TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

DISSERTATION

Algorithmic Techniques for Constraint Satisfaction Problems over Finitely Bounded Homogeneous Structures

ausgeführt zum Zwecke der Erlangung des akademischen Grades
Doktor der Naturwissenschaften

eingereicht an der Technischen Universität Wien,
Fakultät für Mathematik und Geoinformation

von

Tomáš Nagy

Matrikelnummer: 11936034

Betreuung: Associate Prof. Dipl.-Ing. Dr.techn. **Michael Pinsker**
Institut für Diskrete Mathematik
Forschungsbereich Algebra
Technische Universität Wien
Wiedner Hauptstraße 8-10/104, 1040 Wien, Österreich

Begutachtung: Associate Prof. Dipl.-Ing. Dr.techn. **Stefan Hetzl**
Institut für Diskrete Mathematik
Forschungsbereich Computational Logic
Technische Universität Wien
Wiedner Hauptstraße 8-10/104, 1040 Wien, Österreich

Begutachtung: Prof. Dr. **Antoine Mottet**
Forschungsgruppe Theoretische Informatik
Technische Universität Hamburg
Blohmstraße 15, 21079 Hamburg, Deutschland

Wien, im Mai 2023

Acknowledgements

First of all, I would like to thank my supervisor, Michael Pinsker, for all his help during my studies. On the one hand, I am very grateful for the research topics he pointed out to me, for the days and nights he spent reading my secret notes and decrypting all the twisted claims and proofs, and for all the time we spent discussing various mathematical problems. These discussions eventually led to all the results that are presented in this thesis. On the other hand, I am particularly grateful for the freedom he gave me – both in choosing the questions I wanted to work on and in scheduling my work.

I would also like to thank Antoine Mottet for his patience and willingness to discuss mathematical problems, and for the many problems and ideas he shared with me. I am very grateful for all the fruitful collaboration that eventually led to two publications that form a substantial part of this thesis.

I am also very grateful to Michał Wrona. In the first year of my studies, he was willing to think about a problem that I mentioned to him and shared with me his ideas about it which were eventually used in the first publication of my doctoral studies.

I must also mention David Stanovský, the supervisor of my bachelor's and master's theses. He strongly encouraged me to do a PhD abroad and also recommended that I look for a position in Vienna with Michael Pinsker. Without him, my PhD studies might not have even started.

I also have to mention my colleagues, office mates and friends, Albert Vucaj and Jakub Rydval. I am very grateful for all the time we spent together discussing mathematical problems, discussing completely irrelevant problems, gossiping, drinking iced doppio macchiato or bouldering.

I am also very grateful to all the other members of the CSP community. When I went to my first conference in the first year of my studies, I immediately felt welcomed into this community. I am also very grateful for all the time we spent discussing all kinds of problems, climbing, bouldering and hiking together.

I am also very thankful to all the friends I made in Vienna for all the free time we spent together – their support, as well as the opportunity to take a break from everything related to my studies and mathematics in general, certainly helped me to relax and then to approach the mathematical problems I was trying to solve anew.

I am especially grateful to my parents for their unwavering support, no matter what I did or what decision I made. In particular, I am very grateful for all the time I was able to spend at home during the lockdowns that occurred during my studies. Their support and presence helped me to lift my spirits during those difficult times and even to concentrate on my work.

It may sound strange, but I must also mention the mountains – especially the Tatras and the Alps – where I was able to relax, but also where I was able to concentrate better and solve the problems whose solutions form an essential part of this thesis. In fact, most of the ideas for solving these problems came to me while hiking, skitouring or climbing in the mountains.

Finally, I would like to thank the One who is for the light that shines in the darkness, but also – as hard as it may be – for the darkness that is sometimes needed to see the light.

This research was funded in whole or in part by the Austrian Science Fund (FWF) [P 32337, I 5948]. The author has also received funding from the Austrian Science Fund (FWF) through Lise Meitner Grant No M 2555-N35 and from the Czech Science Foundation (Grant No 18-20123S).

Kurzfassung

Ein *Bedingungserfüllungsproblem* über einer relationalen Struktur $\mathbb{A}$ ist das Problem, bei dem eine *Instanz* – d.h. eine endliche Menge von Variablen und Bedingungen – gegeben ist und das Ziel ist zu entscheiden, ob es eine Zuordnung der Werte aus dem Universum von $\mathbb{A}$ zu den Variablen gibt, so dass alle Bedingungen erfüllt sind. In dieser Arbeit befassen wir uns mit Bedingungserfüllungsproblemen über Strukturen, die in endlich beschränkten homogenen Strukturen erster Ordnung definierbar sind (sogenannte *Redukte erster Ordnung* dieser Strukturen). Während uns die Homogenität versichert, dass jede Lösung einer Instanz bis auf Automorphismen durch die Relationen auf ihrem Bild gegeben ist, impliziert die endliche Beschränktheit, dass jede Lösung nur auf Teilmengen mit einer festen Größe überprüft werden muss.

In Kapitel 1 geben wir eine Einführung in die Theorie der Bedingungserfüllungsprobleme und diskutieren verwandte Konzepte aus der Modelltheorie und der universellen Algebra, die im Rest der Arbeit benötigt werden.

Prüfung auf lokale Konsistenz ist eine der wichtigsten algorithmischen Techniken im Gebiet der Bedingungserfüllungsprobleme. In Kapitel 2 beweisen wir eine obere Schranke für die Menge an lokaler Konsistenz, die benötigt wird, um Bedingungserfüllungsprobleme über Redukten erster Ordnung von endlich beschränkten homogenen Strukturen zu lösen, die bestimmte algebraische Bedingungen erfüllen. Als Korollar erhalten wir eine obere Schranke für die Menge an lokaler Konsistenz, die benötigt wird, um Bedingungserfüllungsprobleme über Redukte erster Ordnung von vielen bekannten relationalen Strukturen zu lösen, die durch Prüfung auf lokale Konsistenz lösbar sind. Wir erhalten auch eine Charakterisierung der begrenzten Weite für Redukte erster Ordnung von unären Strukturen und für bestimmte Strukturen, die mit der Logik MMSNP zusammenhängen.

In Kapitel 3 betrachten wir Bedingungserfüllungsprobleme über Redukte erster Ordnung von bestimmten uniformen Hypergraphen. In [73] wurde festgestellt, dass diese Bedingungserfüllungsprobleme nicht mit den Standardmethoden gelöst werden können, die in den meisten bekannten Klassifikationen von Bedingungserfüllungsproblemen über Redukte erster Ordnung von endlich beschränkten homogenen Strukturen verwendet werden. Daher stellen wir einen neuen Algorithmus zur Lösung der betrachteten Bedingungserfüllungsprobleme vor. Dieser Algorithmus verwendet Umformulierungen mehrerer Begriffe, die in dem Algorithmus für die Lösung aller traktablen Bedingungserfüllungsprobleme über endlichen Domänen aus [83, 84] verwendet werden.

In Kapitel 4 betrachten wir eine Klasse von Redukten erster Ordnung von endlich beschränkten homogenen Strukturen, die durch eine stärkere Form der lokalen Konsistenz lösbar sind. Wir beweisen, dass diese Strukturen eine begrenzte Ausdruckkraft in Form von implikationeller Einfachheit haben. Dies impliziert, dass nur eine begrenzte Menge an lokaler Konsistenz erforderlich ist, um ihre Bedingungserfüllungsprobleme zu lösen.

Kapitel 2 entspricht der Veröffentlichung [72] (mit Antoine Mottet, Michael Pinsker und

Michał Wrona), die eine Zeitschriftenversion von [71] ist; Kapitel 3 entspricht der Veröffentlichung [70] (mit Antoine Mottet und Michael Pinsker), Kapitel 4 basiert auf noch nicht veröffentlichten Ergebnissen, die in Zusammenarbeit mit Michael Pinsker erhalten wurden.

Abstract

A *Constraint Satisfaction Problem (CSP)* over a relational structure $\mathbb{A}$ is the problem where one is given an *instance* – i.e., a finite set of variables and constraints, and one has to decide whether there exists an assignment of the values from the domain of $\mathbb{A}$ to the variables such that all constraints are satisfied. In this thesis, we deal with CSPs over structures which are first-order definable in finitely bounded homogeneous structures (so-called *first-order reducts* of these structures). While homogeneity assures us that every solution of a CSP instance is up to automorphisms given by the relations holding on its image, finite boundedness implies that every solution must only be verified locally on subsets of a fixed size.

In Chapter 1, we give an introduction to the theory of CSPs and we discuss related concepts from model theory and universal algebra that are be needed in the rest of the thesis.

Local consistency checking is one of the most important algorithmic techniques in the area of constraint satisfaction. In Chapter 2, we prove an upper bound on the amount of local consistency needed to solve CSPs over first-order reducts of finitely bounded homogeneous structures which satisfy certain algebraic conditions. As a corollary, we obtain a bound on the amount of local consistency needed to solve CSPs over first-order reducts of many well-known relational structures which are solvable by local consistency checking. We also obtain a characterization of bounded width for first-order reducts of unary structures and for certain structures related to the logic MMSNP.

In Chapter 3, we consider CSPs over first-order reducts of certain uniform hypergraphs. It was observed in [73] that these CSPs cannot be solved by standard methods used in most known classifications of CSPs of first-order reducts of finitely bounded homogeneous structures. Therefore, we introduce a new algorithm for solving the CSPs under consideration. This algorithm uses reformulations of several notions used in the algorithm for solving all tractable CSPs over finite domains from [83, 84].

In Chapter 4, we consider a class of first-order reducts of finitely bounded homogeneous structures that are solvable by a stronger form of local consistency. We prove that these structures have limited expressibility in the form of implicational simplicity. This will imply that only a restricted amount of local consistency checking is needed in order to solve their CSPs.

Chapter 2 corresponds to the publication [72] (with Antoine Mottet, Michael Pinsker and Michał Wrona) which is a journal version of [71]; Chapter 3 corresponds to the publication [70] (with Antoine Mottet and Michael Pinsker); Chapter 4 is based on results that have not been published yet and that were obtained in collaboration with Michael Pinsker.

Contents

1 Introduction

In this chapter, we introduce and motivate basic notions connected to the theory of constraint satisfaction problems (CSPs), in particular to CSPs over first-order reducts of finitely bounded homogeneous structures.

1.1 Notation and basic notions from model theory

1.1.1 Notation

All relational structures in this thesis are assumed to be at most countable and over a finite relational signature. We will use the blackboard bold font to denote relational structures. The domain of a relational structure $\mathbb{A}$ will be denoted by A unless stated otherwise. Let $\mathbb{A}$ be a relational structure in a relational signature τ. For any first-order formula ϕ over τ, we will denote by $\phi^{\mathbb{A}}$ its interpretation in $\mathbb{A}$. When the structure $\mathbb{A}$ is clear from the context, we will slightly abuse the notation and use R both for the relational symbol from τ and for its interpretation in $\mathbb{A}$. Let ϕ, ψ be first-order formulas over τ.

For $k \geq 1$, we write $[k]$ for the set $\{1, \ldots, k\}$. Let R be a relation of arity $m \geq 1$, let $n \geq 1$, and let $i_1, \ldots, i_n \in [m]$. We denote by $\text{proj}_{(i_1,\ldots,i_n)}(R)$ the projection of R to the coordinates $(i_1, \ldots, i_n)$, i.e., $\text{proj}_{(i_1,\ldots,i_n)}(R) = \{(b_{i_1}, \ldots, b_{i_n}) \mid (b_1, \ldots, b_n) \in R\}$.

Let A be a set, and let $m \geq 2$. We say that a tuple $\mathbf{t} \in A^m$ is *injective* if its entries are pairwise different. We say that a relation is injective if it contains only injective tuples. We write I_m^A for the relation containing all injective m-tuples of elements of A. When the set A is clear from the context, we omit the upper index and write I_m instead of I_m^A.

1.1.2 Relational structures and permutation groups

A *first-order reduct* of a structure $\mathbb{A}$ is a structure on the same domain whose relations have a first-order definition in $\mathbb{A}$. A *first-order expansion* of $\mathbb{A}$ is a first-order reduct of $\mathbb{A}$ which contains all relations of $\mathbb{A}$ among its relations.

Let $\mathbb{A}$ and $\mathbb{B}$ be relational structures in the same signature τ. We say that a mapping $f\colon A \to B$ is a *homomorphism* from $\mathbb{A}$ to $\mathbb{B}$ if for every relational symbol $R \in \tau$ of arity k and for every $\mathbf{a} = (a_1, \ldots, a_k) \in R^{\mathbb{A}}$, $f(\mathbf{a}) = (f(a_1), \ldots, f(a_k)) \in R^{\mathbb{B}}$. If f is moreover injective and for every $\mathbf{a} \in A^k$, it holds that $\mathbf{a} \in R^{\mathbb{A}}$ if, and only if, $f(\mathbf{a}) \in R^{\mathbb{B}}$, then f is called an *embedding*. An *isomorphism* is a bijective embedding. Two relational structures are *isomorphic* if there exists an isomorphism between them. A homomorphism from a relational structure to itself is an *endomorphism* of the structure and an isomorphism from a structure to itself is an *automorphism*. The set of all endomorphisms of a relational structure $\mathbb{A}$ will be denoted by $\text{End}(\mathbb{A})$ and the set of all its automorphisms by $\text{Aut}(\mathbb{A})$.

Let $\mathbb{A}$ be a relational structure, and let $n \geq 1$. We define the n-th power $\mathbb{A}^n$ of $\mathbb{A}$ as follows. Let R be a relational symbol from the signature of $\mathbb{A}$ of arity $k \geq 1$. Then $R^{\mathbb{A}^n}$ contains all tuples $((a_1^1, \ldots, a_n^1), \ldots, (a_1^k, \ldots, a_n^k))$ such that $(a_i^1, \ldots, a_i^k) \in R^{\mathbb{A}}$ for every $i \in [n]$.

Definition 1.1.1. *Let $\mathscr{G}$ be a permutation group acting on a set A, let $k \geq 1$, and let $\mathbf{a} \in A^k$. The* orbit *of $\mathbf{a}$ under $\mathscr{G}$ is the set $\{g(\mathbf{a}) \mid g \in G\}$. We say that $\mathscr{G}$ is* oligomorphic *if for every $k \geq 1$, $\mathscr{G}$ has only finitely many orbits of k-tuples in its action on A.*

We say that a relational structure $\mathbb{A}$ is ω-categorical *if its automorphism group is oligomorphic.*

It is a well-known fact that the automorphisms of a relational structure $\mathbb{A}$ preserve all relations that are first-order definable in $\mathbb{A}$. More precisely, if R is any such relation, $\alpha \in \mathrm{Aut}(\mathbb{A})$ and $\mathbf{t} \in R$, then $\alpha(\mathbf{t}) \in R$. It follows that if $\mathbb{B}$ is a first-order reduct of $\mathbb{A}$, then $\mathrm{Aut}(\mathbb{A}) \subseteq \mathrm{Aut}(\mathbb{B})$ and in particular, if $\mathbb{A}$ is ω-categorical, then so is $\mathbb{B}$. If $\mathbb{B}$ is a first-order expansion of $\mathbb{A}$, then $\mathrm{Aut}(\mathbb{B}) = \mathrm{Aut}(\mathbb{A})$.

For an ω-categorical relational structure $\mathbb{A}$ and for every $k \geq 1$, it holds that two k-tuples of elements from A are in the same orbit under $\mathrm{Aut}(\mathbb{A})$ if, and only if, they have the same *type*, i.e., if they satisfy the same first-order formulas.

Definition 1.1.2. *Let $\mathscr{G}$ be a permutation group acting on a set A and let $k \geq 1$. We say that $\mathscr{G}$ is* k-transitive *if it has only one orbit in its action on I_k^A. We say that $\mathscr{G}$ is* transitive *if it is 1-transitive.*

We say that $\mathscr{G}$ is k-homogeneous *if for every $\ell \geq k$, the orbit of every ℓ-tuple under $\mathscr{G}$ is uniquely determined by the orbits of its k-subtuples.*

We say that $\mathscr{G}$ has no k-algebraicity *if the only fixed points of any stabilizer of $\mathscr{G}$ by $k-1$ elements are these elements themselves. We say that $\mathscr{G}$ has* no algebraicity *if it has no k-algebraicity, for every $k \geq 1$.*

The canonical k-ary structure of $\mathscr{G}$ *is the relational structure on A that has a relation for every orbit of k-tuples under $\mathscr{G}$.*

The *atomic type* of a tuple of elements in a relational structure $\mathbb{A}$ is the set of all atomic formulas satisfied by this tuple.

Definition 1.1.3. *Let $\mathbb{A}$ be a relational structure. We say that $\mathbb{A}$ is* transitive *if $\mathrm{Aut}(\mathbb{A})$ is. For $k \geq 1$, we say that $\mathbb{A}$ is* k-homogeneous *if $\mathrm{Aut}(\mathbb{A})$ is k-homogeneous; $\mathbb{A}$ is* homogeneous *if any two tuples of the same atomic type belong to the same orbit. We say that $\mathbb{A}$ has* no algebraicity *if $\mathrm{Aut}(\mathbb{A})$ has no algebraicity.*

Homogeneity is often in the literature defined in a slightly different way – a relational structure $\mathbb{A}$ is homogeneous if any isomorphism between any two finite substructures of $\mathbb{A}$ can be extended to an automorphism of $\mathbb{A}$. However, it is easy to see that this definition is equivalent to our definition of homogeneity. Note moreover that if a relational structure $\mathbb{A}$ in a finite relational signature τ is homogeneous and if n is the maximum of the arities of the relational symbols from τ, then $\mathbb{A}$ is k-homogeneous for every $k \geq n$. It is also easy to observe that every homogeneous relational structure in a finite relational signature is ω-categorical.

Definition 1.1.4. *Let $\ell \geq 1$ and let $\mathbb{A}$ be a relational structure in a signature τ. We say that $\mathbb{A}$ is* ℓ-bounded *if for every finite τ-structure $\mathbb{X}$, if all substructures $\mathbb{Y}$ of $\mathbb{X}$ of size at most ℓ embed to $\mathbb{A}$, then $\mathbb{X}$ embeds to $\mathbb{A}$.*

We say that $\mathbb{A}$ is finitely bounded *if there exists a finite set $\mathcal{F}$ of finite τ-structures such that for every finite τ-structure $\mathbb{X}$, it holds that $\mathbb{X}$ embeds to $\mathbb{A}$ if no $\mathbb{F} \in \mathcal{F}$ embeds to $\mathbb{X}$. Let $\mathcal{F}_{\mathbb{A}}$ be such a set of finite τ-structures with the property that the size of the biggest structure contained in $\mathcal{F}_{\mathbb{A}}$ is the smallest possible among all choices of $\mathcal{F}$; we write $b_{\mathbb{A}}$ for this size.*

Note that a finitely bounded relational structure $\mathbb{A}$ in a finite relational signature τ is ℓ-bounded for every $\ell \geq b_{\mathbb{A}}$.

While the finite substructures of a finitely bounded structure $\mathbb{A}$ are defined by not embedding any of the finitely many structures from $\mathcal{F}_{\mathbb{A}}$, it also makes sense to consider structures where the set of finite structures having a homomorphism to them is defined by forbidding finitely many homomorphic images. This is achieved in the following definition.

Definition 1.1.5. *Let $\mathbb{A}$ be a relational structure in a signature τ. We say that $\mathbb{A}$ has a* finite duality *if there exists a finite set $\mathcal{G}$ of finite τ-structures such that for every finite τ-structure $\mathbb{X}$, it holds that $\mathbb{X}$ has a homomorphism to $\mathbb{A}$ if no $\mathbb{G} \in \mathcal{G}$ has a homomorphism to $\mathbb{A}$. Let $\mathcal{G}_{\mathbb{A}}$ be such a set of finite τ-structures with the property that the size of the biggest structure contained in $\mathcal{G}_{\mathbb{A}}$ is the smallest possible among all choices of $\mathcal{G}$; we denote this size by $d_{\mathbb{A}}$.*

Definition 1.1.6. *A relational structure $\mathbb{A}$ is* universal *for a class of finite structures in its signature if it embeds all members of the class.*

There exist up to isomorphism unique countable homogeneous relational structures that are universal for the following classes of finite structures:

- graphs,
- $\mathbb{K}_n$-free graphs (i.e., graphs not containing the complete graph on n vertices as an induced subgraph) for any fixed $n \geq 3$,
- k-uniform hypergraphs for any fixed $k \geq 3$,
- tournaments,
- partial orders.

All of these homogeneous structures are obtained as *Fraïssé limits* of the respective classes of finite structures [52]. Since we will not need the construction in this thesis, we choose not to present the details – the interested reader can however find them e.g. in [58].

1.2 CSP, polymorphisms and identities

In the whole thesis, we will use the shortcut *CSP* for “constraint satisfaction problem”.

1.2.1 Constraint satisfaction problems

Definition 1.2.1. *A* CSP instance *over a set A is a pair $\mathcal{I} = (\mathcal{V}, \mathcal{C})$, where $\mathcal{V}$ is a finite set of variables and $\mathcal{C}$ is a finite set of* constraints *such that each constraint $C \in \mathcal{C}$ is a subset of A^U for some non-empty $U \subseteq \mathcal{V}$; the set U is called the* scope *of C.*

A solution *of a CSP instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ is a mapping $f \colon \mathcal{V} \to A$ such that for every $C \in \mathcal{C}$ with scope U, $f|_U \in C$.*

CSP over a set A can be understood as a computational problem – an instance of this CSP is given as an input and one has to decide whether this instance has a solution or not. This version is known as the *decision version* of the CSP. In this thesis, this version of the CSP is considered.

In order to systematically study CSPs, it is natural to view constraints of a CSP instance over a set A as relations on A. This can be formalized as follows.

Definition 1.2.2. *Let $\mathbb{A}$ be a relational structure. We say that a CSP instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ over the set A is an* instance of CSP($\mathbb{A}$) *if every constraint $C \in \mathcal{C}$ can be viewed as a relation of $\mathbb{A}$ by totally ordering its scope U; more precisely, we require that there exists an enumeration $u_1, \ldots, u_k$ of the elements of U and a k-ary relation R of $\mathbb{A}$ such that for all $f \colon U \to A$ we have $f \in C \Leftrightarrow (f(u_1), \ldots, f(u_k)) \in R$. The relational structure $\mathbb{A}$ is called the* template *of the CSP.*

Different definitions of a CSP over a relational structure appear in the literature. One of the most common alternative approaches to the definition of a CSP over a relational structure $\mathbb{A}$ is to define it as a homomorphism problem, i.e., an instance of CSP($\mathbb{A}$) is a finite structure $\mathbb{I}$ in the signature of $\mathbb{A}$ and a solution of this instance is a homomorphism from $\mathbb{I}$ to $\mathbb{A}$. However, it is easy to see that this definition is equivalent to our definition in the following sense. If we are given a finite structure $\mathbb{I}$ with domain $\mathcal{V}$, then we can define an instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of CSP($\mathbb{A}$) in our sense as follows. For every relational symbol R from the signature of $\mathbb{A}$ and for every tuple $(v_1, \ldots . v_k) \in R^{\mathbb{I}}$, we add the constraint $C := \{f \in A^{\{v_1,\ldots,v_k\}} \mid f(v_1, \ldots, v_k) \in R^{\mathbb{A}}\}$ to $\mathcal{C}$. It immediately follows that a homomorphism from $\mathbb{I}$ to $\mathbb{A}$ is a solution of $\mathcal{I}$ and vice versa. On the other hand, given an instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of CSP($\mathbb{A}$), we can define a structure $\mathbb{I}$ over the set $\mathcal{V}$ as follows. For every relational symbol R from the signature of $\mathbb{A}$, we define $R^{\mathbb{I}}$ as the set of all tuples $(v_1, \ldots, v_k) \in \mathcal{V}^k$ such that there exists a constraint $C \in \mathcal{C}$ with $f \in C \Leftrightarrow (f(v_1), \ldots, f(v_k)) \in R^{\mathbb{A}}$. It again immediately follows that every homomorphism from $\mathbb{I}$ to $\mathbb{A}$ is a solution of $\mathcal{I}$ and there are no other solutions.

There are many examples of constraint satisfaction problems which appear in different branches of mathematics and computer science.

Example 1.2.3. *Let $\mathbb{K}_3 = (\{v_1, v_2, v_3\}; E)$ be the complete graph on 3 vertices. Then* CSP($\mathbb{K}_3$) *corresponds to the problem of deciding whether the vertices of a given finite graph can be coloured by 3 colours such that no pair of vertices which are coloured by the same colour is connected by an edge. The easiest way of seeing this is to understand* CSP($\mathbb{K}_3$) *as a homomorphism problem. An instance of* CSP($\mathbb{K}_3$) *is a finite graph $\mathbb{G}$ and any homomorphism from $\mathbb{G}$ to $\mathbb{K}_3$ corresponds to a valid colouring of the vertices of $\mathbb{G}$ and vice versa. This problem is known to be* NP-complete.

Example 1.2.4. *Let $\mathbb{A} = (\{0,1\}; R)$ where $R \subseteq \{0,1\}^3$ contains the tuples $(1,0,0)$, $(0,1,0)$ and $(0,0,1)$. Then $\mathrm{CSP}(\mathbb{A})$ is the computational problem 1-IN-3 SAT which is known to be NP-complete.*

Example 1.2.5. *Solving a finite set X of linear equations over a field $\mathbb{F}$ can also be understood as a CSP over the relational structure whose domain is F and which contains a relation for every linear equation from X; this relation is defined as the set of all tuples satisfying the corresponding equation. Let, e.g., $\mathbb{F} := \mathbb{Z}_2$ and let X be a set consisting of the single linear equation $x + y + z = 1$. Then the corresponding relation is ternary and contains the tuples $(1,0,0), (0,1,0), (0,0,1)$ and $(1,1,1)$. It is well-known that every such CSP can be solved in polynomial time by the Gaussian elimination.*

The fact that many natural computational problems can be expressed as CSPs led to a big interest in classifying the complexity of CSPs. In 1998, Feder and Vardi conjectured that any CSP whose template has a finite domain is either solvable in polynomial time (*tractable*) or NP-complete. This can be put into a contrast with the theorem of Ladner [65] stating that if P $\neq$ NP, then there are computational problems that are neither solvable in polynomial time nor NP-complete.

1.2.2 Clones and polymorphisms

In 2005, a connection between the computational complexity of CSPs over structures with finite domains and universal algebra was discovered [44]. It turned out that for any relational structure $\mathbb{A}$ with finite domain, the complexity of $\mathrm{CSP}(\mathbb{A})$ depends only on functions *preserving* all constraints of all instances of this CSP.

Let $n, m \geq 1$, let A be a set, let f be an m-ary function on A, and let $g_1, \ldots, g_m$ be n-ary functions on A. The mn-ary function $f(g_1, \ldots, g_m)$ on A is defined as follows.

$$f(g_1, \ldots, g_m)(x_1^1, \ldots, x_n^1, \ldots, x_1^m, \ldots, x_n^m) := f(g_1(x_1^1, \ldots, x_n^1), \ldots, g_m(x_1^m, \ldots, x_n^m))$$

Definition 1.2.6. *Let $n \geq 1$. We say that a constraint C of a* CSP *instance over a set A is* preserved *by a function $f\colon A^n \to A$ if for all $g_1, \ldots, g_n \in C$, we have $f(g_1, \ldots, g_n) \in C$.*

A relation R on a set A of arity $k \geq 1$ is preserved by a function $f\colon A^n \to A$ if the constraint $\{g \in A^{\{v_1,\ldots,v_k\}} \mid g(v_1, \ldots, v_k) \in R\}$ is preserved by f. In this case, we also say that R is *invariant* under f. Let $\mathbb{A}$ be a relational structure. Functions preserving all constraints of all instances of $\mathrm{CSP}(\mathbb{A})$ are called *polymorphisms* of $\mathbb{A}$. Note that it is also possible to define polymorphisms without using the notion of CSP. Indeed, it is easy to see that a function $f\colon A^n \to A$ is a polymorphism of $\mathbb{A}$ if, and only if, it is a homomorphism from $\mathbb{A}^n$ to $\mathbb{A}$. The set of all polymorphisms of a CSP template $\mathbb{A}$ is denoted by $\mathrm{Pol}(\mathbb{A})$; it is easy to observe that $\mathrm{Pol}(\mathbb{A})$ is a *function clone.*

Definition 1.2.7. *A set $\mathscr{C}$ of finitary operations on a fixed set C is called a* function clone *if both of the following hold.*

- *for every $k \geq 1$ and for every $i \in [k]$, $\mathscr{C}$ contains the k-ary* projection *on the i-th coordinate, i.e., the function $\pi_i^k\colon C^k \to C$ defined by $\pi_i^k(x_1, \ldots, x_k) := x_i$,*

- *for all $m, n \geq 1$, for every $f \in \mathscr{C}$ of arity m and for all $g_1, \dots, g_m \in \mathscr{C}$ of arity n, it holds that $f(g_1, \dots, g_m) \in \mathscr{C}$.*

For a function clone $\mathscr{C}$, we denote the domain of its functions by C; we say that $\mathscr{C}$ acts on *C. Let $n \geq 1$. We say that a relation $R \subseteq C^n$ is invariant under $\mathscr{C}$ if it is invariant under all functions from $\mathscr{C}$. The clone $\mathscr{C}$ also naturally acts (componentwise) on C^k for any $k \geq 1$, on any invariant subset S of C (by restriction), and on the classes of any invariant equivalence relation $\sim$ on an invariant subset S of C (by its action on representatives of the classes). We write $\mathscr{C} \curvearrowright C^k$, $\mathscr{C} \curvearrowright S$ and $\mathscr{C} \curvearrowright S/{\sim}$ for these actions. Any action $\mathscr{C} \curvearrowright S/{\sim}$ is called a* subfactor *of $\mathscr{C}$, and we also call the pair $(S, \sim)$ a subfactor. A subfactor $(S, \sim)$ is* minimal *if $\sim$ has at least two classes and no proper subset of S intersecting at least two $\sim$-classes is invariant under $\mathscr{C}$. For a clone $\mathscr{C}$ acting on a set X and $Y \subseteq X$ we write $\langle Y \rangle_{\mathscr{C}}$ for the smallest $\mathscr{C}$-invariant subset of X containing Y.*

The discovery of the connection between CSPs and universal algebra laid the foundations of the *algebraic approach to constraint satisfaction* that eventually led to a confirmation of the Feder-Vardi conjecture independently by Bulatov and Zhuk in 2017. Moreover, it turns out that the border between tractability and NP-completeness is given by the existence of a polymorphism of the template satisfying certain algebraic conditions.

Theorem 1.2.8 ([46, 83, 84])**.** *Let $\mathbb{A}$ be a relational structure over a finite domain and suppose that $P \neq NP$. Then precisely one of the following applies.*

- CSP($\mathbb{A}$) *is tractable, and there exists a 6-ary polymorphism s of $\mathbb{A}$ satisfying for every $a, b, c \in A$*
$$s(a, b, c, b, c, a) = s(b, c, a, a, b, c).$$
- CSP($\mathbb{A}$) *is NP-complete.*

1.2.3 Identities

The polymorphism from the first item of Theorem 1.2.8 is called a *Siggers polymorphism* as it is characterized by satisfying the *Siggers identity*. An *identity* is an equation of terms built from some functional symbols where all variables are implicitly interpreted as being universally quantified. For example, a 6-ary function s on some domain C *satisfies* the Siggers identity $s(x, y, z, y, z, x) = s(y, z, x, x, y, z)$ if the equation $s(a, b, c, b, c, a) = s(b, c, a, a, b, c)$ holds for all elements $a, b, c \in C$.

We list below some identities and operations that will be used in the thesis. Let $k \geq 1$, let C be a set, and let $f \colon A^k \to A$.

- f is *idempotent* if it satisfies the identity $f(x, \dots, x) = x$,
- f is a *weak near-unanimity (WNU)* operation if it satisfies the set of identities containing an equation for each pair of terms in $\{f(y, x, \dots, x), f(x, y, \dots, x), \dots, f(x, \dots, x, y)\}$,
- f is *totally symmetric* if it satisfies all identities of the form $f(x_1, \dots, x_k) = f(y_1, \dots, y_k)$ whenever $\{x_1, \dots, x_k\} = \{y_1, \dots, y_k\}$,

- f is a *quasi near-unanimity (qnu)* operation if it is a WNU operation and moreover, it satisfies $f(x, \ldots, x) = f(y, x, \ldots, x)$,
- f is a *near-unanimity* operation if it is an idempotent quasi near-unanimity operation,
- a ternary operation m is a *minority* if it satisfies $m(x, x, y) = m(x, y, x) = m(y, x, x) = y$,
- f is a *local near-unanimity* operation on a finite subset $F \subseteq C$ if it is a near-unanimity operation when restricted to F,
- a binary operation g on a two-element domain is a *semilattice operation* if it satisfies $g(x, y) = g(y, x) = g(x, x) = x$ for some enumeration $\{x, y\}$ of its domain.

Note that a local near-unanimity operation and a semilattice operation are not characterized by the satisfaction of a certain set of identities since we require them not only to satisfy a certain set of identities but we also make assumptions about the domains on which these identities should be satisfied.

Each set of identities also has a *pseudo*-variant obtained by composing each term appearing in the identities with a distinct unary function symbol. For example, a ternary operation f is a pseudo-WNU operation if there exist unary functions $e_1, \ldots, e_6$ on the same domain such that f together with these functions satisfies the identities: $e_1 \circ f(y, x, x) = e_2 \circ f(x, y, x)$, $e_3 \circ f(y, x, x) = e_4 \circ f(x, x, y)$ and $e_5 \circ f(x, y, x) = e_6 \circ f(x, x, y)$.

So far, we were interested in the satisfaction of some set of identities in the clone of polymorphisms of a particular CSP template. However, sometimes it makes sense to ask whether these identities are satisfied in some particular subclone of the polymorphism clone. This motivates the following notion. A function clone $\mathscr{C}$ *satisfies* a given set of identities if the function symbols which appear in the identities can be assigned functions in $\mathscr{C}$ in such a way that all identities in the set are satisfied by the assigned functions; if $\mathscr{U} \subseteq \mathscr{C}$ is a set of unary functions, then $\mathscr{C}$ *satisfies a set of pseudo-identities modulo* $\mathscr{U}$ if it satisfies the identities in such a way that the unary function symbols are assigned values in $\mathscr{U}$.

Unlike for the other identities, the k-ary idempotency identity is *trivial* for every $k \geq 1$, i.e. it is satisfied in every function clone – indeed, the k-ary projection to the first coordinate is certainly idempotent. Therefore, it makes sense only to ask whether all functions in a given function clone satisfy this identity, i.e. if the clone is *idempotent*. A set of identities that is not satisfied in every function clone is *non-trivial*. Remarkably, for clones over finite domains, the satisfaction of the Siggers identity is equivalent to the satisfaction of any non-trivial set of identities [77].

1.2.4 Infinite-domain CSPs

As we have seen in Section 1.2, many computational problems can be expressed as CSPs over relational structures with finite domains. However, many other natural problems from computer science can be phrased as CSPs only over templates with infinite domains. This is the case for some problems in artificial intelligence but also, e.g., for the *digraph acyclicity problem*.

