak{B})$ is conservative. Then precisely one of the following holds:*

1. *There exist distinct $a, b \in B$ such that for every $f \in \text{Pol}(\mathfrak{B})^{(n)}$ the restriction of f to $\{a, b\}^n$ is a projection. In this case, $\text{CSP}(\mathfrak{B})$ is NP-complete.*
2. *$\text{Pol}(\mathfrak{B})$ contains a Siggers operation; in this case, $\text{CSP}(\mathfrak{B})$ is in P.*

We now discuss fundamental results about the universal-algebraic approach for constraint satisfaction problems of structures with an infinite domain.

Theorem 2.19 (Bodirsky & Nešetřil, 2006). *Let \mathfrak{B} be a homogeneous structure with finite relational signature. Then a relation is preserved by $\text{Pol}(\mathfrak{B})$ if and only if it is primitively positively definable in \mathfrak{B} .*

The following definition is a preparation to formulate the next theorem which is a well-known condition that implies NP-hardness of $\text{CSP}(\mathfrak{B})$ for homogeneous structures with a finite relational signature.

Definition 2.20. *Let \mathcal{K} be a class of algebras. Then we have*

- $\text{H}(\mathcal{K})$ *is the class of homomorphic images of algebras from \mathcal{K} and*
- $\text{S}(\mathcal{K})$ *is the class of subalgebras of algebras from \mathcal{K} .*
- $\text{P}^{\text{fin}}(\mathcal{K})$ *is the class of finite products of algebras from \mathcal{K} .*

An operation clone \mathcal{C} on a set B can also be seen as an algebra \mathbf{B} with domain B whose signature consists of the operations of \mathcal{C} such that $f^{\mathbf{B}} := f$ for all $f \in \mathcal{C}$.

The following is a classical condition for NP-Hardness, see for example Theorem 10 in the survey by Bodirsky (2008).

Theorem 2.21. *Let \mathfrak{B} be a homogeneous structure with finite relational signature. If $\text{HSP}^{\text{fin}}(\{\text{Pol}(\mathfrak{B})\})$ contains a 2-element algebra where all operations are projections, then $\text{CSP}(\mathfrak{B})$ is NP-hard.*

In the following let $\mathbf{A} \in \text{RRA}$ be finite and with normal representation \mathfrak{B} .

Definition 2.22. *Let $a_1, \dots, a_n \in A_0$ be atoms of \mathbf{A} . Then the $2n$ -ary relation $(a_1, \dots, a_n)^{\mathfrak{B}}$ is defined as follows:*

$$(a_1, \dots, a_n)^{\mathfrak{B}} := \{(x_1, \dots, x_n, y_1, \dots, y_n) \in B^{2n} \mid \bigwedge_{i \in \{1, \dots, n\}} a_i^{\mathfrak{B}}(x_i, y_i)\}.$$

An operation $f: B^n \rightarrow B$ is called *edge-conservative* if it satisfies for all $x, y \in B^n$ and all $a_1, \dots, a_n \in A_0$

$$(a_1, \dots, a_n)^{\mathfrak{B}}(x, y) \Rightarrow (f(x), f(y)) \in \bigcup_{i \in \{1, \dots, n\}} a_i^{\mathfrak{B}}.$$

Note that for every $D \subseteq A_0$ the structure \mathfrak{B} contains the relation $\bigcup_{a_i \in D} a_i^{\mathfrak{B}}$. Therefore the next proposition follows immediately since polymorphisms of \mathfrak{B} preserve all relations of \mathfrak{B} .

Proposition 2.23. *All polymorphisms of \mathfrak{B} are edge-conservative.*

Definition 2.24. *Let $X \subseteq A_0$. An operation $f: B^n \rightarrow B$ is called X -canonical (with respect to \mathfrak{B}) if there exists a function $\bar{f}: X^n \rightarrow A_0$ such that for all $a, b \in B^n$ and $O_1, \dots, O_n \in X$, if $(a_i, b_i) \in O_i$ for all $i \in \{1, \dots, n\}$ then $(f(a), f(b)) \in \bar{f}(O_1, \dots, O_n)^{\mathfrak{B}}$. An operation f is called canonical (with respect to \mathfrak{B}) if it is A_0 -canonical. In this case we say that the behaviour \bar{f} is total. If $X \subsetneq A_0$ we call \bar{f} a partial behaviour. The function \bar{f} is called the behaviour of f on X . If $X = A_0$ then \bar{f} is just called the behaviour of f .*

We denote by $\text{Pol}^{\text{can}}(\mathfrak{B})$ the set of all polymorphisms of \mathfrak{B} that are canonical with respect to \mathfrak{B} . It will always be clear from the context what the domain of a behaviour \bar{f} is. An operation $f: S^2 \rightarrow S$ is called *symmetric* if for all $x, y \in S$ it holds that $f(x, y) = f(y, x)$. An X -canonical function f is called *X -symmetric* if the behaviour of f on X is symmetric.

2.6 Ramsey Theory and Canonisation

We avoid giving an introduction to Ramsey theory, since the only usage of the Ramsey property is via Theorem 2.26, and rather refer to the survey by Bodirsky (2015) for an introduction.

Let \mathfrak{A} be a homogeneous τ -structure such that $\text{Age}(\mathfrak{A})$ has the strong amalgamation property. Then the class of all $(\tau \cup \{<\})$ -structures \mathfrak{A} such that $<^{\mathfrak{A}}$ is a linear order and whose τ -reduct (i.e. the structure on the same domain, but only with the relations that are denoted by symbols from τ , see e.g. the book by Hodges (1997)) is from $\text{Age}(\mathfrak{A})$ is a strong amalgamation class, too (see e.g. Bodirsky, 2015). By Theorem 2.12 there exists an up to isomorphism unique countable homogeneous structure of that age, which we denote by $\mathfrak{A}_{<}$. It can be shown by a straightforward back-and-forth argument that $\mathfrak{A}_{<}$ is isomorphic to an expansion of \mathfrak{A} , so we identify the domain of \mathfrak{A} and of $\mathfrak{A}_{<}$ along this isomorphism, and call $\mathfrak{A}_{<}$ the *expansion of \mathfrak{A} by a generic linear order*.

Theorem 2.25 (Nešetřil & Rödl, 1989; Hubička & Nešetřil, 2019). *Let \mathfrak{A} be a relational τ -structure such that $\text{Age}(\mathfrak{A})$ has the free amalgamation property. Then the expansion of \mathfrak{A} by a generic linear order has the Ramsey property.*

The following theorem gives a connection of the Ramsey property with the existence of canonical functions and plays a key role in our analysis.

Theorem 2.26 (Bodirsky & Pinsker, 2021). *Let \mathfrak{B} be a countable homogeneous structure with finite relational signature and the Ramsey property. Let $h: B^k \rightarrow B$ be an operation and let $L := \{(x_1, \dots, x_k) \mapsto \alpha(h(\beta_1(x_1), \dots, \beta_k(x_k))) \mid \alpha, \beta_1, \dots, \beta_k \in \text{Aut}(\mathfrak{B})\}$.*

Then there exists a canonical operation $g: B^k \rightarrow B$ such that for every finite $F \subset B$ there exists $g' \in L$ such that $g'|_{F^k} = g|_{F^k}$.

Remark 2.27. *Let \mathfrak{A} and \mathfrak{B} be homogeneous structures with finite relational signatures. If \mathfrak{A} and \mathfrak{B} have the same domain and the same automorphism group, then \mathfrak{A} has the Ramsey property if and only if \mathfrak{B} has it (see, e.g., Bodirsky, 2015).*

3. Representable Relation Algebras with a Flexible Atom

In this section we define the concept of a flexible atom and show how to reduce the classification problem for the network satisfaction problem for a finite $\mathbf{A} \in \text{RRA}$ with a flexible atom to the situation where \mathbf{A} is additionally integral (Proposition 3.3). A finite representable relation algebra \mathbf{A} is called *integral* if the element Id is an atom of \mathbf{A} (cf. Maddux, 2006b).

Then we show that an integral $\mathbf{A} \in \text{RRA}$ with a flexible atom has a normal representation. Therefore, the universal-algebraic approach is applicable; in particular, we make heavy use of polymorphisms and their connection to primitive positive definability in later sections (cf. Theorem 2.19). Furthermore, we prove that every normal representation of a finite representable relation algebra with a flexible atom has a Ramsey expansion (Section 3.2). Therefore, the tools from Section 2.6 can be applied, too. Finally we give some examples of representable relation algebras with a flexible atom (Section 3.3). We start with the definition of a flexible atom.

Definition 3.1. *Let $\mathbf{A} \in \text{RRA}$ and let $I := \{a \in A \mid a \leq \text{Id}\}$. An atom $s \in A_0 \setminus I$ is called flexible if for all $a, b \in A \setminus I$ it holds that $s \leq a \circ b$.*

This definition can for example be found in the book by Hirsch and Hodkinson (2002, Chapter 11, Exercise 1). Note that this definition does not require the representable relation algebra \mathbf{A} to be integral. This is slightly more general than the definition by Maddux (1994, 2006b). As mentioned before, we show in the following section that it is sufficient for our result to classify the computational complexity of NSPs for finite representable relation algebras with a flexible atom that are additionally integral. This means that readers who prefer this second definition by Maddux (1994, 2006b) (assuming integrality) can perfectly skip the following section and read the article with this other definition in mind. In this case representable relation algebras with a flexible atom are always implicitly integral.

3.1 Integral Representable Relation Algebras

Let $\mathbf{A} \in \text{RRA}$ and let $I := \{a \in A \mid a \leq \text{Id}\}$. The atoms in $I \cap A_0$ are called *identity atoms*. Therefore, \mathbf{A} is integral if and only if \mathbf{A} has exactly one identity atom. The first claim of the following lemma is a well-known fact about atoms of representable relation algebras. This fact is used to prove a second claim about representable relation algebras with a flexible atom. In simple terms, this claim states that a representation of a non-integral representable relation algebra with a flexible atom is only “square” on precisely those elements that are in one certain identity atom.

Lemma 3.2. *Let $\mathbf{A} \in \text{RRA}$ be finite. Then there exists for every atom s a unique $e_1 \in A_0$ with $0 < e_1 \leq \text{Id}$ such that $s = e_1 \circ s$. Furthermore, if s is a flexible atom then for all $e_2 \in A_0$ with $0 < e_2 \leq \text{Id}$ and $e_2 \neq e_1$ we have that $e_2 \circ \overline{\text{Id}} = 0$.*

Proof. Note that $\text{Id} \circ s = s$ by definition and therefore $e \circ s \subseteq s$ for all $e \in A_0$ with $0 < e \leq \text{Id}$. Since s is an atom either $e \circ s = 0$ or $e \circ s = s$. By $\text{Id} = \bigcup \{e \in A_0 \mid 0 < e \leq \text{Id}\}$ and $\text{Id} \circ s = s$ there exists at least one $0 < e \leq \text{Id}$ with $e \circ s = s$. In the next step we prove uniqueness of such an element e . Assume for contradiction that there exist distinct $e_1, e_2 \in A_0$ with $0 < e_1 \leq \text{Id}$ and $0 < e_2 \leq \text{Id}$ such that $e_1 \circ s = s$ and $e_2 \circ s = s$. Note that $e_1 \circ e_2 = 0$ since

e_1 and e_2 are identity atoms. Therefore, we get

$$0 = 0 \circ s = (e_1 \circ e_2) \circ s = e_1 \circ (e_2 \circ s) = e_1 \circ s = s,$$

which is a contradiction since s is an atom. This proves the first statement.

For the second statement assume for contradiction that there exists $e_2 \in A_0 \setminus \{e_1\}$ such that $e_2 \leq \text{Id}$ and $e_2 \circ \overline{\text{Id}} \neq 0$. Let a be an atom with $a \leq e_2 \circ \overline{\text{Id}}$. Since $e_1 \circ e_2 = 0$ we get $e_1 \circ a \leq e_1 \circ (e_2 \circ \overline{\text{Id}}) = (e_1 \circ e_2) \circ \overline{\text{Id}} = 0 \circ \overline{\text{Id}} = 0$. Since s is a flexible atom it holds that $s \leq a \circ a^\smile$ and therefore

$$s = e_1 \circ s \leq e_1 \circ (a \circ a^\smile) = (e_1 \circ a) \circ a^\smile = 0 \circ a^\smile = 0,$$

which is a contradiction. \square

Proposition 3.3. *Let $\mathbf{A} \in \text{RRA}$ be finite and with a flexible atom. Then there exists a finite integral $\mathbf{A}' \in \text{RRA}$ with a flexible atom such that the following statements hold:*

1. *There exists a polynomial-time many-one reduction from $\text{NSP}(\mathbf{A})$ to $\text{NSP}(\mathbf{A}')$.*
2. *There exists a polynomial-time many-one reduction from $\text{NSP}(\mathbf{A}')$ to $\text{NSP}(\mathbf{A})$.*
3. *The atom structure of \mathbf{A} has a polymorphism that satisfies the Siggers identity if and only if the atom structure of \mathbf{A}' has such a polymorphism (see Definition 4.1).*

Proof. If \mathbf{A} is integral there is nothing to be shown. So assume that \mathbf{A} is not integral and let s be a flexible atom. Let \mathfrak{B} be a representation of \mathbf{A} such that $\text{NSP}(\mathbf{A})$ and $\text{CSP}(\mathfrak{B})$ are the same problem up to the translation between \mathbf{A} -networks and A -sentences. Such a representation exists by Theorem 2.9. Let $(x, y) \in \overline{\text{Id}}^{\mathfrak{B}}$ and let $e_1 \in A_0$ be the unique element with $e_1 \leq \text{Id}$ and $s = e_1 \circ s$ that exists by Lemma 3.2. The second statement of Lemma 3.2 implies $e_1 \circ \overline{\text{Id}} = \overline{\text{Id}}$ and therefore we have that $(x, x) \in e_1^{\mathfrak{B}}$ and $(y, y) \in e_1^{\mathfrak{B}}$. Let \mathfrak{C}' be the substructure of \mathfrak{B} on the domain $\{x \in B \mid (x, x) \in e_1^{\mathfrak{B}}\}$. The set of relations of \mathfrak{C}' clearly induces a proper relation algebra which is integral. We denote this representable relation algebra by \mathbf{A}' . Note that we can also consider A' as a subset of A . Let \mathfrak{B}' be the representation of \mathbf{A}' such that $\text{NSP}(\mathbf{A}')$ and $\text{CSP}(\mathfrak{B}')$ are the same problem up to the translation between \mathbf{A}' -networks and A' -sentences. As before, such a representation exists by Theorem 2.9.

Proof of 1.: Note that if a connected instance of $\text{CSP}(\mathfrak{B})$ is satisfiable, then either all variables are mapped to an atom from the subset of A_0 that corresponds to A'_0 or all variables are mapped to one element x with $e^*(x, x)$ and $e^* \in \{e \in A_0 \mid e \leq \text{Id} \text{ and } e \neq e_1\}$. This leads to the following polynomial-time many-one reduction from $\text{CSP}(\mathfrak{B})$ to $\text{CSP}(\mathfrak{B}')$, which proves together with Theorem 2.9 the claim that there exists a polynomial-time many-one reduction from $\text{NSP}(\mathbf{A})$ to $\text{NSP}(\mathbf{A}')$. Consider the following algorithm: For a given primitive positive A -sentence φ it computes the connected components $\varphi_1, \dots, \varphi_n$ (see paragraph after Definition 2.7). Then it checks for every $i \in \{1, \dots, n\}$ whether there exists $e^* \in \{e \in A_0 \mid e \leq \text{Id} \text{ and } e \neq e_1\}$ such that for every conjunct $a(x, y)$ of φ_i it holds that $e^* \leq a$. Let $I \subseteq \{1, \dots, n\}$ be the set of indices for which this is not the case. Then the algorithm defines new A' -sentences φ'_i from φ_i for every $i \in I$ by replacing every conjunct

$a(x, y)$ with $a'(x, y)$ where $a' := a \setminus \{e \in A_0 \mid e \leq \text{Id} \text{ and } e \neq e_1\}$. Let φ' be the conjunction of all φ'_i for $i \in I$.

It is easy to see that φ' is computable from φ in polynomial time and that φ is a satisfiable instance of $\text{CSP}(\mathfrak{B})$ if and only if φ' is a satisfiable instance of $\text{CSP}(\mathfrak{B}')$.

Proof of 2.: Consider now an \mathbf{A}' -network $(V; f')$. We claim that $(V; f')$ is satisfiable as an \mathbf{A}' -network if and only if it is satisfiable as an \mathbf{A} -network. Suppose that $(V; f')$ is satisfiable in a representation \mathfrak{D}' of \mathbf{A}' by an assignment α . Let y_i be fresh elements for every atom $e_i \leq \text{Id}$ with $e_i \neq e_1$. We build the disjoint union of \mathfrak{D}' with one-element $\{e_i\}$ -structures $(\{y_i\}; \{(y_i, y_i)\})$ and then close the structure under union and intersection of binary relations. This results in a representation of \mathbf{A} that satisfies $(V; f')$ again by the assignment α . For the other direction, if $(V; f')$ is satisfiable in a representation \mathfrak{D} of \mathbf{A} we can again consider the substructure on the domain $(x, x) \in e_1^{\mathfrak{D}}$ and get a representation of \mathbf{A}' that satisfies $(V; f')$.

Proof of 3.: Let g be a polymorphism of the atom structure of \mathbf{A} that satisfies the Siggers identity. By assumption g satisfies

$$\forall x_1, \dots, x_6 \in A_0. g(x_1, \dots, x_6) \in \{x_1, \dots, x_6\}$$

and therefore the restriction of g to $(A_0 \cap A')^6$ is a polymorphism of the atom structure of \mathbf{A}' .

For the other direction choose an arbitrary ordering of the atoms $\{e \in A_0 \mid e \leq \text{Id} \text{ and } e \neq e_1\} = \{l_1, \dots, l_j\}$. If g is a Siggers polymorphism of the atom structure of \mathbf{A}' one can extend g to an operation $g^*: A^6 \rightarrow A$ by defining

$$g^*(x_1, \dots, x_6) := \begin{cases} \min(\{l_1, \dots, l_j\} \cap \{x_1, \dots, x_6\}) & \text{if } \{l_1, \dots, l_j\} \cap \{x_1, \dots, x_6\} \neq \emptyset, \\ g(x_1, \dots, x_6) & \text{otherwise.} \end{cases}$$

It is easy to see that this operation satisfies also the Siggers identity. Furthermore, since every atom e from $\{e \in A_0 \mid e \leq \text{Id} \text{ and } e \neq e_1\}$ is only contained in allowed triples of the form (e, e, e) it follows that f^* preserves the allowed triples from \mathbf{A} (see after Definition 4.1). \square

3.2 Normal Representations

Let $\mathbf{A} \in \text{RRA}$ be for the rest of the section finite, integral, and with a flexible atom s . We consider the following subset of A :

$$A - s := \{a \in A \mid s \not\leq a\}.$$

Let (V, g) be an \mathbf{A} -network and let \mathfrak{C} be the corresponding A -structure (see paragraph before Definition 2.5). Let $\mathfrak{C} - s$ be the $(A - s)$ -structure on the same domain V as \mathfrak{C} such that for all $x, y \in V$ and $a \in (A - s) \setminus \{0\}$ we have

$$a^{\mathfrak{C}-s}(x, y) \text{ if and only if } (a^{\mathfrak{C}}(x, y) \vee (a \cup s)^{\mathfrak{C}}(x, y)).$$

We call $\mathfrak{C} - s$ the s -free companion of an \mathbf{A} -network (V, f) .

The next lemma follows directly from the definitions of flexible atoms and s -free companions.

Lemma 3.4. *Let \mathcal{C} be the class of s -free companions of atomic closed \mathbf{A} -networks. Then \mathcal{C} has the free amalgamation property.*

Proof. Let \mathfrak{A} , \mathfrak{B}_1 , and \mathfrak{B}_2 be structures in \mathcal{C} with embeddings $e_1: \mathfrak{A} \rightarrow \mathfrak{B}_1$ and $e_2: \mathfrak{A} \rightarrow \mathfrak{B}_2$. Since s is a flexible atom and \mathcal{C} is a class of s -free companions we get that the structure $\mathfrak{C} := \mathfrak{B}_1 \cup \mathfrak{B}_2$ is in \mathcal{C} . Therefore, the natural embeddings $f_1: \mathfrak{B}_1 \rightarrow \mathfrak{C}$ and $f_2: \mathfrak{B}_2 \rightarrow \mathfrak{C}$ prove the free amalgamation property of \mathcal{C} . \square

As a consequence of this lemma we obtain the following.

Proposition 3.5. *\mathbf{A} has a normal representation \mathfrak{B} .*

Proof. Let \mathcal{C} be the class from Lemma 3.4. This class is closed under taking substructures and isomorphisms. By Lemma 3.4 it also has the amalgamation property and therefore we get by Theorem 2.12 a homogeneous structure \mathfrak{B}' with $\text{Age}(\mathfrak{B}') = \mathcal{C}$. Let \mathfrak{B}'' be the expansion of \mathfrak{B}' by the following relation

$$s(x, y) :\Leftrightarrow \bigwedge_{a \in A_0 \setminus \{s\}} \neg a^{\mathfrak{B}'}(x, y).$$

Let \mathfrak{B} be the (homogeneous) expansion of \mathfrak{B}'' by all Boolean combinations of relations from \mathfrak{B}'' . Then \mathfrak{B} is a representation of $\mathbf{A} \in \text{RRA}$. Since $\text{Age}(\mathfrak{B}')$ is the class of all atomic closed \mathbf{A} -networks, \mathfrak{B} is fully universal. The definition of s witnesses that \mathfrak{B} is a square representation of \mathbf{A} : for all elements $x, y \in B$ there exists an atom $a \in A_0$ such that $a^{\mathfrak{B}}(x, y)$ holds. \square

The next theorem is another consequence of Lemma 3.4.

Theorem 3.6. *Let \mathfrak{B} be a normal representation of \mathbf{A} . Let $\mathfrak{B}_{<}$ be the expansion of \mathfrak{B} by a generic linear order. Then $\mathfrak{B}_{<}$ has the Ramsey property.*

Proof. Let \mathfrak{B}' be the $(A_0 \setminus \{s\})$ -reduct of \mathfrak{B} . The age of this structure has the free amalgamation property by Lemma 3.4. Therefore, Theorem 2.25 implies that the expansion of \mathfrak{B}' by a generic linear order has the Ramsey property. By Remark 2.27 the structure $\mathfrak{B}_{<}$ also has the Ramsey property since $\mathfrak{B}_{<}$ and $(\mathfrak{B}')_{<}$ have the same automorphism group. \square

Remark 3.7. *The binary first-order definable relations of $\mathfrak{B}_{<}$ form a proper relation algebra since $\mathfrak{B}_{<}$ has quantifier-elimination (see Hodges, 1997). By the definition of the generic order the atoms of this proper relation algebra are of the following form*

- $a^{\mathfrak{B}_{<}} \cap <^{\mathfrak{B}_{<}}$ for $a \in A_0 \setminus \{\text{Id}\}$, or
- $a^{\mathfrak{B}_{<}} \cap >^{\mathfrak{B}_{<}}$ for $a \in A_0 \setminus \{\text{Id}\}$, or
- Id ,

where the relation $>^{\mathfrak{B}_{<}}$ consists of all tuples (x, y) such that $(y, x) \in <^{\mathfrak{B}_{<}}$ holds.

◦	Id	<i>a</i>	<i>b</i>
Id	Id	<i>a</i>	<i>b</i>
<i>a</i>	<i>a</i>	{Id, <i>a</i> , <i>b</i> }	{ <i>a</i> , <i>b</i> }
<i>b</i>	<i>b</i>	{ <i>a</i> , <i>b</i> }	{Id, <i>a</i> , <i>b</i> }

◦	Id	<i>a</i>	<i>b</i>
Id	Id	<i>a</i>	<i>b</i>
<i>a</i>	<i>a</i>	{Id, <i>b</i> }	{ <i>a</i> , <i>b</i> }
<i>b</i>	<i>b</i>	{ <i>a</i> , <i>b</i> }	{Id, <i>a</i> , <i>b</i> }

Figure 2: Multiplication tables of representable relation algebras #18 (left) and #17 (right).

3.3 Examples

We give two concrete examples of finite integral symmetric representable relation algebras with a flexible atom (Examples 3.8 and 3.9), and a systematic way of building such algebras from arbitrary representable relation algebras (Example 3.10). The numbering of the algebras in the examples is by Andréka and Maddux (1994).

Example 3.8 (Representable relation algebra #18). *The representable relation algebra #18 has three atoms, namely the identity atom Id and two symmetric atoms a and b. The multiplication table for the atoms is given in Fig. 2. In this representable relation algebra the atoms a and b are flexible. Consider the countable, homogeneous, undirected graph $\mathfrak{R} = (V; E^{\mathfrak{R}})$, whose age is the class of all finite undirected graphs (see, e.g., Hodges, 1997), also called the Random graph. The expansion of \mathfrak{R} by all binary first-order definable relations is a normal representation of the algebra #18. In this representation the atoms a and b are interpreted as the relation $E^{\mathfrak{R}}$ and the relation $N^{\mathfrak{R}}$, where $N^{\mathfrak{R}}$ is defined as $\neg E(x, y) \wedge x \neq y$.*

Example 3.9 (Representable relation algebra #17). *The representable relation algebra #17 also consists of three symmetric atoms. The multiplication table in Fig. 2 shows that in this algebra the element b is a flexible atom. To see that a is not a flexible atom, note that $a \not\leq a \circ a = \{\text{Id}, b\}$. Let $\mathfrak{R} = (V; E^{\mathfrak{R}})$ be the countable, homogeneous, undirected graph, whose age is the class of all finite undirected graphs that do not embed the complete graph on three vertices (see, e.g., Hodges, 1997). This structure is called a Henson graph. If we expand \mathfrak{R} by all binary first-order definable relations we get a normal representation of the algebra #17. To see this note that we interpret a as the relation $E^{\mathfrak{R}}$. That \mathfrak{R} is triangle free, i.e. triangles of $E^{\mathfrak{R}}$ are forbidden, matches with the fact that $a \not\leq a \circ a$ holds in the representable relation algebra.*

Example 3.10. *Consider an arbitrary finite, integral $\mathbf{A} = (A; \cup, \bar{\cdot}, 0, 1, \text{Id}, \smile, \circ) \in \text{RRA}$. Clearly \mathbf{A} does not have a flexible atom s in general. Nevertheless we can expand the domain of \mathbf{A} to implement an “artificial” flexible atom.*

*Let s be some symbol not contained in A. Let us mention that every element in \mathbf{A} can uniquely be written as a union of atoms from A_0 . Let A' be the set of all subsets of $A_0 \cup \{s\}$. The set A' is the domain of our new algebra \mathbf{A}' . Note that on A' there exists the subset-ordering and A' is closed under set-union and complement (in $A_0 \cup \{s\}$) We define s to be symmetric and therefore get the following unary function * in \mathbf{A}' as follows. For an element $x \in A'$ we define*

$$x^* := \begin{cases} y^\smile \cup \{s\} & \text{if } x = y \cup \{s\} \text{ for } y \in A, \\ x^\smile & \text{otherwise.} \end{cases}$$

The new function symbol \circ'_A in \mathbf{A}' is defined on the atoms $A_0 \cup \{s\}$ as follows:

$$x \circ_{A'} y := \begin{cases} A_0 \cup \{s\} & \text{if } \{s\} = \{x, y\}, \\ (A_0 \setminus \{\text{Id}\}) \cup \{s\} & \text{if } \{s, a\} = \{x, y\} \text{ for } a \in A_0 \setminus \{s, \text{Id}\}, \\ \{a\} & \text{if } \{\text{Id}, a\} = \{x, y\} \text{ for } a \in A_0 \cup \{s\}, \\ (x \circ y) \cup \{s\} & \text{otherwise.} \end{cases}$$

One can check that $\mathbf{A}' = (A'; \cup, \bar{\cdot}, \emptyset, A_0 \cup \{s\}, \text{Id}, *, \circ_{A'})$ is a finite integral representable relation algebra with a flexible atom s . Note that the forbidden triples of \mathbf{A}' are exactly those of \mathbf{A} together with triples which are permutations of (s, a, Id) for some $a \in A_0$.

4. Polynomial-time Tractability

In this section we introduce for every finite $\mathbf{A} \in \text{RRA}$ an associated finite structure, called the *atom structure* of \mathbf{A} . Note that it is closely related, but not the same, as the type structure introduced by Bodirsky and Mottet (2016). In the context of relation algebras the atom structure has the advantage that its domain is the set of atoms of \mathbf{A} , rather than the set of 3-types, which would be the domain of the type structure of Bodirsky and Mottet (2016); hence, our domain is smaller and has some advantages on which the main result of this section (Proposition 4.4) is based. Up to some differences in the signature, our atom structure is the same as the atom structure introduced by Lyndon (1950) which was used there for different purposes (see also Maddux, 1982; Hirsch & Hodkinson, 2001; Hirsch, Jackson, & Kowalski, 2019).

Let \mathfrak{B} be a normal representation of a finite $\mathbf{A} \in \text{RRA}$. We will reduce $\text{CSP}(\mathfrak{B})$ to the CSP of the atom structure of \mathbf{A} . This means that if the CSP of the atom structure is in P, then so are $\text{CSP}(\mathfrak{B})$ and $\text{NSP}(\mathbf{A})$. For our main result we will show later that every network satisfaction problem for a finite integral symmetric representable relation algebra with a flexible atom that cannot be solved in polynomial time by this method is NP-complete.

Definition 4.1. *The atom structure of $\mathbf{A} \in \text{RRA}$ is the finite relational structure \mathfrak{D} with domain A_0 and the following relations:*

- for every $x \in A$ the unary relation $x^\mathfrak{D} := \{a \in A_0 \mid a \leq x\}$,
- the binary relation $E^\mathfrak{D} := \{(a_1, a_2) \in A_0^2 \mid a_1 \bar{\cdot} = a_2\}$,
- the ternary relation $H^\mathfrak{D} := \{(a_1, a_2, a_3) \in A_0^3 \mid a_3 \leq a_1 \circ a_2\}$.