Example 1.2.9. *Let us consider the relational structure* $(\mathbb{Q}; <)$. *This structure is easily seen to be universal for the class of all finite linear orders. Note moreover that a finite digraph* $\mathbb{D} = (D; E)$ *is acyclic if, and only if, it is possible to find a linear order* $<^*$ *on* D *such that for every* $(u, v) \in E$, *it holds that* $u <^* v$. *It follows that a finite digraph* $\mathbb{D}$ *is acyclic if, and only if, it has a homomorphism to* $(\mathbb{Q}; <)$ *(here, we identify the relational symbols* E *and* $<$*). Hence,* $\mathrm{CSP}(\mathbb{Q}; <)$ *is equivalent to the digraph acyclicity problem in the sense of the remark below Definition 1.2.2.*

It is therefore natural to study also CSPs over infinite domains. However, every computational problem is polynomial-time Turing-equivalent to a CSP over some template with infinite domain [20]. Hence, it is natural to restrict ourselves only to CSPs with infinite templates that satisfy some additional properties – the first natural assumption on the template is that every instance of its CSP has only finitely many solutions up to some property. This is the case for structures that are ω-categorical: For any instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of CSP over an ω-categorical relational structure $\mathbb{A}$, for any solution $f\colon \mathcal{V} \to A$ of $\mathcal{I}$ and for any $\alpha \in \mathrm{Aut}(\mathbb{A})$, αf is a solution of $\mathcal{I}$ as well. Hence, the number of solutions of $\mathcal{I}$ up to automorphisms of $\mathbb{A}$ is bounded by the number of orbits of $|\mathcal{V}|$-tuples under $\mathrm{Aut}(\mathbb{A})$ and is therefore finite.

Moreover, for an ω-categorical structure $\mathbb{A}$, the complexity of $\mathrm{CSP}(\mathbb{A})$ depends only on the polymorphisms of $\mathbb{A}$ [33]. Unfortunately, the sole assumption of ω-categoricity of the template is still insufficient to assess the complexity of its CSP [55, 56]. This is because even though every instance of such CSP has only finitely many solutions up to automorphisms, the set of possible solutions still does not need to be algorithmically enumerable. A natural assumptions on the relational structure $\mathbb{A}$ that guarantee the algorithmical enumerability of the solution set of its CSP are that $\mathbb{A}$ is homogeneous and finitely bounded. While homogeneity amounts to assuming that any solution is uniquely determined by the relations holding on its image, finite boundedness assures us that in order to decide whether a given map is a possible solution of some instance, it is necessary to verify if it satisfies certain conditions on subsets of the variable set of a fixed size.

Observe moreover that if $\mathbb{A}$ is a first-order reduct of a finitely bounded homogeneous structure $\mathbb{B}$, then the set of possible solutions of instances of $\mathrm{CSP}(\mathbb{A})$ can still be algorithmically enumerated in the sense of the previous paragraph since orbits under $\mathrm{Aut}(\mathbb{A})$ are unions of orbits under $\mathrm{Aut}(\mathbb{B})$. For first-order reducts of finitely bounded homogeneous structures, a generalization of the Feder-Vardi dichotomy conjecture was formulated by Bodirsky and Pinsker in 2011 [38].

Conjecture 1.2.10. *Let* $\mathbb{A}$ *be a first-order reduct of a finitely bounded homogeneous structure* $\mathbb{B}$ *and assume that* $\mathrm{P} \neq \mathrm{NP}$. *Then precisely one of the following applies.*

- $\mathrm{CSP}(\mathbb{A})$ *is* NP*-complete.*
- $\mathrm{CSP}(\mathbb{A})$ *is in P.*

Moreover, $\mathrm{CSP}(\mathbb{A})$ not being NP-complete implies that the polymorphism clone of a certain CSP template related to $\mathbb{A}$ has to satisfy the pseudo-Siggers identity. For more details, see Section 1.3.1. In this section, we will also present a more precise formulation of Conjecture 1.2.10. This conjecture was confirmed for first-order reducts of many finitely bounded homogeneous structures as well as for structures representing several classes of problems from computer science. We list some of the structures for which the conjecture was confirmed below.

- first-order reducts of finitely bounded homogeneous graphs [35, 28],
- first-order reducts of $(\mathbb{Q}; <)$ [23],
- first-order reducts of unary structures [30, 29],
- first-order reducts of the random poset [63],
- first-order reducts of the random tournament [69],
- first-order reducts of the homogeneous branching C-relation [21],
- structures representing all CSPs in the class MMSNP [26, 25],
- CSPs of representations of some relational algebras [24],
- CSPs of ω-categorical monadically stable structures [41].

In Chapter 3, we confirm Conjecture 1.2.10 for first-order reducts of certain k-uniform hypergraphs for every $k \geq 3$.

1.3 Universal algebra and reductions between CSPs

1.3.1 Primitive positive definability and clone homomorphisms

In this section, we present several notions from universal algebra and model theory and their connection to complexity reductions between CSPs.

One of the simplest and oldest reductions between CSPs is connected to the notion of primitive positive definability. A first-order formula is called *primitive positive* (pp) if it is built exclusively from atomic formulae, existential quantifiers, and conjunction. A relation is pp-definable in a relational structure $\mathbb{A}$ if it is first-order definable by a pp-formula. A relation is pp-definable in a relational structure $\mathbb{A}$ if, and only if, it is preserved by all its polymorphisms (see e.g. [78]).

Theorem 1.3.1 ([59]). *Let $\mathbb{A}$ be a relational structure over a finite signature and let $\mathbb{B}$ be a structure over the same domain such that every relation of $\mathbb{B}$ is pp-definable in $\mathbb{A}$. Then* $\mathrm{CSP}(\mathbb{B})$ *can be reduced to* $\mathrm{CSP}(\mathbb{A})$ *in polynomial time.*

Another easy reduction between CSPs comes from the notion of homomorphic equivalence. We say that two relational structures $\mathbb{A}$ and $\mathbb{B}$ in the same signature are *homomorphically equivalent* if there exists a homomorphism from $\mathbb{A}$ to $\mathbb{B}$ and a homomorphism from $\mathbb{B}$ to $\mathbb{A}$.

Theorem 1.3.2. *Let $\mathbb{A}$ and $\mathbb{B}$ be relational structures over the same finite relational signature and suppose that $\mathbb{A}$ and $\mathbb{B}$ are homomorphically equivalent. Then* $\mathrm{CSP}(\mathbb{B})$ *can be reduced to* $\mathrm{CSP}(\mathbb{A})$ *in polynomial time.*

Finally, the third basic reduction comes from the notion of a model-complete core. Let $\mathbb{A}$ be an ω-categorical relational structure. We say that $\mathbb{A}$ is a *model-complete* core if $\mathrm{Aut}(\mathbb{A})$ is dense in $\mathrm{End}(\mathbb{A})$, i.e., if for every $e \in \mathrm{End}(\mathbb{A})$ and for every finite set $F \subseteq A$, there exists $\alpha \in \mathrm{Aut}(\mathbb{A})$ such that $e|_F = \alpha|_F$. In any ω-categorical model-complete core $\mathbb{A}$, all orbits of n-tuples with respect to the automorphism group $\mathrm{Aut}(\mathbb{A})$ are pp-definable for all $n \geq 1$. For every ω-categorical relational structure $\mathbb{A}$, there exists an up to isomorphism unique ω-categorical structure $\mathbb{B}$ which is homomorphically equivalent to $\mathbb{A}$ and which is a model-complete core [16]; $\mathbb{B}$ is called the *model-complete core of* $\mathbb{A}$. Moreover, if we formulate CSP as a homomorphism problem as in Section 1.2.1, $\mathrm{CSP}(\mathbb{A})$ and $\mathrm{CSP}(\mathbb{B})$ are the same computational problems. Note that a relational structure $\mathbb{A}$ over a finite domain is a model-complete core if, and only if, $\mathrm{End}(\mathbb{A}) = \mathrm{Aut}(\mathbb{A})$. In this case, $\mathbb{A}$ is called a *core*. It also makes sense to generalize the notion of model-complete cores for clones – a clone $\mathscr{C}$ is a *model-complete core* if for every unary function $f \in \mathscr{C}$ and for every finite $F \subseteq C$, there exists $\alpha \in \mathrm{Aut}(\mathbb{A})$ such that $f|_F = \alpha|_F$. Note that a relational structure $\mathbb{A}$ is a model-complete core if, and only if, $\mathrm{Pol}(\mathbb{A})$ is.

However, one may wonder if it is possible to find a reduction that generalizes all the reductions above. A positive answer was given in [12] where the notion of pp-construction was introduced. We say that a structure $\mathbb{B}$ has a *pp-construction* in a structure $\mathbb{A}$ if $\mathbb{B}$ is homomorphically equivalent to a structure with domain A^n, for some $n \geq 1$, whose relations are pp-definable in $\mathbb{A}$ (for this purpose, a k-ary relation on A^n is regarded as a kn-ary relation on A).

Moreover, pp-constructibility of a relational structure $\mathbb{B}$ in a relational structure $\mathbb{A}$ can be characterized algebraically by the existence of a certain mapping from $\mathrm{Pol}(\mathbb{A})$ to $\mathrm{Pol}(\mathbb{B})$.

Definition 1.3.3. *Let $\mathscr{C}$ and $\mathscr{D}$ be function clones. A map $\xi\colon \mathscr{C} \to \mathscr{D}$ is a* clone homomorphism *if it satisfies all of the following.*

- *ξ preserves arities, i.e., for every $f \in \mathscr{C}$ of arity $n \geq 1$, $\xi(f)$ has arity n.*
- *ξ preserves projections, i.e., it maps every projection in $\mathscr{C}$ to the corresponding projection in $\mathscr{D}$.*
- *ξ preserves compositions, i.e., for all $n, m \geq 1$, for every $f \in \mathscr{C}$ of arity n and for all $g_1, \ldots, g_n \in \mathscr{C}$ of arity m, $\xi(f(g_1, \ldots, g_m)) = \xi(f)(\xi(g_1), \ldots, \xi(g_n))$.*

We say that ξ is a minion homomorphism *if it preserves arities and compositions with projections, i.e., it satisfies $\xi(f \circ (\pi_1, \ldots, \pi_n)) = \xi(f) \circ (\pi_1, \ldots, \pi_n)$ for all $n, m \geq 1$ and all n-ary $f \in \mathscr{C}$ and m-ary projections $\pi_1, \ldots, \pi_n$.*

Note that every clone homomorphism is a minion homomorphism. Moreover, it is easy to see that clone homomorphisms preserves identities, i.e., if $\xi\colon \mathscr{C} \to \mathscr{D}$ is a clone homomorphism and a particular set of identities is satisfies in $\mathscr{C}$, then the same set of identities is satisfied in $\mathscr{D}$. Indeed, if some functions $f_1, \ldots, f_n \in \mathscr{C}$ witness the satisfaction of a given set of identities in $\mathscr{C}$, the functions $\xi(f_1), \ldots, \xi(f_n)$ witness the satisfaction of the same set of identities in $\mathscr{D}$. Similarly, minion homomorphisms preserve the satisfaction of *identities of height 1* (h1-identities), i.e., identities that do not contain any nesting of functional symbols.

Definition 1.3.4. *Let $\mathscr{C}$ and $\mathscr{D}$ be function clones and suppose that the domain of $\mathscr{D}$ is finite. We say that a minion homomorphism $\xi\colon \mathscr{C} \to \mathscr{D}$ is* uniformly continuous *if for all $n \geq 1$, there exists a finite subset F of C^n such that $\xi(f) = \xi(g)$ for all n-ary functions $f, g \in \mathscr{C}$ which agree on F.*

Clearly, if the domain of $\mathscr{C}$ is finite, then every minion homomorphism from $\mathscr{C}$ is uniformly continuous. Now, we introduced all notions that are needed to state the algebraic characterization of pp-constructability.

Theorem 1.3.5 ([12]). *Let $\mathbb{A}$ be an at most countable ω-categorical relational structure and let $\mathbb{B}$ be a relational structure with a finite domain. Then the following are equivalent.*

- *$\mathbb{B}$ has a pp-construction in $\mathbb{A}$.*
- *There exists a uniformly continuous minion homomorphism $\xi\colon \mathrm{Pol}(\mathbb{A}) \to \mathrm{Pol}(\mathbb{B})$.*

Moreover, if any of the items applies, then CSP($\mathbb{B}$) *can be reduced to* CSP($\mathbb{A}$) *in polynomial time.*

Now, we can formulate the precise statement of Conjecture 1.2.10 based on the recent progress from [8, 7]. We will denote by $\mathscr{P}$ the clone of projections on a two-element domain.

Conjecture 1.3.6. *Let $\mathbb{A}$ be a* CSP *template which is a first-order reduct of a finitely bounded homogeneous structure. Then one of the following applies.*

- Pol($\mathbb{A}$) *has a uniformly continuous minion homomorphism to $\mathscr{P}$, and* CSP($\mathbb{A}$) *is* NP-*complete.*
- Pol($\mathbb{A}$) *has no uniformly continuous minion homomorphism to $\mathscr{P}$, and* CSP($\mathbb{A}$) *is in* P.

Observe that for every CSP template $\mathbb{B}$ over a two-element domain, $\mathscr{P} \subseteq \mathrm{Pol}(\mathbb{B})$. Hence, the first item of Conjecture 1.3.6 corresponds to every CSP template over a two-element domain being pp-constructible in $\mathbb{A}$. Since there are NP-complete CSPs on two-element domain (e.g., the CSP from Example 1.2.4), this implies that if Pol($\mathbb{A}$) has a uniformly continuous minion homomorphism to $\mathscr{P}$, then CSP($\mathbb{A}$) is NP-complete. Hence, in order to prove Conjecture 1.3.6, it remains to prove that if Pol($\mathbb{A}$) has no uniformly continuous minion homomorphism to $\mathscr{P}$, then CSP($\mathbb{A}$) is in P. By the results from [16] mentioned above, it is enough to prove it for the model-complete core $\mathbb{B}$ of $\mathbb{A}$. For an ω-categorical model-complete core $\mathbb{B}$, it is known that if Pol($\mathbb{B}$) does not have a uniformly continuous minion homomorphism to $\mathscr{P}$, then the pseudo-Siggers identity is satisfied in Pol($\mathbb{B}$) [13].

1.3.2 Canonical functions

As we mentioned in Section 1.2.4, Conjecture 1.2.10 has been confirmed for many CSP templates under consideration. In all proofs, a reduction to some finite domain CSP is used. In many of these reductions, so-called canonical functions play a key role.

Definition 1.3.7. *Let $k, n \geq 1$, let $\mathscr{G}$ be a permutation group acting on a set C, and let $f: C^k \to C$. We say that f is* n-canonical *with respect to $\mathscr{G}$ if for all $a_1, \ldots, a_k \in C^n$ and all $\alpha_1, \ldots, \alpha_k \in \mathscr{G}$ there exists $\beta \in \mathscr{G}$ such that $f(a_1, \ldots, a_k) = \beta \circ f(\alpha_1(a_1), \ldots, \alpha_k(a_k))$. A function that is n-canonical with respect to $\mathscr{G}$ for all $n \geq 1$ is called* canonical *with respect to $\mathscr{G}$. We say that a function is* diagonally canonical *if it satisfies the definition of n-canonicity in the case $\alpha_1 = \cdots = \alpha_k$ for every $n \geq 1$.*

In particular, a canonical function f induces an operation on the set $C^n/\mathscr{G}$ of orbits of n-tuples under $\mathscr{G}$. If all functions of a function clone $\mathscr{C}$ are n-canonical with respect to $\mathscr{G}$, then $\mathscr{C}$ acts on $C^n/\mathscr{G}$ and we write $\mathscr{C}^n/\mathscr{G}$ for this action; if the permutation group $\mathscr{G}$ is oligomorphic, then $\mathscr{C}^n/\mathscr{G}$ is a function clone on a finite set.

We write $\mathscr{G}_{\mathscr{C}}$ to denote the largest permutation group contained in a function clone $\mathscr{C}$, and say that $\mathscr{C}$ is oligomorphic if $\mathscr{G}_{\mathscr{C}}$ is oligomorphic. For $n \geq 1$, the *n-canonical (canonical) part* of $\mathscr{C}$ is the clone of those functions of $\mathscr{C}$ which are n-canonical (canonical) with respect to $\mathscr{G}_{\mathscr{C}}$. We write $\mathscr{C}_n^{\mathrm{can}}$ and $\mathscr{C}^{\mathrm{can}}$ for these sets, which form themselves function clones.

Let $k \geq 1$ and let $\mathbb{B}$ be a finitely bounded k-homogeneous structure. For a first-order reduct $\mathbb{A}$ of $\mathbb{B}$, the polymorphisms of $\mathbb{A}$ which are canonical with respect to $\mathbb{B}$ act on the orbits of k-tuples under $\mathrm{Aut}(\mathbb{B})$, for every $k \geq 1$. Moreover, by the k-homogeneity, for any instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of $\mathrm{CSP}(\mathbb{A})$, every solution $f: \mathcal{V} \to A$ of $\mathcal{I}$ is up to $\mathrm{Aut}(\mathbb{B})$ determined by the orbits of the images of k-tuples of variables from f. This means that we can transform every instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of $\mathrm{CSP}(\mathbb{A})$ into an instance having one variable for every k-tuple of variables of $\mathcal{V}$ and such that these variables are meant to take values in the set of orbits of k-tuples under $\mathrm{Aut}(\mathbb{B})$. If we can solve every such finite instance in polynomial time, then we can solve the original CSP in polynomial time. This idea was formalized in [29, 30].

Theorem 1.3.8. *Let $\mathbb{A}$ be a first-order reduct of a finitely bounded homogeneous structure $\mathbb{B}$ and let $\mathscr{C} \subseteq \mathrm{Pol}(\mathbb{A})$ be the clone of all polymorphisms of $\mathbb{A}$ which are canonical with respect to* $\mathrm{Aut}(\mathbb{B})$. *Then both of the following hold.*

- *If $\mathscr{C}$ satisfies the pseudo-Siggers identity modulo* $\mathrm{Aut}(\mathbb{B})$, *then* $\mathrm{CSP}(\mathbb{A})$ *is tractable.*
- *If $\mathscr{C}$ contains quasi-WNU operations modulo* $\mathrm{Aut}(\mathbb{B})$ *of all arities $n \geq 3$, then* $\mathrm{CSP}(\mathbb{A})$ *has bounded width.*

In Chapter 2, we will introduce a variant of the reduction from [29, 30] that will enable us to prove a bound on the relational width of first-order reduct $\mathbb{A}$ of a finitely bounded homogeneous structure $\mathbb{B}$ which has canonical pseudo-WNU operations modulo $\mathrm{Aut}(\mathbb{B})$ of all arities $n \geq 3$ among its polymorphisms.

In order to make use of Theorem 1.3.8, we first need to show that a first-order reduct $\mathbb{A}$ of a finitely bounded homogeneous structure $\mathbb{B}$ under consideration possesses the desired canonical polymorphisms. In order to do that, a *Ramsey expansion* of the structure $\mathbb{B}$ is often used.

Let $\mathbb{A}$ and $\mathbb{B}$ be relational structures in the same signature. We write $\binom{\mathbb{B}}{\mathbb{A}}$ for the set of all embeddings of $\mathbb{A}$ into $\mathbb{B}$. We say that a class $\mathcal{S}$ of finite relational structures in the same signature has the *Ramsey property* if for every $r \geq 1$ and for all $\mathbb{A}, \mathbb{B} \in \mathcal{S}$ there exists $\mathbb{C} \in \mathcal{S}$ such that for every mapping $\chi: \binom{\mathbb{C}}{\mathbb{A}} \to [r]$ there exists $f \in \binom{\mathbb{C}}{\mathbb{B}}$ such that $|\chi(f \circ \binom{\mathbb{B}}{\mathbb{A}})| \leq 1$. We say that a homogeneous structure $\mathbb{B}$ is *Ramsey* if its *age*, i.e., the set of all finite structures that embed into $\mathbb{B}$, has the Ramsey property. For more details on Ramsey structures, see e.g. [17]. It is an open problem whether every finitely bounded homogeneous structure has a first-order expansion which is Ramsey and in general, it is not easy to show that a particular structure has one – the most general results that have proven to be useful in the context of CSP come from [76, 49].

For a set of functions $\mathscr{F}$ over the same fixed set C we write $\overline{\mathscr{F}}$ for the set of those functions g such that for all finite subsets $F \subseteq C$, there exists a function in $\mathscr{F}$ which agrees with g on F. For $k \geq 1$, for k-ary functions f, g and for a permutation group $\mathscr{G}$ such that f, g and $\mathscr{G}$ act on the same domain, we say that f *locally interpolates* g modulo $\mathscr{G}$ if $g \in \overline{\{\beta \circ f(\alpha_1, \dots, \alpha_k) \mid \beta, \alpha_1, \dots, \alpha_k \in \mathscr{G}\}}$. Similarly, we say that f *diagonally interpolates* g modulo $\mathscr{G}$ if f locally interpolates g with $\alpha_1 = \dots = \alpha_k$. If $\mathscr{G}$ is the automorphism group of a *Ramsey structure*, then every function on its domain locally interpolates a canonical function modulo $\mathscr{G}$, and diagonally interpolates a diagonally canonical function modulo $\mathscr{G}$ [39, 37]. Intuitively speaking, this means that Ramsey structures have many canonical polymorphisms. We say that a clone $\mathscr{D}$ locally interpolates a clone $\mathscr{C}$ modulo a permutation group $\mathscr{G}$ if for every $g \in \mathscr{D}$ there exists $f \in \mathscr{C}$ such that g locally interpolates f modulo $\mathscr{G}$.

1.4 Local consistency

Local consistency checking plays a prominent role in the area of constraint satisfaction. In this section, we introduce two different notions of local consistency - *relational width* and *strict width*. First, we need to introduce a few related concepts, in particular the notion of minimality.

A CSP instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ over a set A is *non-trivial* if it does not contain any empty constraint; otherwise, it is *trivial*. Given a constraint $C \subseteq A^U$ and a tuple $\mathbf{v} \in U^k$ for some $k \geq 1$, the *projection of* C *onto* $\mathbf{v}$ is defined by $\text{proj}_{\mathbf{v}}(C) := \{f(\mathbf{v}) \colon f \in C\}$. Let $U \subseteq \mathcal{V}$. We define the *restriction* of $\mathcal{I}$ to U to be an instance $\mathcal{I}|_U = (U, \mathcal{C}|_U)$ where the set of constraints $\mathcal{C}|_U$ contains for every $C \in \mathcal{C}$ the constraint $C|_U = \{g|_U \mid g \in C\}$.

We denote by $\text{CSP}_{\text{Inj}}(\mathbb{A})$ the restriction of $\text{CSP}(\mathbb{A})$ to those instances of $\text{CSP}(\mathbb{A})$ where for every constraint C and for every pair of distinct variables u, v in its scope, $\text{proj}_{(u,v)}(C) \subseteq I_2^A$.

1.4.1 Minimality

Definition 1.4.1. *Let $1 \leq m \leq n$. We say that an instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of* $\text{CSP}(\mathbb{A})$*is* (m, n)-minimal *if both of the following hold:*

- *every at most n-element subset of the variable set $\mathcal{V}$ is contained in the scope of some constraint in $\mathcal{I}$;*
- *for every at most m-element subset of variables $K \subseteq \mathcal{V}$ and for any two constraints $C_1, C_2 \in \mathcal{I}$ whose scopes contain K, the restrictions of C_1 and C_2 to K coincide.*

*We say that an instance $\mathcal{I}$ is m-*minimal *if it is (m, m)-minimal.*

Let $1 \leq m \leq n$. Note that an instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ is (m, n)-minimal if, and only if, for every $\mathbf{u} \in \mathcal{V}^n$, all variables of $\mathbf{u}$ are contained in the scope of some constraint, and for every $\mathbf{v} \in \mathcal{V}^m$

and for all $C_1, C_2 \in \mathcal{C}$ containing all variables of $\mathbf{v}$ in their scopes, $\mathrm{proj}_{\mathbf{v}}(C_1) = \mathrm{proj}_{\mathbf{v}}(C_2)$. It follows that if $\mathcal{I}$ is an m-minimal instance and $\mathbf{v}$ is a tuple of variables of length at most m, then there exists a constraint of $\mathcal{I}$ whose scope contains $\mathbf{v}$, and all the constraints who do have the same projection on $\mathbf{v}$. We write $\mathrm{proj}_{\mathbf{v}}(\mathcal{I})$ for this projection, and call it the *projection of* $\mathcal{I}$ *onto* $\mathbf{v}$.

Let $1 \leq m \leq n$, let $\mathbb{A}$ be a relational structure, and let p denote the maximum of n and the maximal arity of the relations of $\mathbb{A}$. Clearly not every instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of $\mathrm{CSP}(\mathbb{A})$ is (m,n)-minimal. However, every instance $\mathcal{I}$ is *equivalent* to an (m,n)-minimal instance $\mathcal{I}'$ of $\mathrm{CSP}(\mathbb{A}')$ where $\mathbb{A}'$ is the expansion of $\mathbb{A}$ by all at most p-ary relations pp-definable in $\mathbb{A}$ in the sense that $\mathcal{I}$ and $\mathcal{I}'$ have the same set of solutions. In particular we have that if $\mathcal{I}'$ is trivial, then $\mathcal{I}$ has no solutions. Moreover, $\mathrm{CSP}(\mathbb{A}')$ has the same complexity as $\mathrm{CSP}(\mathbb{A})$ by Theorem 1.3.1.

Moreover, if $\mathbb{A}$ is ω-categorical, then the instance $\mathcal{I}'$ can be computed in only polynomially many steps in the number of variables and constraints as follows. First introduce a new constraint A^L for every set $L \subseteq \mathcal{V}$ with at most n elements to satisfy the first condition. Then remove orbits with respect to $\mathrm{Aut}(\mathbb{A})$ from the constraints in the instance as long as the second condition is not satisfied.

To see that $\mathcal{I}'$ is an instance of $\mathrm{CSP}(\mathbb{A}')$, observe that the relation A^L is pp-definable in $\mathbb{A}$ for every $L \subseteq \mathcal{V}$ and hence, the instance created from $\mathcal{I}$ by adding the new constraints is still an instance of $\mathrm{CSP}(\mathbb{A})$. Moreover, when the algorithm removes orbits from some constraint, then the new constraint is obtained by a pp-definition from two old constraints and hence, it is pp-definable in $\mathbb{A}$. Whence, the instance $\mathcal{I}'$ is an instance of $\mathrm{CSP}(\mathbb{A}')$.

The algorithm just described introduces at most $\sum_{i=1}^{n} \binom{|\mathcal{V}|}{i}$ new constraints and it uses only relations of arity at most p. Note that by the ω-categoricity of $\mathbb{A}$, every relation pp-definable in $\mathbb{A}$ is a finite union of orbits of tuples with respect to $\mathrm{Aut}(\mathbb{A})$. Hence, since in every step the algorithm removes at least one orbit from some constraint and since the number of orbits that can be removed is polynomial in $|\mathcal{V}| + |\mathcal{C}|$ by the discussion above, the number of steps of the algorithm is polynomial in $|\mathcal{V}| + |\mathcal{C}|$.

In the following section, CSPs for which every sufficiently minimal instance has a solution will be introduced. However, the above algorithm producing an (m,n)-minimal instance can also be used as a subroutine in more complicated algorithms for solving certain CSPs. This is the case also for the algorithm from Chapter 3.

1.4.2 Relational width

Definition 1.4.2. *Let* $1 \leq m \leq n$. *A relational structure* $\mathbb{A}$ *has* relational width (m,n) *if every non-trivial* (m,n)*-minimal instance equivalent to an instance of* $\mathrm{CSP}(\mathbb{A})$ *has a solution.* $\mathbb{A}$ *has* bounded width *if it has relational width* (m,n) *for some* m,n.

If $\mathbb{A}$ has relational width (m,n), then we will also say that $\mathrm{CSP}(\mathbb{A})$ has relational width (m,n). We say that $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$ has relational width (m,n) if every non-trivial (m,n)-minimal instance equivalent to an instance of $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$ has a solution.

Example 1.4.3. *The relational structure* $\mathbb{A} := (\mathbb{Q}; <)$ *has relational width* $(2,3)$ *but not relational width* $(2,2)$ *or* $(1,k)$ *for any* $k \geq 1$.

To show that $\mathbb{A}$ *does not have relational width* $(1,k)$ *for any* $k \geq 1$*, let us consider the instance* $\mathcal{I} = (\{u,v\},\mathcal{C})$ *of* $\mathrm{CSP}(\mathbb{A})$ *where* $\mathcal{C}$ *contains the following two constraints:* $C_1 := \{f \in \mathbb{Q}^{\{u,v\}} \mid f(u) < f(v)\}$ *and* $C_2 := \{f \in \mathbb{Q}^{\{u,v\}} \mid f(u) > f(v)\}$*. It follows that both* C_1 *and* C_2 *are non-empty,* $\mathrm{proj}_u(C_1) = \mathrm{proj}_u(C_2) = \mathrm{proj}_v(C_1) = \mathrm{proj}_v(C_2) = \mathbb{Q}$ *and hence,* $\mathcal{I}$ *is* $(1,k)$*-minimal for any* $k \geq 1$*. However, it has obviously no solution.*

To show that $\mathbb{A}$ *does not have relational width* $(2,2)$*, let us consider the instance* $\mathcal{I}' = (\{u,v,w\},\mathcal{C}')$ *of* $\mathrm{CSP}(\mathbb{A})$ *such that* $\mathcal{C}'$ *contains for every* $(a,b) \in \{(u,v),(v,w),(w,u)\}$ *the constraint* $C_{a,b} := \{f \in A^{\{a,b\}} \mid f(a) < f(b)\}$*. This instance is easily seen to be* $(2,2)$*-minimal and non-trivial but it does not have any solution.*

Finally, it is easy to see and well-known that $\mathbb{A}$ *has relational width* $(2,3)$*.*

The following statement gives an overview of conditions characterizing relational structures with bounded width that are used in several proofs later. Let $p \geq 2$ be a prime number and let R_0 and R_1 be the relations defined by $R_i := \{(x,y,z) \in \mathbb{Z}_p \mid x+y+z = i \bmod p\}$ for $i \in \{0,1\}$. An *affine clone* is the clone of affine maps over a finite module. We say that a function clone $\mathscr{C}$ is *equationally affine* if it has a clone homomorphism to an affine clone. A *Datalog program* is a set of formulas of the form $\phi_1(\mathbf{x}_1) \wedge \cdots \wedge \phi_n(\mathbf{x}_n) \Rightarrow \phi_0(\mathbf{x}_0)$ where ϕ_i is an atomic formula over free variables from the tuple $\mathbf{x}_i$ for every $i \in \{0,\ldots,n\}$.

Theorem 1.4.4. *Let* $\mathbb{A}$ *be an* ω*-categorical relational structure. All the implications* $(i) \Rightarrow (j)$ *and* $(i) \Leftrightarrow (i')$ *in the list below hold for* $1 \leq i \leq j \leq 3$*.*

1. $\mathbb{A}$ *has relational width* $(2,3)$*.*

2. $\mathbb{A}$ *has bounded width.*

(2') *The class of finite structures that do not have a homomorphism to* $\mathbb{A}$ *is definable by a Datalog program.*

3. $\mathrm{Pol}(\mathbb{A})$ *does not admit a uniformly continuous minion homomorphism to an affine clone.*

(3') $\mathbb{A}$ *does not pp-construct* $(\mathbb{Z}_p; R_0, R_1)$ *for any prime number* p*.*

If $\mathbb{A}$ *is finite, then all these conditions are equivalent, and moreover equivalent to the following two statements.*

4. *The expansion* $\mathbb{B}$ *of the core of* $\mathbb{A}$ *by all unary singleton relations is such that* $\mathrm{Pol}(\mathbb{B})$ *is not equationally affine.*

5. $\mathrm{Pol}(\mathbb{A})$ *contains WNU operations of all arities* $n \geq 3$*.*

The implication from (1) to (2) is trivial, the equivalence of (2) and (2') follows from [53, Theorem 23] and [47, Corollary 1]. (2) implies (3) by [66] and [12], and the equivalence of (3) and (3') is a folklore consequence of [12]. For finite structures, the implication from (4) to (1) was proven in [6], (3) implies (4) by [12], and (4) is equivalent to (5) by the results from [67, 64] (in fact, a slightly weaker statement is formulated there; the precise statement made here is attributed to E. Kiss in [64, Theorem 2.8], and another proof can be found in [85]).

If $\mathbb{A}$ is an ω-categorical model-complete core, then it is known that the clone of polymorphisms of $\mathbb{A}$ which are canonical with respect to $\overline{\mathrm{Aut}(\mathbb{A})}$ is either equationally affine, or it contains pseudo-WNU operations modulo $\overline{\mathrm{Aut}(\mathbb{A})}$ of all arities $n \geq 3$ [38, 69].

The following theorem shows that finite CSP templates which have relational width $(1, 1)$ can be characterized algebraically. In fact, Theorem 1.4.4 together with the results from [53, 50] imply that a finite CSP template with bounded width has always relational width $(1, 1)$ or $(2, 3)$.

Theorem 1.4.5 ([51, 53])**.** *Let $\mathbb{A}$ be a finite relational structure. $\mathbb{A}$ has relational width $(1, 1)$ if, and only if, $\mathrm{Pol}(\mathbb{A})$ contains totally symmetric operations of all arities.*

We have just seen that every finite CSP template with bounded width has relational width $(2, 3)$ and that such templates can be characterized algebraically by Item 5 of Theorem 1.4.4. Unfortunately, none of those is true for first-order reducts of finitely bounded homogeneous structures. In fact, already the first-order reducts of $(\mathbb{Q}, <)$ with bounded width cannot be characterized algebraically [34]. The following example exhibits a class of finitely bounded homogeneous structures with arbitrarily large relational width.

Example 1.4.6. *Let $k \geq 2$, let $\mathbb{H} = (H; E)$ be the universal homogeneous k-uniform hypergraph, and let $\mathbb{A} = (H; E, N)$ be an expansion of $\mathbb{H}$, where $N := I_k^H \backslash E$. We will show that $\mathbb{A}$ does not have relational width $(k - 1, n)$ for any $n \geq k - 1$ but it has relational width $(k, k + 1)$.*

To see that $\mathbb{A}$ does not have relational width $(k - 1, n)$ for any $n \geq k - 1$, let $\mathcal{I} = (\{v_1, \dots, v_k\}, \mathcal{C})$ be the instance of $\mathrm{CSP}(\mathbb{A})$ where $\mathcal{C}$ contains two constraints $C_1 := \{f \in H^{\{v_1,\dots,v_k\}} \mid (f(v_1), \dots, f(v_k)) \in E\}$ and $C_2 := \{f \in H^{\{v_1,\dots,v_l\}} \mid (f(v_1), \dots, f(v_k)) \in N\}$. It immediately follows that $\mathcal{I}$ is $(k - 1, n)$-minimal for every $n \geq k - 1$ but it does not have any solution.

On the other hand, it can be easily seen that $\mathrm{Pol}(\mathbb{A})$ contains pseudo-symmetric operations of all arities which are canonical with respect to $\mathrm{Aut}(\mathbb{H})$ and Theorem 2.1.2 yields that $\mathbb{A}$ has relational width $(k, k + 1)$.

In Chapter 2, we find an upper bound on the relational width of first-order reducts of k-homogeneous ℓ-bounded structures for every $k, \ell \geq 1$ that depends only on k and ℓ under the assumption that these first-order reducts satisfy certain algebraic conditions.

1.4.3 Strict width

While the notion of relational width describes CSPs where locally consistent instances are guaranteed to have a solution, it makes sense to strengthen this assumption and require that for instances that are sufficiently locally consistent, every local solution extends to a global one. This concept is formalized in the notion of *strict width* which was introduced by Feder and Vardi [53]. Let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be an instance of CSP over a set A, and let $\mathcal{V}' \subseteq \mathcal{V}$. We say that a mapping $g\colon \mathcal{V}' \to A$ is a *partial solution* of $\mathcal{I}$, if for every $C \in \mathcal{C}$ with scope $U \subseteq \mathcal{V}$, it holds that $g|_{U \cap \mathcal{V}'} \in C|_{U \cap \mathcal{V}'}$.

Definition 1.4.7. *Let $\mathbb{A}$ be a relational structure and let $m \geq 1$. We say that $\mathbb{A}$ has* strict width *m if there exists $n > m$ such that for every (m, n)-minimal instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of*

CSP($\mathbb{A}$), *for every* $\mathcal{V}' \subseteq \mathcal{V}$, *and for every partial solution* $g\colon \mathcal{V}' \to A$ *of* $\mathcal{I}$, *there exists a solution* $f\colon \mathcal{V} \to A$ *of* $\mathcal{I}$ *such that* $f|_{\mathcal{V}'} = g$. *We say that* $\mathbb{A}$ *has* bounded strict width *if it has strict width* m *for some* $m \geq 1$.