Note that the relation $H^\mathfrak{D}$ consists of the allowed triples of $\mathbf{A} \in \text{RRA}$. We say that an operation *preserves the allowed triples* if it preserves the relation $H^\mathfrak{D}$.

Proposition 4.2. *Let \mathfrak{B} be a fully universal representation of a finite $\mathbf{A} \in \text{RRA}$. There is a polynomial-time reduction from $\text{CSP}(\mathfrak{B})$ to $\text{CSP}(\mathfrak{D})$.*

Proof. Let Ψ be an instance of $\text{CSP}(\mathfrak{B})$ with variable set $X = \{x_1, \dots, x_n\}$. We construct an instance Φ of $\text{CSP}(\mathfrak{D})$ as follows. The variable set Y of Φ is given by $Y := \{(x_i, x_j) \in X^2 \mid i \leq j\}$. The constraints of Φ are of the two kinds:

1. Let $a \in A$ be an element of $\mathbf{A} \in \text{RRA}$ and let $a((x_i, y_j))$ be an atomic formula of Ψ . If $i < j$, then we add the atomic (unary) formula $a((x_i, x_j))$ to Φ ; otherwise we add the atomic formula $a^\smile((x_j, x_i))$. If $j = i$ we additionally add $\text{Id}((x_i, x_j))$.
2. Let $x_i, x_j, x_l \in X$ be such that $i \leq j \leq l$. Then we add the atomic formula $H((x_i, x_j), (x_j, x_l), (x_i, x_l))$ to Φ .

It remains to show that this reduction is correct. Let $s: X \rightarrow B$ be a satisfying assignment for Ψ . This assignment maps every pair of variables x_i and x_j to a unique atom in A_0 and therefore induces a map $s': Y \rightarrow A_0$. The map s' clearly satisfies all atomic formulas introduced by (1.). To see that it also satisfies all formulas introduced by (2.) note that s maps the elements $x_i, x_j, x_l \in X$ to a substructure of \mathfrak{B} , which does not induce a forbidden triple.

For the other direction assume that $s': Y \rightarrow A_0$ is a satisfying assignment for Φ . This induces an A -structure \mathfrak{X} on X (maybe with some identification of variables) as follows: we add (x_i, x_j) to the relation $a^\mathfrak{X}$ if $i \leq j$ and $s'((x_i, x_j)) = a$; if otherwise $j < i$ and $s'((x_j, x_i)) = a$ we add (x_i, x_j) to the relation $(a^\smile)^\mathfrak{X}$. It is clear that no forbidden triple from \mathbf{A} is induced by \mathfrak{X} . Also note that \mathfrak{X} satisfies Ψ by the choice of the (unary) constraints of the first kind. Since \mathfrak{B} is a fully universal representation the structure \mathfrak{X} is a substructure of \mathfrak{B} . Hence, the instance Ψ is satisfiable in \mathfrak{B} . \square

The atom structure has another property which is fundamental for our proof of Theorem 1. Recall that every canonical polymorphism f induces a behaviour $\bar{f}: A_0^n \rightarrow A_0$. In the next proposition we show that then \bar{f} is a polymorphism of \mathfrak{D} . Moreover the other direction also holds. Every $g \in \text{Pol}(\mathfrak{D})$ is the behaviour of a canonical polymorphism of \mathfrak{B} .

Proposition 4.3. *Let \mathfrak{B} be a normal representation of a finite $\mathbf{A} \in \text{RRA}$.*

1. *Let $g \in \text{Pol}(\mathfrak{B})^{(n)}$ be canonical and let $\bar{g}: A_0^n \rightarrow A_0$ be its behaviour. Then $\bar{g} \in \text{Pol}(\mathfrak{D})^{(n)}$.*
2. *Let $f \in \text{Pol}(\mathfrak{D})^{(n)}$. Then there exists a canonical $g \in \text{Pol}(\mathfrak{B})^{(n)}$ whose behaviour equals f .*

Proof. For (1): Let $g \in \text{Pol}(\mathfrak{B})^{(n)}$ be canonical and let $c^1, \dots, c^n \in H^\mathfrak{D}$. Then by the definition of $H^\mathfrak{D}$ there exist tuples $x^1, \dots, x^n \in B^3$ such that for all $i \in \{1, \dots, n\}$ we have

$$c_1^{i\mathfrak{B}}(x_1^i, x_2^i), c_2^{i\mathfrak{B}}(x_2^i, x_3^i), \text{ and } c_3^{i\mathfrak{B}}(x_1^i, x_3^i).$$

We apply the canonical polymorphism g and get $y := g(x^1, \dots, x^n) \in B^3$. Then there exists an allowed triple $(d_1, d_2, d_3) \in A_0^3$ such that

$$d_1^{\mathfrak{B}}(y_1, y_2), d_2^{\mathfrak{B}}(y_2, y_3), \text{ and } d_3^{\mathfrak{B}}(y_1, y_3).$$

We have that $d = (d_1, d_2, d_3) \in H^\mathfrak{D}$ and by the definition of the behaviour of a canonical function we get $\bar{g}(c^1, \dots, c^n) = d$. The other relations in \mathfrak{D} are preserved trivially and therefore $\bar{g} \in \text{Pol}(\mathfrak{D})^{(n)}$.

For (2): Since \mathfrak{B} is fully universal and homogeneous it follows by a compactness argument (see e.g., Lemma 2 by Bodirsky & Dalmau, 2013) that every countable A_0 -structure

which does not induce a forbidden triple and its square has a homomorphism to \mathfrak{B} . It is therefore enough to show that every operation $h: B^n \rightarrow B$ with behaviour f does not induce a forbidden triple in the image. Let $x^1, \dots, x^n \in B^3$ be such that the application of a canonical function with behaviour f on x^1, \dots, x^n would give a tuple $y \in B^3$ with $d = (d_1, d_2, d_3) \in A_0^3$ such that

$$d_1^{\mathfrak{B}}(y_1, y_2), d_2^{\mathfrak{B}}(y_2, y_3), \text{ and } d_3^{\mathfrak{B}}(y_1, y_3).$$

Since f preserves $H^{\mathfrak{D}}$ the triple d is not forbidden. □

Recall from Proposition 2.23 that polymorphisms of \mathfrak{B} are edge-conservative. Note that this implies that polymorphisms of \mathfrak{D} are conservative. In fact, Theorem 2.18 and the Proposition 4.3 imply the following.

Proposition 4.4. *If $\text{Pol}(\mathfrak{B})$ contains a canonical polymorphism s whose behaviour \bar{s} is a Siggers operation in $\text{Pol}(\mathfrak{D})$ then $\text{CSP}(\mathfrak{B})$ is in P .*

We demonstrate how this result can be used to prove polynomial-time tractability of $\text{NSP}(\mathbf{A})$ for a symmetric, integral $\mathbf{A} \in \text{RRA}$ with a flexible atom.

Example 4.5 (Polynomial-time tractability the NSP of representable relation algebra #18).

The polynomial-time tractability of the NSP of the representable relation algebra #18 (see Example 3.8) was first shown by Cristiani and Hirsch (2004) (see also Section 8.4 of Bodirsky & Pinsker, 2015). Here we consider the following function $\bar{s}: \{\text{Id}, a, b\}^6 \rightarrow \{\text{Id}, a, b\}$.

$$\bar{s}(x_1, \dots, x_6) := \begin{cases} a & \text{if } a \in \{x_1, \dots, x_6\}, \\ b & \text{if } b \in \{x_1, \dots, x_6\} \text{ and } a \notin \{x_1, \dots, x_6\}, \\ \text{Id} & \text{otherwise.} \end{cases}$$

Let \mathfrak{R}' be the normal representation of the algebra #18 given in Example 3.8. Note that \bar{s} is the behaviour of an injective, canonical polymorphism of \mathfrak{R}' . The injectivity follows from the last line of the definition; if $\bar{s}(x_1, \dots, x_6) = \text{Id}$ then $\{x_1, \dots, x_6\} = \{\text{Id}\}$. Therefore \bar{s} preserves all allowed triples, since in the algebra #18 the only forbidden triples involve Id. One can check that \bar{s} is a Siggers operation and therefore we get by Proposition 4.4 that $\text{NSP}(\#18)$ is in P .

Example 4.6. *Consider the construction of representable relation algebras with a flexible atom from Example 3.10. It is easy to see that $\text{NSP}(\mathbf{A})$ for a finite integral $\mathbf{A} \in \text{RRA}$ has a polynomial-time reduction to $\text{NSP}(\mathbf{A}')$ where \mathbf{A}' is the representable relation algebra with a flexible atom that is constructed in Example 3.10. We get as a consequence that if a normal representation of \mathbf{A}' satisfies the condition of Proposition 4.4 then $\text{NSP}(\mathbf{A})$ is in P .*

Theorem 2.18 has additionally to Proposition 4.4 another important consequence.

Corollary 4.7. *If $\text{Pol}(\mathfrak{D})$ does not have a Siggers operation then there exist elements $a_1, a_2 \in A_0$ such that the restriction of every operation from $\text{Pol}(\mathfrak{D})^{(n)}$ to $\{a_1, a_2\}^n$ is a projection.*

5. Network Consistency Problems

The purpose of this section is to give an additional perspective on the class of network satisfaction problems of finite symmetric integral representable relation algebras with a flexible atom. Even more, we define these computational problems in this section completely without the use of the relation algebra framework. Our classification result for these problems does not depend on the content of this section and the reader may skip it.

We introduce a class of computational decision problems which we call *network consistency problems (NCPs)*. It is easy to see that NCPs are in a 1-to-1 correspondence with NSPs of finite, symmetric, integral $\mathbf{A} \in \text{RRA}$ with a flexible atom.

Definition 5.1. *Let A be a finite set and $R \subseteq A^3$. Then R is called totally symmetric if for all bijections $\pi: \{1, 2, 3\} \rightarrow \{1, 2, 3\}$ we have*

$$(a_1, a_2, a_3) \in R \Rightarrow (a_{\pi(1)}, a_{\pi(2)}, a_{\pi(3)}) \in R.$$

We call an element $p \in A$ identity element if for all $x, y \in A$ the following holds:

$$(p, x, y) \in R \Leftrightarrow x = y.$$

A structure $(A; R)$ is called a stencil if R is totally symmetric and it contains an identity element.

Definition 5.2. *Let $(G; F)$ be an undirected graph and let Q be a set. We call a map $c: F \rightarrow Q$ an edge Q -coloring of $(G; F)$ if for all $x, y \in G$ with $(x, y) \in F$ it holds that $c((x, y)) = c((y, x))$.*

For each fixed stencil, we define an NCP as follows.

Definition 5.3. *Let (A, R) be a stencil. The network completion problem of (A, R) , denoted by $\text{NCP}(A, R)$, is the following problem. Given a finite undirected graph $(G; F)$ with an edge $\mathcal{P}(A)$ -coloring f the task is to decide whether there exists an edge A -coloring f' of $(G; F)$ such that*

1. *for all $x, y \in G$ with $(x, y) \in F$ it holds that $f'((x, y)) \in f((x, y))$.*
2. *for all $x, y, z \in G$ with $(x, y), (y, z), (x, z) \in F$ we have*

$$(f'((x, y)), f'((y, z)), f'((x, z))) \in R.$$

The following proposition illustrates how NCPs correspond to a certain class of NSPs.

Proposition 5.4. *The class of NCPs and the class of NSPs for finite symmetric integral representable relation algebras with a flexible atom are in a natural 1-to-1 correspondence such that corresponding problems are polynomial-time equivalent.*

Proof. Let (A', R') be a stencil with $p \in A'$ according to (2) in Definition 5.1 and let s be new element with $s \notin A'$. We define a relational structure \mathfrak{D} as follows. The domain of \mathfrak{D} is the set $A' \cup \{s\}$. We assume that \mathfrak{D} has every subset of its domain as a unary relation.

Furthermore \mathfrak{D} contains the binary relation $E^{\mathfrak{D}} := \{(x, x) \mid x \in A' \cup \{s\}\}$ and the ternary relation $H^{\mathfrak{D}}$ which is defined as follows:

$$(x, y, z) \in H^{\mathfrak{D}} :\Leftrightarrow \left((x, y, z) \in R' \vee (s \in \{x, y, z\} \wedge p \notin \{x, y, z\}) \vee ((x, y, z) \in \{(p, s, s), (s, p, s), (s, s, p)\}) \right).$$

One can find a finite symmetric integral algebra $\mathbf{A} \in \text{RRA}$ with domain $\mathcal{P}(A' \cup \{s\})$ and a flexible atom s such that \mathfrak{D} is the atom structure of \mathbf{A} (see, e.g., Theorem 2.2 and Theorem 2.6 in the article by Maddux (1982)).

Furthermore, given a finite symmetric integral algebra $\mathbf{A} \in \text{RRA}$ with a flexible atom s let \mathfrak{D} be the atom structure of \mathbf{A} . Let $(A'; R')$ be the substructure of the $\{R\}$ -reduct of \mathfrak{D} induced by $D \setminus \{s\}$. By the properties of \mathfrak{D} we get that $(A'; R')$ is a stencil.

We show that the instances of $\text{NCP}(A', R')$ and $\text{NSP}(\mathbf{A})$ are in a natural 1-to-1 correspondence that preserves the acceptance condition of the computational problems. Let $(G; F)$ be a finite undirected graph with an edge $\mathcal{P}(A')$ -coloring f . We define an \mathbf{A} -network $(G; g)$ by defining

$$g(x, y) = \begin{cases} f(x, y) & (x, y) \in F, \\ \{p\} & (x, y) \notin F \text{ and } x = y \\ \{s\} & \text{else.} \end{cases}$$

It is easy to see that $(G; F)$ is an accepted instance of $\text{NCP}(A', R')$ if and only if $(G; g)$ is an accepted instance of $\text{NSP}(\mathbf{A})$. Since we can reverse this and find for every \mathbf{A} -network $(G; g)$ a finite undirected graph $(G; F)$ with an edge $\mathcal{P}(A')$ -coloring f such that each of them is an accepted instance if and only if the other one is. These two reductions show that the computational decision problems $\text{NCP}(A', R')$ and $\text{NSP}(\mathbf{A})$ are polynomial-time equivalent. \square

We end this section by providing a rich source of examples for NCPs.

Example 5.5 (“Distance problems”). *Let $A \subset \mathbb{Q}$ be an arbitrary finite set that contains 0. We define the relation $R \subseteq A^3$ of all tuples which satisfy all instantiations of the triangle inequality, i.e.*

$$(a_1, a_2, a_3) \in R :\Leftrightarrow (a_1 \leq a_2 + a_3) \wedge (a_2 \leq a_1 + a_3) \wedge (a_3 \leq a_1 + a_2),$$

where the addition is meant to be the usual one on rational numbers. The relation R is by definition totally symmetric and the element 0 is an identity element. Therefore, $(A; R)$ is a stencil.

Now consider a finite undirected graph $(G; F)$ with an edge $\mathcal{P}(A)$ -coloring f . This can be seen as a labeling of each edge in the graph by a set of possible (or allowed) distances. The computational task of $\text{NCP}(A, R)$ is to decide whether one can choose for each edge (x, y) one of the possible distance such that in the end this choice satisfies on each triangle of edges the triangle inequalities of metric spaces.