Moreover, for any $m \geq 1$, finite CSP templates with strict width m can be characterized algebraically as follows.

Theorem 1.4.8 ([53]). *Let* $\mathbb{A}$ *be a relational structure with finite domain and let* $m \geq 2$. *Then the following are equivalent.*

- $\mathbb{A}$ *has strict width* m.
- $\mathrm{Pol}(\mathbb{A})$ *contains a near-unanimity operation of arity* $m+1$.

However, in contrast with the case of relational width where Theorem 1.4.4 implies that every finite template with bounded width has relational width $(2,3)$, the following example shows that no bound on the amount of strict width needed to ensure that a CSP template has bounded strict width can be found already for templates over a 2-element domain.

Example 1.4.9. *Let* $m \geq 3$ *and let* $\mathbb{A} = (\{0,1\}; R)$ *where* $R := \{0,1\}^m \backslash (1,\ldots,1)$. *We claim that* $\mathbb{A}$ *has strict width* m *but not* $m-1$. *To this end, we will first show that* $\mathbb{A}$ *does not have strict width* $m-1$. *Let us suppose that this is the case and let* $f\colon \{0,1\}^m \to \{0,1\}$ *be a near-unanimity polymorphism of* $\mathbb{A}$ *that exists by Theorem 1.4.8. We have*

$$\begin{pmatrix}0\\1\\1\\\cdots\\1\end{pmatrix}, \begin{pmatrix}1\\0\\1\\\cdots\\1\end{pmatrix}, \ldots, \begin{pmatrix}1\\1\\\cdots\\1\\0\end{pmatrix} \in R.$$

and applying f *to the rows yields that* $(1,\ldots,1) \in R$, *which is a contradiction.*

On the other hand, consider the operation $g\colon \{0,1\}^{m+1} \to \{0,1\}$ *defined by*

$$g(a_1,\ldots,a_{m+1}) := \begin{cases} 1, & \text{if there exists at most one } i \in [m+1] \text{ such that } a_i = 0,\\ 0, & \text{otherwise.}\end{cases}$$

It is easy to see that $g \in \mathrm{Pol}(\mathbb{A})$ *and* g *is a near-unanimity operation of arity* $m+1$. *Theorem 1.4.8 implies that* $\mathbb{A}$ *has strict width* m.

On the other hand, strict width is one of the few concepts from finite-domain CSP where the algebraic characterization can be lifted to CSPs over ω-categorical templates as follows. Let $\mathbb{A}$ be an ω-categorical relational structure. We say that $f \in \mathrm{Pol}(\mathbb{A})$ is *oligopotent* if for every finite subset $B \subseteq A$, there exists $\alpha \in \mathrm{Aut}(\mathbb{A})$ such that $f(b,\ldots,b) = \alpha(b)$ for every $b \in B$.

Theorem 1.4.10 ([19]). *Let* $\mathbb{A}$ *be an* ω*-categorical relational structure and let* $m \geq 2$. *Then the following are equivalent.*

- $\mathbb{A}$ *has strict width* m,
- $\mathrm{Pol}(\mathbb{A})$ *contains an oligopotent quasi near-unanimity operation of arity* $m+1$,
- *for every finite subset* $F \subseteq A$, $\mathrm{Pol}(\mathbb{A})$ *contains a local near-unanimity operation on* F *of arity* $m+1$.

In Chapter 4, we will use this characterization to show a bound on relational width of first-order reducts of certain finitely bounded homogeneous structures with bounded strict width.

2 Smooth approximations and relational width collapses

2.1 Introduction

Local consistency checking is an algorithmic technique that is central in computer science. Intuitively speaking, it consists in propagating local information through a structure so as to infer global information: consider, for example, computing the transitive closure of a relation as deriving global information from local one. Local consistency checking has a prominent role in the area of constraint satisfaction – the local consistency algorithm can be used to decrease the size of the search space efficiently or even to correctly solve some constraint satisfaction problems in polynomial time (for example, 2-SAT or Horn-SAT). However, the use of local consistency methods is not limited to constraint satisfaction. Indeed, local consistency checking is also used for such problems as the graph isomorphism problem, where it is is known as the Weisfeiler-Leman algorithm. Again, the technique can be used to derive implied constraints that an isomorphism between two graphs has to satisfy so as to narrow down the search space, but local consistency is in fact powerful enough to solve the graph isomorphism problem over any non-trivial minor-closed class of graphs [57]. Notably, the best algorithm for graph isomorphism to date also uses local consistency as a subroutine [5]. Finally, local consistency can be used to solve games involved in formal verification such as parity games and mean-payoff games [27].

One of the reasons for the ubiquity of local consistency is that its underlying principles can be described in many different languages, such as the language of category theory [1], in the language of finite model theory (by Spoiler-Duplicator games [62] or by homomorphism duality [3]), and logical definability (in Datalog, or infinitary logics with bounded number of variables). For constraint satisfaction problems over a finite template $\mathbb{A}$, the power of local consistency checking can additionally be characterised algebraically (see Theorem 1.4.4). Moreover, whenever local consistency correctly solves CSP($\mathbb{A}$), then in fact only a very restricted form of local consistency checking is needed. This fact is known as the *collapse of the bounded width hierarchy*, and it has strong consequences both for complexity and logic. On the one hand, the collapse gives efficient algorithms that are able to solve *all* the CSPs that are solvable by local consistency methods, and in fact this gives a polynomial-time algorithm solving instances of the *uniform CSP*, where the template $\mathbb{A}$ is also part of the input. On the other hand, this collapse induces collapses in all the areas mentioned at the beginning of this paragraph.

As we have seen in Section 1.2.4, many natural problems from computer science can only be phrased as CSPs where the template is infinite. This is the case for linear programming, some reasoning problems in artificial intelligence such as ontology-mediated querying, or even problems as simple to formulate as the digraph acyclicity problem. In order to understand the power of local consistency in more generality it is thus necessary to consider its use for infinite-domain CSPs. Infinite-domain CSPs with an ω-categorical template form a very general class of problems for which the algebraic approach from the finite case can be extended, and numerous results in the recent years have shown the power of this approach. An algebraic characterisation of local consistency checking for infinite-domain CSPs is, however, missing. In fact, the negative results of [31, 32], refined in [55], show that no purely algebraic description of local consistency is possible for CSPs with ω-categorical templates; this is even the case for temporal CSPs [34]. These negative results are to be compared with the recent result by Mottet and Pinsker [69] that did provide an algebraic description of local consistency for several subclasses of ω-categorical templates.

In the finite, the algebraic characterisation of local consistency relies on a set of algebraic tools whose development eventually led to the solutions of the Feder-Vardi dichotomy conjecture. Bulatov's proof of the Feder-Vardi conjecture [46] builds on his theory of edge-colored algebras, that were also used in his characterisation of bounded width [45]; Zhuk's proof [83, 84] relies on the concept of absorption, which was developed by Barto and Kozik in their effort to prove the bounded width conjecture [9, 11]. Comparable algebraic tools, or a general theory, are at the moment missing in the theory of infinite-domain CSPs, even with an ω-categorical template. The most general results obtained so far use *canonical operations*, which behave like operations on finite sets, and for which it is sometimes possible to mimic the universal-algebraic approach to finite-domain CSPs. Canonical operations alone do not seem to be sufficient in full generality and a characterisation of their applicability is also missing, but on the positive side their applicability covers a vast majority of the results that were proved in the area. The application of canonical operations to approach the question of local consistency for infinite-domain CSPs has only been started recently [29, 30, 69, 82].

2.1.1 Results

In this section, we focus on applying the theory of canonical functions to study the power of local consistency checking for constraint satisfaction problems over ω-categorical templates. Our objective is two-fold: on the one hand, we wish to obtain generic sufficient conditions that imply that local consistency solves a given CSP, and on the other hand we wish to understand the amount of locality needed for local consistency to solve the CSP, as measured by the so-called *relational width*.

In order to solve the first objective, we build on recent work by Mottet and Pinsker [69] and expand the use of their *smooth approximations* to fully suit *equational (non-)affineness*, which is roughly the algebraic situation imposed by local consistency solvability. The main technical contribution is a new *loop lemma* that exploits deep algebraic tools from the finite [10] and, assuming the use of canonical functions is unfruitful, allows to obtain the existence of polymorphisms of every arity $n \geq 2$ and satisfying certain strong symmetry conditions. Using this loop lemma, we are able to obtain a characterisation of bounded width for particular classes of templates, namely for first-order reducts of unary structures (i.e., structures that have only unary relations) and for certain structures related to Feder and Vardi's logic MMSNP [53].

In particular, we solve the following open problem from [15, 54]. A *Datalog program* is a sentence in the existential positive fragment of least fixpoint logic (see, e.g., [53] for a clear introduction of Datalog in the context of constraint satisfaction). The Datalog-rewritability problem for MMSNP is the problem of deciding, given as input an MMSNP sentence Φ, whether $\neg\Phi$ is equivalent to a Datalog program.

Theorem 2.1.1. *The Datalog-rewritability problem for MMSNP is decidable, and is 2NExpTime-complete.*

In order to solve the second objective, we prove that for infinite CSP templates with a certain finite presentation, sufficiently locally consistent instances can be turned into locally consistent instances of a finite-domain CSP. If this finite-domain CSP has bounded width, then it has relational width $(2,3)$ by [6], which in turn allows us to obtain a bound on the relational width of the original CSP. In this way, we obtain a collapse of the relational width reminiscent of the collapse in the finite case for all structures whose clone of canonical polymorphisms satisfies suitable identities. In particular, it turns out that the relational width of a structure then only depends on two simple parameters of the structure whose automorphism group is considered in the notion of canonicity. These parameters determine the finite presentation of the CSP template, and structures which have such a finite presentation contain the range of the complexity dichotomy conjecture for infinite-domain CSPs [38, 14, 13, 8, 7]. Finally, let us note that while the notion of bounded width makes mathematical sense even outside the scope of ω-categorical templates, its algorithmic meaning there is not guaranteed since the recursion establishing consistency of an instance need not terminate after a finite number of steps.

Theorem 2.1.2. *Let $k, \ell \geq 1$, and let $\mathbb{A}$ be a first-order reduct of a k-homogeneous ℓ-bounded ω-categorical structure $\mathbb{B}$.*

- *If the clone of polymorphisms of $\mathbb{A}$ which are canonical with respect to* $\mathrm{Aut}(\mathbb{B})$ *contains pseudo-WNU operations modulo* $\overline{\mathrm{Aut}(\mathbb{B})}$ *of all arities $n \geq 3$, then $\mathbb{A}$ has relational width* $(2k, \max(3k, \ell))$.
- *If the clone of polymorphisms of $\mathbb{A}$ which are canonical with respect to* $\mathrm{Aut}(\mathbb{B})$ *contains pseudo-totally symmetric operations modulo* $\overline{\mathrm{Aut}(\mathbb{B})}$ *of all arities, then $\mathbb{A}$ has relational width* $(k, \max(k+1, \ell))$.

Note that every finite structure $\mathbb{A}$ with domain $\{a_1, \ldots, a_n\}$ is a first-order reduct of the structure $(\{a_1, \ldots, a_n\}; \{a_1\}, \ldots, \{a_n\})$, which is easily seen to be 1-homogeneous and 2-bounded. Thus the width obtained in Theorem 2.1.2 coincides with the width given by Barto's collapse result from [6].

As a corollary of Theorem 2.1.2, we obtain a collapse of the bounded width hierarchy for first-order reducts of the unary structures mentioned above, as well as of numerous other structures studied in the literature [35, 29, 30, 28, 63].

Corollary 2.1.3. *Let $\mathbb{A}$ be a structure that has bounded width. If $\mathbb{A}$ is a first-order reduct of:*

- *the universal homogeneous graph $\mathbb{G}$ or tournament $\mathbb{T}$, or of a unary structure, then $\mathbb{A}$ has relational width at most* $(4, 6)$*;*

- *the universal homogeneous K_n-free graph $\mathbb{H}_n$, where $n \geq 3$, then at most $(2, n)$;*
- $(\mathbb{N}; =)$*, the countably infinite equivalence relation with infinitely many equivalence classes $\mathbb{C}_\omega^\omega$, or the universal homogeneous partial order $\mathbb{P}$, then at most $(2, 3)$.*

Proof. For the first item, let $\mathbb{A}$ be a first-order reduct of $\mathbb{G}$ or $\mathbb{T}$. Then $\mathbb{A}$ has a model-complete core by [16]; this model-complete core has the same relational width as $\mathbb{A}$. Moreover, by Lemma 17 and Lemma 46 in [69], this model-complete core is again a first-order reduct of $\mathbb{G}$ or $\mathbb{T}$, respectively, or a one-element structure. In the latter case, $\mathbb{A}$ has relational width 1; hence, we may assume that $\mathbb{A}$ is itself a model-complete core. It then follows that $\mathbb{A}$ has bounded width if and only if the algebraic condition in the first item of Theorem 2.1.2 is satisfied [69]. Since both $\mathbb{G}$ and $\mathbb{T}$ are 2-homogeneous and 3-bounded our claim follows.

By Lemma 6.7 in [30], the set of first-order reducts of unary structures is closed under taking model-complete cores. Hence, by appeal to Theorem 2.1.2 and Theorem 2.5.1 our claim holds for this class as well.

First-order reducts of $\mathbb{H}_n$, $(\mathbb{N}; =)$ or $\mathbb{C}_\omega^\omega$ have bounded width if and only if the condition in the second item of Theorem 2.1.2 is satisfied, by [28], [22] and [40], respectively. Since $\mathbb{H}_n$ is 2-homogeneous and n-bounded, and since both $(\mathbb{N}; =)$ and $\mathbb{C}_\omega^\omega$ are 2-homogeneous and 3-bounded, the claimed bound follows.

Finally, a first-order reduct of $\mathbb{P}$ with bounded width is either homomorphically equivalent to a first-order reduct of $(\mathbb{Q}; <)$ or it satisfies the algebraic condition in the second item of Theorem 2.1.2 [63]. In the latter case we are done by Theorem 2.1.2, in the former we appeal to the syntactical characterization of first-order reducts of $(\mathbb{Q}; <)$. Indeed, such a structure has bounded width if and only if it is definable by a conjunction of so-called Ord-Horn clauses [34]. It then follows by [43] that a first-order reduct of $(\mathbb{Q}; <)$ with bounded width has relational width $(2, 3)$. The result for $\mathbb{P}$ follows. ☐

The following example shows that for some of the structures under consideration, the bounds on relational width provided by Corollary 2.1.3 are tight.

Example 2.1.4. *To show the tightness of the bound in the case of the universal homogeneous graph $\mathbb{G} = (G; E)$, we exhibit a first-order reduct $\mathbb{A}$ such that for all $i \leq j$ with $i \leq 4$, if $1 \leq i < 4$ or $1 \leq j < 6$, then there exists a non-trivial, (i, j)-minimal instance equivalent to an instance of* $\mathrm{CSP}(\mathbb{A})$ *that has no solution. Let $N := (G^2 \setminus E) \cap I_2^G$. Consider the first-order reduct $\mathbb{A} := (G; R_=, R_{\neq})$ of $\mathbb{G}$, where $R_= := \{(a, b, c, d) \in G^4 \mid E(a, b) \wedge E(c, d)$ or $N(a, b) \wedge N(c, d)\}$ and $R_{\neq} := \{(a, b, c, d) \in G^4 \mid E(a, b) \wedge N(c, d)$ or $N(a, b) \wedge E(c, d)\}$.*

It can be seen that $\mathbb{A}$ has bounded width, so that Corollary 2.1.3 implies that $\mathbb{A}$ has relational width $(4, 6)$. It is easy to see that the instance $\mathcal{I}_1 = (\{v_1, v_2, v_3, v_4\}, \{C_1, C_2\})$ where $C_1 = \{G^{\{v_1,\ldots,v_4\}} \mid (f(v_1), \ldots, f(v_4)) \in R_=\}$, $C_2 = \{G^{\{v_1,\ldots,v_4\}} \mid (f(v_1), \ldots, f(v_4)) \in R_{\neq}\}$ is non-trivial, (i, j)-minimal for all $i \leq j$ with $1 \leq i < 4$, and has no solution. Moreover, the $(4, 5)$-minimal instance equivalent to the instance $\mathcal{I}_2 = (\{v_1, \ldots, v_6\}, \{C_1, C_2, C_3\})$ where $C_1 = \{G^{\{v_1,\ldots,v_4\}} \mid (f(v_1), \ldots, f(v_4)) \in R_{\neq}\}$, $C_2 = \{G^{\{v_3,\ldots,v_6\}} \mid (f(v_3), \ldots, f(v_6)) \in R_{\neq}\}$ and $C_3 = \{G^{\{v_1,v_2,v_5,v_6\}} \mid (f(v_1), f(v_2), f(v_5), f(v_6)) \in R_{\neq}\}$ is non-trivial and has no solution. It follows that the exact relational width of $\mathbb{A}$ is $(4, 6)$.

The tightness of the bound for first-order reducts of the universal homogeneous tournament can be shown similarly. However, it is an open question whether the bound is tight for first-order reducts of unary structures (for more details, see [73]).

The bound on relational width provided in the second item of Corollary 2.1.3 is tight. Indeed, let $n \geq 3$, let $\mathbb{H}_n := (H_n; E)$ be the universal homogeneous K_n-free graph, let $N := ((H_n)^2 \setminus E) \cap I_2^{H_n}$, and let $\mathbb{A} := (H_n; E, N)$. $\mathbb{A}$ is preserved by canonical pseudo-totally symmetric operations modulo $\overline{\mathrm{Aut}(\mathbb{H}_n)}$ of all arities and therefore has relational width $(2, n)$ by Theorem 2.1.2 (setting $\mathbb{B} := \mathbb{H}_n$). But the non-trivial, $(2, n-1)$-minimal instance $\mathcal{I}_1 = (\{v_1, \ldots, v_n\}, \{C_{i,j} \mid 1 \leq i \neq j \leq n\})$ where $C_{i,j} = \{f \in (H_n)^{\{v_i,v_j\}} \mid (f(v_i), f(v_j)) \in E\}$ has no solution; moreover, the instance $\mathcal{I}_2 = (\{v_1, v_2\}, \{C_1, C_2\})$ where $C_1 = \{f \in (H_n)^{\{v_1,v_2\}} \mid (f(v_1), f(v_2)) \in E\}$, $C_2 = \{f \in (H_n)^{\{v_1,v_2\}} \mid (f(v_1), f(v_2)) \in N\}$ is non-trivial, $(1, j)$-minimal for every $j \geq 1$ and has no solution either.

For the structures from the third item of Corollary 2.1.3, the tightness of the bound can be shown similarly.

2.1.2 Related results

Local consistency for ω-categorical structures was studied for the first time in [19] where basic notions were introduced and some basic results provided. First-order reducts of certain k-homogeneous ℓ-bounded structures with bounded width were characterized in [69, 34].

The articles [82] and [81] give the upper bound $(2, \ell)$ on the relational width for first-order expansions of some classes of 2-homogeneous, ℓ-bounded structures under the stronger assumption of bounded strict width; Corollary 2.1.3 for first-order reducts of $\mathbb{H}_n$ and of $\mathbb{C}_\omega^\omega$ also follows from [82].

2.1.3 Organisation of the present chapter

In Section 2.2 we provide some definitions and basic facts connected to the theory of smooth approximations from [69]. The reduction to the finite using canonical functions which leads to the collapse of the bounded width hierarchy is given in Section 2.3. We then extend the algebraic theory of smooth approximations in Section 2.4 before applying it to first-order reducts of unary structures and MMSNP in Section 2.5.

2.2 Smooth Approximations

We are going to apply the fundamental theorem of smooth approximations [69] to lift an action of a function clone to a larger clone.

Definition 2.2.1. (Smooth approximations) *Let A be a set, $n \geq 1$, and let $\sim$ be an equivalence relation on a subset S of A^n. We say that an equivalence relation η on some set S' with $S \subseteq S'$* approximates *$\sim$ if the restriction of η to S is a (possibly non-proper) refinement of $\sim$. We call η an* approximation *of $\sim$.*

For a permutation group $\mathscr{G}$ acting on A and leaving η as well as each $\sim$-class invariant, we say that the approximation η is smooth *if each equivalence class C of $\sim$ intersects some equivalence class C' of η such that $C \cap C'$ contains a $\mathscr{G}$-orbit.*

Theorem 2.2.2 (The fundamental theorem of smooth approximations [69])**.** *Let $\mathscr{C} \subseteq \mathscr{D}$ be function clones on a set A, and let $\mathscr{G}$ be a permutation group on A such that $\mathscr{D}$ locally* *interpolates $\mathscr{C}$ modulo $\mathscr{G}$. Let $\sim$ be a $\mathscr{C}$-invariant equivalence relation on $S \subseteq A$ with $\mathscr{G}$-invariant classes and finite index, and η be a $\mathscr{D}$-invariant smooth approximation of $\sim$ with respect to $\mathscr{G}$. Then there exists a uniformly continuous minion homomorphism from $\mathscr{D}$ to $\mathscr{C} \curvearrowright S/\sim$.*

2.3 Collapses in the Relational Width Hierarchy

Definition 2.3.1. *Let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be a CSP instance over a set A. Let $\mathscr{G}$ be a permutation group on A, let $k \geq 1$, and for every $K \in \binom{\mathcal{V}}{k}$, let $\mathcal{O}^K_{\mathscr{G}}$ be the set of K-*orbits*, i.e., orbits of maps $f\colon K \to A$ under the natural action of $\mathscr{G}$. Let $\mathcal{O}^{\mathcal{V}}_{\mathscr{G},k} := \bigcup_{K \in \binom{\mathcal{V}}{k}} \mathcal{O}^K_{\mathscr{G}}$ and let $\mathcal{I}_{\mathscr{G},k}$ be the following instance over $\mathcal{O}^{\mathcal{V}}_{\mathscr{G},k}$:*

- *The variable set of $\mathcal{I}_{\mathscr{G},k}$ is the set $\binom{\mathcal{V}}{k}$ of k-element subsets of $\mathcal{V}$. Every variable K of $\mathcal{I}_{\mathscr{G},k}$ is meant to take a value in $\mathcal{O}^K_{\mathscr{G}}$.*
- *For every constraint $C \subseteq A^U$ in $\mathcal{I}$, $\mathcal{I}_{\mathscr{G},k}$ contains the constraint $C_{\mathscr{G},k} \subseteq (\mathcal{O}^{\mathcal{V}}_{\mathscr{G},k})^{\binom{U}{k}}$ defined by*

$$C_{\mathscr{G},k} := \left\{ g\colon \binom{U}{k} \to \mathcal{O}^{\mathcal{V}}_{\mathscr{G},k} \mid \exists f \in C \;\; \forall K \in \binom{U}{k} \;\; (f|_K \in g(K)) \right\}.$$

  That is, the set of k-element subsets of U is the scope of a constraint which allows precisely those assignments of orbits to these subsets which are naturally induced by the assignments allowed by C for the variables in U. Note that for every $K \in \binom{\mathcal{V}}{k}$, for every constraint C whose scope contains K and for every $g \in C_{\mathscr{G},k}$ we have $g(K) \in \mathcal{O}^K_{\mathscr{G}}$.

Observe that if $\mathcal{I}$ is non-trivial, then so is $\mathcal{I}_{\mathscr{G},k}$.

Lemma 2.3.2. *Let $1 \leq a \leq b$. If $\mathcal{I}$ is (ak, bk)-minimal, then $\mathcal{I}_{\mathscr{G},k}$ is (a, b)-minimal.*

Proof. Let $K_1, \ldots, K_b \in \binom{\mathcal{V}}{k}$. Note that $U := \bigcup_i K_i$ has size at most bk, and therefore there exists a set W with $U \subseteq W \subseteq \mathcal{V}$ and a constraint $C \subseteq A^W$ in $\mathcal{I}$ since $\mathcal{I}$ is (ak, bk)-minimal. The scope of the associated constraint $C_{\mathscr{G},k}$ is $\binom{W}{k}$, which contains $K_1, \ldots, K_b$. Hence, the first item of Definition 1.4.1 is satisfied.

For the second item, let $K := \{K_1, \ldots, K_a\}$ and let $C_{\mathscr{G},k} \subseteq (\mathcal{O}^{\mathcal{V}}_{\mathscr{G},k})^{\binom{U}{k}}, D_{\mathscr{G},k} \subseteq (\mathcal{O}^{\mathcal{V}}_{\mathscr{G},k})^{\binom{W}{k}}$ be two constraints whose scopes contain K. Then $\bigcup K$ is contained in the scope of the associated $C \subseteq A^U$ and $D \subseteq A^W$ and has size at most ak, so that by (ak, bk)-minimality of $\mathcal{I}$, the restrictions of C and D to $\bigcup K$ coincide. Thus for every $g \in C_{\mathscr{G},k}$, there exists by definition an $f \in C$ such that $f|_{K_i} \in g(K_i)$ for all i, and by the previous sentence there exists $f' \in D$ such that $f'|_{K_i} \in g(K_i)$ for all i. Thus, $g|_K$ is in the restriction of $D_{\mathscr{G},k}$ to K. The argument is symmetric, showing that the restrictions of $C_{\mathscr{G},k}$ and $D_{\mathscr{G},k}$ to K coincide. ☐

Note that for every solution h of $\mathcal{I}$, the map $\chi_h\colon \binom{\mathcal{V}}{k} \to \mathcal{O}^{\mathcal{V}}_{\mathscr{G},k}$ defined by $K \mapsto \{\alpha h|_K \mid \alpha \in \mathscr{G}\}$ defines a solution to $\mathcal{I}_{\mathscr{G},k}$. The next lemma proves that every solution to $\mathcal{I}_{\mathscr{G},k}$ is of the form χ_h for some solution h of $\mathcal{I}$, provided that $\mathcal{I}$ is (k, ℓ)-minimal and that $\mathscr{G} = \mathrm{Aut}(\mathbb{B})$ for some k-homogeneous ℓ-bounded structure $\mathbb{B}$.

Lemma 2.3.3. *Let $1 \leq k < \ell$. Let $\mathbb{B}$ be a k-homogeneous ℓ-bounded structure, let $\mathbb{A}$ be a first-order reduct of $\mathbb{B}$, and let $\mathcal{I}$ be a (k,ℓ)-minimal instance equivalent to an instance of* $\mathrm{CSP}(\mathbb{A})$*. Then every solution to $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$ lifts to a solution of $\mathcal{I}$.*

Proof. Let $h\colon \binom{\mathcal{V}}{k} \to \mathcal{O}^{\mathcal{V}}_{\mathscr{G},k}$ be a solution to $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$. Recall that $h(K)$ is a K-orbit for any $K \in \binom{\mathcal{V}}{k}$, and one can therefore restrict $h(K)$ to any $L \subseteq K$ by setting $h(K)|_L := \{f|_L \mid f \in h(K)\}$. Note that since $\mathcal{I}$ is k-minimal, we have $h(K)|_{K\cap K'} = h(K')|_{K\cap K'}$ for all $K, K' \in \binom{\mathcal{V}}{k}$.

We now define an equivalence relation $\sim$ on $\mathcal{V}$. Suppose first that $k = 1$. Then every orbit of elements of $\mathbb{B}$ under the action of $\mathrm{Aut}(\mathbb{B})$ must be a singleton (for any orbit with two distinct elements a, b, the pairs (a,a) and (a,b) would not be in the same orbit but their subtuples of length one would be, so that $\mathbb{B}$ would not be 1-homogeneous). In that case, we identify $\mathcal{O}^{\mathcal{V}}_{\mathscr{G},k}$ with the domain B itself, and set $x \sim y$ if and only if $h(\{x\}) = h(\{y\})$; that is, $\sim$ is essentially the kernel of h.

Suppose next that $k \geq 2$, and set $x \sim y$ if and only if there is $K \in \binom{\mathcal{V}}{k}$ containing x, y such that $h(K)|_{\{x,y\}}$ consists of constant maps. One could equivalently ask that this holds for *all* K containing x, y by 2-minimality, and it then follows that this is indeed an equivalence relation by $(2,3)$-minimality of $\mathcal{I}$. Moreover, h descends to $\binom{\mathcal{V}/\sim}{k}$: if $K' = \{[v_1]_\sim, \ldots, [v_k]_\sim\}$ is a k-element set, define $\tilde{h}(K') := h(\{v_1, \ldots, v_k\})$. The definition of $\tilde{h}$ does not depend on the choice of representatives, by the very definition of $\sim$.

Define a finite structure $\mathbb{C}$ with domain $\mathcal{V}/\sim$ in the signature of $\mathbb{B}$ as follows. Let $\mathcal{V}/\sim = \{[v_1]_\sim, \ldots, [v_n]_\sim\}$. We define $\mathbb{C}$ such that the relations holding on the tuple $([v_1]_\sim, \ldots, [v_n]_\sim)$ in $\mathbb{C}$ are the same as the relations holding on an arbitrary tuple $(b_1, \ldots, b_n) \in B^n$ that satisfies the following. For every $m \leq k$ and for all $[v_{i_1}]_\sim, \ldots [v_{i_m}]_\sim$ pairwise different, the atomic types of $(b_{i_1}, \ldots, b_{i_m})$ and $(g([v_{i_1}]_\sim), \ldots, g([v_{i_m}]_\sim))$ agree for arbitrary $g \in \tilde{h}(K)$ and $K \supseteq \{v_{i_1}, \ldots, v_{i_m}\}$. This construction does not depend on the choice of the tuple $(b_1, \ldots, b_n)$ by the k-homogeneity of $\mathbb{B}$ and is well-defined by the consistency of the assignment given by h.

Finally, note that all substructures of $\mathbb{C}$ of size at most ℓ embed into $\mathbb{B}$. Indeed, let $m \leq \ell$ and let $\mathbb{L}$ be an m-element substructure of $\mathbb{C}$, and let $L' \subseteq V$ be an m-element set containing one representative for each element of $\mathbb{L}$. By (k,ℓ)-minimality of $\mathcal{I}$, there exists $C \subseteq A^{L'}$ in $\mathcal{I}$, and a corresponding constraint $C_{\mathrm{Aut}(\mathbb{B}),k}$ of $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$. Thus, $h|_{\binom{L'}{k}} \in C_{\mathrm{Aut}(\mathbb{B}),k}$, so that there exists $g \in C$ such that for all $K \in \binom{L'}{k}$, $g|_K \in h(K)$. Thus g corresponds to an embedding of every k-element substructure of $\mathbb{L}$ into $\mathbb{B}$, and since $\mathbb{B}$ is k-homogeneous, g is an embedding of $\mathbb{L}$ into $\mathbb{B}$. Finally, since $\mathbb{B}$ is ℓ-bounded, it follows that there exists an embedding e of $\mathbb{C}$ into $\mathbb{B}$.

It remains to check that the composition of e with the canonical projection $\mathcal{V} \to \mathcal{V}/\sim$ is a solution to $\mathcal{I}$, which is trivial since the relations of $\mathbb{A}$ are unions of orbits under $\mathrm{Aut}(\mathbb{B})$. □

For every finite set $\mathcal{V}$ and for every $K \in \binom{\mathcal{V}}{k}$, every operation f that is canonical with respect to a permutation group $\mathscr{G}$ induces an operation on the set orbits of K-tuples under $\mathscr{G}$. We denote this operation by $f^K_{\mathscr{G}}$. Finally, we denote by $f^{\mathcal{V}}_{\mathscr{G},k}$ the union of $f^K_{\mathscr{G}}$ for all $K \in \binom{\mathcal{V}}{k}$ and we call it a *multisorted* operation, i.e., this operation is defined only on tuples where all elements belong to the same $\mathcal{O}^K_{\mathscr{G}}$ for some fixed $K \in \binom{\mathcal{V}}{k}$. We say that a multisorted operation $f^{\mathcal{V}}_{\mathscr{G},k}$ is a *multisorted WNU* if $f^K_{\mathscr{G}}$ satisfies the WNU identities for every $K \in \binom{\mathcal{V}}{k}$.

We remark that we could have avoided the use of multisorted functions by using k-tuples of elements of $\mathcal{V}$ instead of k-element subsets of $\mathcal{V}$ in Definition 2.3.1; however, encoding the original instance this way would have added considerable redundancy and made the proof more technical.

Note that in the above situation, for every constraint C of an instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ it makes formally sense to ask whether $f^{\mathcal{V}}_{\mathscr{G},k}$ preserves $C_{\mathscr{G},k}$. Indeed, let f be n-ary. Whenever K is a k-element subset of the scope of C, and hence a variable in the scope of $C_{\mathscr{G},k}$, then for all $g_1, \dots, g_n \in C_{\mathscr{G},k}$ we have $g_1(K), \dots, g_n(K) \in \mathcal{O}^K_{\mathscr{G}}$; hence $f^{\mathcal{V}}_{\mathscr{G},k}$ can be applied to these values, and doing so for all variables in the scope of $C_{\mathscr{G},k}$ altogether yields a function from this scope to $\mathcal{O}^{\mathcal{V}}_{\mathscr{G},k}$ which can be an element of $C_{\mathscr{G},k}$ or not.

Lemma 2.3.4. *Let f be a polymorphism of $\mathbb{A}$ that is canonical with respect to $\mathscr{G}$. Every constraint $C_{\mathscr{G},k}$ in $\mathcal{I}_{\mathscr{G},k}$ is preserved by $f^{\mathcal{V}}_{\mathscr{G},k}$.*

Proof. Let n be the arity of f and let $C \subseteq A^U$ be a constraint in $\mathcal{I}$. In particular since $\mathcal{I}$ is an instance of CSP($\mathbb{A}$), and f is a polymorphism of $\mathbb{A}$, we have that C is preserved by f.

Let $g_1, \dots, g_n \in C_{\mathscr{G},k}$. By definition, for every $i \in \{1, \dots, n\}$ there is $g'_i \in C$ such that for all $K \in \binom{U}{k}$, $g'_i|_K \in g_i(K)$. Note that $f(g'_1|_K, \dots, g'_n|_K) = f(g'_1, \dots, g'_n)|_K$, so that $f(g'_1, \dots, g'_n)|_K \in f^{\mathcal{V}}_{\mathscr{G},k}(g_1(K), \dots, g_n(K))$. Since $f(g'_1, \dots, g'_n) \in C$, it follows that $f^{\mathcal{V}}_{\mathscr{G},k}(g_1, \dots, g_n)$ is in $C_{\mathscr{G},k}$. □

Finally, this allows us to prove Theorem 2.1.2 from the introduction.

Proof of Theorem 2.1.2. Suppose that the assumption of the first item of Theorem 2.1.2 is satisfied. Let $\mathcal{I}$ be a non-trivial $(2k, \max(3k, \ell))$-minimal instance equivalent to an instance of CSP($\mathbb{A}$), and let $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$ be the associated instance from Definition 2.3.1. Note that $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$ is a CSP instance over a finite set by the ω-categoricity of $\mathbb{B}$. Note moreover that any $(2k, \max(3k, \ell))$-minimal non-trivial instance with less than k variables admits a solution. Hence, we may assume that $\mathcal{I}$ has at least k variables. By Lemma 2.3.2, $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$ is a $(2, 3)$-minimal instance, and it is non-trivial by definition and since $\mathcal{I}$ has at least k variables. The constraints of $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$ are preserved by multisorted WNUs of all arities $m \geq 3$ (2.3.4). By an easy corollary of the equivalence of Item 1 and Item 5 in Theorem 1.4.4 for multisorted WNUs, $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$ admits a solution. Since $\mathcal{I}$ is $(2k, \max(3k, \ell))$-minimal, it is also (k, ℓ)-minimal, and hence this solution lifts to a solution of $\mathcal{I}$ by Lemma 2.3.3. Thus, $\mathbb{A}$ has relational width $(2k, \max(3k, \ell))$.

Suppose now that the assumption in the second item is satisfied. By the same reasoning but using Theorem 1.4.5 instead of Theorem 1.4.4, given a $(k, \max(k+1, \ell))$-minimal instance $\mathcal{I}$, the associated instance $\mathcal{I}_{\mathrm{Aut}(\mathbb{B}),k}$ is $(1, 1)$-minimal and therefore has a solution. Since $\mathcal{I}$ is $(k, \max(k+1, \ell))$-minimal, this solution lifts to a solution of $\mathcal{I}$. □

2.4 A New Loop Lemma for Smooth Approximations

We refine the algebraic theory of smooth approximations from [69]. Building on deep algebraic results from [10] on finite idempotent algebras that are equationally non-trivial, we lift some of the theory from binary symmetric relations to cyclic relations of arbitrary arity.

2.4.1 The loop lemma

Definition 2.4.1. *The* linkedness congruence *of a binary relation $R \subseteq A \times B$ is the equivalence relation λ_R on $\mathrm{proj}_{(2)}(R)$ defined by $(b, b') \in \lambda_R$ if there are $k \geq 0$ and $a_0, \dots, a_{k-1} \in A$ and $b = b_0, \dots, b_k = b' \in B$ such that $(a_i, b_i) \in R$ and $(a_i, b_{i+1}) \in R$ for all $i \in \{0, \dots, k-1\}$. We say that R is* linked *if it is non-empty and λ_R relates any two elements of $\mathrm{proj}_{(2)}(R)$.*

If A is a set and $m \geq 2$, then we call a relation $R \subseteq A^m$ cyclic *if it is invariant under cyclic permutations of the components of its tuples. The* support *of R is its projection on any argument. We apply the same terminology as above to any cyclic R, viewing R as a binary relation between $\mathrm{proj}_{(1,\dots,m-1)}(R)$ and $\mathrm{proj}_{(m)}(R)$.*

If R is invariant under an oligomorphic group action on $A \times B$, then there is an upper bound on the length k to witness $(b, b') \in \lambda_R$, and therefore λ_R is pp-definable from R; in particular, it is invariant under any function clone acting on $A \times B$ and preserving R.

Definition 2.4.2. *Let $\mathscr{G}$ be a permutation group acting on a set A. A* pseudo-loop with respect to *$\mathscr{G}$ is a tuple of elements of A all of whose components belong to the same $\mathscr{G}$-orbit [74, 13, 14]. If $\mathscr{G}$ contains only the identity function, then a pseudo-loop is called a* loop.

Our next goal is the proof of Theorem 2.4.4. In order to refer to results on finite algebras more easily, we will use the language of algebras rather than clones in the proof. We need the following definitions.

Definition 2.4.3. *Let $\mathbf{A}$ be an algebra and let B be a subuniverse of $\mathbf{A}$, i.e., a subset of the universe of $\mathbf{A}$ closed under all operations of $\mathbf{A}$. We say that B is an* absorbing subuniverse *of $\mathbf{A}$ (or B absorbs $\mathbf{A}$) if there exists a term operation f of arity $n \geq 2$ such that for any $j \in \{1, \dots, n\}$ and for any $(a_1, \dots, a_n) \in A^n$ with $a_i \in B$ for all $i \neq j$, $t(a_1, \dots, a_n) \in B$.*

If B is an absorbing subuniverse of $\mathbf{A}$ and no proper subuniverse of B absorbs $\mathbf{A}$, we call B a minimal absorbing subuniverse *of $\mathbf{A}$ and write $B \lll \mathbf{A}$.*

Theorem 2.4.4 (Consequence of the proof of Theorem 4.2 in [10])**.** *Let $\mathscr{C}$ be an idempotent function clone on a finite domain that is equationally non-trivial. Then any $\mathscr{C}$-invariant cyclic linked relation on its domain contains a loop.*

Proof. Let R as in the statement be given, and denote its arity by m; we may assume $m \geq 2$. Given $1 \leq i \leq m$, we set $R_i := \mathrm{proj}_{(1,\dots,i)}(R)$; moreover, we set $R_{(i,j)} := \mathrm{proj}_{(i,j)}(R)$ for all distinct i, j with $1 \leq i, j \leq m$.

We denote the support of R by A. Note that for all $i \in \{1, \dots, m-1\}$ we have that R_{i+1} is linked when viewed as a binary relation between R_i and A. We give the short argument showing linkedness for the convenience of the reader. Let $a, b \in A$ be arbitrary; since they are linked in R_m, there exist $k \geq 1$ and $a = c_0, \dots, c_{2k} = b$ such that $(c_{2j+1}, c_{2j}) \in R_m$ and $(c_{2j+1}, c_{2j+2}) \in R_m$ for all $0 \leq j < k$. For all such j, we have $c_{2j+1} \in R_{m-1}$; moreover, $\mathrm{proj}_{(\ell,\dots,m-1)}(c_{2j+1}) \in R_{m-\ell}$ for all $1 \leq \ell \leq m-1$ by the cyclicity of R. Hence, these elements prove linkedness of a, b in R_{i+1} for all $1 \leq i \leq m-1$.

Observe that in particular, $\mathrm{proj}_{(1,i)}(R)$ is linked for all $i \in \{2, \dots, m\}$.

Let $\mathbf{A}$ be the algebra whose domain is the support A of R and whose fundamental operations are the restrictions of the functions of the clone $\mathscr{C}$ to that domain; note that this is well-defined since A, as the projection of R to any coordinate, is invariant under the functions of $\mathscr{C}$. Let $\mathbf{R}$ be the subalgebra of $\mathbf{A}^m$ with domain R. Similarly, we write $\mathbf{R}_i$ for the subalgebra of $\mathbf{A}^i$ on the domain R_i, for all $1 \leq i \leq m$. Following the proof of Theorem 4.2 in [10], we prove by induction on $i \in \{1, \dots, m\}$ that the following properties hold:

1. There exists $I \lll \mathbf{A}$ such that $I^i \lll \mathbf{R}_i$.
2. For all $I_1, \dots, I_i \lll \mathbf{A}$ such that $R_i \cap (I_1 \times \cdots \times I_i) \neq \emptyset$, we have $R_i \cap (I_1 \times \cdots \times I_i) \lll \mathbf{R}_i$.

The case $i = 1$ is trivial, since $R_1 = A$. The case $i = m$ gives us a constant tuple in R.

We prove property (2) for $i+1$. Let $I_1, \dots, I_{i+1} \lll \mathbf{A}$ be such that $R_{i+1} \cap (I_1 \times \cdots \times I_{i+1}) \neq \emptyset$. By the induction hypothesis, $R_i \cap (I_1 \times \cdots \times I_i) \lll \mathbf{R}_i$. By the argument above, R_{i+1} is linked as a relation between R_i and A. Thus, by Proposition 2.15 in [10], we have $R_{i+1} \cap (I_1 \times \cdots \times I_{i+1}) \lll \mathbf{R}_{i+1}$.

We prove property (1) for $i + 1$. Define a directed graph H on R_i by setting

$$H := \{((a_1, \dots, a_i), (a_2, \dots, a_{i+1})) \mid (a_1, \dots, a_{i+1}) \in R_{i+1}\}.$$

Let $I \lll \mathbf{A}$ be such that $I^i \lll \mathbf{R}_i$, which exists by the induction hypothesis. We show that:

- I^i is a subset of a (weak) connected component of H;
- this connected component has algebraic length 1, that is, it contains a path from some element to itself whose numbers of forward arcs and backward arcs differ by 1.

For the first item, let $X := \{x \mid \exists a_1, \dots, a_i \in I \ ((a_1, \dots, a_i, x) \in R_{i+1})\}$, which is an absorbing subuniverse of $\mathbf{A}$. Let $X_1 \subseteq X$ be a minimal absorbing subuniverse of $\mathbf{A}$. Then since $(I^i \times X_1) \cap R_{i+1} \neq \emptyset$ property (2) gives us $(I^i \times X_1) \subseteq R_{i+1}$. Reiterating this idea, we find minimal absorbing subuniverses $X_2, \dots, X_i$ of $\mathbf{A}$ such that for all $1 \leq j \leq i$ we have that $I^{i-j+1} \times X_1 \times \cdots \times X_j$ is contained in R_{i+1}. Now pick an arbitrary tuple $(a_1, \dots, a_i) \in I^i$, and an arbitrary tuple $(x_1, \dots, x_i) \in X_1 \times \cdots \times X_i$. Then there is a path in H from $(a_1, \dots, a_i)$ to $(x_1, \dots, x_i)$ by the above, proving the first item.

For the second item, let

$$E := \{(x, y) \mid \exists c_2, \dots, c_i \in I \ ((x, c_2, \dots, c_i, y) \in R_{i+1})\}.$$

Let V_1, V_2 be the projection of E onto its first and second coordinate, respectively; then E is a relation between V_1 and V_2. We have that E is an absorbing subuniverse of the algebra $\mathbf{R}_{(1,i+1)}$ induced by $R_{(1,i+1)}$ in $\mathbf{A}^2$. By assumption on R, we have that $R_{(1,i+1)}$ is linked. Therefore, E is linked, and V_1 and V_2 are absorbing subuniverses of $\mathbf{A}$. Note that $I \subseteq V_1$ and $I \subseteq V_2$. Let $b \in I$ be arbitrary. Then there exist $k \geq 0$ and $c_0, \dots, c_{2k+1}$ such that $(c_{2j}, c_{2j+1}) \in E$ for all $0 \leq j \leq k$ and $(c_{2j+2}, c_{2j+1}) \in E$ for all $0 \leq j \leq k-1$ and such that $c_0 = b = c_{2k+1}$. These elements can be assumed to lie in minimal absorbing subuniverses of $\mathbf{A}$ by Proposition 2.15 in [10]. We then have, by property (2), $(c_{2j}, b, \dots, b, c_{2j+1}) \in R_{i+1}$ for all $0 \leq j \leq k$ and $(c_{2j+2}, b, \dots, b, c_{2j+1}) \in R_{i+1}$ for all $0 \leq j \leq k-1$. This gives a path of algebraic length 1 in H from $(b, \dots, b)$ to itself. This proves the second item.

By Theorem 3.6 in [10] there exists a loop in H which lies in a minimal absorbing subuniverse K of $\mathbf{R}_i$; by the definition of H, this loop is a constant tuple $(a, \ldots, a)$. By projecting K on the first component, we obtain a minimal absorbing subuniverse $J \lll \mathbf{A}$; since $a \in J$, we have that $J^{i+1} \cap R_{i+1} \neq \emptyset$. By property (2), we get that $J^{i+1} \lll \mathbf{R}_{i+1}$, so that property (1) holds. □

The following is a generalization of [69, Theorem 11] from binary symmetric relations to arbitrary cyclic relations.

Proposition 2.4.5. *Let $n \geq 1$, and let $\mathscr{D}$ be an oligomorphic function clone on a set A which is a model-complete core. Let $\mathscr{C} \subseteq \mathscr{D}_n^{\text{can}}$ be such that $\mathscr{C}^n/\mathscr{G}_{\mathscr{D}}$ is equationally non-trivial. Let $(S, \sim)$ be a minimal subfactor of the action $\mathscr{C}^n$ with $\mathscr{G}_{\mathscr{D}}$-invariant $\sim$-classes. Then for every $\mathscr{D}$-invariant cyclic relation R with support $\langle S \rangle_{\mathscr{D}}$ one of the following holds:*

1. *The linkedness congruence of R is a $\mathscr{D}$-invariant approximation of $\sim$.*

2. *R contains a pseudo-loop with respect to $\mathscr{G}_{\mathscr{D}}$.*

Proof. Let R be given, and denote its arity by m. Assuming that (1) does not hold, we prove (2).

Denote by $\mathcal{O}$ the set of orbits of n-tuples under the action of $\mathscr{G}_{\mathscr{D}}$ thereon. Let R' be the relation obtained by considering R as a relation on $\mathcal{O}$, i.e.,

$$R' := \{(O_1, \ldots, O_m) \in \mathcal{O}^m \mid R \cap (O_1 \times \cdots \times O_m) \neq \emptyset\}.$$

Thus, R' is an m-ary cyclic relation with support $S' \subseteq \mathcal{O}$, and R' contains a loop if and only if R satisfies (2).

By assumption, the action $\mathscr{C}^n/\mathscr{G}_{\mathscr{D}}$ is equationally non-trivial; moreover, it is idempotent since $\mathscr{D}$ is a model-complete core. Note also that R', and in particular S', are preserved by this action. It is therefore sufficient to show that R' is linked and apply Theorem 2.4.4.

Recall that we consider R also as a binary relation between $\text{proj}_{m-1}(R)$ and $\langle S \rangle_{\mathscr{D}}$; similarly, we consider R' as a binary relation between $\text{proj}_{m-1}(R')$ and S'. By the oligomorphicity of $\mathscr{D}$, the linkedness congruence λ_R of R is invariant under $\mathscr{D}$.

By our assumption that (1) does not hold, there exist $c, d \in S$ which are not $\sim$-equivalent and such that $\lambda_R(c, d)$ holds; otherwise, λ_R would be an approximation of $\sim$. This implies that the orbits O_c, O_d of c, d are related via $\lambda_{R'}$. By the minimality of $(S, \sim)$, we have that $\langle S \rangle_{\mathscr{D}} = \langle \{c, d\} \rangle_{\mathscr{D}}$. Since $\mathscr{D}$ is a model-complete core, it preserves the $\mathscr{G}_{\mathscr{D}}$-orbits, and it follows that any tuple in $\langle S \rangle_{\mathscr{D}} = \langle \{c, d\} \rangle_{\mathscr{D}}$ is λ_R-related to a tuple in the orbit of c. Hence, $\lambda_{R'} = (S')^2$, and thus R' is linked. Theorem 2.4.4 therefore implies that R' contains a loop, and hence R contains a pseudo-loop with respect to $\mathscr{G}_{\mathscr{D}}$, which is what we had to show. □

The following is a generalization of Lemma 14 in [69] from binary relations and functions to relations and functions of higher arity.

Lemma 2.4.6. *Let $n \geq 1$, and let $\mathscr{D}$ be an oligomorphic polymorphism clone on a set A that is a model-complete core. Let $\sim$ be an equivalence relation on a set $S \subseteq A^n$ with $\mathscr{G}_{\mathscr{D}}$-invariant classes. Let $m \geq 1$, and let P be an m-ary relation on $\langle S \rangle_{\mathscr{D}}$. Suppose that every*

m-ary $\mathscr{D}$-invariant cyclic relation R on $\langle S\rangle_{\mathscr{D}}$ which contains a tuple in P with components in at least two $\sim$-classes contains a pseudo-loop with respect to $\mathscr{G}_{\mathscr{D}}$.

Then there exists an m-ary $f \in \mathscr{D}$ such that for all $a_1, \ldots, a_m \in A^n$ we have that if the tuple $(f(a_1, \ldots, a_m), f(a_2, \ldots, a_m, a_1), \ldots, f(a_m, a_1, \ldots, a_{m-1}))$ is in P, then it intersects at most one $\sim$-class.

Proof. The proof is similar to the proof of Lemma 14 in [69]. Fix $m \geq 2$. Call an m-tuple $(a_1, \ldots, a_m)$ of elements of A^n *troublesome* if there exists an m-ary $r \in \mathscr{D}$ such that

$$(b_1, \ldots, b_m) := (r(a_1, \ldots, a_m), r(a_2, \ldots, a_m, a_1), \ldots, r(a_m, a_1, \ldots, a_{m-1}))$$

is in P and has components in at least two $\sim$-classes. Note that if $(a_1, \ldots, a_m)$ is not troublesome and $h \in \mathscr{D}$, then $(h(a_1, \ldots, a_m), \ldots, h(a_m, a_1, \ldots, a_{m-1}))$ is not troublesome either. Moreover, if $a_1, \ldots, a_m$ are in the same orbit under $\mathscr{G}_{\mathscr{D}}$, then $(a_1, \ldots, a_m)$ is not troublesome: since $\mathscr{D}$ is a model-complete core, $h(a_1, \ldots, a_m), \ldots, h(a_m, a_1, \ldots, a_{m-1})$ are also in the same orbit under $\mathscr{G}_{\mathscr{D}}$ for any m-ary $h \in \mathscr{D}$; since the classes of $\sim$ are closed under $\mathscr{G}_{\mathscr{D}}$, our claim follows.

For each troublesome tuple $(a_1, \ldots, a_m)$, the smallest $\mathscr{D}$-invariant relation containing the set

$$\{(b_1, \ldots, b_m), \ldots, (b_m, b_1, \ldots, b_{m-1})\},$$

where $(b_1, \ldots, b_m)$ is defined via a witnessing function r as above, is cyclic, contains a tuple in P with components in at least two $\sim$-classes, and its support is contained in $\langle S\rangle_{\mathscr{D}}$. Hence, it contains a pseudo-loop with respect to $\mathscr{G}_{\mathscr{D}}$ by our assumptions. This implies that there exists an m-ary $g \in \mathscr{D}$ such that the entries of the tuple

$$(g(b_1, \ldots, b_m), \ldots, g(b_m, b_1, \ldots, b_{m-1}))$$

all belong to the same $\mathscr{G}_{\mathscr{D}}$-orbit. The function

$$s(x_1, \ldots, x_m) := g(r(x_1, \ldots, x_m), \ldots, r(x_m, x_1, \ldots, x_{m-1}))$$

thus has the property that the entries of

$$(s(a_1, \ldots, a_m), \ldots, s(a_m, a_1, \ldots, a_{m-1}))$$

all lie in the same $\mathscr{G}_{\mathscr{D}}$-orbit.

In conclusion, for every tuple $(a_1, \ldots, a_m)$ of elements of A^n – troublesome or not – there exists an m-ary function d in $\mathscr{D}$ such that $(d(a_1, \ldots, a_m), \ldots, d(a_m, a_1, \ldots, a_{m-1}))$ is not troublesome: if $(a_1, \ldots, a_m)$ is troublesome one can take d to be the operation s we just described; if $(a_1, \ldots, a_m)$ is not troublesome then the first projection works.

Let $((a_1^i, \ldots, a_m^i))_{i\in\omega}$ be an enumeration of all m-tuples of elements of A^n. We build by induction on $i \in \omega$ an operation $f^i \in \mathscr{D}$ such that

$$(f^i(a_1^j, \ldots, a_m^j), \ldots, f^i(a_m^j, a_1^j, \ldots, a_{m-1}^j))$$

is not troublesome for any $j < i$. For $i = 0$ there is nothing to show, so suppose that f^i is built. Let $d \in \mathscr{D}$ be an m-ary operation making the tuple

$$(f^i(a_1^i, \ldots, a_m^i), \ldots, f^i(a_m^i, a_1^i, \ldots, a_{m-1}^i))$$

not troublesome, in the manner of the previous paragraph. Then setting

$$f^{i+1}(x_1, \ldots, x_m) := d(f^i(x_1, \ldots, x_m), \ldots, f^i(x_m, x_1, \ldots, x_{m-1}))$$

clearly has the desired property for the tuple $(a_1^i, \ldots, a_m^i)$, and also for all tuples $(a_1^j, \ldots, a_m^j)$ with $j < i$ by the remark in the first paragraph.

By a standard compactness argument using the oligomorphicity of $\mathscr{G}_{\mathscr{D}}$ (essentially from [36]) and the fact that the polymorphism clone is topologically closed, we may assume that the sequence $(f^i)_{i \in \omega}$ converges to a function f. This function f satisfies the claim of the lemma. □

2.5 Applications: Collapses of the bounded width hierachies for some classes of infinite structures

We now apply the algebraic results of Section 2.4 and the theory of smooth approximations to obtain a characterisation of bounded width for CSPs of first-order reducts of unary structures and for CSPs in MMSNP. Moreover, we obtain a collapse of the bounded width hierarchy for such CSPs.

2.5.1 Unary Structures

We are going to prove the following characterization of bounded width for first-order reducts of unary structures.

Theorem 2.5.1. *Let $\mathbb{A}$ be a first-order reduct of a unary structure, and assume that $\mathbb{A}$ is a model-complete core. Then one of the following holds:*

- $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ *is not equationally affine, or equivalently, it contains pseudo-WNUs modulo* $\mathrm{Aut}(\mathbb{A})$ *of all arities $n \geq 3$;*
- $\mathrm{Pol}(\mathbb{A})$ *has a uniformly continuous minion homomorphism to an affine clone.*

In the first case, the stated equivalence follows from Section 1.4.2 and $\mathbb{A}$ has relational width $(4, 6)$ by Theorem 2.1.2, and in the second case it does not have bounded width by Theorem 1.4.4. We remark that Theorem 2.5.1 gives a characterization of bounded width for *all* first-order reducts of unary structures, since this class is closed under taking model-complete cores by Lemma 6.7 in [30].

The two items of Theorem 2.5.1 are invariant under expansions of $\mathbb{A}$ by a finite number of constants (for the first item, this follows exactly as the preservation of the pseudo-Siggers identity under adding constants in [13, 14], for the second it follows directly from [12]). Thus, by Proposition 6.8 in [30], one can assume that $\mathbb{A}$ is a first-order expansion of $(\mathbb{N}; V_1, \ldots, V_r)$ where $V_1, \ldots, V_r$ form a partition of $\mathbb{N}$ in which every set is either a singleton or infinite.

Such partitions were called *stabilized partitions* in [30], and we shall also call the structure $(\mathbb{N}; V_1, \ldots, V_r)$ a stabilized partition.

We will use the following fact which states that $\mathrm{Pol}(\mathbb{A})$ locally interpolates $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$.

Lemma 2.5.2 (Proposition 6.5 in [30]). *Let $\mathbb{A}$ be a first-order expansion of a stabilized partition $(\mathbb{N}; V_1, \ldots, V_r)$. For every $f \in \mathrm{Pol}(\mathbb{A})$ there exists $g \in \mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ which is locally interpolated by f modulo $\mathrm{Aut}(\mathbb{A})$.*

Proposition 2.5.3. *Let $\mathbb{A}$ be a first-order expansion of a stabilized partition $(\mathbb{N}; V_1, \ldots, V_r)$, and assume it is a model-complete core. Suppose that $\mathrm{Pol}(\mathbb{A})$ contains operations of all arities whose restrictions to V_i are injective for all $1 \leq i \leq r$. Then the following are equivalent:*

- $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ *is equationally affine;*
- $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}} \curvearrowright \mathbb{N}/\mathrm{Aut}(\mathbb{A})$ *is equationally affine.*

Proof. Trivially, if $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}} \curvearrowright \mathbb{N}/\mathrm{Aut}(\mathbb{A})$ is equationally affine, then so is $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$.

For the other direction, assume that $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}} \curvearrowright \mathbb{N}/\mathrm{Aut}(\mathbb{A})$ is not equationally affine. Then for all $k \geq 3$, $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ contains a k-ary operation g_k whose action on $\mathbb{N}/\mathrm{Aut}(\mathbb{A})$ is a WNU operation by the implication from Item 4 to Item 5 in Theorem 1.4.4 and since $\mathbb{N}/\mathrm{Aut}(\mathbb{A})$ is a finite set. Fix for all such $k \geq 3$ an operation f_k of arity k whose restriction to V_i is injective for all $i \in \{1, \ldots, r\}$, and consider the operation

$$h_k(x_1, \ldots, x_k) := f_k(g_k(x_1, \ldots, x_k), g_k(x_2, \ldots, x_k, x_1), \ldots, g_k(x_k, x_1, \ldots, x_{k-1})).$$

Then evaluating $h(a, b, \ldots, b)$ for arbitrary $a, b \in \mathbb{N}$, all the arguments of f_k belong to the same set V_i, by the fact that g_k acts on $\mathbb{N}/\mathrm{Aut}(\mathbb{A})$ as a weak near-unanimity operation. Since $\mathbb{A}$ is an expansion of $(\mathbb{N}; V_1, \ldots, V_r)$, we obtain that $h_k(a, b, \ldots, b)$ belongs to V_i. The same is true for any permutation of the tuple $(a, b, \ldots, b)$, so that h_k acts as a WNU operation on $\mathbb{N}/\mathrm{Aut}(\mathbb{A})$. By the injectivity of f_k when restricted to V_i, it also follows that h_k acts as a WNU operation on $\mathbb{N}^2/\mathrm{Aut}(\mathbb{A})$; hence, it is a pseudo-WNU operation since it preserves the equivalence of orbits of 2-tuples under $\mathrm{Aut}(\mathbb{N}; V_1, \ldots, V_r)$ and since the stabilized partition is 2-homogeneous. Taking for every h_k an operation in $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ locally interpolated by h_k modulo $\mathrm{Aut}(\mathbb{A})$ which exists by Lemma 2.5.2, we see that $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ contains pseudo-WNU operations of all arities ≥ 3, and hence it is not equationally affine. □

The following will allow us to assume, in most proofs, the presence of functions in $\mathrm{Pol}(\mathbb{A})$ which are injective on every set in the stabilized partition. This is the analogue to the efforts to obtain binary injections in [69].

Lemma 2.5.4 (Subset of the proof of Proposition 6.6 [30]). *Let $\mathbb{A}$ be a first-order expansion of a stabilized partition $(\mathbb{N}; V_1, \ldots, V_r)$, and assume it is a model-complete core. If $\mathrm{Pol}(\mathbb{A})$ has no continuous clone homomorphism to $\mathscr{P}$, then it contains operations of all arities whose restrictions to V_i are injective for all $1 \leq i \leq r$.*

Proof. We show by induction that for all $1 \leq j \leq r$, there exists a binary operation in $\mathrm{Pol}(\mathbb{A})$ whose restriction to each of $V_1, \ldots, V_j$ is injective; higher arity functions with the same property are then obtained by nesting the binary operation. For the base case $j = 1$, observe that the disequality relation $\neq$ is preserved on V_1 since $\mathbb{A}$ is a model-complete core; together with the restriction of $\mathrm{Pol}(\mathbb{A})$ to V_1 being equationally non-trivial, we then obtain an operation which acts as as an essential function on V_1. This in turn easily yields a function that acts as a binary injection on V_1 – see e.g. [22]. For the induction step, assuming the statement holds for $1 \leq j < r$, we show the same for $j + 1$. By the induction hypothesis, there exist binary functions $f, g \in \mathrm{Pol}(\mathbb{A})$ such that the restriction of f to each of the sets $V_1, \ldots, V_j$ is injective, and the restriction of g to V_{j+1} is injective. If the restriction of f to V_{j+1} depends only on its first or only on its second variable, then it is injective in that variable since disequality is preserved on V_{j+1}, and hence either the function $f(g(x, y), f(x, y))$ or the function $f(f(x, y), g(x, y))$ has the desired property. If on the other hand the restriction of f to V_{j+1} depends on both variables, then the same argument as in the base case yields a function which is injective on V_{j+1}, and this function is still injective on each of the sets $V_1, \ldots, V_j$. □

Proof of Theorem 2.5.1. Let $\mathbb{A}$ as in Theorem 2.5.1 be given; by the remark following that theorem, we may without loss of generality assume that $\mathbb{A}$ is a first-order expansion of a stabilized partition $(\mathbb{N}; V_1, \ldots, V_r)$. Assume henceforth that $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ is equationally affine; we show that $\mathrm{Pol}(\mathbb{A})$ has a uniformly continuous minion homomorphism to an affine clone.

If $\mathrm{Pol}(\mathbb{A})$ has a continuous clone homomorphism to $\mathscr{P}$, then we are done. Assume therefore the contrary; then by Lemma 2.5.4, $\mathrm{Pol}(\mathbb{A})$ contains for all $k \geq 2$ a k-ary operation whose restriction to V_i is injective for all $1 \leq i \leq r$. In particular, Proposition 2.5.3 applies, and thus $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}} \curvearrowright \mathbb{N}/\mathrm{Aut}(\mathbb{A})$ is equationally affine. Let $(S, \sim)$ be a minimal subfactor of $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ such that $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ acts on the $\sim$-classes as an affine clone; the fact that this exists is well-known (see, e.g., Proposition 3.1 in [80]).

Let R be any $\mathrm{Pol}(\mathbb{A})$-invariant cyclic relation with support $\langle S\rangle_{\mathrm{Pol}(\mathbb{A})}$, containing a tuple with components in pairwise distinct $\mathrm{Aut}(\mathbb{A})$-orbits and which intersects at least two $\sim$-classes. By Proposition 2.4.5, R either gives rise to a $\mathrm{Pol}(\mathbb{A})$-invariant approximation of $\sim$, or it contains a pseudo-loop with respect to $\mathrm{Aut}(\mathbb{A})$. In the first case, the presence of the tuple required above implies smoothness of the approximation: if $t \in R$ is such a tuple, $c \in \langle S\rangle_{\mathrm{Pol}(\mathbb{A})}$ appears in t, and $d \in \langle S\rangle_{\mathrm{Pol}(\mathbb{A})}$ belongs to the same $\mathrm{Aut}(\mathbb{A})$-orbit as c, then there exists an element of $\mathrm{Aut}(\mathbb{A})$ which sends c to d and fixes all other elements of t. Hence, c and d are linked in R, and the entire $\mathrm{Aut}(\mathbb{A})$-orbit of c is contained in a class of the linkedness relation of R. Thus, $\mathrm{Pol}(\mathbb{A})$ admits a uniformly continuous minion homomorphism to an affine clone by Theorem 2.2.2.

Hence we may assume that for any R as above the second case holds. We are now going to show that this leads to a contradiction, finishing the proof of Theorem 2.5.1. By Lemma 2.4.6 applied with any $m \geq 2$ and P the set of m-tuples with entries in pairwise distinct $\mathrm{Aut}(\mathbb{A})$-orbits within $\langle S\rangle_{\mathrm{Pol}(\mathbb{A})}$, we obtain an m-ary function $f \in \mathrm{Pol}(\mathbb{A})$ with the property that the tuple $(f(a_0, \ldots, a_{m-1}), \ldots, f(a_1, \ldots, a_{m-1}, a_0))$ intersects at most one $\sim$-class whenever it has entries in pairwise distinct $\mathrm{Aut}(\mathbb{A})$-orbits, for all $a_0, \ldots, a_{m-1} \in S$. Let $(\mathbb{A}, <)$ be the expansion of $\mathbb{A}$ by a linear order that is convex with respect to the partition $V_1, \ldots, V_r$ and dense and without endpoints on every infinite set of the partition. The structure $(\mathbb{A}, <)$ can be seen to be a *Ramsey structure*, since $\mathrm{Aut}(\mathbb{A}, <)$ is isomorphic as a permutation group to the action of the product $\prod_{i=1}^{r} \mathrm{Aut}(V_i; <)$, and each of the groups of the product is either trivial or the automorphism group of a Ramsey structure [60]. By diagonal interpolation we may assume that f is diagonally canonical with respect to $\mathrm{Aut}(\mathbb{A}, <)$. Let $a, a' \in A^m$ be so that a_i, a'_i belong to the same orbit with respect to $\mathrm{Aut}(\mathbb{A})$ for all $1 \leq i \leq m$. Then there exists $\alpha \in \mathrm{Aut}(\mathbb{A}, <)$ such that $\alpha(a) = a'$, and hence $f(a)$ and $f(a')$ lie in the same $\mathrm{Aut}(\mathbb{A})$-orbit by diagonal canonicity; hence f is 1-canonical with respect to $\mathrm{Aut}(\mathbb{A})$. Applying Lemma 2.5.2, we obtain a canonical function $g \in \mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ which acts like f on $\mathbb{N}/\mathrm{Aut}(\mathbb{A})$. The property of f stated above then implies for g that $g(a_0, \ldots, a_{m-1}) \sim g(a_1, \ldots, a_{m-1}, a_0)$ for all $a_0, \ldots, a_{m-1} \in S$ such that the values $g(a_0, \ldots, a_{m-1}), \ldots, g(a_{m-1}, a_0, \ldots, a_{m-2})$ lie in pairwise distinct $\mathrm{Aut}(\mathbb{A})$-orbits.

By the choice of $(S, \sim)$ we have that $\mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ acts on $S/{\sim}$ by affine functions over a finite module. We use the symbols $+, \cdot$ for the addition and multiplication in the corresponding ring, and also $+$ for the addition in the module and $\cdot$ for multiplication of elements of the module with elements of the ring. We denote by 1 the multiplicative identity of the ring, by -1 its additive inverse, and identify their powers in the additive group with the non-zero integers. The domain of the module is $S/{\sim}$, and we denote the identity element of its additive group by $[a_0]_\sim$. Pick an arbitrary element $[a_1]_\sim \neq [a_0]_\sim$ from $S/{\sim}$, and let $m \geq 2$ be its order in the additive group of the module, i.e., the minimal positive number such that $m \cdot [a_1]_\sim = [a_0]_\sim$. For $i \in \{2, \ldots, m-1\}$, let a_i be an arbitrary element such that $[a_i]_\sim = i \cdot [a_1]_\sim$. Let $g \in \mathrm{Pol}(\mathbb{A})^{\mathrm{can}}$ be the m-ary operation obtained in the preceding paragraph. If the values $g(a_0, \ldots, a_{m-1}), \ldots, g(a_{m-1}, a_0, \ldots, a_{m-2})$ lie in pairwise distinct $\mathrm{Aut}(\mathbb{A})$-orbits, then (computing indices modulo m) we have that $g([a_0]_\sim, \ldots, [a_{m-1}]_\sim), \ldots, g([a_{m-1}]_\sim, \ldots, [a_{m+m-1}]_\sim)$ are all equal. If on the other hand they do not, then $g([a_k]_\sim, \ldots, [a_{k+m-1}]_\sim) = g([a_{k+j}]_\sim, \ldots, [a_{k+j+m-1}]_\sim)$ for some $0 \leq k < m$ and $1 \leq j < m$. Hence, in either case we may assume the latter equation holds. By assumption, g acts on $S/{\sim}$ as an affine map, i.e., as a map of the form $(x_0, \ldots, x_{m-1}) \mapsto \sum_{i=0}^{m-1} c_i \cdot x_i$, where $c_0, \ldots, c_{m-1}$ are elements of the ring which sum up to 1. We compute (with indices to be read modulo m)

$$\begin{aligned}
[a_0]_\sim &= g([a_{k+j}]_\sim, \ldots, [a_{k+j+m-1}]_\sim) + (-1) \cdot g([a_k]_\sim, \ldots, [a_{k+m-1}]_\sim) \\
&= \sum_{i=0}^{m-1} c_i \cdot [a_{k+j+i}]_\sim + (-1) \cdot \sum_{i=0}^{m-1} c_i \cdot [a_{k+i}]_\sim \\
&= \sum_{i=0}^{m-1} c_i \cdot (k+i+j) \cdot [a_1]_\sim + (-1) \cdot \sum_{i=0}^{m-1} c_i \cdot (k+i) \cdot [a_1]_\sim \\
&= \left(\sum_{i=0}^{m-1} c_i\right) \cdot j \cdot [a_1]_\sim = j \cdot [a_1]_\sim.
\end{aligned}$$

But $j \cdot [a_1]_\sim \neq [a_0]_\sim$ since the order of $[a_1]_\sim$ equals $m > j$, a contradiction. □

2.5.2 MMSNP

MMSNP is a fragment of existential second order logic that was discovered by Feder and Vardi in their seminal paper [53]. MMSNP is defined as the class of formulas of the form

$$\Phi := \exists M_1 \dots \exists M_r \forall x_1 \dots \forall x_k \bigwedge_i \neg(\alpha_i \wedge \beta_i)$$

where $M_1, \dots, M_r$ are unary predicates, each α_i is a conjunction of positive atoms in a signature τ not containing the equality (called the *input signature*) and each β_i is a conjunction of positive or negative existentially quantified unary predicates. We say that Φ is *connected* if every conjunction α_i is connected, that is, no α_i can be written as the conjunction of two non-empty formulas which do not share any variables. A simple syntactic procedure shows that every Φ is equivalent to a disjunction of connected MMSNP sentences, which can be computed in exponential time.