By Proposition 5.4 there exists a finite symmetric integral algebra $\mathbf{A}' \in \text{RRA}$ with a flexible atom s such that $\text{NCP}(A, R)$ and $\text{NSP}(\mathbf{A}')$ are polynomial-time equivalent. By the

proof of Proposition 5.4 we have that the domain of \mathbf{A}' is equal to $\mathcal{P}(A \cup \{s\})$. It is easy to observe that $A \cup \{s\}$ is the set of atoms A'_0 from \mathbf{A}' .

We define an operation $f: A_0^6 \rightarrow A_0$ as follows:

$$f(x_1, \dots, x_6) = \begin{cases} s & s \in \{x_1, \dots, x_6\}, \\ \max\{x_1, \dots, x_6\} & \text{otherwise,} \end{cases}$$

where the max operation is the usual from in \mathbb{Q} .

The allowed triples of \mathbf{A}' are, up to triples that involve the flexible atom s , those which arise from valid triangle inequalities. For this reason the operation f preserves the allowed triples of \mathbf{A}' . Moreover, one can check that f satisfies the Siggers identity. This implies that all these “distance problems” satisfy the first condition in Theorem 1.1 and are therefore solvable by a polynomial-time algorithm. Furthermore, if the normal representation of \mathbf{A}' induces a proper relation algebra \mathbf{A} on the set A , then \mathbf{A} is a representable relation algebra with a normal representation. This follows from Proposition 2.7.4 in the thesis of Conant (2015) (see also Delhommé, Laflamme, Pouzet, & Sauer, 2007), since the composition operation of \mathbf{A} is associative. We get that if the representable relation algebra \mathbf{A} exists then the argument from Example 3.10 implies that $\text{NSP}(\mathbf{A})$ is also in P .

6. Binary Injective Polymorphisms

We give in this section a proof of the following proposition.

Proposition 6.1. *Let \mathfrak{B} be a normal representation of a finite, symmetric, integral $\mathbf{A} \in \text{RRA}$ with a flexible atom s . If $\text{HSP}^{\text{fin}}(\{\text{Pol}(\mathfrak{B})\})$ does not contain a 2-element algebra where all operations are projections, then \mathfrak{B} has a binary injective polymorphism.*

This statement is a consequence of well known results that can be found in the book by Bodirsky (2021) and results by Mottet and Pinsker (2022) applied to the class of normal representations of finite, symmetric, integral $\mathbf{A} \in \text{RRA}$ with a flexible atom s . An operation $f \in \text{Pol}(\mathfrak{B})$ is called *essentially unary* if it depends on at most one of its variables and f is called *essential* otherwise.

Following the terminology of Mottet and Pinsker (2022), we now define *free 2-orbits*. The existence of a free 2-orbit appeared under the name ‘*orbital extension property*’ for example in the habilitation thesis of Bodirsky (2012).

Definition 6.2. *Let \mathfrak{B} be a structure. A 2-orbit O of $\text{Aut}(\mathfrak{B})$ is called free if for all elements $x, y \in B$ there exists $z \in B$ with $(z, x) \in O$ and $(z, y) \in O$.*

Note that if $\text{Aut}(\mathfrak{B})$ has a free 2-orbit then it is transitive. The following theorem generalises a fact that was first proved for first-order reducts of $(\mathbb{Q}; <)$ by Bodirsky and Kára (2009).

Proposition 6.3 (Lemma 5.3.10 in Bodirsky, 2012). *Let \mathfrak{B} be a structure such that \mathfrak{B} has a free 2-orbit. If $\text{Pol}(\mathfrak{B})$ contains an essential operation then it contains a binary essential operation.*

The following is essentially taken from the article by Mottet and Pinsker (2022).

Definition 6.4. Let \mathfrak{B} be a structure. Then the canonical binary structure of \mathfrak{B} is the structure with domain B and a binary relation for each 2-orbit O of $\text{Aut}(\mathfrak{B})$ such that $(x, y) \in O$ implies $x \neq y$.

Definition 6.5. A τ -structure \mathfrak{B} has finite duality if there exists a finite set \mathcal{F} of finite τ -structures such that a τ -structure \mathfrak{J} has a homomorphism to \mathfrak{B} if and only if no element of \mathcal{F} has a homomorphism to \mathfrak{J} .

We establish finite duality for the class of structures which is important for our classification purposes.

Lemma 6.6. Let \mathfrak{B} be a normal representation of a finite, integral $\mathbf{A} \in \text{RRA}$ with a flexible atom s . Then the canonical binary structure of \mathfrak{B} has finite duality.

Proof. Let $\tau := A_0 \setminus \{\text{Id}\}$. Note that since \mathbf{A} is integral and \mathfrak{B} is homogeneous, the canonical binary structure \mathfrak{C} of $\text{Aut}(\mathfrak{B})$ is precisely the τ -reduct of \mathfrak{B} . Let \mathcal{F} be the set of all τ -structures with domain $\{1, 2, 3\}$ that do not have a homomorphism to \mathfrak{C} . We show that the set \mathcal{F} witnesses the finite duality of the canonical binary structure \mathfrak{C} of \mathfrak{B} . Let \mathfrak{J} be a τ -structure with a homomorphism to \mathfrak{C} . If there exists $\mathfrak{F} \in \mathcal{F}$ with a homomorphism to \mathfrak{J} , then \mathfrak{F} also has a homomorphism to \mathfrak{C} , contradicting the choice of \mathcal{F} . For the other direction assume that no element from \mathcal{F} has a homomorphism to \mathfrak{J} . Let \mathfrak{J}' be the τ -expansion of the $(\tau \setminus \{s\})$ -reduct of \mathfrak{J} where the relation $s^{\mathfrak{J}'}$ is defined by

$$s^{\mathfrak{J}'} := \{(x, y) \in I^2 \mid (x, y) \in s^{\mathfrak{J}} \vee (x \neq y \wedge \forall a \in \tau \setminus \{s\}. (x, y) \notin a^{\mathfrak{J}})\}.$$

By the definition of the flexible atom s it follows that no element from \mathcal{F} has a homomorphism to \mathfrak{J}' . This implies that for all distinct elements x, y from \mathfrak{J}' the tuple (x, y) is in at most one relation from τ . The definition of \mathfrak{J}' ensures that (x, y) is in at least one relation from τ . Recall that Proposition 2.10 states that \mathfrak{C} is finitely bounded by τ -structures of size at most three. Assume that one of those bounds \mathfrak{N} embeds into \mathfrak{J}' . This implies by what we noted before that all elements of \mathfrak{N} are in precisely one relation from τ . On the other hand \mathfrak{N} is not in \mathcal{F} and therefore has a homomorphism to \mathfrak{C} . Since all elements of \mathfrak{N} are related by precisely one relation from τ , this homomorphism needs to be an embedding, contradicting our assumption on \mathfrak{N} to be a bound. Therefore, none of the bounds of \mathfrak{C} embeds into \mathfrak{J}' , which means that \mathfrak{J}' is a substructure of \mathfrak{C} . Clearly, there exists a homomorphism from \mathfrak{J} to \mathfrak{J}' which proves the lemma. \square

The following proposition about the existence of injective operations is from Mottet and Pinsker (2022), building on ideas of Bodirsky and Pinsker (2015) and Bodirsky, Martin, Pinsker, and Pongrácz (2019).

Proposition 6.7 (Mottet & Pinsker, 2022). Let \mathfrak{B} be a homogeneous structure such that $\text{Aut}(\mathfrak{B})$ is transitive and such that the canonical binary structure of $\text{Aut}(\mathfrak{B})$ has finite duality. If $\text{Pol}(\mathfrak{B})$ contains a binary essential operation that preserves \neq then it contains a binary injective operation.

We are now able to prove the main result of this section.

Proof of Proposition 6.1. Note that since \mathbf{A} is integral and \mathfrak{B} is homogeneous, the flexible atom s is a free 2-orbit of $\text{Aut}(\mathfrak{B})$. Furthermore, $\text{Aut}(\mathfrak{B})$ is transitive. Suppose that $\text{HSP}^{\text{fin}}(\{\text{Pol}(\mathfrak{B})\})$ does not contain a 2-element algebra where all operations are projections. Since all operations of $\text{Pol}(\mathfrak{B})$ are edge conservative, it follows that $\text{Pol}(\mathfrak{B})$ contains an operation that does not behave as a projection on $\{s, \text{Id}\}$. This implies that $\text{Pol}(\mathfrak{B})$ contains an essential operation. By Proposition 6.3, $\text{Pol}(\mathfrak{B})$ must also contain a binary essential operation. Since the canonical binary structure of \mathfrak{B} has finite duality by Lemma 6.6 we can apply Proposition 6.7 and get that $\text{Pol}(\mathfrak{B})$ contains a binary injective operation. \square

The following shows how to use Proposition 6.1 to obtain a hardness result for the concrete $\mathbf{A} \in \text{RRA}$ from Example 3.9.

Example 6.8 (Hardness of representable relation algebra #17, see Bodirsky et al., 2019; Bodirsky & Knäuer, 2020). *Let \mathfrak{N}' be the normal representation of the algebra #17 mentioned in Example 3.9. We claim that the structure \mathfrak{N}' does not have a binary injective polymorphism. To see this, consider a substructure of \mathfrak{N}'^2 on elements $x, y, z \in V^2$ such that $(E, =)(x, y)$, $(=, E)(y, x)$, and $(E, E)(x, z)$. Assume \mathfrak{N}' has a binary injective polymorphism f . This means that $\bar{f}(E, \text{Id}) = E = \bar{f}(\text{Id}, E)$ holds. Then we get that $E(f(x), f(y))$, $E(f(y), f(z))$, and $E(f(x), f(z))$ hold in \mathfrak{N}' , which is a contradiction, since in \mathfrak{N}' triangles of this form are forbidden. By the contraposition of Proposition 6.1 it follows that $\text{HSP}^{\text{fin}}(\{\text{Pol}(\mathfrak{B})\})$ contains a 2-element algebra where all operations are projections. We conclude with Theorem 2.21 that $\text{NSP}(\#17)$ is an NP-hard problem.*

7. From Partial to Total Canonical Behaviour

In this section we prove that in many cases the existence of a polymorphism with a certain partial behaviour implies the existence of a canonical polymorphism with the same partial behaviour. Following this idea we start in Section 7.1 with the proof that the existence of an injective polymorphism implies the existence of a canonical injective polymorphism. In some cases the existence of an $\{a, b\}$ -canonical polymorphism implies the existence of a canonical polymorphism with the same behaviour on $\{a, b\}$. We prove this separately for binary (Section 7.2) and ternary (Section 7.3) operations, making use of the binary injective polymorphism that exists by the results from Section 6 and Section 7.1.

Let us remark that most proofs of this section would fail if the representable relation algebra \mathbf{A} was not symmetric. Indeed, every representable relation algebra that contains a non-symmetric atom a normal representation would not satisfy Proposition 7.2 below, since the stated behaviour is not well-defined.

We assume for this section that \mathfrak{B} is a normal representation of a finite, symmetric, integral $\mathbf{A} \in \text{RRA}$ with a flexible atom s . Let furthermore $\mathfrak{B}_{<}$ be the expansion of \mathfrak{B} by the generic linear order. The structure $\mathfrak{B}_{<}$ exists by the observations in Section 3.2.

Proposition 7.1. *Let $f \in \text{Pol}(\mathfrak{B})^{(n)}$ be injective. Then there exists a polymorphism $f_{<}$ of $\mathfrak{B}_{<}$ and an injective endomorphism e of \mathfrak{B} such that*

$$f = e \circ f_{<}$$

as mappings from B^n to B .

Proof. Let $U := f(B^n)$ and consider the substructure \mathfrak{U} induced by \mathfrak{B} on U . There exists a linear ordering on B^n , namely the lexicographic order given by the linear order of $\mathfrak{B}_<$ on each coordinate.

Let $\mathfrak{U}_<$ be the expansion of \mathfrak{U} by the linear order that is induced by the lexicographic linear order of $\mathfrak{B}_<$ on the preimage. This is well defined since f is injective. By the definition of $\mathfrak{B}_<$ and a compactness argument the structure $\mathfrak{U}_<$ embeds into $\mathfrak{B}_<$. In this way we obtain a homomorphism $f_<$ from $\mathfrak{B}_<^n$ to $\mathfrak{B}_<$. Again by a compactness argument also an endomorphism e with the desired properties exists. \square

7.1 Canonical Binary Injective Polymorphisms

We prove in this section that the existence of an injective polymorphism implies the existence of a canonical injective polymorphism. We say that a polymorphism f of $\mathfrak{B}_<$ is *canonical with respect to $\mathfrak{B}_<$* if f satisfies Definition 2.24, where the underlying representable relation algebra is the proper relation algebra induced by the binary first-order definable relations (i.e., unions of 2-orbits) of $\mathfrak{B}_<$. Note that in a normal representation the set of 2-orbits equals the set of the interpretations of the atoms of the representable relation algebra.

Proposition 7.2. *Let f be a binary polymorphism of $\mathfrak{B}_<$ that is canonical with respect to $\mathfrak{B}_<$. Let $h: A_0^2 \rightarrow A_0$ be the map such that for all $x, y, z \in A_0$*

$$h(x, y) = z \Leftrightarrow \bar{f}(x^{\mathfrak{B}_< \cap \leq^{\mathfrak{B}_<}}, y^{\mathfrak{B}_< \cap \leq^{\mathfrak{B}_<}}) = z^{\mathfrak{B}_< \cap \leq^{\mathfrak{B}_<}},$$

(cf. Remark 3.7). Then h is well defined and there exists a canonical binary polymorphism of \mathfrak{B} with behaviour h .

Proof. The function h is well defined since all atoms are symmetric. We show that there exists a canonical polymorphism of \mathfrak{B} that has h as a behaviour. Consider the following structure \mathfrak{A} on the domain B^2 . Let $x, y \in B^2$ and let $a, a_1, a_2 \in A_0$ be atoms of \mathbf{A} with $a_1^{\mathfrak{B}}(x_1, y_1)$ and $a_2^{\mathfrak{B}}(x_2, y_2)$. Then we define that $a^{\mathfrak{A}}(x, y)$ holds if and only if $h(a_1, a_2) = a$.

We show in the following that \mathfrak{A} has a homomorphism to \mathfrak{B} . This is enough to prove the statement, because a homomorphism from \mathfrak{A} to \mathfrak{B} is a canonical polymorphism of \mathfrak{B} . Since \mathfrak{B} is homogeneous it suffices to show that every finite substructure of \mathfrak{A} homomorphically maps to \mathfrak{B} . Let \mathfrak{F} be a finite substructure of \mathfrak{A} and assume for contradiction that \mathfrak{F} does not homomorphically map to \mathfrak{B} . We can view \mathfrak{F} as an atomic \mathbf{A} -network. Since \mathfrak{B} is fully universal \mathfrak{F} is not closed. There must exist elements $b^1, b^2, b^3 \in B^2$ of \mathfrak{F} and atoms $a_1, a_2, a_3 \in A_0$ such that $a_1 \not\leq a_2 \circ a_3$ holds in \mathbf{A} and

$$a_1^{\mathfrak{F}}(b^1, b^3), a_2^{\mathfrak{F}}(b^1, b^2), \text{ and } a_3^{\mathfrak{F}}(b^2, b^3).$$

This means that the substructure induced on the elements b^1, b^2, b^3 by \mathfrak{F} contains a forbidden triple.