As for graphs, we call a structure *connected* if it is not the disjoint union of two of its proper substructures. Let τ be a relational signature, let σ be a unary signature whose relations are called the *colors*, and let $\mathcal{F}$ be a finite set of finite connected $(\tau \cup \sigma)$-structures, with the property that every element of any structure in $\mathcal{F}$ has exactly one color. We call $\mathcal{F}$ a *colored obstruction set* in the following. The problem FPP($\mathcal{F}$) takes as input a τ-structure $\mathbb{X}$ and asks whether there exists a $(\tau \cup \sigma)$-expansion $\mathbb{X}^*$ of $\mathbb{X}$ whose vertices are all colored with exactly one color and such that for every $\mathbb{F} \in \mathcal{F}$, there exists no homomorphism from $\mathbb{F}$ to $\mathbb{X}^*$. We say in this case that the coloring $\mathbb{X}^*$ is *$\mathcal{F}$-free* or *obstruction-free*. It can be seen that connected MMSNP and FPP are equivalent: for every connected MMSNP sentence Φ, there exists a set $\mathcal{F}$ such that $\mathbb{X} \models \Phi$ if and only if $\mathbb{X}$ is a yes-instance to FPP($\mathcal{F}$), for any τ-structure $\mathbb{X}$. Conversely, given any $\mathcal{F}$, there exists a connected MMSNP sentence Φ as above. We call $\mathcal{F}$ a *colored obstruction set associated with* Φ.

Every connected MMSNP sentence Φ has an equivalent *normal form* Ψ that can be computed from $\mathcal{F}$ in double exponential time [25, Lemma 4.4]. It is shown in [25, Definition 4.12] that for every MMSNP sentence Φ in normal form, there exists an ω-categorical structure $\mathbb{A}_\Phi$ (denoted in [25] by $\mathfrak{C}^\tau_\Phi$) such that for any τ-structure $\mathbb{X}$ it is true that $\mathbb{X} \models \Phi$ if and only if $\mathbb{X}$ admits a homomorphism to $\mathbb{A}_\Phi$ (which can equivalently be taken to be injective) [25, Lemma 4.13].

Additionally, Φ is in *strong* normal form if any identification of two existentially quantified predicates in Φ yields an inequivalent sentence. Finally, we say that Φ is *precolored* if for every symbol $M \in \sigma$, there is an associated unary symbol $P_M \in \tau$, and Φ contains the conjunct $\neg(P_M(x) \wedge M'(x))$ for every color $M' \in \sigma \setminus \{M\}$. Note that any precolored sentence in normal form is automatically in strong normal form (identifying two unary predicates M, M' would yield a sentence Ψ containing the conjunct $\neg(P_M(x) \wedge M(x))$, and Φ and Ψ can be separated by a 1-element structure whose sole vertex is in P_M). Any sentence Φ has a *standard precoloration* obtained by adding the necessary predicates and conjuncts. The colored obstruction set of the standard precoloration of Φ consists of the obstruction set for Φ together with one-element obstructions $\mathbb{F}$ whose sole vertex belongs to P_M and M' in $\mathbb{F}$.

A useful (but imprecise) parallel between MMSNP and finite-domain CSPs is the following. Connected MMSNP sentences in normal form, in strong normal form, and in precolored normal form have the same relationship as do finite structures, finite cores, and expansions of cores by all singleton unary relations.

When Φ is in strong normal formal or precolored, $\mathbb{A}_\Phi$ can additionally be chosen to have the following properties:

(i) If Φ is precolored, then the orbits of the elements of $\mathbb{A}_\Phi$ under $\mathrm{Aut}(\mathbb{A}_\Phi)$ correspond to the colors of Φ and to the corresponding predicates in τ [25, Lemma 6.2]. In particular, the action of $\mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}}$ on $\mathrm{Aut}(\mathbb{A}_\Phi)$-orbits of elements is idempotent [25, Proposition 7.2].

(ii) The expansion of $\mathbb{A}_\Phi$ by a generic linear order $<$ is a Ramsey structure [25, Corollary 5.9]. In particular, every $f \in \mathrm{Pol}(\mathbb{A}_\Phi)$ locally interpolates an operation $g \in \mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}}$, and every f diagonally interpolates an operation f' that is diagonally canonical with respect to $\mathrm{Aut}(\mathbb{A}_\Phi, <)$.

We finally solve the Datalog-rewritability problem for MMSNP and prove that a precolored connected sentence Φ in normal form is equivalent to a Datalog program if and only if the action of $\mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}}$ on $\mathrm{Aut}(\mathbb{A}_\Phi)$-orbits of elements is not equationally affine.

The following proposition is proved in [25, Lemma 7.5] in the case where $m = 2$. We give the proof for the convenience of the reader.

Proposition 2.5.5. *Let Φ be a precolored MMSNP sentence in normal form and let $m \geq 1$. There exist self-embeddings $e_1, \dots, e_m$ of $\mathbb{A}_\Phi$ such that the tuples $(e_{i_1}(a_1), \dots, e_{i_m}(a_m))$ and $(e_{j_1}(b_1), \dots, e_{j_m}(b_m))$ are in the same orbit under $\mathrm{Aut}(\mathbb{A}_\Phi, <)$ provided that:*

- *a_k and b_k are in the same color for all $k \in \{1, \dots, m\}$*
- *a_k and a_ℓ are in distinct colors for all $k \neq \ell$,*
- *$\{i_1, \dots, i_m\} = \{j_1, \dots, j_m\} = \{1, \dots, m\}$.*

Proof. Let $\mathbb{A}^*$ be the expansion of $\mathbb{A}_\Phi$ by all pp-definable relations. Let $(\mathbb{H}, <)$ be the countably infinite homogeneous structure whose finite substructures are exactly the finite substructures of structures satisfying Φ and expanded by all pp-definable relations and by an arbitrary linear order. By [49], such $(\mathbb{H}, <)$ exists and it is a Ramsey structure. There exists a homomorphism $h\colon \mathbb{H} \to \mathbb{A}^*$ since the reduct of $\mathbb{H}$ to the input signature of Φ satisfies Φ by definition and hence, it admits a homomorphism to $\mathbb{A}_\Phi$. The homomorphism h can moreover be assumed to be *canonical* from $(\mathbb{H}, <)$ to $(\mathbb{A}_\Phi, <)$, i.e., sending tuples of the same type to tuples of the same type, by [37].

Let $\mathbb{B}$ be $\{1, \dots, m\} \times \mathbb{A}^*$, the disjoint union of m copies of $\mathbb{A}^*$. Endow $\mathbb{B}$ with a linear order that is convex with respect to the colors. Then there exists an embedding e' of $(\mathbb{B}, <)$ into $(\mathbb{H}, <)$ since $\mathbb{A}_\Phi$ satisfies Φ and hence, also $\mathbb{B}$ satisfies Φ. All the finite substructures of $(\mathbb{B}, <)$ then embed into $(\mathbb{H}, <)$, and by compactness $(\mathbb{B}, <)$ itself embeds into $(\mathbb{H}, <)$. Let $e_i(x) := (h \circ e')(i, x)$.

To check that these self-embeddings satisfy the required properties, let $a_1, b_1, \dots, a_m, b_m$ and $i_1, j_1, \dots, i_m, j_m$ be as in the statement. Note that since a_k and b_k are in the same color for all k, they are in the same orbit in $\mathbb{A}_\Phi$ by the fact that Φ is precolored (i). They therefore satisfy the same formulas, meaning that e' maps (i_k, a_k) and (j_k, b_k) to elements of $\mathbb{H}$ that are in the same orbit ($\mathbb{H}$ is homogeneous). Note moreover that by definition of the order on $\mathbb{B}$, $(i_k, a_k) < (i_\ell, a_\ell)$ if and only if $(j_k, b_k) < (j_\ell, b_\ell)$, and that no other atomic relation holds within the tuples $((i_1, a_1), \dots, (i_m, a_m))$ and $((j_1, b_1), \dots, (j_m, b_m))$. Thus, the required tuples are in the same orbit in $(\mathbb{H}, <)$, by the homogeneity of $(\mathbb{H}, <)$. Since h is canonical, we obtain that their image under h is in the same orbit in $(\mathbb{A}_\Phi, <)$. □

Lemma 2.5.6. *Let* $(S, \sim)$ *be a subfactor of* $\mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}}$ *with* $\mathrm{Aut}(\mathbb{A}_\Phi)$*-invariant* $\sim$*-classes. Let* $m \geq 2$*, and let* $f \in \mathrm{Pol}(\mathbb{A}_\Phi)$ *be such that for all* $a_1, \dots, a_m \in A_\Phi$*, if the entries of the tuple* $(f(a_1, \dots, a_m), f(a_2, \dots, a_m, a_1), \dots, f(a_m, a_1, \dots, a_{m-1}))$ *all belong to different colors, then the tuple intersects at most one* $\sim$*-class. Let* $O_0, \dots, O_{m-1} \in S$ *be pairwise distinct orbits under* $\mathrm{Aut}(\mathbb{A}_\Phi)$*. There exists* $g \in \mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}}$ *that is locally interpolated by* f *and that satisfies*

$$g(O_k, \dots, O_{k+m-1}) \sim g(O_{j+k}, \dots, O_{j+k+m-1}) \qquad (\star)$$

for some $0 \leq k < m$ *and* $1 \leq j < m$ *(where the indices are computed modulo* m*).*

Proof. Recall that by (ii) the expansion of $\mathbb{A}_\Phi$ by a generic linear order is a Ramsey structure. Thus, f diagonally interpolates a function $g \in \mathrm{Pol}(\mathbb{A}_\Phi)$ with the same properties and which is diagonally canonical with respect to $\mathrm{Aut}(\mathbb{A}_\Phi, <)$, and without loss of generality we can therefore assume that f is itself diagonally canonical.

Let $e_0, \dots, e_{m-1}$ be self-embeddings of $\mathbb{A}_\Phi$ with the properties stated in Proposition 2.5.5. Consider $f'(x_0, \dots, x_{m-1}) := f(e_0 x_0, \dots, e_{m-1} x_{m-1})$, and note that f' is 1-canonical when restricted to m-tuples where all entries are in pairwise distinct orbits. Let $g \in \mathrm{Pol}(\mathbb{A}_\Phi)$ be a function that is diagonally interpolated by f' and which is diagonally canonical with respect to $\mathrm{Aut}(\mathbb{A}_\Phi, <)$. In particular $g \in \mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}}$ and $g(O_k, \dots, O_{k+m-1})$ and $f'(O_k, \dots, O_{k+m-1})$ are in S and $\sim$-equivalent for all k.

As in the proof of Theorem 2.5.1, there are suitable $0 \leq k < m$ and $1 \leq j < m$ such that

$$f(e_k O_k, \dots, e_{k+m-1} O_{k+m-1}) \sim f(e_{k+j} O_{k+j}, \dots, e_{k+j+m-1} O_{j+k+m-1})$$

holds, where indices are computed modulo m. Then

$$\begin{aligned} g(O_k, \dots, O_{k+m-1}) &\sim f(e_0 O_k, \dots, e_{m-1} O_{k+m-1}) \\ &\sim f(e_k O_k, \dots, e_{k+m-1} O_{k+m-1}) && (\ddagger) \\ &\sim f(e_{k+j} O_{k+j}, \dots, e_{k+j+m-1} O_{k+j+m-1}) \\ &\sim f(e_0 O_{k+j}, \dots, e_{m-1} O_{k+j+m-1}) && (\ddagger) \\ &\sim g(O_{k+j}, \dots, O_{k+j+m-1}), \end{aligned}$$

where the equivalences marked (‡) hold by the fact that f is diagonally canonical with respect to $\mathrm{Aut}(\mathbb{A}_\Phi, <)$ and by Proposition 2.5.5. □

Let Φ be an MMSNP sentence in normal form and let $\mathcal{F}$ be the associated colored obstruction set. Note that $\mathbb{A}_\Phi$ satisfies Φ, so that there exists a coloring of $\mathbb{A}_\Phi$ that is $\mathcal{F}$-free. Fix such a coloring $\mathbb{A}_\Phi^*$. Let $\mathbb{B}_\Phi$ be the finite structure whose domain is the set of colors of Φ, with one k-ary relation $R_\mathbb{F}$ for each obstruction $\mathbb{F} \in \mathcal{F}$ of size k, and where the interpretation of $R_\mathbb{F}$ in $\mathbb{B}_\Phi$ contains all tuples $(C_1, \dots, C_k)$ of colors such that there exist $a_1, \dots, a_k$ in $\mathbb{A}_\Phi$ inducing a homomorphic image of the τ-reduct of $\mathbb{F}$, and such that a_i is in the color C_i in $\mathbb{A}_\Phi^*$ for all i.

The following theorem gives a characterization of Datalog-rewritability for precolored normal forms. The proof is similar to that of Theorem 2.5.1.

Theorem 2.5.7. *Let Φ be a precolored connected MMSNP τ-sentence in normal form, and let ℓ be the maximum size of a structure in a minimal colored obstruction set associated with Φ. The following are equivalent:*

1. *$\neg\Phi$ is equivalent to a Datalog program;*

2. *$\mathrm{Pol}(\mathbb{A}_\Phi)$ does not have a uniformly continuous minion homomorphism to an affine clone;*

3. *The action of $\mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}}$ on $\mathrm{Aut}(\mathbb{A}_\Phi)$-orbits of elements is not equationally affine;*

4. *$\mathrm{Pol}(\mathbb{B}_\Phi)$ is not equationally affine;*

5. *$\mathbb{A}_\Phi$ has relational width $(2, \max(3, \ell))$.*

Proof. (1) implies (2) by the implication from Item 2' to Item 3 in Theorem 1.4.4.

(2) implies (3). We do the proof by contraposition. The proof is essentially the same as in the case of reducts of unary structures (Theorem 2.5.1). Suppose that $\mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}} \curvearrowright \mathbb{A}_\Phi/\mathrm{Aut}(\mathbb{A}_\Phi)$ is equationally affine and let $(S, \sim)$ be a minimal module for this action.

Let $m \geq 2$ and let R be an m-ary cyclic relation invariant under $\mathrm{Pol}(\mathbb{A}_\Phi)$ and containing a tuple $(a_1, \ldots, a_m)$ whose entries are pairwise distinct. By Proposition 2.4.5, either the linkedness congruence of R defines an approximation of $\sim$, or R contains a pseudoloop modulo $\mathrm{Aut}(\mathbb{A}_\Phi)$. In the first case, the approximation is smooth by the same argument as in the proof of Theorem 2.5.1 and we obtain a uniformly continuous minion homomorphism from $\mathrm{Pol}(\mathbb{A}_\Phi)$ to an affine clone by Theorem 2.2.2.

So let us assume that for all $m \geq 2$, every such relation R contains a pseudoloop. By applying Lemma 2.4.6 with P being the set of m-tuples whose entries belong to pairwise distinct colors, we obtain a polymorphism f such that for all $a_1, \ldots, a_m$, if the elements $f(a_1, \ldots, a_m), \ldots, f(a_m, a_0, \ldots, a_{m-1})$ are in pairwise distinct colors, then they intersect at most one $\sim$-class. As in the proof of Theorem 2.5.1, pick an arbitrary $a_1 \in S$ such that $[a_1]_\sim$ is not the zero element of the module $S/{\sim}$. Let $m \geq 2$ be its order, and let O_i be the orbit of $i \cdot [a_1]_\sim$, for $i \in \{0, 1, \ldots, m-1\}$. By Lemma 2.5.6, we obtain $g \in \mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}}$ such that

$$g(O_k, \ldots, O_{k+m-1}) \sim g(O_{j+k}, \ldots, O_{j+k+m-1})$$

for some $k \in \{0, \ldots, m-1\}$ and $j \in \{1, \ldots, m-1\}$. The same computation as in Theorem 2.5.1 then gives a contradiction and concludes the proof.

(3) implies (4). Since Φ is precolored, orbits under $\mathrm{Aut}(\mathbb{A}_\Phi)$ corresponds to the colors of Φ, i.e., to the domain of $\mathbb{B}_\Phi$ (by (i)). Then under the bijection between colors and orbits, $\mathrm{Pol}(\mathbb{A}_\Phi)_1^{\mathrm{can}} \curvearrowright \mathbb{A}_\Phi/\mathrm{Aut}(\mathbb{A}_\Phi)$ is a subset of $\mathrm{Pol}(\mathbb{B}_\Phi)$. Thus if the smaller clone is not equationally affine, the larger clone is also not.

(4) implies (5). Let $\mathcal{F}$ be a minimal colored obstruction set associated with Φ. Suppose that $\mathrm{Pol}(\mathbb{B}_\Phi)$ is not equationally affine. Then $\mathbb{B}_\Phi$ does not have a minion homomorphism to an affine clone and by the implication from Item 3 to Item 1 in Theorem 1.4.4, $\mathbb{B}_\Phi$ has relational width $(2, 3)$. It is proven in [25] that there exists a weak reduction from $\mathrm{CSP}(\mathbb{A}_\Phi)$ to $\mathrm{CSP}(\mathbb{B}_\Phi)$ (using the terminology of [4], this is a *positive quantifier-free reduction without*

parameters). Since the maximal arity of a relation in $\mathbb{B}_\Phi$ is ℓ, this reduction implies the claimed bound on the relational width of $\mathbb{A}_\Phi$.

(5) implies (1). By the implication from Item 2 to Item 2' in Theorem 1.4.4, (5) implies that the class of finite structures that do not have a homomorphism to $\mathbb{A}_\Phi$ is definable by a Datalog program. This class is by construction the class of finite models of $\neg\Phi$, which proves (1). □

Note that the characterization in Theorem 2.5.7 only holds for precolored sentences in normal form. We show below how to characterize Datalog also for arbitrary MMSNP sentences in normal form. The steps involved in the proof to go from normal form to precolored normal form reflect a similar situation as for finite domain CSPs, when going from arbitrary finite structures to their cores: given a finite structure $\mathbb{A}$ and its core $\mathbb{B}$, we have that $\mathrm{CSP}(\mathbb{A})$ is in Datalog if and only if $\mathrm{Pol}(\mathbb{B})$ is not equationally affine if and only if $\mathrm{Pol}(\mathbb{A})$ does not admit a minion homomorphism to an affine clone (see Theorem 1.4.4). As for finite CSPs, those steps involve the use of pp-constructions.

The following proposition shows that for the question of Datalog-rewritability, one can go from strong normal form to precolored normal form without loss of generality. The same proposition was shown in [26, 25] for the P/NP-complete dichotomy, with $\mathscr{P}$ replacing affine clones in the statement.

Proposition 2.5.8. *Let Φ be an MMSNP sentence in strong normal form and let Ψ be its standard precoloration. There is a uniformly continuous minion homomorphism from $\mathrm{Pol}(\mathbb{A}_\Psi)$ to an affine clone if and only if there is a uniformly continuous minion homomorphism from $\mathrm{Pol}(\mathbb{A}_\Phi)$ to an affine clone.*

Proof. It is shown in [25, Theorem 6.9] that $\mathrm{Pol}(\mathbb{A}_\Psi)$ has a uniformly continuous minion homomorphism to $\mathrm{Pol}(\mathbb{A}_\Phi)$ and that $\mathrm{Pol}(\mathbb{A}_\Phi, \neq)$ has a uniformly continuous minion homomorphism to $\mathrm{Pol}(\mathbb{A}_\Psi)$. Thus, it suffices to show that if $\mathrm{Pol}(\mathbb{A}_\Phi, \neq)$ has a uniformly continuous minion homomorphism to an affine clone, then so does $\mathrm{Pol}(\mathbb{A}_\Phi)$.

Let $p \geq 2$ be prime and let R_0 and R_1 be the relations from Item 3' in Theorem 1.4.4. By the equivalence of Item 3 and Item 3' in Theorem 1.4.4, it is enough to show that if $(\mathbb{A}_\Phi, \neq)$ pp-constructs $(\mathbb{Z}_p; R_0, R_1)$, then so does $\mathbb{A}_\Phi$.

Suppose that $(\mathbb{Z}_p; R_0, R_1)$ has a pp-construction in $(\mathbb{A}_\Phi, \neq)$. Thus, there is $n \geq 1$ and pp-formulas $\phi_0(x, y, z), \phi_1(x, y, z)$ defining relations S_0, S_1 such that $(A^n; S_0, S_1)$ and $(\mathbb{Z}_p; R_0, R_1)$ are homomorphically equivalent; we take n to be minimal with the property that such pp-formulas exist. Let $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n), z = (z_1, \dots, z_n)$. Since R_0 and R_1 are totally symmetric relations (i.e., the order of the entries in a tuple does not affect its membership into any of R_0 or R_1), we can assume that S_0 and S_1 are, too, and that the formulas pp-defining them are syntactically invariant under permutation of the blocks of variables x, y, and z.

We first claim that ϕ_i does not contain any inequality atom $x_r \neq y_r$ for $r \in \{1, \dots, n\}$ (so that by symmetry, also $y_r \neq z_r$ and $x_r \neq z_r$ do not appear). Let $h \colon (\mathbb{Z}_p; R_0, R_1) \to (A^n; S_0, S_1)$ be a homomorphism. Since $(0, 0, 0) \in R_0$, we have that $(h(0), h(0), h(0))$ satisfies ϕ_0, and therefore the listed inequality atoms cannot appear. The same holds for ϕ_1, by considering $(h(0), h(0), h(1))$ and its permutations.

Moreover, the only possible equality atoms in ϕ_0 are of the form $u_r = v_r$ for some $r \in \{1, \ldots, n\}$ and $u, v \in \{x, y, z\}$. Indeed, if ϕ_0 contains $u_r = v_s$ for $r \neq s$, then by symmetry of ϕ_0 we obtain by transitivity that ϕ_0 entails $x_r = x_s$. One could obtain a pp-construction with smaller dimension by adding all the equalities $x_r = x_s, y_r = y_s, z_r = z_s$ to ϕ_0 and ϕ_1 and existentially quantifying the sth coordinate of the two formulas. The same reasoning applies to ϕ_1.

Let ψ_i be the formula obtained from ϕ_i by removing the possible inequality literals, and let T_i be defined by ψ_i in $\mathbb{A}_\Phi$. We claim that $(A^n; T_0, T_1)$ and $(A^n; S_0, S_1)$ are homomorphically equivalent, which concludes the proof. Since ϕ_i implies ψ_i, we have that $(A^n; S_0, S_1)$ is a (non-induced) substructure of $(A^n; T_0, T_1)$, and therefore it homomorphically maps to $(A^n; T_0, T_1)$ by the identity map. For the other direction, we prove the result by compactness and show that every finite substructure $\mathbb{B}$ of $(A^n; T_0, T_1)$ has a homomorphism to $(A^n; S_0, S_1)$. Let $b^1, \ldots, b^m$ be the elements of $\mathbb{B}$, where $b^i = (b^i_1, \ldots, b^i_n)$. Let $\mathbb{C}$ be the τ-structure over at most $n \cdot m$ elements $\{c^i_r \mid i \in \{1, \ldots, m\}, r \in \{1, \ldots, n\}\}$ such that:

- c^i_r and c^j_s are taken to be equal if, and only if, b^i_r and b^j_s are connected by a sequence of equalities coming from ϕ_0 and ϕ_1. By the claims above, this is only possible if $r = s$.
- the relations of $\mathbb{C}$ are defined by pulling back the relations from $\mathbb{A}_\Phi$ under the map $\pi\colon c^i_j \mapsto b^i_j$.

Note that π is a homomorphism $\mathbb{C} \to \mathbb{A}_\Phi$, and therefore $\mathbb{C}$ admits an injective homomorphism g to $\mathbb{A}_\Phi$. Let $c^i = (c^i_1, \ldots, c^i_n)$ for $i \in \{1, \ldots, n\}$. We claim that if $(b^i, b^j, b^k) \in T_0$ then $(g(c^i), g(c^j), g(c^k)) \in S_0$. Indeed, suppose that (b^i, b^j, b^k) satisfies ψ_0. Then by construction (c^i, c^j, c^k) satisfies ψ_0 in $\mathbb{C}$, and thus $(g(c^i), g(c^j), g(c^k))$ satisfies ψ_0 in $\mathbb{A}_\Phi$. Moreover, by injectivity of g, we have $g(c^i_r) \neq g(c^j_s)$ as long as $r \neq s$. Consider any inequality atom in ϕ_0. By our first claim, it can only be of the form $x_r \neq y_s$ for some $r \neq s$, and therefore it is satisfied by $(g(c^i), g(c^j), g(c^k))$. Thus, $(g(c^i), g(c^j), g(c^k))$ satisfies ϕ_0. The same reasoning for ϕ_1 shows that g induces a homomorphism $\mathbb{B} \to (A^n; S_0, S_1)$ by mapping b^i to $g(c^i)$. □

Proposition 2.5.9. *Let Φ be an MMSNP sentence in normal form. Then $\neg\Phi$ is equivalent to a Datalog program if, and only if,* $\mathrm{Pol}(\mathbb{B}_\Phi)$ *does not admit a minion homomorphism to an affine clone.*

Proof. Let Ψ be an MMSNP sentence in strong normal form equivalent to Φ, and let Θ be the standard precoloration of Ψ.

We prove that $\mathbb{B}_\Theta$ is the expansion of the core of $\mathbb{B}_\Phi$ by all unary singleton relations. Note that since Φ and Ψ are equivalent, the structures $\mathbb{A}_\Phi$ and $\mathbb{A}_\Psi$ are homomorphically equivalent. The homomorphisms realizing the equivalence can be taken to be 1-canonical without loss of generality (by (ii)), and thus they induce a homomorphic equivalence between $\mathbb{B}_\Phi$ and $\mathbb{B}_\Psi$. Moreover, $\mathbb{B}_\Psi$ is a core: if h is an endomorphism of $\mathbb{B}_\Psi$ such that $h(M) = h(M')$ for some colors $M \neq M'$, one can identify in Ψ the colors according to h and obtain an equivalent sentence, contradicting the fact that Ψ is in strong normal form. Recall that the standard precoloration of Ψ is obtained by adding 1-element obstructions for each pair of colors $M \neq M'$. Since in $\mathbb{A}_\Theta$ the relation P_M equals the color M (by (i)), these obstructions yield in $\mathbb{B}_\Theta$ unary relations containing a single element. Thus, $\mathbb{B}_\Theta$ is the expansion of the core of $\mathbb{B}_\Phi$ by all unary singleton relations.

If $\mathrm{Pol}(\mathbb{B}_\Phi)$ does not admit a minion homomorphism to an affine clone, then $\mathrm{Pol}(\mathbb{B}_\Theta)$ is not equationally affine by the implication from Item 3 to Item 4 in Theorem 1.4.4, and by Theorem 2.5.7, $\neg\Theta$ is equivalent to a Datalog program. Note that if $\mathbb{X}$ is a structure in the signature of Φ, it satisfies Φ if and only if its expansion by the relations P_M for every color M interpreted as empty relations satisfies Θ. Thus, we obtain that $\neg\Phi$ is equivalent to a Datalog program (it suffices to take a program for $\neg\Theta$ and remove all rules involving the extra predicates P_M).

Conversely, suppose that $\mathrm{Pol}(\mathbb{B}_\Phi)$ has a minion homomorphism to an affine clone. Then $\mathrm{Pol}(\mathbb{B}_\Theta)$ is equationally affine by the implication from Item 3 to Item 4 in Theorem 1.4.4, therefore by Theorem 2.5.7 there is a uniformly continuous minion homomorphism from $\mathrm{Pol}(\mathbb{A}_\Theta)$ to an affine clone. By Proposition 2.5.8, there is a uniformly continuous minion homomorphism from $\mathrm{Pol}(\mathbb{A}_\Psi)$ to an affine clone. Since $\mathbb{A}_\Psi$ and $\mathbb{A}_\Phi$ are homomorphically equivalent, we obtain a uniformly continuous minion homomorphism from $\mathrm{Pol}(\mathbb{A}_\Phi)$ to an affine clone. Finally, by the implication from Item 2 to Item 3 in Theorem 1.4.4, this implies that $\mathbb{A}_\Phi$ does not have bounded width, i.e., $\neg\Phi$ is not equivalent to a Datalog program by the same reasoning as in the proof of (5) implies (1) in Theorem 2.5.7. □

This finally allows us to obtain Theorem 2.1.1 from the introduction.

Theorem 2.1.1. *The Datalog-rewritability problem for MMSNP is decidable, and is 2NExpTime-complete.*

Proof. Let Φ be an MMSNP sentence, which is equivalent to a disjunction $\Phi_1 \vee \cdots \vee \Phi_p$ of connected MMSNP sentences, and this decomposition can be computed in exponential time [26, Proposition 3.2]. Each Φ_i has size polynomial in Φ. Moreover, if p is minimal then $\neg\Phi$ is equivalent to a Datalog program if and only if every $\neg\Phi_i$ is equivalent to a Datalog program (see, e.g., Proposition 3.3 in [25], for a proof of a similar fact). Such a minimal set $\{\Phi_1, \ldots, \Phi_p\}$ of sentences can be computed in exponential time, given Φ as input. After having computed any set $\{\Phi_1, \ldots, \Phi_p\}$ whose disjunction is equivalent to Φ, it suffices to iterate the following procedure: for any $i, j \in \{1, \ldots, p\}$, check whether Φ_i implies Φ_j (by Theorem 5.15 in [25], this problem is in NP). If Φ_i implies Φ_j for some $i \neq j$, then remove Φ_i and continue. Otherwise we claim $\{\Phi_1, \ldots, \Phi_p\}$ is minimal. Indeed, no Φ_i disjunct from $\Phi_{j_1}, \ldots, \Phi_{j_k}$ implies any disjunction $\Phi_{j_1} \vee \cdots \vee \Phi_{j_k}$: by taking for each $j \in \{j_1, \ldots, j_k\}$ a structure $\mathbb{X}_j$ witnessing that Φ_i does not imply Φ_j (i.e., $\mathbb{X}_j$ satisfies Φ_i but not Φ_j), then the disjoint union $\mathbb{X}_{j_1} \cup \cdots \cup \mathbb{X}_{j_k}$ witnesses that Φ_i does not imply the disjunction.

Fix $i \in \{1, \ldots, p\}$. Let Ψ_i be a normal form associated with Φ_i, which can be computed in double exponential time, and where Ψ itself can be of size doubly exponential in the size of Φ_i. By Proposition 2.5.9, $\neg\Psi_i$ is equivalent to a Datalog program if, and only if, $\mathrm{Pol}(\mathbb{B}_{\Psi_i})$ does not admit a minion homomorphism to an affine clone. Deciding this property is in NP [48, Corollary 6.8]. We obtain overall a 2NExpTime algorithm. The complexity lower bound is Theorem 18 in [42]. □

3 A new algorithm for infinite-domain CSPs

3.1 Introduction

As we mentioned in Section 1.2.4, Conjecture 1.3.6 has been confirmed for many subclasses. There are two regimes in which these complexity classifications were proven. In the first regime, the aim is to show that already the clone of *canonical* polymorphisms admits a rich algebraic structure, which yields an efficient many-one reduction to a tractable finite-domain CSP [29, 30].

In order to prove that the structures under consideration have canonical polymorphisms enabling an efficient reduction of their CSPs to tractable finite-domain CSPs, a demanding case distinction had to be done until recently. Today, at least two general approaches that avoid this case distinction are available: the theory of *smooth approximations* of Mottet and Pinsker [69] and the work of Bodirsky and Bodor on the *unique interpolation property* [18].

In the second regime, the standard action of the canonical polymorphisms of the structure on orbits of tuples is trivial, and therefore ad hoc algorithms that do not use the standard reduction to the finite-domain CSP mentioned above need to be given to confirm Conjecture 1.3.6. In [69], a non-standard action of the canonical polymorphisms is used to explain the known algorithms in the case of temporal CSPs, but the method does not seem to generalize. It was remarked in [18] that the boundary between the two regimes is roughly drawn by whether or not the template has the *strict order property* (SOP) (for details on the SOP, see [61]). Indeed, in all of the above-mentioned complexity classifications, it was possible to confirm Conjecture 1.3.6 for first-order reducts of structures without the SOP within the first regime.

However, in [73, Example 1], a worrying example of a first-order reduct $\mathbb{A}$ of the universal homogeneous 3-uniform hypergraph $\mathbb{H}$ is given. Even though $\mathbb{H}$ does not have the SOP, the structure $\mathbb{A}$ should have a polynomial-time solvable CSP according to Conjecture 1.3.6 but it does not possess canonical polymorphisms that would allow us to use the standard reduction to a tractable finite-domain CSP. The polymorphisms of $\mathbb{A}$ depend on a linear order on the domain of $\mathbb{H}$ even though the hypergraph $\mathbb{H}$ is not ordered. This surprising discovery implies that new algorithmic techniques are needed to solve CSPs of first-order reducts of finitely bounded homogeneous hypergraphs.

3.1.1 Results

In this chapter, we will use the above-mentioned theory of smooth approximations to confirm Conjecture 1.3.6 for first-order reducts of certain finitely bounded homogeneous ℓ-

uniform hypergraphs (where $\ell \geq 3$). In particular, we show that the templates from [73, Example 1] have polynomial-time solvable CSPs. For the precise definitions of the mentioned concepts that are not defined in Chapter 1, see Section 3.2.

Theorem 3.1.1. *Let $\ell \geq 3$, let $\mathbb{H}$ be a finitely bounded homogeneous ℓ-uniform hypergraph whose expansion with a freely added linear order $<$ is a Ramsey structure and whose automorphism group is n-"primitive" for every $n \geq 1$ and let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$. Then precisely one of the following applies.*

1. *The clone* $\mathrm{Pol}(\mathbb{A})$ *has a uniformly continuous minion homomorphism to the clone of projections* $\mathscr{P}$*, and* $\mathrm{CSP}(\mathbb{A})$ *is* NP*-complete.*

2. *The clone* $\mathrm{Pol}(\mathbb{A})$ *has no uniformly continuous minion homomorphism to the clone of projections* $\mathscr{P}$*, and* $\mathrm{CSP}(\mathbb{A})$ *is in* P*.*

The complexity classification from Theorem 3.1.1 is of particular interest for the following reasons:

- New algorithms are needed to prove Theorem 3.1.1. Since proving Conjecture 1.3.6 would in particular give an algorithm solving all tractable finite-domain CSPs, it seems likely that the methods existing in the finite either have to be used as a black box or have to be adapted in the infinite setting. The black box method from [30] does not work for the class of structures considered in this chapter, so we resort to the second option and introduce algorithmic techniques inspired by Zhuk's algorithm for finite-domain CSPs. These techniques are then coupled with the classical reduction to finite-domain CSPs, resulting in an intriguing interplay between infinitary and finitary methods.

- The result depends only on some general properties of the automorphism group of the base structure (n-*"primitivity"*) and on the fact that the base structure has a particular Ramsey expansion. Moreover, the base structures, i.e., ℓ-uniform hypergraphs satisfying the properties from Theorem 3.1.1, are not classified and a complete classification seems to be demanding if not hopeless (some 3-uniform hypergraphs satisfying our assumptions were classified in [2]). This is the first classification where no structural results about the base structures are known.

- In [69], the scalability of the theory of smooth approximations, i.e., the fact that this theory does not require us to analyze all first-order reducts of the particular structure, was claimed to be one of the main contributions of this theory compared to the archaic case-distinction method from the early literature on the subject. Theorem 3.1.1 provides us with the first complexity classification using smooth approximations that truly stands by this promise. By [79], even for the universal homogeneous ℓ-uniform hypergraph, the number of first-order reducts of this hypergraph grows with ℓ, putting an exhaustive case analysis out of reach.

Finally, in Section 3.8, we obtain as an easy consequence of the proof of Theorem 3.1.1 a classification of first-order expansions of the finitely bounded homogeneous hypergraphs under consideration whose CSP is solvable by local consistency methods.

3.2 Hypergraphs and clones

3.2.1 Hypergraphs and model-theoretic notions

Let $\ell \geq 2$. A structure $\mathbb{H} = (H; E)$ is an ℓ-uniform hypergraph if the relation E (called a *hyperedge*) is of arity ℓ, contains only injective tuples and is *fully symmetric*, i.e., for any tuple in E, all tuples obtained by permuting its components are in E as well.