Now we consider the substructures that are induced on b_1^1, b_1^2, b_1^3 and b_2^1, b_2^2, b_2^3 by \mathfrak{B} . Our goal is to order these elements such that for all $i, j \in \{1, 2, 3\}$

$$\neg(b_1^i < b_1^j \wedge b_2^i > b_2^j). \tag{1}$$

If we achieve this we know that there exist elements in $\mathfrak{B}_<$ that induce isomorphic copies of the induced structures of the elements b_1^1, b_1^2, b_1^3 and b_2^1, b_2^2, b_2^3 with the additional ordering.

Now the application of the polymorphism f on these elements results in a structure whose A_0 -reduct is isomorphic to the substructure induced by b^1, b^2 and b^3 on \mathfrak{F} by the definition of the canonical behaviour h . This contradicts our assumption because a polymorphism can not have a forbidden substructure in its image.

It remains to show that we can choose orderings on the elements b_1^1, b_1^2, b_1^3 and b_2^1, b_2^2, b_2^3 such that (1) holds. Without loss of generality we can assume that $\{b_1^1, b_1^2, b_1^3\} \cap \{b_2^1, b_2^2, b_2^3\} = \emptyset$ holds. Now consider the following cases:

1. $|\{b_1^1, b_1^2, b_1^3\}| = 3$ and $|\{b_2^1, b_2^2, b_2^3\}| = 3$.

We can obviously choose linear orders on both sets such that (1) holds.

2. $|\{b_1^1, b_1^2, b_1^3\}| = 2$ and $|\{b_2^1, b_2^2, b_2^3\}| = 3$.

Assume that $\text{Id}^{\mathfrak{B}}(b_1^1, b_1^2)$ holds then the possible orders are

$$b_1^1 = b_1^2 < b_1^3 \text{ and } b_2^1 < b_2^2 < b_2^3.$$

3. $|\{b_1^1, b_1^2, b_1^3\}| = 2$ and $|\{b_2^1, b_2^2, b_2^3\}| = 2$.

First consider the case that $\text{Id}^{\mathfrak{B}}(b_1^1, b_1^2)$ and $\text{Id}^{\mathfrak{B}}(b_2^1, b_2^2)$ hold. Then we choose as orders

$$b_1^1 = b_1^2 < b_1^3 \text{ and } b_2^1 = b_2^2 < b_2^3.$$

In the second possible case we can assume without loss of generality that $\text{Id}^{\mathfrak{B}}(b_1^1, b_1^2)$ and $\text{Id}^{\mathfrak{B}}(b_2^2, b_2^3)$ hold. Note that otherwise we could change the role of two of the tuples b^1, b^2 and b^3 and get this case. The compatible order is then

$$b_1^1 = b_1^2 < b_1^3 \text{ and } b_2^1 < b_2^2 = b_2^3.$$

4. $|\{b_1^1, b_1^2, b_1^3\}| = 1$ and $|\{b_2^1, b_2^2, b_2^3\}| = 3$.

In this case we choose the order

$$b_1^1 = b_1^2 = b_1^3 \text{ and } b_2^1 < b_2^2 < b_2^3.$$

5. $|\{b_1^1, b_1^2, b_1^3\}| = 1$ and $|\{b_2^1, b_2^2, b_2^3\}| = 2$.

Assume that $\text{Id}^{\mathfrak{B}}(b_2^1, b_2^2)$ holds and that we have

$$b_1^1 = b_1^2 = b_1^3 \text{ and } b_2^1 = b_2^2 < b_2^3.$$

6. $|\{b_1^1, b_1^2, b_1^3\}| = 1$ and $|\{b_2^1, b_2^2, b_2^3\}| = 1$.

For this case we trivially get

$$b_1^1 = b_1^2 = b_1^3 \text{ and } b_2^1 = b_2^2 = b_2^3.$$

Note that up to the symmetry of the arguments for both coordinates these are all the possible cases. This completes the proof of the proposition. \square

Corollary 7.3. *Suppose that \mathfrak{B} has a binary injective polymorphism. Then \mathfrak{B} also has a canonical binary injective polymorphism.*

Proof. By Proposition 7.1 we may assume that there exists also an injective polymorphism of $\mathfrak{B}_<$. The structure $\mathfrak{B}_<$ has the Ramsey property by Theorem 3.6. Therefore, Theorem 2.26 implies that there also exists an injective canonical polymorphism g of $\mathfrak{B}_<$. According to Proposition 7.2 the restriction of the behaviour \bar{g} to the 2-orbits that satisfy $x \leq y$ induces the behaviour of a canonical polymorphism of \mathfrak{B} which is also injective. \square

7.2 Canonical $\{a, b\}$ -symmetric Polymorphisms

We will now use the results about binary injective polymorphism from Section 7.1 to show the existence of a canonical $\{a, b\}$ -symmetric polymorphism in case there exists an $\{a, b\}$ -symmetric polymorphism.

Lemma 7.4. *Let $a, b \in A_0 \setminus \{\text{Id}\}$ be atoms. Then every binary $\{a, b\}$ -symmetric polymorphism of \mathfrak{B} is injective.*

Proof. Let f be an $\{a, b\}$ -symmetric polymorphism. Without loss of generality $\bar{f}(a, b) = a = \bar{f}(b, a)$. Assume for contradiction that f is not injective. This means that there exist $c \in A_0$ and $x, y \in B^2$ with $(c, \text{Id})(x, y)$ (for the notation see Definition 2.22) such that $\text{Id}(f(x), f(y))$ holds.

Case 1: $s \notin \{a, b\}$. Since s is a flexible atom we may choose $z \in B^2$ such that $(a, b)(z, x)$ and $(s, b)(z, y)$ hold. By the choice of the polymorphism f we get $a(f(z), f(x))$ and $(s \cup b)(f(z), f(y))$ which induces either the forbidden triple (Id, s, a) or the forbidden triple (Id, b, a) on $f(x), f(y)$, and $f(z)$.

Case 2: $s = a$. We choose $z \in B^2$ such that $(a, b)(z, x)$ and $(b, b)(z, y)$. This is possible since a is the flexible atom. We obtain $a(f(z), f(x))$ and $b(f(z), f(y))$ which again induces a forbidden triple on $f(x), f(y)$, and $f(z)$.

Case 3: $s = b$. We choose $z \in B^2$ such that $(a, b)(z, x)$ and $(b, b)(z, y)$. This is possible since a is the flexible atom. We obtain $a(f(z), f(x))$ and $b(f(z), f(y))$ which again induces a forbidden triple on $f(x), f(y)$, and $f(z)$.

Since we obtained in all cases a contradiction we conclude that f is injective. \square

Proposition 7.5. *Let $a, b \in A_0 \setminus \{\text{Id}\}$ be atoms. If \mathfrak{B} has a binary $\{a, b\}$ -symmetric polymorphism, then \mathfrak{B} has also a binary canonical $\{a, b\}$ -symmetric polymorphism.*

Proof. Let f be the binary $\{a, b\}$ -symmetric polymorphism. By Lemma 7.4 we know that f is injective. By Proposition 7.1 it induces a polymorphism $f_<$ on $\mathfrak{B}_<$. The structure $\mathfrak{B}_<$ has the Ramsey property by Theorem 3.6. Let g be the canonization of $f_<$ that exists by Theorem 2.26. The restriction of the behaviour \bar{g} to the 2-orbits that satisfy $x \leq y$ induces by Proposition 7.2 the behaviour of a canonical polymorphism h of \mathfrak{B} . The way we obtained h ensures that h is $\{a, b\}$ -symmetric with the same behaviour on $\{a, b\}$ as f . \square

The following is an easy observation about $\{a, \text{Id}\}$ -symmetric polymorphisms that we will use several times.

Observation 7.6. *Let $a \not\leq \text{Id}$ be an atom and f an $\{a, \text{Id}\}$ -symmetric polymorphism of \mathfrak{B} . Then $\bar{f}(a, \text{Id}) = a = \bar{f}(\text{Id}, a)$.*

Proof. Suppose for contradiction that $\bar{f}(a, \text{Id}) = \text{Id} = \bar{f}(\text{Id}, a)$. Let $x, y, z \in B^2$ be such that

$$(a, \text{Id})(x, y), (\text{Id}, a)(y, z), \text{ and } (a, a)(x, z)$$

and consider the substructure of \mathfrak{B} that is induced by $f(x), f(y)$ and $f(z)$. This structure induces a forbidden triple $(\text{Id}, \text{Id}, a)$ which contradicts the assumption $\mathbf{A} \in \text{RRA}$. \square

In the remainder of the section we combine canonical $\{a, b\}$ -symmetric polymorphisms to obtain a single “maximal-symmetric” polymorphism.

Definition 7.7. We call a subset $\{a, b\} \subseteq A_0$ an edge of $\text{Pol}^{\text{can}}(\mathfrak{B})$ and we call the elements in

$$Q := \{\{a, b\} \subseteq A_0 \mid \exists g \in \text{Pol}^{\text{can}}(\mathfrak{B}) \text{ such that } \bar{g} \text{ is symmetric on } \{a, b\}\}$$

the red edges of $\text{Pol}^{\text{can}}(\mathfrak{B})$.

The terminology of the colored edges as well as the following lemma are from Bulatov (2016).

Lemma 7.8. *There exists a binary canonical polymorphism that is symmetric on all red edges and behaves on each non-red edge like a projection. We call this function maximal-symmetric.*

Proof. For each $\{a, b\} \in Q$ let $f_{a,b}$ be a canonical polymorphism such that its behaviour is symmetric on $\{a, b\}$. We prove the lemma by an induction on the size of subsets of Q , i.e., we show that for every subset $F \subseteq Q$ of size n there exists a polymorphism $f_F \in \text{Pol}^{\text{can}}(\mathfrak{B})$ that is symmetric on all edges from F . For each subset $\{a\}$ of Q of size one, there exists by the definition of red edges a canonical polymorphism $f_{a,a}$ with a behaviour that is symmetric on $\{a\}$. Let $F \subseteq Q$ and suppose there exists a canonical polymorphism g with symmetric behaviour on elements from F . Let $\{a_1, a_2\} \in Q \setminus F$. We want to show that there exists a canonical polymorphism with a behaviour that is symmetric on all elements from $F \cup \{a_1, a_2\}$. We may assume that this does not hold for g , otherwise we are done. Therefore, and since g is edge-conservative, we have

$$\bar{g}(a_1, a_2) \neq \bar{g}(a_2, a_1) \text{ and } \bar{g}(a_1, a_2), \bar{g}(a_2, a_1) \in \{a_1, a_2\}.$$

With this it is easy to see that $f_{a_1, a_2}(g(x, y), g(y, x))$ is a polymorphism with a behaviour that is symmetric on all elements from $F \cup \{a_1, a_2\}$.

This proves the first part of the statement. For the second part note that for a binary canonical edge-conservative polymorphism there are only 4 possibilities for the behaviour on a set $\{a_1, a_2\}$. If $\{a_1, a_2\}$ is not a red edge then every binary canonical edge-conservative polymorphism behaves like a projection on $\{a_1, a_2\}$. \square

7.3 Canonical Ternary Polymorphisms

We obtain in this section a result that states that the existence of a ternary $\{a, b\}$ -canonical polymorphism f implies the existence of a canonical polymorphism with the same behaviour on $\{a, b\}$ as f (Corollary 7.11). This is done similarly as in Section 7.2.

Lemma 7.9. *Let $s' \in \text{Pol}^{\text{can}}(\mathfrak{B})$ be a binary, injective, maximal-symmetric polymorphism. Then the function $s^*: \mathfrak{B}^3 \rightarrow \mathfrak{B}^3$ where $s^*(x_1, x_2, x_3)$ is defined by*

$$(s'(s'(x_1, x_2), s'(x_2, x_3)), s'(s'(x_2, x_3), s'(x_3, x_1)), s'(s'(x_3, x_1), s'(x_1, x_2)))$$

is a homomorphism. Moreover, for all $x, y \in B^3$ with $x \neq y$ it holds that

$$\overline{\text{Id}}^{\mathfrak{B}^3}(s^*(x), s^*(y)).$$

Note that this means that two distinct tuples in the image of s^* have distinct entries in each coordinate.

Proof. Let $x, y \in B^3$ and suppose that $a^{\mathfrak{B}^3}(x, y)$ holds for $a \in A$. By the definition of the product structure $a^{\mathfrak{B}}(x_i, y_i)$ holds for all $i \in \{1, 2, 3\}$. Since s' is a polymorphism clearly $a^{\mathfrak{B}}(s^*(x)_i, s^*(y)_i)$ holds by the definition of s^* . Now we use again the definition of a product structure and get $a^{\mathfrak{B}^3}(s^*(x), s^*(y))$ which shows that s^* is a homomorphism.

For the second part of the statement let $x, y \in B^3$ distinct. Suppose that $a_1^{\mathfrak{B}}(x_1, y_1)$, $a_2^{\mathfrak{B}}(x_2, y_2)$ and $a_3^{\mathfrak{B}}(x_3, y_3)$ hold for some $a_1, a_2, a_3 \in A_0$, where at least one atom is different from Id . Since s' is injective we have

$$\overline{\text{Id}}^{\mathfrak{B}}(s'(x_1, x_2), s'(y_1, y_2)) \text{ or } \overline{\text{Id}}^{\mathfrak{B}}(s'(x_2, x_3), s'(y_2, y_3)).$$

Again by the injectivity of s' we get

$$\overline{\text{Id}}^{\mathfrak{B}}(s'(s'(x_1, x_2), s'(x_2, x_3)), s'(s'(y_1, y_2), s'(y_2, y_3))).$$

By the definition of s^* this shows that $\overline{\text{Id}}^{\mathfrak{B}^3}(s^*(x)_1, s^*(y)_1)$ holds. It is easy to see by analogous arguments that the same is true for the other coordinates. Therefore, the statement follows. \square

Proposition 7.10. *Let $a \not\leq \text{Id}$ and $b \not\leq \text{Id}$ be atoms of \mathbf{A} such that $\{a, b\} \notin Q$. Let $m \in \text{Pol}(\mathfrak{B})$ be ternary $\{a, b\}$ -canonical and $s' \in \text{Pol}(\mathfrak{B})$ be injective, maximal-symmetric. Then there exists a canonical $m' \in \text{Pol}(\mathfrak{B})$ with the same behaviour as m on $\{a, b\}$.*

Proof. By Lemma 7.8 we may assume that s' behaves on $\{a, b\}$ like the projection to the first coordinate since $\{a, b\} \notin Q$. Let s^* be the function defined in Lemma 7.9 and consider the function $m': B^3 \rightarrow B$ which is defined by $m^*(x) := m(s^*(x))$.

Claim 1: m^* is injective. Let $x, y \in B^3$ be two distinct elements. By Lemma 7.9 we know that $\overline{\text{Id}}^{\mathfrak{B}^3}(s^*(x), s^*(y))$ holds. Since m is a polymorphism of \mathfrak{B} we directly get that $\overline{\text{Id}}^{\mathfrak{B}}(m^*(x), m^*(y))$ holds, which proves the injectivity of m^* .

Claim 2: m^* is $\{a, b\}$ -canonical and behaves on $\{a, b\}$ like m .