Let $\ell \geq 2$. The *universal homogeneous ℓ-uniform hypergraph* is the up to isomorphism unique countably infinite homogeneous structure that is *universal* for the class of all finite ℓ-uniform hypergraphs.The universal homogeneous ℓ-uniform hypergraph is not a Ramsey structure for any ℓ but it has a finitely bounded homogeneous *Ramsey expansion* – if we add a linear order *freely* (i.e., so that the new age consists of the structures from the old age ordered in all possible ways and the resulting structure is homogeneous), the resulting structure will be homogeneous, finitely bounded and Ramsey – this follows, e.g., from the Nešetřil-Rödl theorem [76].

Definition 3.2.1 ("Primitivity")**.** *Let A be a set and $n \geq 1$. A permutation group $\mathscr{G}$ acting on A is n-*"primitive" *if for every orbit $O \subseteq A^n$ of $\mathscr{G}$, every $\mathscr{G}$-invariant equivalence relation on O containing some pair (a, b) with a, b disjoint is full.*

Example 3.2.2. *The automorphism group of the universal homogeneous ℓ-uniform hypergraph $\mathbb{H}$ is n-"primitive" for any $n \geq 1$. Indeed, let $n \geq 1$, let O be an orbit of n-tuples under $\mathrm{Aut}(\mathbb{H})$ and let $\sim$ be an equivalence relation on O containing (a, b) such that a, b are disjoint. Let $c, d \in O$ be arbitrary. We define $\mathbb{X}$ to be an ℓ-uniform hypergraph over $3n$ elements $\{x_i^j \mid i \in [n], j \in [3]\}$ such that the following holds. The hypergraph induced by $(x_1^j, \dots, x_n^j, x_1^{j+1}, \dots, x_n^{j+1})$ is isomorphic to the structure induced by (a, b) in $\mathbb{H}$ for every $j \in [2]$ and the hypergraph induced by $(x_1^1, \dots, x_n^1, x_1^3, \dots, x_n^3)$ is isomorphic to the structure induced by (c, d) in $\mathbb{H}$. By the universality of $\mathbb{H}$, $\mathbb{X}$ embeds to $\mathbb{H}$. By the homogeneity of $\mathbb{H}$, we can assume that the embedding maps $(x_1^1, \dots, x_n^1, x_1^3, \dots, x_n^3)$ to (c, d). By the transitivity of $\sim$, $c \sim d$.*

In the whole chapter, we fix $\ell \geq 3$ and a finitely bounded homogeneous ℓ-uniform hypergraph $\mathbb{H}$ whose expansion with a freely added linear order $<$ is a Ramsey structure and whose automorphism group is n-"primitive" for every $n \geq 1$.

Note that not every finitely bounded homogeneous ℓ-uniform hypergraph satisfies our assumptions. However, the universal homogeneous ℓ-uniform hypergraph does satisfy the assumptions for every $\ell \geq 3$. Additionally, for every fixed $\ell \geq 3$ and $n > \ell$, there exists a homogeneous ℓ-uniform hypergraph that is universal for the class of K_n^ℓ-free hypergraphs, where K_n^ℓ is the ℓ-hypergraph on n vertices whose every ℓ-element subset forms a hyperedge; these hypergraphs also satisfy our assumptions.

In the whole chapter, for $n \geq 1$, we denote by I_n the set I_n^H and we write $I := I_\ell$. We write N for the complement of the hyperedge relation E in I and we call it the *non-hyperedge* relation.

Note that $\mathbb{H}$ has no algebraicity. To see this, suppose that there exists $a_0 \in H$ that is first-order definable using elements $a_1, \dots, a_n \in H$ as parameters. Let us define ordered

ℓ-uniform hypergraphs $\mathbb{X} = (\{x_0, \dots, x_n\}, E^{\mathbb{X}}, <^{\mathbb{X}})$ and $\mathbb{Y} = (\{y_0, \dots, y_n\}, E^{\mathbb{Y}}, <^{\mathbb{Y}})$ such that both $(x_{i_1}, \dots, x_{i_\ell}) \in E^{\mathbb{X}}$ and $(y_{i_1}, \dots, y_{i_\ell}) \in E^{\mathbb{Y}}$ if, and only if, $(a_{i_1}, \dots, a_{i_\ell}) \in E$ for every $(i_1, \dots, i_\ell) \in \{0, \dots, n\}^\ell$. Let moreover both $x_i <^{\mathbb{X}} x_j$ and $y_i <^{\mathbb{Y}} y_j$ if, and only if $a_i < a_j$ for every $i, j \in [n]$ and let $x_0 < x_i$ and $y_0 > y_i$ for every $i \in [n]$. It follows that $\mathbb{X}$ and $\mathbb{Y}$ embeds to $(\mathbb{H}, <)$ by embeddings e_X and e_Y and by the homogeneity, we may suppose that $e_X(a_i) = e_Y(b_i) = x_i$ for every $i \in [n]$. Hence, $e_X(x_0), e_Y(y_0)$ and a_0 satisfy the same first-order formulas over $(\mathbb{H}, \{\{a_i\} \mid i \in [n]\})$ but $e_X(x_0) \neq e_Y(y_0)$, a contradiction.

3.2.2 Universal algebra

For a function f of arity $n \geq 1$ and for $i \in [n]$, we say that the i-th variable of f is *essential* if there exist $a_1, \dots, a_n, a_i'$ from the domain of f such that $f(a_1, \dots, a_n) \neq f(a_1, \dots, a_{i-1}, a_i', a_{i+1}, \dots, a_n)$. A function is called *essentially unary* if at most one of its variables is essential, otherwise the function is *essential*.

We say that a function clone $\mathscr{C}$ is *equationally trivial* if it has a clone homomorphism to the clone $\mathscr{P}$, and *equationally non-trivial* otherwise.

For a first-order reduct $\mathbb{A}$ of $\mathbb{H}$, we will consider the following two subclones of $\mathrm{Pol}(\mathbb{A})$:

- $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ is the clone of those polymorphisms of $\mathbb{A}$ which preserve the equivalence of orbits of injective tuples under $\mathrm{Aut}(\mathbb{H})$.
- $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\mathrm{inj}}$ is the clone of those polymorphisms of $\mathbb{A}$ which preserve the equivalence of orbits of injective tuples under $\mathrm{Aut}(\mathbb{A})$.

3.3 Overview of the proof of Theorem 3.1.1

Let $\mathbb{A}$ be a model-complete core of a first-order reduct $\mathbb{A}'$ of $\mathbb{H}$. By [12], it is enough to prove that Theorem 3.1.1 holds for $\mathbb{A}$ since there exists a uniformly continuous minion homomorphism from $\mathrm{Pol}(\mathbb{A}')$ to $\mathrm{Pol}(\mathbb{A})$ and from $\mathrm{Pol}(\mathbb{A})$ to $\mathrm{Pol}(\mathbb{A}')$. By Proposition 3.4.2, we can assume that $\mathbb{A}$ is itself a first-order reduct of $\mathbb{H}$.

By applying some folklore results, as well as a compactness argument, we finally obtain in Proposition 3.4.5 that $\mathrm{Pol}(\mathbb{A})$ contains an injective operation with certain properties unless $\mathrm{Pol}(\mathbb{A})$ admits a uniformly continuous clone homomorphism to $\mathscr{P}$. This polymorphism witnesses that I is a *binary absorbing subuniverse* of H^ℓ. In the rest of the chapter, we will prove that assuming that I is absorbing, then $\mathrm{CSP}(\mathbb{A})$ is tractable if, and only if, $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally non-trivial.

Let us therefore first suppose that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally non-trivial. In this case, $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$ can be solved by the reduction to a tractable finite domain CSP from [29, 30]. We will show that $\mathrm{CSP}(\mathbb{A})$ can be reduced to $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$.

By the classification of clones on a two-element domain from [75], $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is either equationally non-affine, or it consists of affine maps over $\mathbb{Z}_2$. In the first case, $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$ has relational width $(2\ell, \max(3\ell, b_{\mathbb{H}}))$ by an easy modification of the proof of Theorem 2.1.2. Now, Theorem 3.1.1 follows from Corollary 3.5.2 (details in Section 3.5.1). In the second case, $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$ amounts to solving linear equations over $\mathbb{Z}_2$. In this situation, we can apply the following algorithm – for more details, see Section 3.5.2.

Let $\mathcal{I}$ be an instance of CSP($\mathbb{A}$). Our algorithm transforms $\mathcal{I}$ into an equi-satisfiable instance $\mathcal{I}'$ that is sufficiently minimal, such that the solution set of a certain relaxation of $\mathcal{I}'$ is subdirect on all projections to an ℓ-tuple of pairwise distinct variables, and that additionally satisfies a condition which we call *inj-irreducibility*, inspired by Zhuk's notion of irreducibility [83, 84]. We then prove that any satisfiable instance satisfying those properties has an injective solution. This step is to be compared with the case of *absorbing reductions* in Zhuk's algorithm, and in particular with Theorem 5.5 in [84], in which it is proved that any sufficiently minimal and irreducible instance that has a solution also has a solution where an arbitrary variable is constrained to belong to an absorbing subuniverse. Since in our setting I is an absorbing subuniverse of H^ℓ, this fully establishes a parallel between the present work and [84].

If $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}} \curvearrowright \{E, N\}$ is equationally trivial, our goal is to prove that CSP($\mathbb{A}$) is NP-hard. The first step is to establish that $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}}$ is equationally trivial as well and that $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}} \subseteq \mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}}$ (Lemma 3.6.1). Moreover, Proposition 36 in [69] implies that there exists $k \geq \ell$ such that the action of $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}}$ on orbits of injective k-tuples under Aut($\mathbb{A}$) is equationally trivial. Therefore, there exists a *naked set* $(S, \sim)$ for this action. A naked set of $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}} \curvearrowright I_k/\text{Aut}(\mathbb{A})$ consists of an invariant subset $S \subseteq I_k/\text{Aut}(\mathbb{A})$ and an invariant equivalence relation $\sim$ on S such that $\sim$ has at least two equivalence classes and such that $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}} \curvearrowright I_k/\text{Aut}(\mathbb{A})$ acts on $S/\sim$ by projections. By classical results in finite clone theory, the existence of such a naked set is equivalent to the existence of a clone homomorphism from $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}} \curvearrowright I_k/\text{Aut}(\mathbb{A})$ to $\mathscr{P}$, which extends to a uniformly continuous clone homomorphism $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}} \to \mathscr{P}$.

Now, we can employ the theory of smooth approximations to extend this homomorphism further and obtain a uniformly continuous clone homomorphism Pol($\mathbb{A}$) $\to \mathscr{P}$. We recall below the relevant definitions from the theory of smooth approximations.

Definition 3.3.1 (Smooth approximations)**.** *Let A be a set and let $\sim$ be an equivalence relation on $S \subseteq A$. We say that an equivalence relation η on a set S' with $S \subseteq S' \subseteq A$* approximates *$\sim$ if the restriction of η to S is a refinement of $\sim$. η is called an* approximation *of $\sim$.*

For a permutation group $\mathscr{G}$ acting on A and leaving the $\sim$-classes invariant as well as η, we say that the approximation η is smooth *if every equivalence class C of $\sim$ intersects some equivalence class C' of η such that $C \cap C'$ contains a $\mathscr{G}$-orbit. η is* very smooth *if orbit-equivalence with respect to $\mathscr{G}$ is a refinement of η on S.*

Recall that $\mathbb{H}$ has no algebraicity and hence, the hypotheses of the loop lemma of smooth approximations [69, Theorem 10] are met and we may use the following reformulation of the lemma to our situation.

Theorem 3.3.2. *Let $k \geq 1$ and suppose that $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},inj} \curvearrowright I_k/\text{Aut}(\mathbb{A})$ is equationally trivial. Then there exists a naked set $(S, \sim)$ of $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},inj} \curvearrowright I_k/\text{Aut}(\mathbb{A})$ with* Aut($\mathbb{A}$)*-invariant $\sim$-classes such that one of the following holds:*

- *$\sim$ is approximated by a* Pol($\mathbb{A}$)*-invariant equivalence relation that is very smooth with respect to* Aut($\mathbb{A}$)*;*
- *every* Pol($\mathbb{A}$)*-invariant binary symmetric relation $R \subseteq (I_k)^2$ that contains a pair $(a, b) \in S^2$ such that $a \neq b$ and such that $a \not\sim b$ contains a pseudo-loop modulo* Aut($\mathbb{A}$)*, i.e., a pair (c, c') where c, c' belong to the same orbit under* Aut($\mathbb{A}$)*.*

In [69, Theorem 10], the first item of the statement gives an approximation $\sim$ that is not very smooth, but *presmooth*. By [69, Lemma 8], under the assumption that $\text{Pol}(\mathbb{A})$ preserves I_2 and that $\text{Aut}(\mathbb{A})$ is n-"primitive", every presmooth approximation is very smooth; this justifies our reformulation above.

Suppose that the first case of Theorem 3.3.2 applies. By a minor modification of the smooth approximation toolbox (Lemma 3.6.2), this implies that $\text{Pol}(\mathbb{A})$ has a uniformly continuous clone homomorphism to $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}} \curvearrowright I_k/\text{Aut}(\mathbb{A})$ and hence to the clone of projection.

Suppose now that the second case of Theorem 3.3.2 applies. By [69, Lemma 13], $\text{Pol}(\mathbb{A})$ contains a *weakly commutative function*, i.e., a binary operation f with the property that $f(a,b) \sim f(b,a)$ holds for all $a, b \in I_k$ such that $f(a,b)$ and $f(b,a)$ are in S and disjoint. It follows from a fairly involved compactness argument in Lemma 3.7.2 that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}} \curvearrowright \{E, N\}$ contains a semilattice operation and in particular, is equationally non-trivial, which is a contradiction.

3.4 Model-Complete Cores and Injective Polymorphisms

In this section, we first prove some basic facts about model-complete cores of first-order reducts of $\mathbb{H}$. In the second part, we prove that such first-order reducts that are model-complete cores and that are equationally non-trivial have binary injective polymorphisms acting as a projection or as a semilattice operation on $\{E, N\}$. These binary injections will play an important in the algorithm for tractable CSPs in Section 3.5.

3.4.1 Model-complete cores

Let $\mathscr{G}$ be a permutation group and let $g\colon G \to G$ be a function. We say that g is *range-rigid* with respect to $\mathscr{G}$ if all orbits of tuples under $\mathscr{G}$ that intersect the range of g are invariant under g. We will use the following theorem to understand the model-complete cores of first-order reducts of $\mathbb{A}$.

Theorem 3.4.1 ([68]). *Let $\mathbb{A}$ be a first-order reduct of a homogeneous Ramsey structure $\mathbb{B}$ and let $\mathbb{A}'$ be its model-complete core. Then $\mathbb{A}'$ is a first-order reduct of a homogeneous Ramsey substructure $\mathbb{B}'$ of $\mathbb{B}$.*

Moreover, there exists $g \in \text{End}(\mathbb{A})$ that is range-rigid with respect to $\text{Aut}(\mathbb{B})$ and such that the age of $\mathbb{B}'$ is equal to the age of the structure induced by the range of g in $\mathbb{B}$.

Proposition 3.4.2. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$. Then the model-complete core of $\mathbb{A}$ is a one-element structure or a first-order reduct of $\mathbb{H}$. Moreover, if $\mathbb{A}$ is a model-complete core that is a first-order reduct of $\mathbb{H}$ and not of $(H;=)$, then the range of every $f \in \text{End}(\mathbb{A})$ intersects every orbit under $\text{Aut}(\mathbb{H},<)$.*

Proof. Using Theorem 3.4.1, we obtain that the model-complete core $\mathbb{A}'$ of $\mathbb{A}$ is a first-order reduct of a homogeneous Ramsey substructure $\mathbb{B}'$ of $(\mathbb{H},<)$. Moreover, there exists $g \in \text{End}(\mathbb{A})$ which is range-rigid with respect to $\text{Aut}(\mathbb{H},<)$ and such that the age of the structure induced by the range of g in $(\mathbb{H},<)$ is equal to the range of $\mathbb{B}'$.

If the range of g contains just one element, $\mathbb{B}'$ is a one-element structure. Otherwise, the range contains at least two elements $a < b$ and by the range-rigidity, it then contains infinitely many elements. In particular, it contains a hyperedge or a non-hyperedge.

If the range contains only hyperedges, then g sends every ℓ-tuple in N to an ℓ-tuple ordered in the same way as the original tuple that lies in the hyperedge relation E. It follows that for all injective tuples a, b of the same length, there exists an automorphism α of $\mathbb{H}$ and an embedding e from the range of g into $\mathbb{B}'$ such that $e \circ g \circ \alpha(a) = b$. It follows that $\mathrm{Aut}(\mathbb{A}')$ is the full symmetric group on the domain of $\mathbb{A}'$ and hence, $\mathbb{A}'$ is a first-order reduct of $(A', =)$ which is isomorphic to $(H; =)$. If the range of g contains only ℓ-tuples in N, $\mathbb{A}'$ is a first-order reduct of $(A', =)$ by the same argument where the roles of E and N are switched. Finally, if the range of g contains both a hyperedge as well as a non-hyperedge, it follows from the range-rigidity of g that $\mathbb{B}'$ is isomorphic to $(\mathbb{H}, <)$ and $\mathbb{A}'$ is isomorphic to $\mathbb{A}$. In particular, $\mathbb{A}'$ is a first-order reduct of $\mathbb{H}$.

Suppose now that $\mathbb{A}$ is a model-complete core that is a first-order reduct of $\mathbb{H}$ but not of $(H; =)$ and let $f \in \mathrm{End}(\mathbb{A})$. Suppose that the range of f does not intersect every orbit under $\mathrm{Aut}(\mathbb{H}, <)$. By Lemma 15 and Lemma 11 in [68] applied to $\mathrm{End}(\mathbb{A})$, the range-rigid function g does not intersect every orbit under $\mathrm{Aut}(\mathbb{H}, <)$ either and we obtain a contradiction with the previous paragraph. □

3.4.2 Injective binary polymorphisms

Lemma 3.4.3. *Let $\mathbb{A}$ be a first-order reduct of a finitely bounded homogeneous hypergraph $\mathbb{H}$ that is a model-complete core. If $\mathrm{Pol}(\mathbb{A})$ does not have a uniformly continuous clone homomorphism to $\mathscr{P}$, then it contains a binary essential operation.*

Proof. It follows from Corollary 6.9 in [14] that $\mathrm{Pol}(\mathbb{A})$ contains a ternary essential operation. Moreover, the binary relation $O := \{(a, b) \mid a \neq b \in H\}$ is an orbit under $\mathrm{Aut}(\mathbb{H})$ that is *free*, i.e., for every $(c, d) \in H^2$ there exists $a \in H$ such that $(a, c), (a, d) \in O$. Now, the lemma follows directly from Proposition 22 in [69]. □

Note that for every binary injective operation f on H, there exists an embedding e of the substructure induced by the range of f into $\mathbb{H}$ such that $f' := e \circ f$ acts lexicographically on the order, i.e., $f'(x, y) < f'(x', y')$ if $x < x'$ or $x = x', y < y'$. To see this, let us define an ordered hypergraph $(\mathbb{Y}, E^{\mathbb{Y}}, <^{\mathbb{Y}})$ on H^2 such that $((x_1, y_1) \ldots, (x_\ell, y_\ell)) \in E^{\mathbb{Y}}$ if $(f(x_1, y_1), \ldots, f(x_n, y_n)) \in E$ and such that $(x, y) <^{\mathbb{Y}} (x', y')$ if $x < x'$ or $x = x', y < y'$. It is easy to see that there exists an isomorphism i from the substructure of $\mathbb{H}$ induced by the range of f in $\mathbb{H}$ to $\mathbb{Y}$. Since $(\mathbb{H}, <)$ is an expansion of $\mathbb{H}$ by a linear order that is added freely, it is universal for the class of all l-uniform linearly ordered hypergraphs $(\mathbb{X}, <^{\mathbb{X}})$ such that $\mathbb{X}$ embeds to $\mathbb{H}$. It follows that $\mathbb{Y}$ embeds to $(\mathbb{H}, <)$ by an embedding e'. Finally, by setting $e := e' \circ i$, we get the desired embedding e.

Lemma 3.4.4. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ that is a model-complete core. Suppose that $\mathrm{Pol}(\mathbb{A})$ contains a binary essential operation. Then $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj}$ contains a binary injection f such that one of the following holds:*

- *f acts like a semilattice operation on $\{E, N\}$, or*

- *f acts like a projection on $\{E, N\}$ and is canonical with respect to $(\mathbb{H}, <)$.*

Proof. The binary disequality relation $\neq$ is clearly an orbit under $\mathrm{Aut}(\mathbb{A})$ and since $\mathbb{A}$ is a model-complete core, $\mathbb{A}$ pp-defines $\neq$. $\mathbb{H}$ is clearly transitive, i.e., $\mathrm{Aut}(\mathbb{H})$ has only one orbit in its action on H. Moreover, $\mathrm{Aut}(\mathbb{H})$ has only one orbit of injective pairs in its action on H^2, namely the orbit of $\{(a,b) \mid a \neq b \in H\}$ and the structure $(H; \neq)$ clearly has finite duality. Proposition 25 in [69] then implies that $\mathrm{Pol}(\mathbb{A})$ contains a binary injection.

Note that if $\mathbb{A}$ is a reduct of $(H; =)$, then this binary injection can be composed with a unary injection so that the result acts like a semilattice operation on $\{E, N\}$, so that we can assume below that $\mathbb{A}$ is not a first-order reduct of $(H; =)$.

Since $(\mathbb{H}, <)$ is a homogeneous Ramsey structure and since any function interpolated by an injective function is injective, $\mathrm{Pol}(\mathbb{A})$ contains a binary injection f which is canonical with respect to $(\mathbb{H}, <)$. Since $\mathbb{A}$ is a model-complete core and not a first-order reduct of $(H; =)$, we can assume that the action of f on orbits of ℓ-tuples under $\mathrm{Aut}(\mathbb{H}, <)$ is idempotent by Proposition 3.4.2 – the range of the endomorphism e defined by $e(x) := f(x, x)$ intersects every orbit under $\mathrm{Aut}(\mathbb{H}, <)$ and hence, if f does not act idempotently on the orbits of ℓ-tuples under $\mathrm{Aut}(\mathbb{H}, <)$, we may compose f finitely many times with itself until it does. Here, by composition of a binary function g with f we mean the function $g(f(x, y), f(x, y))$. Moreover, by the remark above, we may assume that f acts lexicographically on the order $<$.

Let α be a permutation of $\{1, \ldots, \ell\}$. If O is an orbit of injective ℓ-tuples under $\mathrm{Aut}(\mathbb{H}, <)$, then we write $\alpha(O)$ for the orbit obtained by changing the order of the tuples in O according to α (not by permuting the tuples in O). That way, α acts naturally on orbits of injective ℓ-tuples. We also apply this notation to unions of orbits of injective tuples under $\mathrm{Aut}(\mathbb{H}, <)$. Let J be the set of strictly increasing ℓ-tuples with respect to $<$. For a fixed permutation α, the restriction of the natural action of f on ℓ-tuples to $J \times \alpha(J)$ acts idempotently on $\{E, N\}$; depending on what this action is, we say that f *behaves like p_1 on input α* if it behaves like the first projection on $\{E, N\}$ and we say that it *behaves like p_2* if it behaves like the second projection in this action on input α; in both cases, we say that f *behaves "like a projection" on input α*. Finally, we say that f *behaves like $\vee$ on input α* if $f(E, N) = f(N, E) = f(E, E) = E, f(N, N) = N$ and that f *behaves like $\wedge$ on input α* if $f(E, N) = f(N, E) = f(N, N) = N, f(E, E) = E$; in both cases, we say that f *behaves "like a semilattice" on input α*. Since the action of f on orbits is idempotent, it cannot behave like a constant in the above situation.

We write id for the identity on $\{1, \ldots, \ell\}$. If f behaves like p_2 on input id, then the operation $f(f(x, y), f(y, x))$ behaves like p_1 on id while still satisfying the behaviour with respect to the order we assumed in the beginning, so we assume that this is not the case.

If f behaves like p_1 on id, then $f(x, f(x, y))$ behaves like p_1 on $\{E, N\}$ on all inputs, and hence it acts on $\{E, N\}$ and is canonical on injective tuples with respect to $\mathbb{H}$.

Suppose that f behaves like a semilattice on $\{E, N\}$ on input id; without loss of generality it behaves like $\vee$. If f behaves differently on some other input α, then $f(f(x, y), f(x, \alpha^{-1}(y))$ behaves like $\vee$ on input α as well as on all other inputs where f behaves like $\vee$. Hence, repeating this argument for all α, we can assume that f behaves like $\vee$ on $\{E, N\}$ on any input. Note, however, that the expression $f(f(x, y), f(x, \alpha^{-1}(y))$ only makes sense for the action on orbits unless α is the identity or reverses the order; that is, only in those cases there exists a canonical function on $(\mathbb{H},<)$ whose action on orbits is that of α. The global argument has to be made local in the other cases, as follows. Let $(s,t) \in J \times \alpha(J)$. Let $\overline{\alpha^{-1}}$ be an automorphism of $\mathbb{H}$ which changes the order on t as would α^{-1}. Then $f(f(x,y), f(x, \overline{\alpha^{-1}}(y))$ acts like $\vee$ on (s,t), as well on all pairs of injective ℓ-tuples on which it previously acted like $\vee$. Repeating this procedure for all pairs of ℓ-tuples and applying a standard compactness argument, we obtain a function which acts like $\vee$ on all pairs of injective tuples, and hence on all inputs. □

A combination of Lemma 3.4.3 and Lemma 3.4.4 immediately yields the following proposition.

Proposition 3.4.5. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ that is a model-complete core. If* $\operatorname{Pol}(\mathbb{A})$ *does not have a uniformly continuous clone homomorphism to $\mathscr{P}$, then $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj}$ contains a binary injection f such that one of the following holds:*

- *f acts like a semilattice operation on $\{E,N\}$, or*
- *f acts like a projection on $\{E,N\}$ and is canonical with respect to $(\mathbb{H},<)$.*

3.5 The tractable case

In this section, we will prove that if $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E,N\}$ is equationally non-trivial, then $\operatorname{CSP}(\mathbb{A})$ is in P. We distinguish two cases based on whether $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E,N\}$ is equationally affine.

3.5.1 $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathbf{inj}} \curvearrowright \{E,N\}$ is equationally non-affine

Let $\mathbb{A}$ be a relational structure, let $n \geq 1$, let S,R with $S \subseteq R$ be n-ary relations pp-definable in $\mathbb{A}$ and let $f \in \operatorname{Pol}(\mathbb{A})$ be binary. We say that S *is a binary absorbing subuniverse of* R if for every $s \in S, r \in R$, we have that $f(s,r), f(r,s) \in S$. In this case, we write $S \trianglelefteq_2 R$ and we say that f *witnesses* the binary absorption.

Lemma 3.5.1. *Let $\mathbb{A}$ be an ω-categorical structure, let $1 \leq k \leq m \leq n$, let S,R with $S \subseteq R \subseteq A^k$ be relations pp-definable in $\mathbb{A}$ and suppose that $S \trianglelefteq_2 R$. Let $\mathcal{I} = (\mathcal{V},\mathcal{C})$ be a non-trivial, (m,n)-minimal instance equivalent to an instance of* $\operatorname{CSP}(\mathbb{A})$. *Let $\mathcal{I}'$ be an instance obtained from $\mathcal{I}$ by adding for every $\mathbf{v} = (v_1,\ldots,v_k) \in \mathcal{V}^k$ with $\operatorname{proj}_{\mathbf{v}}(\mathcal{I}) \subseteq R$, $S \cap \operatorname{proj}_{\mathbf{v}}(\mathcal{I}) \neq \emptyset$ and such that $S \cap \operatorname{proj}_{\mathbf{v}}(\mathcal{I}) \trianglelefteq_2 R$ a constraint $\{c \in A^{\{v_1,\ldots,v_k\}} \mid c(\mathbf{v}) \in S\}$. Then the (m,n)-minimal instance equivalent to $\mathcal{I}'$ is non-trivial.*

Proof. Let $\mathcal{I}''$ be the (m,n)-minimal instance equivalent to $\mathcal{I}'$. We will find a non-trivial, (m,n)-minimal instance $\mathcal{J}$ of CSP over the set A with the same set of variables as $\mathcal{I}''$ and such that every constraint of $\mathcal{J}$ is a subset of a constraint of $\mathcal{I}''$. Then it follows that $\mathcal{I}''$ is non-trivial.

Let f be a binary polymorphism of $\mathbb{A}$ witnessing $S \trianglelefteq_2 R$, let $F := \{\alpha f(\beta,\gamma) \mid \alpha,\beta,\gamma \in \operatorname{Aut}(\mathbb{A})\}$ and let $T := \{t \in \operatorname{Pol}(\mathbb{A}) \mid t \text{ is a term in } F\}$. Let C be constraint of $\mathcal{I}$ with scope $\{u_1,\ldots,u_q\}$ and let $c_1,\ldots,c_p \in C$ be such that $(c_1(u_1),\ldots,c_1(u_q)),\ldots,(c_p(u_1),\ldots,c_p(u_q))$ are in pairwise different orbits of q-tuples under $\operatorname{Aut}(\mathbb{A})$ and such that for every $d \in C$

there exists $i \in [p]$ such that $(d(u_1), \dots, d(u_q))$ is in the same orbit under $\mathrm{Aut}(\mathbb{A})$ as $(c_i(u_1), \dots, c_i(u_q))$.

Set for every $C \in \mathcal{C}$,

$$C' := \{t(c_1, \dots, c_p, a_1, \dots, a_r) \mid t \in T, a_i \in C, \text{ all variables of } t \text{ essential}\}.$$

Set now $\mathcal{J}$ to be the instance obtained from $\mathcal{I}$ by replacing every constraint C by C'. Clearly, $\mathcal{J}$ is non-trivial, for every k-tuple of variables $\mathbf{v}$ it holds that if $\mathrm{proj}_{\mathbf{v}}(\mathcal{I}) \subseteq R$ and $S \cap \mathrm{proj}_{\mathbf{v}}(\mathcal{I}) \neq \emptyset$ then $S \cap \mathrm{proj}_{\mathbf{v}}(\mathcal{J}) \trianglelefteq_2 R$ and $C' \subseteq C$ for every $C \in \mathcal{C}$. It remains to show that $\mathcal{J}$ is (m, n)-minimal.

To this end, let $1 \leq m' \leq m$ and let $\mathbf{v} \in \mathcal{V}^{m'}$, let $C, D \in \mathcal{C}$ be such that $\mathrm{proj}_{\mathbf{v}}(C) = \mathrm{proj}_{\mathbf{v}}(D)$ and let $c \in C'$. We need to find $d \in D'$ such that $c(\mathbf{v}) = d(\mathbf{v})$. We may suppose that there exist $d_1, d_2 \in D$ such that $d_1(\mathbf{v}), d_2(\mathbf{v})$ lie in different orbits under $\mathrm{Aut}(\mathbb{A})$ since otherwise, $d(\mathbf{v})$ is in the same orbit of m'-tuples under $\mathrm{Aut}(\mathbb{A})$ as $c(\mathbf{v})$ for any $d \in D$ by the m'-minimality of $\mathcal{I}$, and hence, for any $d \in D'$ we may find $\alpha \in \mathrm{Aut}(\mathbb{A})$ such that $\alpha(d(\mathbf{v})) = c(\mathbf{v})$. By the definition of D', it follows that $\alpha d \in D'$.

We have $c = t(c_1, \dots, c_p, a_1, \dots, a_r)$ for some $a_1, \dots, a_r \in C, t \in T$. Suppose that D' is defined using $d_1, \dots, d_{p'} \in D$ and let $s \in T$ be of arity $p' \geq 2$ (by our assumption from the previous paragraph). Then $\alpha_j \circ s(d_1, \dots, d_{p'})(\mathbf{v}) = c_j(\mathbf{v})$ for some $j \in [p]$ and some $\alpha_j \in \mathrm{Aut}(\mathbb{A})$. Find $x_i \in D$ with $c_i(\mathbf{v}) = x_i(\mathbf{v})$ for every $i \in [p]$ and $y_i \in D$ with $a_i(\mathbf{v}) = y_i(\mathbf{v})$ for every $i \in [r]$. Set now $d := t(x_1, \dots, x_{j-1}, \alpha_j(s(d_1, \dots, d_{p'})), x_{j+1}, \dots, x_p, y_1, \dots, y_r)$. It is easy to see that $d \in D'$.

It follows that $c(\mathbf{v}) = t(c_1, \dots, c_p, a_1, \dots, a_r)(\mathbf{v}) = d(\mathbf{v})$ as desired. □

Lemma 3.5.1 together with the results from [71, 72] give the following corollary.

Corollary 3.5.2. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ which is a model-complete core. Suppose that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj} \curvearrowright \{E, N\}$ is equationally non-affine. Then every non-trivial $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal instance equivalent to an instance of* $\mathrm{CSP}(\mathbb{A})$ *has a solution. Moreover, such a solution s exists where $s(x) \neq s(y)$ for all variables x, y such that* $\mathrm{proj}_{(x,y)}(\mathcal{I}) \cap I_2 \neq \emptyset$.

Proof. It follows from [67, 64] that for every $k \geq 3$, $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ contains an operation of arity k that acts as a WNU operation on $\{E, N\}$. An easy modification of the proof of Theorem 2.1.2 yields that $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$ has relational width $(2\ell, \max(3\ell, b_{\mathbb{H}}))$. Moreover, by Proposition 3.4.5, there exists a binary injection $f \in \mathrm{Pol}(\mathbb{A})$. Therefore, for all $a, b \in A^2$ such that b is injective, $f(a, b), f(b, a) \in I_2$ and hence, $I_2 \triangleleft_2 A^2$. Finally, Lemma 3.5.1 yields the result. □

3.5.2 $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathbf{inj}} \curvearrowright \{E, N\}$ is equationally affine

Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ which is a model-complete core. If $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally affine but equationally non-trivial, then $\mathbb{A}$ is not a first-order reduct of $(H; =)$ and it follows from Proposition 3.4.5 that $\mathrm{Pol}(\mathbb{A})$ contains a binary injection p_1 canonical with respect to $(\mathbb{H}, <)$ that acts as the first projection on $\{E, N\}$.

Moreover, we can suppose that the function $f(x, y) := p_1(y, x)$ acts lexicographically on the order by the remark above Lemma 3.4.4. It follows from [75] and from [29] that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ contains a ternary function m' that acts idempotently and as a minority on $\{E, N\}$.

We may moreover suppose that m' is canonical with respect to $(\mathbb{H},<)$ since any function locally interpolated by a function that acts idempotently and as a minority on $\{E,N\}$ acts idempotently and as a minority on $\{E,N\}$. It can be easily seen that the function $m(x,y,z) = m'(p_1(x,p_1(y,z)),p_1(y,p_1(z,x)),p_1(z,p_1(x,y)))$ is a ternary injection canonical with respect to $(\mathbb{H},<)$ that acts like a minority on $\{E,N\}$.

For $\mathbf{a} \in H^\ell$, we write $O(\mathbf{a})$ for the orbit of $\mathbf{a}$ under $\mathrm{Aut}(\mathbb{H})$ and $O_<(\mathbf{a})$ for the orbit of $\mathbf{a}$ under $\mathrm{Aut}(\mathbb{H},<)$. We say that a non-injective orbit O of ℓ-tuples under $\mathrm{Aut}(\mathbb{H})$ is:

- *deterministic* if for every $\mathbf{a} \in O$, there exists $\alpha \in \mathrm{Aut}(\mathbb{H})$ such that $p_1(O_<(\alpha(\mathbf{a})),E) = p_1(O_<(\alpha(\mathbf{a})),N)$, where the ordering on the second coordinate is strictly increasing,
- *non-deterministic* otherwise.

For a tuple $\mathbf{a}$ contained in a deterministic orbit O, we call any $\alpha \in \mathrm{Aut}(\mathbb{H})$ with the property that $p_1(O_<(\alpha(\mathbf{a})),E) = p_1(O_<(\alpha(\mathbf{a})),N)$ for the strictly increasing ordering in the second coordinate *deterministic for* $\mathbf{a}$. Note that for any $\beta \in \mathrm{Aut}(\mathbb{H},<)$, $\beta\alpha \in Aut(\mathbb{H})$ is deterministic for $\mathbf{a}$ as well since p_1 is canonical with respect to $(\mathbb{H},<)$.