Let $x, y \in B^3$ with $(q_1, q_2, q_3)(x, y)$ such $q_1, q_2, q_3 \in \{a, b\}$. Since s behaves like the first projection on $\{a, b\}$ it follows that $(q_1, q_2, q_3)(s^*(x), s^*(y))$. Together with the $\{a, b\}$ -canonicity of m this proves Claim 2.

Since m^* is injective there exists by Proposition 7.1 a polymorphism $m^*_<$ of $\mathfrak{B}_<$. Since $\mathfrak{B}_<$ is a Ramsey structure we can apply Theorem 2.26 to $m^*_<$. Let g be the resulting

polymorphism that is canonical with respect to $\mathfrak{B}_<$. Note that if we consider g as a polymorphism of \mathfrak{B} it behaves on $\{a, b\}$ like m^* and therefore like m . Now we consider the induced behaviour of g on all 2-orbits that satisfy $x < y$. Since all atoms of \mathbf{A} are symmetric and \bar{g} is conservative this induces a function $h: (A_0 \setminus \{\text{Id}\})^3 \rightarrow A_0 \setminus \{\text{Id}\}$.

Claim 3: The partial behaviour h does not induce a forbidden triple.

We would like to note that we prove this claim with similar arguments as in the proof of Proposition 7.2. Assume for contradiction that there exist $x, y, z \in B^3$ such that the application of an operation with behaviour h would induce a forbidden triple. Without loss of generality we can order the elements of each coordinate of x, y, z strictly with $x_i < y_i < z_i$ for $i \in \{1, 2, 3\}$. Note that if on some coordinate there would be the relation Id then we are out of the domain of the behaviour h .

If we choose such an order we can find isomorphic copies \mathfrak{A} of this structure (with the order) in $\mathfrak{B}_<$. If we apply the polymorphism g to this copy and forget the order of the structure $g(\mathfrak{A})$ we get a structure that is by definition isomorphic to the forbidden triple. This proves Claim 3.

To finish the proof of the lemma note that the composition of s^* with the projection to the i -th coordinate for $i \in \{1, 2, 3\}$ is a canonical, injective polymorphism (for injectivity see Lemma 7.9) and therefore induces a behaviour $f_i: A_0^3 \rightarrow A_0 \setminus \{\text{Id}\}$. We define $f: A_0^3 \rightarrow (A_0 \setminus \{\text{Id}\})^3$ by $f(a_1, a_2, a_3) := (f_1(a_1), f_2(a_2), f_3(a_3))$. The composition $h \circ f: A_0^3 \rightarrow A_0$ is the behaviour of a canonical function of \mathfrak{B} . If $h \circ f$ would induce a forbidden triple then also h would induce a forbidden triple, which contradicts Claim 3. \square

Corollary 7.11. *Let \mathfrak{B} have a binary injective polymorphism. Let $a, b \in A_0$ be such that no $\{a, b\}$ -symmetric polymorphism exists. Let m be a ternary $\{a, b\}$ -canonical polymorphism. Then there exists a canonical polymorphism m' with the same behaviour on $\{a, b\}$ as m .*

Proof. By assumption there is no canonical $\{a, b\}$ -symmetric polymorphism and therefore $\{a, b\} \notin Q$. By Corollary 7.3 \mathfrak{B} has a canonical binary injective polymorphism. This polymorphism is a witness that $\{c, \text{Id}\} \in Q$ for all $c \in A_0 \setminus \{\text{Id}\}$. With Lemma 7.8 we get a maximal symmetric polymorphism h . Since $\{c, \text{Id}\} \in Q$ we get that h is $\{c, \text{Id}\}$ -symmetric for all $c \in A_0$. By Observation 7.6 it follows $\bar{h}(c, \text{Id}) = c = \bar{h}(\text{Id}, c)$, which implies that h is injective. Now Proposition 7.10 implies the statement. \square

8. The Independence Lemma and How To Use It

The central result of this section is Proposition 8.2 which states that the absence of an $\{a, b\}$ -symmetric polymorphism implies that all polymorphism are canonical on $\{a, b\}$. The main ingredients of our proof of this proposition are the fact that $\mathbf{A} \in \text{RRA}$ has a flexible atom and the following “Independence Lemma” (Lemma 8.1).

8.1 The Independence Lemma

The following lemma transfers the absence of a special partially canonical polymorphism to the existence of certain relations of arity 4 that are primitively positively definable in \mathfrak{B} . A lemma of a similar type appeared as Lemma 42 in an article by Bodirsky and Pinsker (2014).

Lemma 8.1 (Independence Lemma). *Let \mathfrak{B} be a homogeneous structure with finite relational signature. Let a and b be 2-orbits of $\text{Aut}(\mathfrak{B})$ such that a , b , and $(a \cup b)$ are primitively positively definable in \mathfrak{B} . Then the following are equivalent:*

1. \mathfrak{B} has an $\{a, b\}$ -canonical polymorphism g that is $\{a, b\}$ -symmetric with $\bar{g}(a, b) = \bar{g}(b, a) = a$.
2. For every primitive positive formula φ such that $\varphi \wedge a(x_1, x_2) \wedge b(y_1, y_2)$ and $\varphi \wedge b(x_1, x_2) \wedge a(y_1, y_2)$ are satisfiable over \mathfrak{B} , the formula $\varphi \wedge a(x_1, x_2) \wedge a(y_1, y_2)$ is also satisfiable over \mathfrak{B} .
3. For every finite $F \subset B^2$ there exists a homomorphism h_F from the substructure of \mathfrak{B}^2 induced by F to \mathfrak{B} that is $\{a, b\}$ -canonical with $\overline{h_F}(a, b) = \overline{h_F}(b, a) = a$.

Proof. The implication from (1) to (2) follows directly by applying the symmetric polymorphisms to tuples from the relation defined by φ .

For the implication from (2) to (3) let F be a finite subset of B^2 . Let $\{e_1, \dots, e_n\}$ with $n \in \mathbb{N}$ be an enumeration of F . To construct h_F consider the formula φ_0 with variables $x_{i,j}$ for $1 \leq i, j \leq n$ that is the conjunction of all atomic formulas $R(x_{i_1, j_1}, \dots, x_{i_k, j_k})$ such that $R(e_{i_1}, \dots, e_{i_k})$ and $R(e_{j_1}, \dots, e_{j_k})$ hold in \mathfrak{B} . Note that this formula states exactly which relations hold on F in \mathfrak{B}^2 . Let P be the set of pairs $((i_1, i_2), (j_1, j_2))$ such that

$$\begin{aligned} & (a \cup b)(e_{i_1}, e_{i_2}) \\ & \text{and } (a \cup b)(e_{j_1}, e_{j_2}) \\ & \text{and } (a(e_{i_1}, e_{i_2}) \vee a(e_{j_1}, e_{j_2})) \\ & \text{and } (b(e_{i_1}, e_{i_2}) \vee b(e_{j_1}, e_{j_2})). \end{aligned}$$

If we show that the formula

$$\psi := \varphi_0 \wedge \bigwedge_{((i_1, i_2), (j_1, j_2)) \in P} a(x_{i_1, j_1}, x_{i_2, j_2})$$

is satisfiable by an assignment α , we get the desired homomorphism by setting $h_F(e_i, e_j) := \alpha(x_{i,j})$. We prove the satisfiability of ψ by induction over the size of subsets I of P . For the inductive beginning consider an element $((i_1, i_2), (j_1, j_2)) \in P$. Without loss of generality we have that $a(i_1, i_2)$ holds. Therefore the assignment $\alpha(x_{i,j}) := e_i$ witnesses the satisfiability of the formula $\varphi_0 \wedge a(x_{i_1, j_1}, x_{i_2, j_2})$. For the inductive step let $I \subseteq P$ be of size m and assume that the statement is true for subsets of size $m - 1$. Let $p_1 = ((u_1, u_2), (v_1, v_2))$ and $p_2 = ((u'_1, u'_2), (v'_1, v'_2))$ be two elements from I . We define the following formula

$$\psi_0 := \varphi_0 \wedge \bigwedge_{((i_1, i_2), (j_1, j_2)) \in I \setminus \{p_1, p_2\}} a(x_{i_1, j_1}, x_{i_2, j_2}).$$

Then by the inductive assumption the formulas $\psi_0 \wedge a(x_{u_1, v_1}, x_{u_2, v_2})$ and $\psi_0 \wedge a(x_{u'_1, v'_1}, x_{u'_2, v'_2})$ are satisfiable. The assumptions on the elements in P give us that also

$$\psi_0 \wedge a(x_{u_1, v_1}, x_{u_2, v_2}) \wedge b(x_{u'_1, v'_1}, x_{u'_2, v'_2})$$

and

$$\psi_0 \wedge b(x_{u_1, v_1}, x_{u_2, v_2}) \wedge a(x_{u'_1, v'_1}, x_{u'_2, v'_2})$$

are satisfiable; since $a \cup b$ is a primitive positive definable relation we are done otherwise. But then we can apply the assumption of (2) and get that also $\psi_0 \wedge a(x_{u_1, v_1}, x_{u_2, v_2}) \wedge a(x_{u'_1, v'_1}, x_{u'_2, v'_2})$ is satisfiable, which proves the inductive step.

The direction from (3) to (1) is a standard application of König's tree lemma. For a reference see for example Lemma 42 in the article by Bodirsky and Pinsker (2014). \square

8.2 Absence of $\{a, b\}$ -symmetric Polymorphisms

We are now able to prove the main result of this section, which will be a corner stone in the proof of Theorem 1.1. Our proof of this proposition makes use of a 4-ary relation $E_{a,b}$ with the following first-order definition:

$$(x_1, x_2, x_3, x_4) \in E_{a,b} :\Leftrightarrow ((a \cup b)(x_1, x_2) \wedge (a \cup b)(x_3, x_4) \wedge a(x_1, x_2) \leftrightarrow a(x_3, x_4)).$$

It is an easy observation that the $\{a, b\}$ -canonical polymorphisms of \mathfrak{B} are precisely those that preserve the relation $E_{a,b}$. By Theorem 2.19 we get that whenever $E_{a,b}$ is primitively positively definable in \mathfrak{B} then all polymorphisms of \mathfrak{B} preserve $E_{a,b}$ and are therefore $\{a, b\}$ -canonical. In the following proof we use the 4-ary relations that are provided by the second item of the Independence Lemma 8.1 to provide a primitive positive definition of $E_{a,b}$.

Proposition 8.2. *Let \mathfrak{B} be a normal representation of a finite integral symmetric relation algebra with a flexible atom s . Suppose that \mathfrak{B} has a binary injective polymorphism. Let $a \not\leq \text{Id}$ and $b \not\leq \text{Id}$ be two atoms such that \mathfrak{B} has no $\{a, b\}$ -symmetric polymorphism. Then all polymorphisms are canonical on $\{a, b\}$.*

Proof. By Corollary 7.3 there exists a canonical binary injective polymorphism of \mathfrak{B} . Therefore, for every $a' \in A_0$ the edge $\{a', \text{Id}\}$ is red and the maximal symmetric polymorphism t that exists by Lemma 7.8 is symmetric on all these edges. Observation 7.6 implies that t is injective. Note that t behaves like a projection on $\{a, b\}$ since there exists no $\{a, b\}$ -symmetric polymorphism.

Let ψ be the formula defined as follows:

$$\psi(x_1, x_2, y_1, y_2) := \overline{\text{Id}}(x_1, y_1) \wedge \overline{\text{Id}}(x_1, y_2) \wedge \overline{\text{Id}}(x_2, y_1) \wedge \overline{\text{Id}}(x_2, y_2).$$

We use ψ to formulate and prove the following claim:

Claim 1: a) There exists a formula $\varphi_a(x_1, x_2, y_1, y_2)$ such that

$$\begin{aligned} \varphi_a \wedge \psi(x_1, x_2, y_1, y_2) \wedge a(x_1, x_2) \wedge b(y_1, y_2) & \text{ is satisfiable in } \mathfrak{B}, \\ \varphi_a \wedge \psi(x_1, x_2, y_1, y_2) \wedge b(x_1, x_2) \wedge a(y_1, y_2) & \text{ is satisfiable in } \mathfrak{B}, \\ \varphi_a \wedge a(x_1, x_2) \wedge a(y_1, y_2) & \text{ is not satisfiable in } \mathfrak{B}. \end{aligned}$$

b) There exists a formula $\varphi_b(x_1, x_2, y_1, y_2)$ that has the same property with a and b in exchanged roles.

Proof of Claim 1. There exists a formula φ_a'' that witnesses the negation of (2) in the Independence Lemma (Lemma 8.1) since \mathfrak{B} does not have an $\{a, b\}$ -symmetric polymorphism h




Figure 3: Application of the injective operation t on tuples u and v . Orange and black edges correspond to atoms c and d , the dotted edge to atom Id , and the dashed lines denote $\overline{\text{Id}}$.

with $\bar{h}(a, b) = a$. Let φ''_b be the formula that witnesses in the same way the non-existence of an $\{a, b\}$ -symmetric polymorphism h of \mathfrak{B} with $\bar{h}(a, b) = b$. We define for $c \in \{a, b\}$ the formula φ'_c as follows:

$$\varphi'_c(x_1, x_2, y_1, y_2) := \varphi''_c(x_1, x_2, y_1, y_2) \wedge (a \cup b)(x_1, x_2) \wedge (a \cup b)(y_1, y_2). \quad (2)$$

If φ'_a and φ'_b witness a) and b) in Claim 1, we are done. So suppose that they do not. Note that if we have $c \in \{a, b\}$ and $d \in \{a, b\} \setminus \{c\}$ such that

$$\varphi'_c(x_1, x_2, y_1, y_2) \wedge c(x_1, x_2) \wedge d(y_1, y_2) \wedge \text{Id}(x_2, y_1) \text{ is satisfiable in } \mathfrak{B} \text{ and} \quad (3)$$

$$\varphi'_c(x_1, x_2, y_1, y_2) \wedge d(x_1, x_2) \wedge c(y_1, y_2) \wedge \overline{\text{Id}}(x_2, y_1) \text{ is satisfiable in } \mathfrak{B}, \quad (4)$$

then φ'_c would satisfy the statement about φ_c in Claim 1, a) or in Claim 1, b). To see this note that we can apply the injective, maximal symmetric polymorphism t that behaves like a projection on $\{a, b\}$ to the tuples u and v that witness (3) and (4). The first tuple satisfies $\overline{\text{Id}}(x_1, y_1)$, $\overline{\text{Id}}(x_2, y_2)$ and $\overline{\text{Id}}(x_1, y_2)$ since $\text{Id}(x_2, y_1)$ and $c \neq d$. The second tuple satisfies $\overline{\text{Id}}(x_2, y_1)$. Then the injectivity of t ensures that the tuples $t(u, v)$ and $t(v, u)$ witness Claim 1, a) or Claim 1, b). Figure 3 illustrates this situation.

By our assumption that Claim 1 is not satisfied by φ'_a and φ'_b we conclude that for at least one $c \in \{a, b\}$ it holds that for $d \in \{a, b\} \setminus \{c\}$

$$\varphi'_c(x_1, x_2, y_1, y_2) \wedge c(x_1, x_2) \wedge d(y_1, y_2) \wedge \text{Id}(x_2, y_1) \text{ is satisfiable in } \mathfrak{B} \text{ and} \quad (5)$$

$$\varphi'_c(x_1, x_2, y_1, y_2) \wedge d(x_1, x_2) \wedge c(y_1, y_2) \wedge \text{Id}(x_2, y_1) \text{ is satisfiable in } \mathfrak{B}. \quad (6)$$

We distinguish the following different cases.