3.5.3 Injectivisation of instances

Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$. Let $\mathcal{I} = (\mathcal{V},\mathcal{C})$ be an instance of $\mathrm{CSP}(\mathbb{A})$. In this section, we always assume that the variable set $\mathcal{V}$ is equipped with an arbitrary linear order; this assumption is inessential and is used to formulate the statements and proofs in a more concise way. We denote by $[\mathcal{V}]^\ell$ the set of injective increasing ℓ-tuples of variables from $\mathcal{V}$. Given any instance $\mathcal{I}$ of $\mathrm{CSP}(\mathbb{A})$, consider the following CSP instance $\mathcal{I}_{\mathrm{fin}}$ over the set $\mathcal{O}$ of orbits of ℓ-tuples under $\mathrm{Aut}(\mathbb{H})$:

- The variable set of $\mathcal{I}_{\mathrm{fin}}$ is the set $[\mathcal{V}]^\ell$.
- For every constraint $C \subseteq A^U$ in $\mathcal{I}$, $\mathcal{I}_{\mathrm{fin}}$ contains the constraint C' containing the maps $g\colon [U]^\ell \to \mathcal{O}$ such that there exists $f \in C$ satisfying $f(\mathbf{v}) \in g(\mathbf{v})$ for every $\mathbf{v} \in [U]^\ell$.

We introduced this instance already in Definition 2.3.1 in a more general setting, where it is denoted by $\mathcal{I}_{\mathrm{Aut}(\mathbb{H}),\ell}$. In the original definition, ℓ-element subsets of $\mathcal{V}$ were used as variables and the domain consisted of orbits of maps. However, the translation between the two definitions is straightforward.

Let $\mathcal{J} = (S,\mathcal{C})$ be an instance over the set $\mathcal{O}$ of orbit of ℓ-tuples under $\mathrm{Aut}(\mathbb{H})$, e.g., $\mathcal{J} = \mathcal{I}_{\mathrm{fin}}$ for some $\mathcal{I}$. The *injectivisation* of $\mathcal{J}$, denoted by $\mathcal{J}^{(\mathrm{inj})}$, is the instance obtained by replacing every constraint $C \in \mathcal{C}$ with scope $U \subseteq S$ by $\{g \in C \mid g(s) \in \{E,N\}$ for every $s \in U\}$.

Let $\mathcal{I} = (\mathcal{V},\mathcal{C})$ be an instance of $\mathrm{CSP}(\mathbb{A})$ and let $S \subseteq [\mathcal{V}]^\ell$. The *finite injectivisation of* $\mathcal{I}$ *on* S is the instance $\mathcal{I}_{\mathrm{fin}}^{(\mathrm{inj})}|_S$. If $S = [\mathcal{V}]^\ell$, we call the finite injectivisation of $\mathcal{I}$ on S just the finite injectivisation of $\mathcal{I}$. For any constraint $C \in \mathcal{C}$, the corresponding constraint in the finite injectivisation of $\mathcal{I}$ is called the finite injectivisation of C. Note that if $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E,N\}$ is equationally non-trivial, $\mathcal{I} = (\mathcal{V},\mathcal{C})$ is an instance of $\mathrm{CSP}(\mathbb{A})$ and $S \subseteq [\mathcal{V}]^\ell$, the finite injectivisation of $\mathcal{I}$ on S is solvable in polynomial time by Lemma 2.3.4 and by the dichotomy theorem for finite-domain CSPs [83, 84, 46].

Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$, let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be an instance of $\mathrm{CSP}(\mathbb{A})$ and let $W \subseteq \mathcal{V}$. Let $g \colon W \to A$ and let $h \colon \mathcal{V} \to A$, we say that g is *ordered by* h if for all $u, v \in W$, we have $g(u) < g(v)$ if, and only if, $h(u) < h(v)$.

Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ preserved by m and by p_1. We can assume that $\mathbb{A}$ has among its relations all unions of orbits of ℓ-tuples under $\mathrm{Aut}(\mathbb{H})$ that are preserved by p_1 and by the ternary injection m. Otherwise, we expand $\mathbb{A}$ by these finitely many relations and we prove that the CSP of this expanded structure is solvable in polynomial time. Note that in particular, every orbit of ℓ-tuples under $\mathrm{Aut}(\mathbb{H})$ is a relation of $\mathbb{A}$. Let n be the maximal arity of a relation from the signature of $\mathbb{A}$. We will moreover suppose that $\mathbb{A}$ has among its relations all relations of arity at most $\max(3\ell, b_{\mathbb{H}}, n)$ that are pp-definable in $\mathbb{A}$. Hence, for every instance $\mathcal{I}$ of $\mathrm{CSP}(\mathbb{A})$, the $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal instance equivalent to $\mathcal{I}$ is again an instance of $\mathrm{CSP}(\mathbb{A})$.

Let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be an instance of $\mathrm{CSP}(\mathbb{A})$ and let $C \in \mathcal{C}$. Since $\mathbb{A}$ is preserved by m, there exists a set of linear equations over $\mathbb{Z}_2$ associated with the finite injectivisation of C. By abuse of notation, we write every linear equation as $\sum_{\mathbf{v} \in S} X_{\mathbf{v}} = P$, where $P \in \{E, N\}$ and S is a set of injective ℓ-tuples of variables from the scope of C. In these linear equations, we identify E with 1 and N with 0, so that e.g. $E + E = N$ and $N + E = E$.

We may assume that no equation $\sum_{\mathbf{v} \in S} X_{\mathbf{v}} = P$ associated with the finite injectivisation of any constraint $C \in \mathcal{C}$ splits into two equations $\sum_{\mathbf{v} \in S_1} X_{\mathbf{v}} = P_1$ and $\sum_{\mathbf{v} \in S_2} X_{\mathbf{v}} = P_2$ where $S_1 \cap S_2 = \emptyset$, $S_1 \cup S_2 = S$, $P_1 + P_2 = P$ and such that every element of the finite injectivisation of C satisfies both these equations. If this is not the case, let us consider only equations satisfying this assumption – it is clear that this new set of equations is satisfiable if, and only if, the original set of equations is satisfiable. We will call every equation satisfying this assumption *unsplittable*.

For an instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of $\mathrm{CSP}(\mathbb{A})$, we define an instance $\mathcal{I}_{\mathrm{eq}} = (\mathcal{V}, \mathcal{C}_{\mathrm{eq}})$ of the equality-CSP over the same base set H corresponding to the closure of the constraints under the full symmetric group on H. Formally, for every constraint $C \in \mathcal{C}$, the corresponding constraint $C_{\mathrm{eq}} \in \mathcal{C}_{\mathrm{eq}}$ contains all functions αh for all $h \in C$ and $\alpha \in \mathrm{Sym}(H)$. Since $\mathbb{A}$ is preserved by a binary injection, $\mathcal{I}_{\mathrm{eq}}$ is preserved by the same or indeed any binary injection and hence, its CSP has relational width $(2, 3)$ by the classification of equality CSPs [22].

Let $\mathcal{I}$ be an ℓ-minimal instance of $\mathrm{CSP}(\mathbb{A})$, let $\mathbf{v} \in [\mathcal{V}]^{\ell}$ and let $R \subseteq \mathrm{proj}_{\mathbf{v}}(\mathcal{I})$ be an ℓ-ary relation from the signature of $\mathbb{A}$. Let $\mathcal{I}^{\mathbf{v} \in R}$ be the instance obtained from $\mathcal{I}$ by replacing every constraint C containing all variables from $\mathbf{v}$ by $\{g \in C \mid g(\mathbf{v}) \in R\}$.

We call an ℓ-minimal instance of $\mathrm{CSP}(\mathbb{A})$ *eq-subdirect* if for every $\mathbf{v} \in [\mathcal{V}]^{\ell}$ and for every non-injective orbit $O \subseteq \mathrm{proj}_{\mathbf{v}}(\mathcal{I})$ under $\mathrm{Aut}(\mathbb{H})$, the instance $(\mathcal{I}^{\mathbf{v} \in O})_{\mathrm{eq}}$ has a solution. Note that by ℓ-minimality and since the instance is preserved by a binary injection, the instance $(\mathcal{I}^{\mathbf{v} \in O})_{\mathrm{eq}}$ has a solution for every injective orbit $O \in \mathrm{proj}_{\mathbf{v}}(\mathcal{I})$ under $\mathrm{Aut}(\mathbb{H})$.

It is clear that we can obtain an eq-subdirect instance out of an ℓ-minimal instance in polynomial time by the algorithm in Figure 3.1.

Lemma 3.5.3. *The instance $\mathcal{I}'$ outputted by the algorithm in Figure 3.1 is ℓ-minimal, eq-subdirect, and it is an instance of* $\mathrm{CSP}(\mathbb{A})$.

Proof. It is easy to see that $\mathcal{I}'$ is ℓ-minimal since in every run of the repeat... until not

```
INPUT: ℓ-minimal instance I = (V, C) of CSP(A);
OUTPUT: eq-subdirect instance I';
repeat
    changed:=false;
    for v ∈ [V]^ℓ do
        if not changed then
            P := proj_v(I) ∩ I;
            for O ⊆ proj_v(I) non-injective orbit do
                if (I^{v∈O})_eq has a solution then
                    P := P ∪ O;
                end if
            end for
            for C ∈ C containing all variables of v in its scope do
                replace C by {g ∈ C | g(v) ∈ P};
            end for
            if I is not ℓ-minimal then
                make I ℓ-minimal;
                changed:=true;
            end if
        end if
    end for
until not changed
return I
```

Figure 3.1: Procedure Eq-Subdirect

changed-loop, we check ℓ-minimality of the produced instance.

To prove that $\mathcal{I}'$ is an instance of CSP($\mathbb{A}$), it is sufficient to show that for every $\mathbf{v} \in [\mathcal{V}]^\ell$ the last P in the first for-loop is preserved by the binary injection p_1 as well as by the ternary injection m. To this end, let $\mathbf{v} \in [\mathcal{V}]^\ell$ be injective, let P be as in the algorithm, let O_1, O_2, O_3 be orbits of ℓ-tuples from P and let $g_i, i \in \{1,2,3\}$ be solutions to $(\mathcal{I}^{\mathbf{v} \in O_i})_{\text{eq}}$, $i \in \{1,2,3\}$.

Let $C \in \mathcal{C}$ be arbitrary and let U be its scope. By the definition of $\mathcal{I}_{\text{eq}}$, there exist $h_1, h_2, h_3 \in C$ and $\alpha_1, \alpha_2, \alpha_3 \in \text{Sym}(H)$ such that $g_i|_U = \alpha_i h_i$ for every $i \in \{1,2,3\}$. Since every binary injection is canonical in its action on $\{=, \neq\}$, it follows that there exist $\beta_1, \beta_2 \in \text{Sym}(H)$ such that $p_1(g_1, g_2)|_U = p_1(\alpha_1 h_1, \alpha_2 h_2) = \beta_1 p_1(h_1, h_2) \in C$ and $m(g_1, g_2, g_3)|_U = m(\alpha_1 h_1, \alpha_2 h_2, \alpha_3 h_3) = \beta_2 m(h_1, h_2, h_3) \in C$. Hence, $p_1(g_1, g_2)|_U, m(g_1, g_2, g_3)|_U \in C$ and since C was chosen arbitrarily, both $p_1(g_1, g_2)$ and $m(g_1, g_2, g_3)$ are solutions to $(\mathcal{I}^{\mathbf{v} \in O_i})_{\text{eq}}$.

Let $\mathbf{v} \in [\mathcal{V}]^\ell$. If C contains all variables from $\mathbf{v}$ in its scope, then either $p_1(g_1, g_2)(\mathbf{v})$ is injective and $p_1(g_1, g_2)(\mathbf{v}) \in P$ or $p_1(g_1, g_2)(\mathbf{v})$ and $p_1(h_1, h_2)(\mathbf{v})$ are non-injective and contained in the same orbit under Aut($\mathbb{H}$) and it follows that $p_1(h_1, h_2)(\mathbf{v}) \in P$. Similar argument holds for $m(h_1, h_2, h_3)$ as well. Therefore, P is preserved by p_1 and m as desired. $\square$

Note that the eq-subdirect instance $\mathcal{I}'$ outputted by the algorithm has the same solution set as $\mathcal{I}$. Moreover, for any $1 \leq m \leq n$ and any instance $\mathcal{I}$ of CSP($\mathbb{A}$), we can compute an instance that is both eq-subdirect and (m,n)-minimal and that has the same solution set as $\mathcal{I}$ in polynomial time. Indeed, it is enough to repeat the above-mentioned algorithm and the (m,n)-minimality algorithm until no orbits under Aut($\mathbb{H}$) are removed from any constraint.

3.5.4 Inj-irreducibility

Let $\mathcal{J} = (\mathcal{V}, \mathcal{C})$ be a CSP instance over a set B. A *path* in $\mathcal{J}$ is a sequence of the form $v_1, C_1, v_2, \ldots, C_k, v_{k+1}$, where $k \geq 1$, $v_i \in \mathcal{V}$ for every $i \in [k+1]$, $C_i \in \mathcal{C}$ for every $i \in [k]$ and v_i, v_{i+1} are contained in the scope of C_i for every $i \in [k]$. We say that two elements $a, b \in B$ are *connected* by a path $v_1, C_1, v_2, \ldots, C_k, v_{k+1}$ if there exists a tuple $(c_1, \ldots, c_{k+1}) \in B^{k+1}$ such that $c_1 = a, c_{k+1} = b$ and such that $(c_i, c_{i+1}) \in \text{proj}_{(v_i, v_{i+1})}(C_i)$ for every $i \in [k]$. Let $\mathcal{J} = (\mathcal{V}, \mathcal{C})$ be a 1-minimal instance over the set B and let $v \in \mathcal{V}$. The *linkedness congruence* on $\text{proj}_v(\mathcal{J})$ is the equivalence relation λ on $\text{proj}_v(\mathcal{J})$ defined by $(a,b) \in \lambda$ if there exists a path from a to b in $\mathcal{J}$. Note that for a finite relational structure $\mathbb{B}$, for a $(2,3)$-minimal instance $\mathcal{J} = (\mathcal{V}, \mathcal{C})$ of CSP($\mathbb{B}$) and for any $v \in \mathcal{V}$, the linkedness congruence λ on $\text{proj}_v(\mathcal{J})$ is a relation pp-definable in $\mathbb{B}$. Indeed, it is easy to see that for every $k \geq 1$, the binary relation containing precisely the pairs $(a,b) \in B^2$ that are connected by a particular path in $\mathcal{I}$ is pp-definable in $\mathbb{B}$. If we concatenate all paths that connect two elements $(a,b) \in \lambda$, the resulting path will connect every pair $(a,b) \in \lambda$ since by the $(2,3)$-minimality of $\mathcal{J}$, every path from v to v connects c to c for every $c \in \text{proj}_v(\mathcal{J})$. It follows that λ is pp-definable.

Definition 3.5.4. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ and let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be a non-trivial ℓ-minimal instance of* CSP($\mathbb{A}$). *We call $\mathcal{I}$* inj-irreducible *if for every set $S \subseteq [\mathcal{V}]^\ell$, one of the following holds for the instance $\mathcal{J} = \mathcal{I}_{\text{fin}}|_S$:*

- *$\mathcal{J}^{(inj)}$ has a solution,*

- *for some* $\mathbf{v} \in S$, $\text{proj}_{\mathbf{v}}(\mathcal{J})$ *contains the two injective orbits and the linkedness congruence on* $\text{proj}_{\mathbf{v}}(\mathcal{J})$ *does not connect them,*

- *for some* $\mathbf{v} \in S$, *the linkedness congruence on* $\text{proj}_{\mathbf{v}}(\mathcal{J})$ *links an injective orbit to a non-deterministic orbit.*

Lemma 3.5.5. *Let* $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ *be a non-trivial* 2*-minimal instance of* $\text{CSP}(\mathbb{A})$ *such that* $\text{proj}_{(u,v)}(\mathcal{I}) \cap I_2 \neq \emptyset$ *for every* $u \neq v \in \mathcal{V}$. *Let* $S \subseteq [\mathcal{V}]^\ell$ *be a set of variables appearing together in an unsplittable linear equation associated with the finite injectivisation of some constraint* $C \in \mathcal{C}$. *Then for every* $g \in C$ *either* $g(\mathbf{s})$ *is in an injective or deterministic orbit for all* $\mathbf{s} \in S$ *or* $g(\mathbf{s})$ *is in a non-deterministic orbit for all* $\mathbf{s} \in S$.

Proof. Suppose not and let $g \in C$ be a counterexample. Let $\emptyset \neq S' \subsetneq S$ be the set of all $\mathbf{s} \in S$ such that $g(\mathbf{s})$ is in an injective or deterministic orbit. Let $\sum_{\mathbf{s} \in S} X_{\mathbf{s}} = P$ be an unsplittable linear equation associated with the finite injectivisation of C and containing all the variables from S.

Let $\{\mathbf{s}_1, \ldots, \mathbf{s}_n\}$ be the set of all $\mathbf{s}_i \in S'$ such that $g(\mathbf{s}_i)$ is in a deterministic orbit. For every $i \in [n]$, let $\alpha_i \in \text{Aut}(\mathbb{H})$ be deterministic for $g(\mathbf{s}_i)$. Let

$$p'_1(x, y) := p_1(\alpha_1 x, p_1(\alpha_2 x, p_1(\ldots, p_1(\alpha_n x, y) \ldots))).$$

It is easy to see that for all $\mathbf{s} \in S$, we have that $p'_1(g(\mathbf{s}), \mathbf{a})$ and $p'_1(g(\mathbf{s}), \mathbf{b})$ are contained in the same injective orbit $O_{\mathbf{s}}$ under $\text{Aut}(\mathbb{H})$ for all increasing $\mathbf{a} \in E, \mathbf{b} \in N$. For all $\mathbf{s} \in S \backslash S'$, let $P_{\mathbf{s}}$ be N if $p'_1(g(\mathbf{s}), \mathbf{a}) \in E$ for every increasing $\mathbf{a} \in E$ and let $P_{\mathbf{s}}$ be E otherwise.

Since every binary projection of $\mathcal{I}$ (and hence of C) has a non-empty intersection with I_2 and since $\mathbb{A}$ is preserved by a binary injection, it follows that there exists a monotone injective $g' \in C$. Moreover, since $S' \subsetneq S$ and the equation under consideration is unsplittable, one can assume that $\sum_{\mathbf{s} \in S'} O(g'(\mathbf{s})) = \sum_{\mathbf{s} \in S'} O_{\mathbf{s}} + \sum_{\mathbf{s} \in S \backslash S'} P_{\mathbf{s}} + E$. Let us consider $h = p'_1(g, g') \in C$. h is clearly injective and we obtain:

$$\begin{aligned}
\sum_{\mathbf{s} \in S} O(h(\mathbf{s})) &= \sum_{\mathbf{s} \in S'} O(h(\mathbf{s})) + \sum_{\mathbf{s} \in S \backslash S'} O(h(\mathbf{s})) = \sum_{\mathbf{s} \in S'} O_{\mathbf{s}} + \sum_{\mathbf{s} \in S \backslash S'} O(g'(\mathbf{s})) + \sum_{\mathbf{s} \in S \backslash S'} P_{\mathbf{s}} \\
&= \sum_{\mathbf{s} \in S'} O(g'(\mathbf{s})) + E + \sum_{\mathbf{s} \in S \backslash S'} O(g'(\mathbf{s})) = \sum_{\mathbf{s} \in S} O(g'(\mathbf{s})) + E \\
&= P + E
\end{aligned}$$

Hence, the mapping $h' \colon S \to \{E, N\}$ defined by $h'(\mathbf{s}) := O(h(\mathbf{s}))$ is a solution to the finite injectivisation of C on S but it does not satisfy the unsplittable linear equation, a contradiction. □

Theorem 3.5.6. *Let* $\mathbb{A}$ *be a first-order reduct of* $\mathbb{H}$ *which is a model-complete core and suppose that* $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj} \curvearrowright \{E, N\}$ *is equationally non-trivial. Let* $\mathcal{I}$ *be a* $(2\ell, \max(3\ell, b_{\mathbb{H}}))$*-minimal, eq-subdirect, inj-irreducible instance of* $\text{CSP}(\mathbb{A})$ *with variables* $\mathcal{V}$ *such that for every distinct* $u, v \in \mathcal{V}$, $\text{proj}_{(u,v)}(\mathcal{I}) \cap I_2 \neq \emptyset$. *Then* $\mathcal{I}$ *has a solution if, and only if, it has an injective one.*

Proof. In the case when $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally non-affine, the statement follows directly from Lemma 3.5.1 and Corollary 3.5.2, where only the $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimality of $\mathcal{I}$ and the fact that for every distinct $u, v \in \mathcal{V}$, $\mathrm{proj}_{(u,v)}(\mathcal{I}) \cap I_2 \neq \emptyset$ are needed.

Suppose henceforth that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally affine (and equationally non-trivial by assumption). Note that if $\mathcal{I}$ has less than ℓ variables, it has an injective solution by the assumption on binary projections of $\mathcal{I}$ and since it is preserved by a binary injection by Proposition 3.4.5. Let us therefore suppose that $\mathcal{I}$ has at least ℓ variables. Let us assume for the sake of contradiction that $\mathcal{I}$ does not have an injective solution. Let $\mathcal{J}$ be $\mathcal{I}_{\mathrm{fin}}$ and let $\mathcal{C}$ be the set of its constraints. By assumption, $\mathcal{J}^{(\mathrm{inj})}$ does not have a solution. Note that $\mathcal{J}^{(\mathrm{inj})}$ corresponds to a system of linear equations over $\mathbb{Z}_2$, which is therefore unsatisfiable. In case this system can be written as a block matrix, there exists a set $S \subseteq [\mathcal{V}]^{\ell}$ of variables such that the system of equations associated with the injectivisation of $\mathcal{K} := \mathcal{J}|_S = (S, \mathcal{C}')$ corresponds to a minimal block, and is therefore unsatisfiable. By definition, this means that $\mathcal{K}^{(\mathrm{inj})}$ is unsatisfiable. The instance $\mathcal{K}$ has the property that for every constraint $C \in \mathcal{C}'$, and every non-trivial partition of the scope of C into parts S_1, S_2, some unsplittable equation associated with C contains a variable from S_1 and a variable from S_2.

Since $\mathcal{I}$ is inj-irreducible, there exists $\mathbf{v} \in S$ such that the two injective orbits are elements of $\mathrm{proj}_{\mathbf{v}}(\mathcal{K})$ and are not linked, or some injective orbit in $\mathrm{proj}_{\mathbf{v}}(\mathcal{K})$ is linked to a non-deterministic orbit in $\mathrm{proj}_{\mathbf{v}}(\mathcal{K})$.

In the first case, we note that for *all* $\mathbf{w} \in S$ such that $\mathrm{proj}_{\mathbf{w}}(\mathcal{K})$ contains the two injective orbits, the two injective orbits are not linked. Indeed, suppose that there exists $\mathbf{w} \in S$ such that $E, N \in \mathrm{proj}_{\mathbf{w}}(\mathcal{K})$ are linked, i.e., there exists a path $\mathbf{v}_1 = \mathbf{w}, C_1, \ldots, C_k, \mathbf{v}_{k+1} = \mathbf{w}$ in $\mathcal{K}$ connecting E and N. Since $\mathcal{I}$ is $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal, Lemma 2.3.2 yields that $\mathcal{J}$ and hence also $\mathcal{K}$ is $(2,3)$-minimal. In particular, there exists $C \in \mathcal{C}'$ containing in its scope both $\mathbf{v}$ and $\mathbf{w}$. Let $O_1, O_2 \in \{E, N\}$ be disjoint such that there exist $g_1, g_2 \in C$ with $g_1(\mathbf{v}) \in O_1, g_1(\mathbf{w}) \in E, g_2(\mathbf{v}) \in O_2, g_2(\mathbf{w}) \in N$. It follows that the path $\mathbf{v}, C, \mathbf{w} = \mathbf{v}_1, C_1, \ldots, C_k, \mathbf{v}_{k+1} = \mathbf{w}, C, \mathbf{v}$ connects O_1 with O_2 in $\mathrm{proj}_{\mathbf{v}}(\mathcal{J})$, a contradiction. Let $g \colon S \to \{E, N\}$ be defined as follows. For a fixed $\mathbf{v} \in S$, let $g(\mathbf{v})$ be an arbitrary element of $\mathrm{proj}_{\mathbf{v}}(\mathcal{K}^{(\mathrm{inj})})$. Next, for $\mathbf{w} \in S$, define $g(\mathbf{w})$ to be the unique injective orbit O such that there exists a path in $\mathcal{K}$ from $\mathbf{v}$ to $\mathbf{w}$ connecting $g(\mathbf{v})$ to O. This g is a solution to $\mathcal{K}^{(\mathrm{inj})}$, a contradiction.

Thus, it must be that a non-deterministic orbit in $\mathrm{proj}_{\mathbf{v}}(\mathcal{K})$ is linked to an injective orbit in $\mathrm{proj}_{\mathbf{v}}(\mathcal{K})$. Hence, there exists a path in $\mathcal{K}$ from $\mathbf{v}$ to $\mathbf{v}$ and connecting an injective orbit to a non-deterministic one. Moreover, up to composing this path with additional constraints, one can assume that this path goes through all the variables in S. This follows by the $(2,3)$-minimality of $\mathcal{J}$. Define a partition of S where $\mathbf{w} \in S_1$ if the first time that $\mathbf{w}$ appears in the path, the element associated with $\mathbf{w}$ is in an injective orbit, and $\mathbf{w} \in S_2$ otherwise. Since the system of unsplittable equations associated with $\mathcal{K}^{(\mathrm{inj})}$ cannot be decomposed as a block matrix, some constraint $C \in \mathcal{C}'$ gives an equation in that system containing $\mathbf{u}_1 \in S_1$ and $\mathbf{u}_2 \in S_2$. Thus, there exists $g \in C$ with $g(\mathbf{u}_1)$ injective, and $g(\mathbf{u}_2)$ non-deterministic. This contradicts Lemma 3.5.5. □

INPUT: instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ of CSP($\mathbb{A}$);
OUTPUT: instance $\mathcal{I}'$ that is either inj-irreducible or trivial and that has a solution if, and only, if $\mathcal{I}$ has a solution;
repeat
  changed := false;
  $\mathcal{J} := (2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal eq-subdirect instance with the same solution set as $\mathcal{I}$;
  $\mathcal{I} := \mathcal{J}$;
  $\mathcal{I} :=$ IdentifyAllEqual($\mathcal{I}$);
  for $\mathbf{u} \in [\mathcal{V}]^\ell$ **do**
    for $\{E^1_{\mathbf{u}}, \ldots, E^m_{\mathbf{u}}\}$ partition on $\text{proj}_{\mathbf{u}}(\mathcal{I})$ with pp-definable classes such that $\text{proj}_{\mathbf{u}}(\mathcal{I}) \cap I \subseteq E^1_{\mathbf{u}}$ and $E^1_{\mathbf{u}}$ contains no non-deterministic orbit **do**
      $S, \{E^i_{\mathbf{w}} \mid i \in [m], \mathbf{w} \in S\} :=$ ExtendPartition($\mathcal{I}, \{\mathbf{u}\}, \{E^i_{\mathbf{u}} \mid i \in [m]\}$);
      solve the finite injectivisation of $\mathcal{I}$ on S;
      if it does not have a solution and not changed **then**
        changed := true;
        for $\mathbf{v} \in S$ **do**
          for $C \in \mathcal{C}$ containing all variables of $\mathbf{v}$ in its scope **do**
            replace C by $\{f \in C \mid f(\mathbf{v}) \notin E^1_{\mathbf{v}}\}$;
          end for
        end for
      end if
    end for
  end for
until not changed;
return $\mathcal{I}$

Figure 3.2: Procedure InjIrreducibility

3.5.5 Establishing inj-irreducibility

We show that the algorithm in Figure 3.2 produces, given an instance $\mathcal{I}$ of CSP($\mathbb{A}$), an instance $\mathcal{I}'$ of CSP($\mathbb{A}$) that is either inj-irreducible or trivial and that has a solution if, and only if, $\mathcal{I}$ has a solution. It uses the fact that the finite injectivisation of an instance $\mathcal{I}$ of CSP($\mathbb{A}$) on S is solvable in polynomial time for any set $S \subseteq [\mathcal{V}]^\ell$. This follows from the fact that the finite injectivisation of $\mathcal{I}$ is preserved by a ternary minority by Lemma 2.3.4.

The algorithm uses the subroutines IdentifyAllEqual and ExtendPartition. The subroutine IdentifyAllEqual takes as an input a 2-minimal instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ and returns an instance where all variables $u, v \in \mathcal{V}$ with $\text{proj}_{(u,v)}(\mathcal{I}) = \{(a,a) \mid a \in H\}$ are identified. It is clear that this can be implemented in polynomial time. Note moreover that for a $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal instance $\mathcal{I}$, the resulting instance IdentifyAllEqual($\mathcal{I}$) is again $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal. The subroutine ExtendPartition is given in Figure 3.3. We note that if $(S, \{E^i_{\mathbf{w}}\})$ is a set of partitions returned by ExtendPartition, then for every $\mathbf{w} \in S$, the linkedness congruence on $\text{proj}_{\mathbf{w}}(\mathcal{I}_{\text{fin}})$ defined by the instance $\mathcal{I}_{\text{fin}}|_S$ is a refinement of the partition $\{E^i_{\mathbf{w}}\}$.

First of all, note that the instance $\mathcal{I}'$ outputted by this algorithm is $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-

INPUT: a tuple $(\mathcal{I}, S, \{E^i_{\mathbf{w}} \mid i \in [m], \mathbf{w} \in S\})$ where

- $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ is a $((2\ell, \max(3\ell, b_{\mathbb{H}})))$-minimal instance of $\mathrm{CSP}(\mathbb{A})$,
- $S \subseteq [\mathcal{V}]^\ell$,
- $\{E^1_{\mathbf{w}}, \ldots, E^m_{\mathbf{w}}\}$ is a partition on $\mathrm{proj}_{\mathbf{w}}(\mathcal{I})$ with pp-definable classes for every $\mathbf{w} \in S$ such that $E^1_{\mathbf{w}}$ contains no non-deterministic orbit and such that $E^i_{\mathbf{w}} := \{\mathbf{a} \in H^\ell \mid \exists \mathbf{b} \in H^\ell \colon (\mathbf{b}, \mathbf{a}) \in \mathrm{proj}_{(\mathbf{v},\mathbf{w})}(\mathcal{I})\}$ for every $\mathbf{v}, \mathbf{w} \in S$ and every $i \in [m]$;

OUTPUT: $S, \{E^i_{\mathbf{w}} : i \in [m], \mathbf{w} \in S\}$ as above such that no tuple $\mathbf{w}$ can be added to S where a partition on $\mathrm{proj}_{\mathbf{w}}(\mathcal{I})$ as above exists;

repeat
  added := false;
  for $\mathbf{v} \in S, \mathbf{w} \in [\mathcal{V}]^\ell$ with $\mathbf{w} \notin S$ **do**
    $D := \mathrm{proj}_{(\mathbf{v},\mathbf{w})}(\mathcal{I})$;
    for $t = 1, \ldots, m$ **do**
      $E^t_{\mathbf{w}} := \{\mathbf{a} \in H^\ell \mid \exists \mathbf{b} \in E^t_{\mathbf{v}} \colon (\mathbf{b}, \mathbf{a}) \in D\}$;
    end for
    if $E^1_{\mathbf{w}}, \ldots, E^m_{\mathbf{w}}$ are disjoint and $E^1_{\mathbf{w}}$ contains no non-deterministic orbit **then**
      $S := S \cup \{\mathbf{w}\}$;
      added := true;
    end if
  end for
until not added
return $S, \{E^i_{\mathbf{w}} : i \in [m], \mathbf{w} \in S\}$

Figure 3.3: Procedure ExtendPartition

minimal and for every $u \neq v \in \mathcal{V}$, $\mathcal{I}'_{(u,v)} \cap I_2 \neq \emptyset$. We will show in Lemma 3.5.8 that $\mathcal{I}'$ is an instance of CSP($\mathbb{A}$). Then it follows that if $\mathcal{I}'$ is non-trivial, then every constraint C' of $\mathcal{I}'$ contains an injective tuple since $\mathcal{I}'$ is preserved by the binary injection p_1.

Let $S \subseteq [\mathcal{V}]^\ell$ be the set outputted by ExtendPartition($\mathcal{I}, \{\mathbf{u}\}, \{E^i_{\mathbf{u}} \mid i \in [m]\}$) for some $\mathbf{u} \in [\mathcal{V}]^\ell$ as in the algorithm for an instance $\mathcal{I}$ with $\mathcal{I} =$ IdentifyAllEqual($\mathcal{I}$) – note that this subroutine is called in the algorithm only for instances satisfying this assumption. Note that for every $\mathbf{w} \in S$, $E^1_{\mathbf{w}}$ contains all injective orbits under Aut($\mathbb{H}$) contained in $\text{proj}_{\mathbf{w}}(\mathcal{I})$. This can be easily shown by induction of the size of S. For S containing just the tuple $\mathbf{u}$, there is nothing to prove. Suppose that it holds for S, let $\mathbf{w} \in \mathcal{V}^\ell$ be a tuple added to S and let $\mathbf{v} \in S$ be such that $E^1_{\mathbf{w}} = \{\mathbf{a} \in H^\ell \mid \exists \mathbf{b} \in H^\ell \colon (\mathbf{b}, \mathbf{a}) \in \text{proj}_{(\mathbf{v},\mathbf{w})}(\mathcal{I})\}$. Let $O \subseteq \text{proj}_{\mathbf{w}}(\mathcal{I})$ be an injective orbit under Aut($\mathbb{H}$) and let $C \in \mathcal{C}$ be a constraint containing all variables from the tuples $\mathbf{v}$ and $\mathbf{w}$ in its scope. Let $g_1 \in C$ be such that $g_1(\mathbf{w}) \in O$ and let $g_2 \in C$ be injective. Then $p_1(g_1, g_2) \in C$ is injective and hence, $p_1(g_1(\mathbf{v}), g_2(\mathbf{v})) \in E^1_{\mathbf{v}}$ by the induction assumption. Moreover, $p_1(g_1(\mathbf{w}), g_2(\mathbf{w})) \in O$ and $p_1(g_1(\mathbf{w}), g_2(\mathbf{w})) \in E^1_{\mathbf{w}}$ by the definition of $E^1_{\mathbf{w}}$.

It is easy to see that the algorithm runs in polynomial time: the algorithm of checking eq-subdirectness as well as the $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimality algorithm both run in polynomial time. Moreover, the number of runs of every for- or repeat...until-loop in the main algorithm as well as in the subroutine ExtendPartition is bounded by $c^c(d + \ell^{\max(3\ell, b_{\mathbb{H}}, n)})\binom{|\mathcal{V}|}{\ell}^2$, where n is maximum of the arities of relations in the signature of $\mathbb{A}$, c is the number of orbits of ℓ-tuples under Aut($\mathbb{H}$) and d is the number of constraints of the original instance $\mathcal{I}$.

Before proving that $\mathcal{I}'$ is an inj-irreducible instance of CSP($\mathbb{A}$) and that it has a solution if, and only if, $\mathcal{I}$ has a solution, we will show a few auxiliary statements.