1. φ'_a satisfies a) in Claim 1 and (5) and (6) hold for $c = b$ and $d = a$.
2. (5) and (6) hold for $c = a$ and $d = b$ and φ'_b satisfies b) in Claim 1.
3. (5) and (6) hold for $c = a$ and $d = b$ as well as for $c = b$ and $d = a$.




Figure 4: φ_b build from φ'_b (red), φ'_a (blue), the atom b (black), the atom a (dashed) and the flexible atom s (dotted). The roles of a and b can be changed.

Case 1: Consider the following formula φ_b with

$$\varphi_b(x_1, x_2, y_1, y_2) := \exists z_1, z_2. (\varphi'_b(x_1, x_2, x_2, z_1) \wedge \varphi'_a(x_2, z_1, z_2, y_1) \wedge \varphi'_b(z_2, y_1, y_1, y_2)).$$

We claim that φ_b satisfies b) in Claim 1. This proves Claim 1, because φ'_a satisfies a) in Claim 1. To see that φ_b fulfills the two satisfiability statements in Claim 1, b) we can first amalgamate the structure \mathfrak{B}_1 induced (as a substructure of \mathfrak{B}) by the elements of a satisfying tuple for φ'_b with the structure \mathfrak{B}_2 induced by the elements of a satisfying tuple for φ'_a (see Figure 4 for an illustration). We amalgamate these two structures over their common substructure \mathfrak{A} induced by the variables x_2 and z_1 , with the variable names from the definition of φ_b . In this amalgamation step all missing edges are filled with the flexible atom s . In a second amalgamation step we amalgamate the resulting structure with another copy of the structure \mathfrak{B}_1 , but now with the common substructure on the variables z_1 and y_1 (again with refer to the names used in the definition of φ_b). As before the missing edges are filled with the flexible atom s . Figure 4 illustrates the situation. It follows from the choice of φ'_a and φ'_b and the definition of φ_b that $\varphi_b \wedge b(x_1, x_2) \wedge b(y_1, y_2)$ is not satisfiable in \mathfrak{B} .

Case 2: This case is analogous to Case 1.

Case 3: Consider the formula φ_a with

$$\varphi_a(x_1, x_2, y_1, y_2) := \exists z. (\varphi'_a(x_1, x_2, x_2, z) \wedge \varphi'_b(x_2, z, z, y_1) \wedge \varphi'_a(z, y_1, y_1, y_2)).$$

We show that $\varphi_a \wedge \psi(x_1, x_2, y_1, y_2) \wedge a(x_1, x_2) \wedge b(y_1, y_2)$ is satisfiable in \mathfrak{B} . Since (5) holds for $c = a$ and $d = b$ and since a and b are distinct, there exists $p_1 \in A_0 \setminus \{\text{Id}\}$ such that

$$\varphi'_a(x_1, x_2, y_1, y_2) \wedge a(x_1, x_2) \wedge b(y_1, y_2) \wedge \text{Id}(x_2, y_1) \wedge p_1(x_1, y_2)$$

is satisfiable in \mathfrak{B} . Similarly, since (5) holds for $c = b$ and $d = a$, there exists $p_2 \in A_0 \setminus \{\text{Id}\}$ such that

$$\text{and } \varphi'_b(x_1, x_2, y_1, y_2) \wedge b(x_1, x_2) \wedge a(y_1, y_2) \wedge \text{Id}(x_2, y_1) \wedge p_2(x_1, y_2).$$




Figure 5: The formula δ build from φ'_a (red), φ'_b (blue), and the flexible atom s (dotted).

Note that there are $u_1, \dots, u_5 \in B$ such that the following atomic formulas hold:

$$\begin{aligned} a(u_1, u_2), p_1(u_1, u_3), s(u_1, u_4), s(u_1, u_5), \\ b(u_2, u_3), p_2(u_2, u_4), s(u_2, u_5), \\ a(u_3, u_4), p_1(u_3, u_5), \\ b(u_4, u_5). \end{aligned}$$

If we choose for the existentially quantified variable z in the definition of φ_a the element u_3 then the tuple (u_1, u_2, u_4, u_5) satisfies the formula

$$\varphi_a \wedge \psi(x_1, x_2, y_1, y_2) \wedge a(x_1, x_2) \wedge b(y_1, y_2).$$

By an analogous argument also $\varphi_a \wedge \psi(x_1, x_2, y_1, y_2) \wedge b(x_1, x_2) \wedge a(y_1, y_2)$ is satisfiable. It follows again from the choice of φ'_a and φ'_b and the definition of φ_a that $\varphi_a \wedge a(x_1, x_2) \wedge a(y_1, y_2)$ is not satisfiable in \mathfrak{B} . By an analogous definition we can find a formula φ_b that satisfies b) in Claim 1. Therefore, we are done with Case 3. Altogether this proves Claim 1.

Let φ_a and φ_b be the two formulas that exist by Claim 1. We define the following formulas

$$\begin{aligned} \varphi'_a(x_1, x_2, y_1, y_2) &:= \varphi_a(x_1, x_2, y_1, y_2) \wedge (a \cup b)(x_1, x_2) \wedge (a \cup b)(y_1, y_2) \\ \varphi'_b(x_1, x_2, y_1, y_2) &:= \varphi_b(x_1, x_2, y_1, y_2) \wedge (a \cup b)(x_1, x_2) \wedge (a \cup b)(y_1, y_2). \end{aligned}$$

We also define a formula δ as follows (see also Figure 5):

$$\begin{aligned} \delta(x_1, x_2, x_3, x_4) := & s(x_1, x_3) \wedge s(x_1, x_4) \wedge s(x_2, x_3) \wedge s(x_2, x_4) \\ & \wedge \exists y_1, y_2, y_3, y_4. (\varphi'_a(x_1, x_2, y_1, y_2) \wedge \varphi'_b(y_1, y_2, y_3, y_4) \\ & \wedge \varphi'_a(y_3, y_4, x_3, x_4)) \\ & \wedge \exists z_1, z_2, z_3, z_4. (\varphi'_b(x_1, x_2, z_1, z_2) \wedge \varphi'_a(z_1, z_2, z_3, z_4) \\ & \wedge \varphi'_b(z_3, z_4, x_3, x_4)). \end{aligned}$$

Analogously to Case 1, an amalgam of the structures that are induced by tuples that satisfy φ'_a and φ'_b shows that the formulas $\delta \wedge a(x_1, x_2) \wedge b(x_3, x_4)$ and $\delta \wedge b(x_1, x_2) \wedge a(x_3, x_4)$ are satisfiable in \mathfrak{B} . Note that this is possible since we ensured in Claim 1 that there exist tuples that additionally satisfy ψ . It also holds that a tuple x that satisfies δ also satisfies

$$(a(x_1, x_2) \wedge b(x_3, x_4)) \vee (b(x_1, x_2) \wedge a(x_3, x_4)). \quad (7)$$

Assume that for a tuple x that satisfies δ it holds that $a(x_1, x_2) \wedge a(x_3, x_4)$. Then there exist y_1, y_2, y_3, y_4 such that $\varphi'_a(x_1, x_2, y_1, y_2) \wedge \varphi'_a(y_3, y_4, x_3, x_4)$ holds. But this is by the definition of φ'_a only possible if $b(y_1, y_2)$ and $b(y_3, y_4)$ hold, in contradiction to $\varphi'_b(y_1, y_2, y_3, y_4)$. The same argument works for proving that $\neg(b(x_1, x_2) \wedge b(x_3, x_4))$ holds.

We complete the proof with a primitive positive definition of $E_{a,b}$. We have the following primitive positive formula

$$\delta'(x_1, x_2, x_3, x_4) := \exists y_1, y_2. (\delta(x_1, x_2, y_1, y_2) \wedge \delta(y_1, y_2, x_3, x_4)),$$

and define $E_{a,b}$ by

$$(x_1, x_2, x_3, x_4) \in E_{a,b} \Leftrightarrow \delta'(x_1, x_2, x_3, x_4).$$

For the forward direction of this equivalence, let x be a tuple from $E_{a,b}$ such that $c(x_1, x_2)$ holds for $c \in \{a, b\}$. Let $d \in \{a, b\} \setminus \{c\}$ and let y_1 and y_2 be two elements from \mathfrak{B} such that $d(y_1, y_2)$ and $s(x_i, y_j)$ for every $i \in \{1, \dots, 4\}$ and every $j \in \{1, 2\}$ holds. Such elements exist since in the substructure of \mathfrak{B} that is induced by $\{x_1, x_2, x_3, x_4, y_1, y_2\}$ all appearing triangles are allowed by the definition of the flexible atom s . The elements y_1 and y_2 witness that x satisfies the formula δ' .

For the other direction assume that a tuple x satisfies δ' . Then there exist y_1 and y_2 such that $\delta(x_1, x_2, y_1, y_2) \wedge \delta(y_1, y_2, x_3, x_4)$ is satisfied. Since we observed that the tuples (x_1, x_2, y_1, y_2) and (y_1, y_2, x_3, x_4) both satisfy (7) we can assume that $c(x_1, x_2)$ holds for $c \in \{a, b\}$. It follows also from (7) that $d(y_1, y_2)$ holds for $d \in \{a, b\} \setminus \{c\}$ and then (7) implies that $c(x_3, x_4)$ holds, which proves the backward direction of the stated equivalence. \square

9. Proof of the Result

In this section we prove the main results of this article. We first obtain a dichotomy theorem for a class of CSPs (Theorem 9.1). This is used in combination with the observations in Section 3 to conclude the proof of Theorem 1.1.

Theorem 9.1. *Let $\mathbf{A} \in \text{RRA}$ be finite integral symmetric and with a flexible atom s and let A_0 be the set of atoms of \mathbf{A} . Then either*

- there exists an operation $f: A_0^6 \rightarrow A_0$ that preserves the allowed triples of \mathbf{A} , satisfies

$$\forall x_1, \dots, x_6 \in A_0. f(x_1, \dots, x_6) \in \{x_1, \dots, x_6\}$$

and satisfies the Siggers identity

$$\forall x, y, z \in A_0. f(x, x, y, y, z, z) = f(y, z, x, z, x, y);$$

in this case, $\text{CSP}(\mathfrak{B})$ is in P , or

- $\text{HSP}^{\text{fin}}(\{\text{Pol}(\mathfrak{B})\})$ contains a 2-element algebra where all operations are projections; in this case, $\text{CSP}(\mathfrak{B})$ is NP-complete.

Proof. Let \mathfrak{B} be a normal representation of \mathbf{A} that exists by Proposition 3.5. The finite boundedness of \mathfrak{B} implies that $\text{CSP}(\mathfrak{B})$ is in NP. Let \mathfrak{D} be the atom structure of \mathbf{A} . We can assume that \mathfrak{B} has a binary injective polymorphism, because otherwise Proposition 6.1 would directly imply the second item. The existence of a binary injective polymorphism implies by Corollary 7.3 the existence of a canonical binary injective polymorphism g .

If the first item of the theorem is satisfied then the operation f is a Siggers polymorphism of \mathfrak{D} and the statement follows by the Propositions 4.3 and 4.4.

Assume therefore that the first item in the theorem does not hold. By Corollary 4.7 there exist elements $a, b \in A_0$ such that the subalgebra of $\text{Pol}(\mathfrak{D})$ on $\{a, b\}$ contains only projections. It holds that $\text{Id} \notin \{a, b\}$, since g is a witness that Id can not be in the domain of a subalgebra that contains only projections. Since all operations from $\text{Pol}(\mathfrak{D})$ are projections on $\{a, b\}$ there exists no canonical polymorphism of \mathfrak{B} that is $\{a, b\}$ -symmetric. By Proposition 7.5 there exists also no $\{a, b\}$ -symmetric polymorphism of \mathfrak{B} . Since \mathfrak{B} has a binary injective polymorphism we can apply Proposition 8.2 and get that all polymorphisms of \mathfrak{B} are $\{a, b\}$ -canonical. The last step is to show that all polymorphisms of \mathfrak{B} behave like projections on $\{a, b\}$.

Assume for contradiction that there exists a ternary, $\{a, b\}$ -canonical polymorphism m that behaves on $\{a, b\}$ like a majority or like a minority. By Corollary 7.11 there exists a canonical polymorphism that is also a majority or minority on $\{a, b\}$ (here we use again the existence of an injective polymorphism). This contradicts our assumption that $\text{Pol}(\mathfrak{D})$ is trivial on $\{a, b\}$. We get that every polymorphism of \mathfrak{B} does not behave on $\{a, b\}$ as an operation from Post's theorem (Theorem 2.15) and therefore must behave as a projection on $\{a, b\}$ by Theorem 2.15. Thus, $\text{HSP}^{\text{fin}}(\{\text{Pol}(\mathfrak{B})\})$ contains a 2-element algebra whose operations are projections and $\text{CSP}(\mathfrak{B})$ is NP-hard, according to Theorem 2.21. \square

We can prove the main result.

Proof of Theorem 1.1. Let \mathbf{A} be as in the assumptions of the theorem and let \mathbf{A}' be the finite symmetric integral representable relation algebra that exists by Proposition 3.3. Suppose that \mathbf{A} satisfies the first condition of Theorem 1.1. By Item 3) in Proposition 3.3, we get that then \mathbf{A}' satisfies the first condition in Theorem 9.1 and therefore $\text{CSP}(\mathfrak{B})$ for the normal representation \mathfrak{B} of \mathbf{A}' is in P. By Section 2.3 we know that $\text{NSP}(\mathbf{A}')$ and $\text{CSP}(\mathfrak{B})$ are polynomial-time equivalent. This, together with the Turing reduction from $\text{NSP}(\mathbf{A})$ to $\text{NSP}(\mathbf{A}')$ by Item 1) in Proposition 3.3 implies that $\text{NSP}(\mathbf{A})$ is in P. This proves the first part of the theorem.

Assume that \mathbf{A} does not satisfy the first condition of Theorem 1.1. Item 3) in Proposition 3.3 again implies that \mathbf{A}' does not satisfy the first condition in Theorem 9.1 and therefore $\text{CSP}(\mathfrak{B})$ for the normal representation \mathfrak{B} of \mathbf{A}' is NP-complete. As before we get that $\text{NSP}(\mathbf{A}')$ is NP-complete and the many-one reduction from $\text{NSP}(\mathbf{A})$ to $\text{NSP}(\mathbf{A}')$ by Item 2) in Proposition 3.3 implies the NP-hardness of $\text{NSP}(\mathbf{A})$. The containment in NP follows by Item 1) in Proposition 3.3. This concludes the proof of Theorem 1.1. \square

Corollary 9.2. *For a given finite, symmetric $\mathbf{A} \in \text{RRA}$ with a flexible atom it is decidable which of the two items in Theorem 1.1 holds. In particular, it is decidable whether $\text{NSP}(\mathbf{A})$ is solvable in polynomial time.*

Proof. Since A_0 is a finite set one can go through all possible operations $f: A_0^6 \rightarrow A_0$ that preserve all the allowed triples of \mathbf{A} and check whether the Siggers identity is satisfied. If $\text{P} \neq \text{NP}$, it follows from Theorem 1.1 that it is decidable whether $\text{NSP}(\mathbf{A})$ is in P. In the case of $\text{P} = \text{NP}$ this is a trivial task. \square

The problem of deciding whether certain identities hold in the polymorphism clone of a structure is well known problem in the study of CSPs. The computational complexity is known to be in NP (Chen & Larose, 2017). The precise complexity of deciding whether a polymorphism clone of an explicitly given structure has a conservative operation that satisfies the Siggers identity is open (Chen & Larose, 2017).

10. Connection to Smooth Approximations

We discuss the relationship of our results to the techniques developed by Mottet and Pinsker (2022). The main invention of Mottet and Pinsker (2022) are *smooth approximations* which are equivalence relations on sets of n -tuples. The purpose of these equivalence relations is to *approximate* the prominent *orbit-equivalence relation*. We have seen the importance of these relations in the present article: the polymorphisms which preserve the orbit-equivalence relation are precisely the canonical ones and they store the information about possible finite-domain algorithms that can be used to solve the infinite-domain CSP (cf. Section 4).

We start by rearranging the results of Section 8 such that we get the following theorem. In order to repeat the key steps of our main proof with a focus on the similarities to Mottet and Pinsker (2022) the assumptions in this theorem are a natural starting point.

Theorem 10.1. *Let \mathfrak{B} be a normal representation of a finite integral symmetric relation algebra with a flexible atom s and let \mathfrak{D} be the atom structure of \mathfrak{B} . Suppose that $\text{Pol}(\mathfrak{B})$ contains a binary injective polymorphism and $\text{Pol}(\mathfrak{D})$ does not have a Siggers operation. Then there exists two atoms $a \not\leq \text{Id}$ and $b \not\leq \text{Id}$ such that one of the following holds:*

1. *The orbit-equivalence relation $E_{a,b}$ is primitively positively definable in \mathfrak{B} .*
2. *For all $x, y \in \{a, b\}$ and every primitive positive formula φ such that $\varphi \wedge x(x_1, x_2) \wedge y(y_1, y_2)$ and $\varphi \wedge y(x_1, x_2) \wedge x(y_1, y_2)$ are satisfiable over \mathfrak{B} , the formula $\varphi \wedge x(x_1, x_2) \wedge x(y_1, y_2)$ is also satisfiable over \mathfrak{B} .*

This theorem is similar to the “Loop lemma of approximations” (Theorem 10 Mottet & Pinsker, 2022), even though it does not make proper use of the approximation idea

($E_{a,b}$ approximates itself). The first case in the theorem leads to the hardness condition that $\text{HSP}^{\text{fin}}(\{\text{Pol}(\mathfrak{B})\})$ contains a 2-element algebra whose operations are projections and therefore $\text{CSP}(\mathfrak{B})$ is NP-complete.

However, the second case is “stronger” than the second case in Theorem 10 by Mottet and Pinsker (2022). The strength in the statement relies on the special class of problems in our article. To see what we mean by this, we continue with the strategy of Mottet and Pinsker (2022). As a next step we can restrict the “Independence Lemma” (Lemma 8.1) as follows:

Lemma 10.2. *Let \mathfrak{B} be a homogeneous structure with finite relational signature. Let a and b be 2-orbits of $\text{Aut}(\mathfrak{B})$ such that a , b , and $(a \cup b)$ are primitively positively definable in \mathfrak{B} .*

Assume that for every primitive positive formula φ such that $\varphi \wedge a(x_1, x_2) \wedge b(y_1, y_2)$ and $\varphi \wedge b(x_1, x_2) \wedge a(y_1, y_2)$ are satisfiable over \mathfrak{B} , the formula $\varphi \wedge a(x_1, x_2) \wedge a(y_1, y_2)$ is also satisfiable over \mathfrak{B} . Then \mathfrak{B} has an $\{a, b\}$ -canonical polymorphism g that is $\{a, b\}$ -symmetric with $\bar{g}(a, b) = \bar{g}(b, a) = a$.

Note that the assumptions in this lemma are precisely what we get from case 2) in Theorem 10.1. A lemma of similar style can be also found as Lemma 13 in the article of Mottet and Pinsker (2022). They obtain in this lemma a *weakly commutative* operation. The $\{a, b\}$ -symmetric operations from our article are a special case of weakly commutative operations. The difference between these two properties seems crucial for the next step of our proof: while $\{a, b\}$ -symmetric operations can often be “lifted” to canonical $\{a, b\}$ -symmetric operations this seems not clear in general for weakly commutative operations. This last step is necessary in order to obtain a contradiction to our assumption that $\text{Pol}(\mathfrak{D})$ does not have a Siggers operation.

In order to apply the results of Mottet and Pinsker (2022) directly to obtain the dichotomy results of our article one would have to find a way to canonize the weakly commutative operations in a suitable way. It seems that this can be done by a similar proof as for Lemma 56 of Mottet and Pinsker (2022) were the authors used among other things arguments which are also incorporated in our proof of the loop lemma (Theorem 10.1).

11. Conclusion

We classified the computational complexity of the network satisfaction problem for finite symmetric $\mathbf{A} \in \text{RRA}$ with a flexible atom and obtained a P versus NP-complete dichotomy. We gave a decidable criterion for \mathbf{A} that is a sufficient condition for the membership of $\text{NSP}(\mathbf{A})$ in P, which is also necessary unless $\text{P}=\text{NP}$. We want to mention that if we drop the assumptions on \mathbf{A} to be symmetric and to have a flexible atom then the statement of Theorem 1.1 is false. An example for this is the Point Algebra; even though the NSP of this representable relation algebra is in P (Vilain et al., 1990), the first condition of Theorem 1.1 does not apply. However, if we only drop the symmetry assumption we conjecture that Theorem 1.1 still holds; yet it is not clear how to obtain the results of Section 7 in this case. Similarly, if we only drop the flexible atom assumption we conjecture that the statement also remains true. For this generalization it would be necessary to obtain Ramsey-type results like Theorem 3.6 without the assumption of the existence of a flexible atom.

Acknowledgments

The first author has received funding from the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 681988, CSP-Infinity). The second author is supported by DFG Graduiertenkolleg 1763 (QuantLA).

References

- Allen, J. F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11), 832–843.
- Alm, J. F., Maddux, R. D., & Manske, J. (2008). Chromatic graphs, Ramsey numbers and the flexible atom conjecture. *The Electronic Journal of Combinatorics*, 15(1).
- Andréka, H., & Maddux, R. D. (1994). Representations for small relation algebras. *Notre Dame Journal of Formal Logic*, 35(4), 550–562.
- Barto, L. (2011). The dichotomy for conservative constraint satisfaction problems revisited. In *Proceedings of the Symposium on Logic in Computer Science (LICS)*, Toronto, Canada.
- Bodirsky, M. (2008). Constraint satisfaction problems with infinite templates. In Vollmer, H. (Ed.), *Complexity of Constraints (a collection of survey articles)*, Vol. 5250 of *Lecture Notes in Computer Science*, pp. 196–228. Springer.
- Bodirsky, M. (2012). Complexity classification in infinite-domain constraint satisfaction. Mémoire d’habilitation à diriger des recherches, Université Diderot – Paris 7. Available at arXiv:1201.0856v8.
- Bodirsky, M. (2015). Ramsey classes: Examples and constructions. In *Surveys in Combinatorics. London Mathematical Society Lecture Note Series 424*. Cambridge University Press. Invited survey article for the British Combinatorial Conference; ArXiv:1502.05146.
- Bodirsky, M. (2018). 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*, pp. 3–17.
- Bodirsky, M. (2021). *Complexity of Infinite-Domain Constraint Satisfaction*. Lecture Notes in Logic. Cambridge University Press.
- Bodirsky, M., & Dalmau, V. (2013). Datalog and constraint satisfaction with infinite templates. *Journal on Computer and System Sciences*, 79, 79–100. A preliminary version appeared in the proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS’05).
- Bodirsky, M., & Jonsson, P. (2017). A model-theoretic view on qualitative constraint reasoning. *Journal of Artificial Intelligence Research*, 58, 339–385.

- Bodirsky, M., & Kára, J. (2009). The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2), 1–41. An extended abstract appeared in the Proceedings of the Symposium on Theory of Computing (STOC).
- Bodirsky, M., & Knäuer, S. (2020). Hardness of network satisfaction for relation algebras with normal representations. In *Relational and Algebraic Methods in Computer Science*, pp. 31–46. Springer International Publishing.
- Bodirsky, M., & Knäuer, S. (2021). Network satisfaction for symmetric relation algebras with a flexible atom. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7), 6218–6226.
- Bodirsky, M., & Kutz, M. (2007). Determining the consistency of partial tree descriptions. *Artificial Intelligence*, 171, 185–196.
- Bodirsky, M., Martin, B., Pinsker, M., & Pongrácz, A. (2019). Constraint satisfaction problems for reducts of homogeneous graphs. *SIAM Journal on Computing*, 48(4), 1224–1264. A conference version appeared in the Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, pages 119:1–119:14.
- Bodirsky, M., & Mottet, A. (2016). Reducts of finitely bounded homogeneous structures, and lifting tractability from finite-domain constraint satisfaction. In *Proceedings of the 31th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pp. 623–632. Preprint available at ArXiv:1601.04520.
- Bodirsky, M., & Nešetřil, J. (2006). Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3), 359–373.
- Bodirsky, M., & Pinsker, M. (2014). Minimal functions on the random graph. *Israel Journal of Mathematics*, 200(1), 251–296.
- Bodirsky, M., & Pinsker, M. (2015). Schaefer’s theorem for graphs. *Journal of the ACM*, 62(3), 52 pages (article number 19). A conference version appeared in the Proceedings of STOC 2011, pages 655–664.
- Bodirsky, M., & Pinsker, M. (2021). Canonical functions: a proof via topological dynamics. *Contributions Discret. Math.*, 16(2), 36–45.
- Bodirsky, M., Pinsker, M., & Pongrácz, A. (2019). Projective clone homomorphisms. *The Journal of Symbolic Logic*, 86(1), 148–161.
- Bulatov, A. A. (2003). Tractable conservative constraint satisfaction problems. In *Proceedings of the Symposium on Logic in Computer Science (LICS)*, pp. 321–330, Ottawa, Canada.
- Bulatov, A. A. (2016). Conservative constraint satisfaction re-revisited. *Journal Computer and System Sciences*, 82(2), 347–356. ArXiv:1408.3690.
- Bulatov, A. A. (2017). A dichotomy theorem for nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pp. 319–330.
- Chen, H., & Larose, B. (2017). Asking the metaquestions in constraint tractability. *TOCT*, 9(3), 11:1–11:27.

- Comer, S. D. (1984). Combinatorial aspects of relations. *Algebra Universalis*, 18(1), 77–94.
- Conant, G. J. (2015). *Model Theory and Combinatorics of Homogeneous Metric Spaces*. Ph.D. thesis, University of Illinois at Chicago.
- Cristiani, M., & Hirsch, R. (2004). The complexity of the constraint satisfaction problem for small relation algebras. *Artificial Intelligence Journal*, 156, 177–196.
- Delhommé, C., Laflamme, C., Pouzet, M., & Sauer, N. (2007). Divisibility of countable metric spaces. *European Journal of Combinatorics*, 28(6), 1746–1769.
- Düntsch, I. (2005). Relation algebras and their application in temporal and spatial reasoning. *Artificial Intelligence Review*, 23, 315–357.
- Hirsch, R., & Hodkinson, I. (2001). Strongly representable atom structures of relation algebras. *Transactions of the American Mathematical Society*, 130(6), 1819–1831.
- Hirsch, R. (1996). Relation algebras of intervals. *Artificial Intelligence Journal*, 83, 1–29.
- Hirsch, R. (1997). Expressive power and complexity in algebraic logic. *Journal of Logic and Computation*, 7(3), 309 – 351.
- Hirsch, R. (1999). A finite relation algebra with undecidable network satisfaction problem. *Logic Journal of the IGPL*, 7(4), 547–554.
- Hirsch, R., & Hodkinson, I. (2002). *Relation Algebras by Games*. North Holland.
- Hirsch, R., Jackson, M., & Kowalski, T. (2019). Algebraic foundations for qualitative calculi and networks. *Theor. Comput. Sci.*, 768, 99–116.
- Hodges, W. (1997). *A shorter model theory*. Cambridge University Press, Cambridge.
- Hubička, J., & Nešetřil, J. (2019). All those Ramsey classes (Ramsey classes with closures and forbidden homomorphisms). *Advances in Mathematics*, 356, 106791.
- Ladkin, P. B., & Maddux, R. D. (1994). On binary constraint problems. *Journal of the Association for Computing Machinery*, 41(3), 435–469.
- Lyndon, R. (1950). The representation of relational algebras. *Annals of Mathematics*, 51(3), 707–729.
- Macpherson, D. (2011). A survey of homogeneous structures. *Discrete Mathematics*, 311(15), 1599–1634.
- Maddux, R. (1982). Some varieties containing relation algebras. *Transactions of the American Mathematical Society*, 272(2), 501–501.
- Maddux, R. D. (1994). A perspective on the theory of relation algebras. *Algebra Universalis*, 31(3), 456–465.
- Maddux, R. D. (2006a). Finite symmetric integral relation algebras with no 3-cycles. In Schmidt, R. A. (Ed.), *Relations and Kleene Algebra in Computer Science, 9th International Conference on Relational Methods in Computer Science and 4th International Workshop on Applications of Kleene Algebra, RelMiCS/AKA 2006, Manchester, UK, August 29–September 2, 2006, Proceedings*, Vol. 4136 of *Lecture Notes in Computer Science*, pp. 2–29. Springer.

- Maddux, R. D. (2006b). *Relation Algebras: Volume 150*. Studies in logic and the foundations of mathematics. Elsevier Science, London, England.
- Mottet, A., & Pinsker, M. (2022). Smooth approximations and cpsps over finitely bounded homogeneous structures. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '22*, New York, NY, USA. Association for Computing Machinery.
- Nešetřil, J., & Rödl, V. (1989). The partite construction and Ramsey set systems. *Discrete Mathematics*, 75(1-3), 327–334.
- Post, E. L. (1941). The two-valued iterative systems of mathematical logic. *Annals of Mathematics Studies*, 5.
- Renz, J., & Nebel, B. (1999). On the complexity of qualitative spatial reasoning: A maximal tractable fragment of the region connection calculus. *Artificial Intelligence*, 108(1-2), 69–123.
- Renz, J., & Nebel, B. (2007). Qualitative spatial reasoning using constraint calculi. In Aiello, M., Pratt-Hartmann, I., & van Benthem, J. (Eds.), *Handbook of Spatial Logics*, pp. 161–215. Springer Verlag, Berlin.
- Siggers, M. H. (2010). A strong Mal'cev condition for varieties omitting the unary type. *Algebra Universalis*, 64(1), 15–20.
- Vilain, M., Kautz, H., & van Beek, P. (1990). Constraint propagation algorithms for temporal reasoning: A revised report. In Weld, D. S., & de Kleer, J. (Eds.), *Readings in Qualitative Reasoning About Physical Systems*, pp. 373–381. Morgan Kaufmann.
- Zhuk, D. N. (2017). A proof of CSP dichotomy conjecture. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pp. 331–342.
- Zhuk, D. N. (2020). A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5).