Lemma 3.5.7. *Let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be an ℓ-minimal non-trivial instance of* CSP($\mathbb{A}$) *such that* $\text{proj}_{(u,v)}(\mathcal{I}) \cap I_2 \neq \emptyset$ *for every $u \neq v \in \mathcal{V}$. Let $S \subseteq [\mathcal{V}]^\ell$ be such that* $\text{proj}_{\mathbf{v}}(\mathcal{I})$ *contains some injective orbit and no non-deterministic orbit for every $\mathbf{v} \in S$. Suppose that $\mathcal{I}$ has a solution. Then the finite injectivisation of $\mathcal{I}$ on S has a solution as well.*

Proof. Let $g \colon \mathcal{V} \to H$ be a solution of $\mathcal{I}$ such that $g(\mathbf{s})$ is not injective for some $\mathbf{s} \in S$. Let $\{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ be the set of all $\mathbf{s}_i \in S'$ such that $g(\mathbf{s}_i)$ is in a non-injective orbit under Aut($\mathbb{H}$). For every $i \in [n]$, let $\alpha_i \in$ Aut($\mathbb{H}$) be deterministic for $g(\mathbf{s}_i)$. Let $p'_1(x, y) := p_1(\alpha_1 x, p_1(\alpha_2 x, p_1(\dots, p_1(\alpha_n x, y))))$.

It is easy to see that for all $\mathbf{s} \in S$, we have that $p'_1(g(\mathbf{s}), \mathbf{a})$ and $p'_1(g(\mathbf{s}), \mathbf{b})$ are contained in the same injective orbit $O_{\mathbf{s}}$ under Aut($\mathbb{H}$) for all increasing $\mathbf{a} \in E, \mathbf{b} \in N$.

We claim that $f \colon S \to \{E, N\}$ defined by $f(\mathbf{s}) = O_{\mathbf{s}}$ is a solution to the finite injectivisation of $\mathcal{I}$ on S. To this end, let $C \in \mathcal{C}$ and let $g' \in C$ be monotone and injective. It follows that $p'_1(g, g') \in C$ and moreover, $O(p'_1(g(\mathbf{s}), g'(\mathbf{s})) = O_{\mathbf{s}}$ for all $\mathbf{s} \in S$ such that all variables from $\mathbf{s}$ are in the scope of C and the lemma follows. □

Lemma 3.5.8. *Let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be a $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal, non-trivial, eq-subdirect instance of* CSP($\mathbb{A}$) *such that* $\text{proj}_{(u,v)}(\mathcal{I}) \cap I_2 \neq \emptyset$ *for every $u \neq v \in \mathcal{V}$. Let $S \subseteq [\mathcal{V}]^\ell$ and let $\{E^i_{\mathbf{s}} \mid \mathbf{s} \in S, i \in [m]\}$ be sets of classes of the partitions from the algorithm. Suppose that there exists $\mathbf{s}' \in S$ such that the relation $\bigcup_{i \in \{2,\dots,m\}} E^i_{\mathbf{s}'}$ is not preserved by p_1 or by m. Then the finite injectivisation of $\mathcal{I}$ on S has a solution.*

Proof. Suppose that the relation $R := \bigcup_{i\in\{2,\dots,m\}} E^i_{\mathbf{s}'}$ is not preserved by p_1. Let $\mathbf{a}_1, \mathbf{a}_2 \in R$ be such that $p_1(\mathbf{a}_1, \mathbf{a}_2) \in E^1_{\mathbf{s}'}$.

For $j \in [2]$, let O^j be the orbit of $\mathbf{a}_j$ under $\mathrm{Aut}(\mathbb{H})$. Let $h^j \colon \mathcal{V} \to H$ be a solution to $(\mathcal{I}^{\mathbf{s}'\in O^j})_{\mathrm{eq}}$ such that $h^j(\mathbf{s}') = \mathbf{a}_j$. Let $h := p_1(h^1, h^2)$. It follows that $h(\mathbf{s}') \in E^1_{\mathbf{s}'}$ and, by the way how the partitions are obtained, $h(\mathbf{s}) \in E^1_{\mathbf{s}}$ for every $\mathbf{s} \in S$. In particular, $h(\mathbf{s})$ lies in an injective or deterministic under $\mathrm{Aut}(\mathbb{H})$ for every $\mathbf{s} \in S$.

Let $\mathbf{s}_1, \dots, \mathbf{s}_n$ be all tuples in S such that $h(\mathbf{s}_i)$ is not injective. For every $i \in [n]$, let $\beta_i \in \mathrm{Aut}(\mathbb{H})$ be such that β_i is deterministic for $h(\mathbf{s}_i)$. Let now

$$p_1''(x, y) := p_1(\beta_1 x, p_1(\beta_2 x, p_1(\dots, p_1(\beta_n x, y)) \dots)).$$

Observe that for all $\mathbf{s} \in S$, we have that $p_1''(h(\mathbf{s}), \mathbf{a})$ and $p_1''(h(\mathbf{s}), \mathbf{b})$ are contained in the same injective orbit $O_{\mathbf{s}}$ under $\mathrm{Aut}(\mathbb{H})$ for all increasing $\mathbf{a} \in E, \mathbf{b} \in N$. Let $f \colon S \to \{E, N\}$ be defined as $f(\mathbf{s}) = O_{\mathbf{s}}$. We claim that f is a solution of the finite injectivisation of $\mathcal{I}$ on S.

To this end, let $C \in \mathcal{C}$. By the choice of h^1, h^2, there exist $g^1, g^2 \in C$ ordered by h^1 and h^2, respectively, It follows that $g := p_1(g^1, g^2)$ is ordered by h and moreover, $g \in C$. For every $i \in [n]$, let $\alpha_i \in \mathrm{Aut}(\mathbb{H})$ be such that $\alpha_i g$ is ordered by $\beta_i h$. Let $p_1'(x, y) := p_1(\alpha_1 x, p_1(\alpha_2 x, p_1(\dots, p_1(\alpha_n x, y)) \dots))$ and let $g' \in C$ be monotone and injective. It follows that for every $\mathbf{s} \in S$ such that all variables from $\mathbf{s}$ are contained in the scope of C and for every $i \in [n]$, $O_<(\alpha_i g(\mathbf{s})) = O_<(\beta_i h(\mathbf{s}))$. In particular, $p_1'(g(\mathbf{s}), g'(\mathbf{s})) \in f(\mathbf{s})$ as desired.

The case when R is not preserved by m is similar. □

Using the auxiliary statements above, we are able to prove the correctness of the main algorithm:

Theorem 3.5.9. *The instance $\mathcal{I}'$ produced by the algorithm in Figure 3.2 is an instance of* $\mathrm{CSP}(\mathbb{A})$ *and it has a solution if, and only if, the original instance has a solution. Moreover, $\mathcal{I}'$ is either trivial or inj-irreducible.*

Proof. It is easy to see that $\mathcal{I}'$ is an instance of $\mathrm{CSP}(\mathbb{A})$ by the algorithm – the only thing that needs to be proven is that if the algorithm constrains an injective tuple $\mathbf{v} \in \mathcal{V}^\ell$ by $\mathrm{proj}_{\mathbf{v}}(\mathcal{I}) \setminus E^1_{\mathbf{v}}$, where $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ is an eq-subdirect, $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal instance obtained during the run of the algorithm, then $\mathrm{proj}_{\mathbf{v}}(\mathcal{I}) \setminus E^1_{\mathbf{v}}$ is preserved by p_1 and by m. But this follows directly from Lemma 3.5.8.

We will now show that if the original instance has a solution, the instance $\mathcal{I}'$ has a solution as well. Suppose that in the repeat... until not changed loop, we get an instance $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ and we removed a solution of $\mathcal{I}$ during one run of this loop. Clearly, we didn't remove any solutions by making the instance eq-subdirect, $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal and by identifying variables in the IdentifyAllEqual subroutine. Hence, the only problem could occur when removing the classes $E^1_{\mathbf{w}}$.

Suppose therefore that there is a set $S \subseteq [\mathcal{V}]^\ell$ and a partition $\{E^j_{\mathbf{v}} \mid j \in \{1, \dots, m\}, \mathbf{v} \in S\}$ on the projections of $\mathcal{I}$ to the tuples from S. Suppose moreover that $\mathcal{I}$ had a solution $s \colon \mathcal{V} \to H$ such that $s(\mathbf{w}) \in E^1_{\mathbf{w}}$ for some $\mathbf{w} \in S$. Then $s(\mathbf{v}) \in E^1_{\mathbf{v}}$ for every $\mathbf{v} \in S$ by the definition of the partition.

Recall that $E^1_{\mathbf{v}}$ contains all injective orbits under $\mathrm{Aut}(\mathbb{H})$ that are contained in $\mathrm{proj}_{\mathbf{v}}(\mathcal{I})$ for all tuples $\mathbf{v} \in S$ by the reasoning below the algorithm. Let $\mathcal{I}''$ be an ℓ-minimal instance equivalent to the instance obtained from $\mathcal{I}$ by adding for every $\mathbf{v} = (v_1, \dots, v_\ell) \in S$ the constraint $\{g\colon \{v_1, \dots, v_\ell\} \to H \mid g(\mathbf{v}) \in E^1_{\mathbf{v}}\}$. Since $E^1_{\mathbf{v}}$ is pp-definable in $\mathbb{A}$, it is by assumption one of the relations of $\mathbb{A}$ and therefore $\mathcal{I}''$ is an instance of $\mathrm{CSP}(\mathbb{A})$. Since s is a solution to $\mathcal{I}$ and hence also to $\mathcal{I}''$, $\mathcal{I}''$ is non-trivial. Since the classes $\{E^1_{\mathbf{v}} \mid \mathbf{v} \in S\}$ were removed, the finite injectivisation of $\mathcal{I}$ on S does not have a solution and it follows that the finite injectivisation of $\mathcal{I}''$ on S does not have a solution as well. Hence, s is not a solution to $\mathcal{I}''$ by Lemma 3.5.7, a contradiction.

We finally prove that if $\mathcal{I}'$ is non-trivial, it is inj-irreducible. Note that by the algorithm, $\mathcal{I}'$ is $(2\ell, \max(3\ell, b_{\mathbb{H}}))$-minimal and if it is non-trivial, then any projection of $\mathcal{I}'$ to a pair of disjoint variables has a non-empty intersection with I_2. In particular, for every injective tuple $\mathbf{v}$ of variables of $\mathcal{I}'$, $\mathrm{proj}_{\mathbf{v}}(\mathcal{I}')$ contains an injective orbit under $\mathrm{Aut}(\mathbb{H})$. Let S be a subset of variables of $\mathcal{I}'_{\mathrm{fin}}$.

Suppose that for some variable $\mathbf{w}$ in S, the linkedness congruence on $\mathrm{proj}_{\mathbf{w}}(\mathcal{I}'_{\mathrm{fin}}|_S)$ links the injective orbits within this set, and separates the injective orbits from the non-deterministic ones.

Each block $B^i_{\mathbf{w}}$ of the linkedness congruence on $\mathrm{proj}_{\mathbf{w}}(\mathcal{I}'_{\mathrm{fin}}|_S)$ defines a subset of $\mathrm{proj}_{\mathbf{w}}(\mathcal{I}')$ by taking $E^i_{\mathbf{w}} := \bigcup B^i_{\mathbf{w}}$. By assumption, this partition of $\mathrm{proj}_{\mathbf{w}}(\mathcal{I}')$ extends to a partition $\{E^1_{\mathbf{v}}, \dots, E^m_{\mathbf{v}}\}$ of $\mathrm{proj}_{\mathbf{v}}(\mathcal{I}')$ for $\mathbf{v} \in S$. Each $E^i_{\mathbf{v}}$ is pp-definable: for an arbitrary orbit O in $B^i_{\mathbf{w}}$, $E^i_{\mathbf{v}}$ consists of the tuples that are reachable from some $\mathbf{a} \in O$ by a suitable path of constraints. Since $\mathbb{A}$ contains all orbits of ℓ-tuples under $\mathrm{Aut}(\mathbb{H})$ by our assumptions, O is a relation of $\mathbb{A}$. Thus, the algorithm has checked that the finite injectivisation of $\mathcal{I}'$ on S has a solution, showing the inj-irreducibility of $\mathcal{I}'$. $\square$

By combining Theorems 3.5.6 and 3.5.9, we obtain the desired result.

Corollary 3.5.10. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ that is a model-complete core and such that $\mathscr{C}^{\mathbb{H},inj}_{\mathbb{A}} \curvearrowright \{E, N\}$ is equationally non-trivial. Then $\mathrm{CSP}(\mathbb{A})$ is in P.*

3.6 The NP-hard case

In this section, we prove that if the first case of Theorem 3.3.2 applies, $\mathrm{CSP}(\mathbb{A})$ is NP-hard.

Lemma 3.6.1. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ that is a model-complete core. Suppose that $\mathscr{C}^{\mathbb{H},inj}_{\mathbb{A}} \curvearrowright \{E, N\}$ is equationally trivial. Then so is $\mathscr{C}^{\mathbb{A},inj}_{\mathbb{A}}$. Moreover, $\mathscr{C}^{\mathbb{H},inj}_{\mathbb{A}} \subseteq \mathscr{C}^{\mathbb{A},inj}_{\mathbb{A}}$.*

Proof. Suppose that $\mathscr{C}^{\mathbb{A},\mathrm{inj}}_{\mathbb{A}}$ is equationally non-trivial. Then $\mathrm{Pol}(\mathbb{A})$ contains a binary injection by Lemma 3.4.4.

We claim that $\mathscr{C}^{\mathbb{A},\mathrm{inj}}_{\mathbb{A}}$ has an operation s which satisfies the pseudo-Siggers identity modulo $\mathrm{Aut}(\mathbb{A})$ on injective tuples. To see this, consider the actions of $\mathscr{C}^{\mathbb{A},\mathrm{inj}}_{\mathbb{A}}$ on $I_n/\mathrm{Aut}(\mathbb{A})$ for all $n \geq 1$. Each of these actions has a Siggers operation s_n, and we may assume that the sequence $(s_n)_{n\geq 1}$ converges in $\mathscr{C}^{\mathbb{A},\mathrm{inj}}_{\mathbb{A}}$ to a function s, which has the desired property.

Let s' be a polymorphism canonical with respect to $(\mathbb{H}, <)$ that is locally interpolated by s modulo $\mathrm{Aut}(\mathbb{H}, <)$. Then on all injective tuples x, y, z of the same length, s' satisfies the pseudo-Siggers equation on x, y, z modulo $\mathrm{Aut}(\mathbb{A})$, by a standard computation (see e.g. the proof of Lemma 35 in [69]). Now for all $1 \leq i \leq 6$ let p_i be a 6-ary function which is canonical with respect to $(\mathbb{H}, <)$, behaves like the i-th projection on $\{E, N\}$, and which acts lexicographically on the order – such functions can easily be obtained from the binary injections from Proposition 3.4.5 and by the remark above Lemma 3.4.4. Set $s'' := s'(p_1(x_1,\dots,x_6),\dots,p_6(x_1,\dots,x_6))$. Then $s'' \in \mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ since application of $p_1,\dots,p_6$ makes forget the order. Say without loss of generality that s'' behaves like the first projection on $\{E,N\}$; the case where it flips E and N of the first coordinate is similar. Let a,b be injective increasing tuples in distinct orbits under $\mathrm{Aut}(\mathbb{A})$ (such tuples only exist if $\mathbb{A}$ is not reduct of $(H;=)$, but we already know from Lemmas 3.4.3 and 3.4.4 that in those cases, $\mathrm{Pol}(\mathbb{A})$ consists of essentially unary operations only and thus the result is also true in this case). Then

$$a \sim_{\mathrm{Aut}(\mathbb{H})} s''(a,b,\dots) \sim_{\mathrm{Aut}(\mathbb{A})} s''(b,a,\dots) \sim_{\mathrm{Aut}(\mathbb{H})} b\,,$$

a contradiction.

The proof that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \subseteq \mathscr{C}_{\mathbb{A}}^{\mathbb{A},\mathrm{inj}}$ is almost identical as the proof of Lemma 32 in [69] but we provide it for the convenience of the reader. Observe that the action $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E,N\}$ is by essentially unary functions by [75]. Moreover, the action $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E,N\}$ determines the action of $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ on orbits of injective tuples under $\mathrm{Aut}(\mathbb{H})$ since $\mathbb{H}$ is homogeneous in an ℓ-ary language. Hence, for every $f \in \mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ of arity $n \geq 1$, there exists $1 \leq i \leq n$ such that the orbit of $f(a_1,\dots,a_n)$ under $\mathrm{Aut}(\mathbb{H})$ is either equal to the orbit of a_i under $\mathrm{Aut}(\mathbb{H})$ for all injective tuples $a_1,\dots,a_n$ of the same length, or the orbit of $f(a_1,\dots,a_n)$ under $\mathrm{Aut}(\mathbb{H})$ is equal to the orbit under $\mathrm{Aut}(\mathbb{H})$ of a tuple obtained from a_i by changing hyperedges to non-hyperedges and vice versa. In the first case, $f \in \mathscr{C}_{\mathbb{A}}^{\mathbb{A},\mathrm{inj}}$ since the orbits under $\mathrm{Aut}(\mathbb{H})$ are refinements of the orbits under $\mathrm{Aut}(\mathbb{A})$. In the second case, $\mathrm{Aut}(\mathbb{A})$ contains a function α changing hyperedges into non-hyperedges and vice versa and hence, $\alpha \circ f \in \mathscr{C}_{\mathbb{A}}^{\mathbb{A},\mathrm{inj}}$ and $f \in \mathscr{C}_{\mathbb{A}}^{\mathbb{A},\mathrm{inj}}$ follows. □

The following is an adjusted version of the fundamental theorem of smooth approximations [69, Theorem 12]. The original version of this theorem does not apply to our case since for a first-order reduct $\mathbb{A}$ of $\mathbb{H}$, we do not have that $\mathrm{Pol}(\mathbb{A})$ locally interpolates $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\mathrm{inj}}$ modulo $\mathrm{Aut}(\mathbb{H})$.

Lemma 3.6.2. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ that is a model-complete core. Assume that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj} \subseteq \mathscr{C}_{\mathbb{A}}^{\mathbb{A},inj}$ and that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj}$ does not contain a binary function that acts as a semilattice operation on $\{E,N\}$. Let $k \geq 1$ be so that $I_k/\mathrm{Aut}(\mathbb{A})$ has at least two elements and let $(S,\sim)$ be a subfactor of $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},inj} \curvearrowright I_k/\mathrm{Aut}(\mathbb{A})$ with $\mathrm{Aut}(\mathbb{A})$-invariant equivalence classes. Assume moreover that η is a very smooth approximation of $\sim$ which is invariant under $\mathrm{Pol}(\mathbb{A})$. Then $\mathrm{Pol}(\mathbb{A})$ has a uniformly continuous clone homomorphism to $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},inj} \curvearrowright S/\sim$.*

Proof. We may assume that $\mathrm{Pol}(\mathbb{A})$ contains an essential operation – otherwise, $\mathrm{Pol}(\mathbb{A})$ has a uniformly continuous clone homomorphism to the clone of projections and hence, to any clone. For all $n \geq 1$ and all $1 \leq i \leq n$, let $p_i^n \in \mathrm{Pol}(\mathbb{A})$ be an n-ary injection which is canonical with respect to $\mathrm{Aut}(\mathbb{H},<)$, acts like the i-th projection on $\{E,N\}$, and follows the first coordinate for the order. These functions are easily obtained from the binary injection acting as a projection on $\{E,N\}$ from Lemma 3.4.4. We define a map ϕ from $\mathrm{Pol}(\mathbb{A})$ to $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\mathrm{inj}} \curvearrowright S/\sim$ as follows. For every $n \geq 1$ and every n-ary $f \in \mathrm{Pol}(\mathbb{A})$, we take any function f' canonical on $(\mathbb{H},<)$ which is locally interpolated by f modulo $\mathrm{Aut}(\mathbb{H},<)$, and then set $f'' := f'(p_1^n,\dots,p_n^n)$; this function is an element of $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$, and hence also of $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\mathrm{inj}}$ by assumption. We set $\phi(f)$ to be f'' acting on $S/\sim$.

Our first claim is that ϕ is well-defined, i.e., it does not depend on the choice of f'. Let $u_1, \ldots, u_n \in S$. Then

$$f(u_1, \ldots, u_n)\ (\eta \circ \sim)\ f'(u_1, \ldots, u_n) \sim f''(u_1, \ldots, u_n)\,.$$

Hence, the action of f'' on $S/\sim$ is completely determined by f, and consequently independent of f'.

It remains to show that ϕ is a uniformly continuous clone homomorphism. ϕ clearly preserves arities, let us show that it preserves compositions as well. Let $n, m \geq 1$, let $f \in \mathrm{Pol}(\mathbb{A})$ be m-ary, let $g_1, \ldots, g_m$ be n-ary, and let $u_1, \ldots, u_n \in S$. Pick, for all $1 \leq i \leq m$, any $v_i \in S$ such that $v_i\ \eta\ g_i(u_1, \ldots, u_n)$. Then

$$\begin{aligned} f(g_1, \ldots, g_m)(u_1, \ldots, u_n)\ \eta\ & f(v_1, \ldots, v_m) \\ (\eta \circ \sim)\ & f''(v_1, \ldots, v_m) \\ (\eta \circ \sim)\ & f''(g_1'', \ldots, g_m'')(u_1, \ldots, u_n), \end{aligned}$$

the last equivalence holding since $g_i(u_1, \ldots, u_n)\ (\eta \circ \sim)\ g''(u_1, \ldots, u_n)$ for all $i \in [n]$. Hence, $\phi(f(g_1, \ldots, g_m)) = \phi(f)(\phi(g_1), \ldots, \phi(g_m))$. It follows that ϕ is a clone homomorphism. Moreover, ϕ is uniformly continuous since for any fixed set of representatives of the classes of $\sim$ and for any $f \in \mathrm{Pol}(\mathbb{A})$, the action of f on these representatives determines $\phi(f)$ by the above proof showing that ϕ is well-defined. □

3.7 The impossible case: few canonical functions and a weakly commutative function

In this section, we prove that if $\mathbb{A}$ is a first-order reduct of $\mathbb{H}$ which has among its polymorphisms a weakly commutative function then $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ contains a binary function that acts as a semilattice operation on $\{E, N\}$. This will allow us to derive a contradiction if the second case of Theorem 3.3.2 applies. To prove this, we will need the following consequence from [18] of the fact that $(\mathbb{H}, <)$ satisfies the Ramsey property.

Let $\mathbb{C}$ be a relational structure and let $A, B \subseteq C$. We say that A *is independent from* B *in* $\mathbb{C}$ if for every $n \geq 1$ and for all $a_1, a_2 \in A^n$ such that a_1, a_2 satisfy the same first-order formulas over $\mathbb{C}$, a_1 and a_2 satisfy the same first-order formulas over $\mathbb{C}$ with parameters from B. We say that two substructures $\mathbb{A}$ and $\mathbb{B}$ of $\mathbb{C}$ are *independent* if A is independent from B in $\mathbb{C}$ and B is independent from A in $\mathbb{C}$. Finally, we say that a substructure $\mathbb{A}$ of $\mathbb{C}$ is *elementary* if the identity mapping from A to C preserves the truth of all first-order formulas in the signature of $\mathbb{C}$.

Theorem 3.7.1 (Theorem 4.4. in [18])**.** *Let $\mathbb{C}$ be a countable homogeneous ω-categorical Ramsey structure. Then $\mathbb{C}$ contains two independent elementary substructures.*

Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ which is a model-complete core and so that $\mathrm{Pol}(\mathbb{A})$ is equationally non-trivial and let us assume that the second case of Proposition 3.4.5 applies, i.e., $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ contains a binary injection g_1 which acts as a projection on $\{E, N\}$ and which is canonical with respect to $(\mathbb{H}, <)$. We will show that in this case, there exist $p_1, p_2 \in \mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}}$ with the following properties.

- p_1, p_2 are both canonical with respect to $(\mathbb{H}, <)$,
- p_i acts like the i-th projection on $\{E, N\}$ for all $i \in \{1, 2\}$,
- p_i acts lexicographically on the order for all $i \in \{1, 2\}$,
- for all $x, y \in H$, if $x < y$, then $p_1(x, y) < p_2(x, y)$, and if $y < x$, then $p_1(x, y) > p_2(x, y)$,
- the ranges of p_1 and p_2 are disjoint and independent as substructures of $\mathbb{H}$.

To show the existence of p_1, p_2, assume without loss of generality that g_1 acts as the first projection on $\{E, N\}$. Then the function $g_2(x, y) := g_1(y, x)$ acts as the second projection on $\{E, N\}$. By the remark above Lemma 3.4.4, we may assume that both g_1 and g_2 act lexicographically on the order $<$. Finally, by Theorem 3.7.1, we may compose g_1 and g_2 with automorphisms of $(\mathbb{H}, <)$ and obtain binary injections g_1', g_2' that still satisfy all the assumptions above and whose ranges are disjoint and induce in $(\mathbb{H}, <)$ substructures that are independent.

Let us define a linear order $<^*$ on $U := \text{im}(g_1') \cup \text{im}(g_2')$ as follows. We set $u <^* v$ if one of the following holds.

- $u < v$ and $u, v \in \text{im}(g_i')$ for some $i \in [2]$, or
- $u = g_i'(x_1, y_1), v = g_j'(x_2, y_2)$ for some $i \neq j \in [2]$, $x_1, x_2, y_1, y_2 \in H$ and one of the following holds
  - $i = 1, j = 2, x_2 \leq y_2$ and $u \leq g_1'(x_2, y_2)$, or
  - $i = 1, j = 2, x_1 \leq y_1$ and $g_2'(x_1, y_1) \leq v$, or
  - $i = 2, j = 1, x_2 > y_2$ and $u \leq g_2'(x_2, y_2)$, or
  - $i = 2, j = 1, x_1 > y_1$ and $g_1'(x_1, y_1) \leq v$.

It is easy to verify that $<^*$ is a linear order on U. Let us define a hyperedge relation E^* on U as the restriction of the relation E to U. It follows that (U, E^*) is isomorphic to the structure induced by the union of the ranges of g_1' and g_2' in $\mathbb{H}$ and hence, $(U, E^*, <^*)$ embeds to $(\mathbb{H}, <)$ by an embedding e since $(\mathbb{H}, <)$ is universal for the class of all ℓ-uniform linearly ordered hypergraphs $(\mathbb{X}, <^{\mathbb{X}})$ such that $\mathbb{X}$ embeds to $\mathbb{H}$. Finally, we obtain the desired functions p_1 and p_2 by setting $p_1 := e \circ g_1'$ and $p_2 := e \circ g_2'$.

Lemma 3.7.2. *Let $\mathbb{A}$ be a first-order reduct of $\mathbb{H}$ which is a model-complete core. Suppose that* $\text{Pol}(\mathbb{A})$ *does not have a uniformly continuous clone homomorphism to $\mathscr{P}$ and contains a binary function f such that there exist $k \geq 1$, $S \subseteq I_k$, and an equivalence relation $\sim$ on S with at least two* $\text{Aut}(\mathbb{A})$*-invariant classes such that $f(a, b) \sim f(b, a)$ for all disjoint injective tuples $a, b \in H^k$ with $f(a, b), f(b, a) \in S$. Then $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj} \curvearrowright \{E, N\}$ contains a semilattice operation.*

Proof. Since $\mathbb{A}$ is a model-complete core, we have $f(a,a) \in S$ for all $a \in S$, and hence picking any $b \in S$ which is not $\sim$-equivalent to a, we see on the values $f(a,a), f(b,b), f(a,b)$, and $f(b,a)$ that f must be essential. It follows from Proposition 3.4.5, from Theorem 3.7.1 and from the discussion above that $\mathrm{Pol}(\mathbb{A})$ contains functions p_1, p_2 as above, or else the conclusion of the lemma follows immediately. Moreover, since its distinctive property is stable under diagonal interpolation modulo $\mathrm{Aut}(\mathbb{H})$, we can assume that f is diagonally canonical with respect to $\mathrm{Aut}(\mathbb{H}, <)$.

It is easy to see that for every $U', V' \subseteq H$ such that $u < v$ for every $u \in U', v \in V'$, f is canonical with respect to $(\mathbb{H}, <)$ on $U' \times V'$. Let $c_1, c_2 \in H$ be arbitrary, let $U := \{p_1(d, c_2) \mid d \in H, d < c_1\}$ and let $V := \{p_2(c_1, d) \mid d \in H, c_2 < d\}$. Hence, for every $u \in U, v \in V$, it holds that $u < v$. Note that for every $m \geq 1$ and for every $a \in H^m$, there exist $b \in U^m, c \in V^m$ such that a, b, c are in the same orbit under $\mathrm{Aut}(\mathbb{H}, <)$. Note moreover that if $m \geq 1$ and (a, b) is any pair of m-tuples such that $(a_i, b_j) \in U \times V$ for all $i, j \in [m]$, then $p_1(a_i, b_j) < p_2(a_i, b_j)$ for every $i, j \in [m]$. Setting $g(x, y) := f(p_1(x, y), p_2(x, y))$, we obtain a function which is canonical with respect to $\mathbb{H}$ on $U \times V$. Hence, the restriction of g to $U \times V$ naturally acts on $\{E, N\}$, and we may assume it does so as an essentially unary function as otherwise we are done. A similar statement holds for its restriction to $V \times U$. If one of the two mentioned essentially unary functions on $\{E, N\}$ is not a projection, then $\mathrm{Aut}(\mathbb{A})$ contains a function flipping E and N. It follows that whenever $a, b \in S$ are so that $(a_i, b_j) \in U \times V$ for all $i, j \in \{1, \ldots, k\}$, then $g(a, b)$ and $g(b, a)$ are elements of S. If a, b are moreover both increasing, then $g(a, b)$ and $f(a, b)$, as well as $g(b, a)$ and $f(b, a)$, belong to the same orbit under $\mathrm{Aut}(\mathbb{A})$. Since $f(a, b) \sim f(b, a)$, the two essentially unary functions mentioned above cannot depend on the same argument, as witnessed by choosing $a, b \in S$ from distinct $\sim$-classes. Hence, we may assume that the one from $U \times V$ depends on the first argument and the other one on the second.

More generally, since $g(x, y) = f(p_1(x, y), p_2(x, y))$, the function g has the property that for all injective ℓ-tuples a, b where a is increasing, the type of their image under this function in $\mathbb{H}$ only depends on the relations of $\mathbb{H}$ on each of a and on b, respectively, and on the order relation on pairs (a_i, b_j), where $i, j \in \{1, \ldots, \ell\}$; this type is precisely the same type obtained when applying the function to a, b', where b' is obtained from b by changing its order to be increasing.

In the following, for any pair (a, b) of injective increasing ℓ-tuples, we shall consider the set B of all pairs (a', b') of injective increasing ℓ-tuples such that the order relation on (a_i, b_j) and that on (a'_i, b'_j) agree for all $i, j \in \{1, \ldots, \ell\}$; we call this set an *increasing diagonal order type*. We then have by the above that g acts naturally on $\{E, N\}$ within each such set B.

Set $h(x, y) := g(e_1 \circ g(x, y), e_2 \circ g(x, y))$, where e_1, e_2 are self-embeddings of $(\mathbb{H}, <)$ which ensure that for all injective increasing a, b, the diagonal order type of (a, b) is equal to that of $(e_1 \circ g(a, b), e_2 \circ g(a, b))$. Then h acts idempotently or as a constant function in its action on $\{E, N\}$ within each increasing diagonal order type B as above. In particular, h then acts like the first projection on $\{E, N\}$ when restricted to $U \times V$, and like the second when restricted to $V \times U$.

We will now rule out the possibility that h acts like a constant function on $\{E, N\}$ within some increasing diagonal order type of injective ℓ-tuples. In the following, we assign to every increasing diagonal order type B of pairs of increasing injective ℓ-tuples a number n_B as follows: if $((x_1, \ldots, x_\ell), (y_1, \ldots, y_\ell)) \in B$, then n_B is the minimal number in $\{0, \ldots, \ell\}$ such that for all $n_B < i \leq \ell$ it holds that $x_i < y_i$, and if $i > 1$, then moreover $y_{i-1} < x_i$.

Suppose that within some increasing diagonal order type B, we have that h acts as a constant function on $\{E, N\}$. Pick such B such that n_B is minimal. Assume without loss of generality that the constant value of h on B is E. We claim that $\mathbb{A}$ has an endomorphism onto a clique, which contradicts the conjunction of the assumptions that $\mathbb{A}$ is a model-complete core and not a reduct of $(H; =)$. We prove this claim by showing that any finite injective tuple can be mapped to a clique by an endomorphism of $\mathbb{A}$. Suppose that $m \geq 1$ and that $(a_1, \ldots, a_m)$ is an injective m-tuple of elements of H which does not induce a clique, i.e., there is a subtuple of length ℓ which is not an element of E. Then $m \geq \ell$, and we may assume without loss of generality that $(a_1, \ldots, a_\ell) \notin E$. By applying a self-embedding e_1 of $\mathbb{H}$, we obtain an increasing tuple $(x_1, \ldots, x_m)$ such that $(x_1, \ldots, x_\ell) \in N$. Applying an appropriate self-embedding e_2 of $\mathbb{H}$ to $(a_1, \ldots, a_m)$, we moreover obtain an increasing tuple $(y_1, \ldots, y_m)$ such that $((x_1, \ldots, x_\ell), (y_1, \ldots, y_\ell)) \in B$ and such that $y_{i-1} < x_i < y_i$ for all $i > \ell$. Then the increasing diagonal order type C for any pair $((x_{i_1}, \ldots, x_{i_\ell}), (y_{i_1}, \ldots, y_{i_\ell}))$ of subtuples, both increasing, has the property that $n_C \leq n_B$ (note that in order to compute n_C, the entries of both tuples receive the indices 1 to ℓ). If $n_C < n_B$, then h acts idempotently on $\{E, N\}$ within C, by the minimality of n_B; if on the other hand $n_C = n_B$, then $B = C$ and h acts as a constant with value E. It follows that applying the endomorphism $h(e_1(x), e_2(x))$ to $(a_1, \ldots, a_m)$, one obtains a tuple which has strictly more subtuples in E than $(a_1, \ldots, a_m)$. The claim follows.

We may thus henceforth assume that within each increasing diagonal order type B, we have that h acts in an idempotent fashion on $\{E, N\}$. Thus, within each such type, h acts as a semilattice operation or as a projection on $\{E, N\}$.

Let $0 \leq j \leq \ell$, and consider the diagonal order type T given by a pair (c, d) of increasing injective j-tuples of elements of H; we call j the *length* of T. In the following, we shall say that T is *categorical* if for all pairs (a, b) of increasing injective ℓ-tuples which extend (c, d) (we mean any extension, not just end-extension) the corresponding diagonal order type is one where h behaves like the first projection on, or if a similar statement holds for the second projection, or for the semilattice behaviour (both semilattice behaviours are considered the same here). Note that for length $j = 0$ every T is non-categorical, by the behaviours on $U \times V$ and $V \times U$. Note also that for length $j = \ell$ every T is trivially categorical. We claim that there exists T of length $j = \ell - 1$ which is not categorical. Suppose otherwise, and take any T which is categorical and implies the behaviour as the second projection; this exists by the behaviour on $V \times U$. Let (c, d) be a pair of injective j-tuples which are ordered such that they represent the diagonal order type T. Let (a_1, b_1) be obtained from (c, d) by extending both increasing tuples by a single element c' and d' at the end, respectively, in such a way that $c' < d'$. Let (c_1, d_1) be tuple obtained from (a_1, b_1) by taking away the first components. Then the order type represented by (c_1, d_1) is categorical but not for the first projection. We continue in this fashion until we arrive at a pair (a_ℓ, b_ℓ) in $U \times V$, a contradiction.

In the following, we assume that there exists T of length $j = \ell - 1$ which extends to diagonal order types where h behaves like different projections; the other case (projection + semilattice) is handled similarly. We show by induction on $m \geq 1$ that for all tuples $a, b \in I^m$ such that for every $i \in \{1, \ldots, m\}$ at most one of the tuples a_i, b_i is in N there exists $u \in \mathrm{Pol}(\mathbb{A})$ such that $u(a, b) \in E^m$. A standard compactness argument then implies that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}}$ contains a function such that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is a semilattice operation.

The base case $m = 1$ is clearly achieved by applying an appropriate projection. For the induction step, let $a, b \in I^m$ for some $m \geq 2$. Since $\mathrm{Pol}(\mathbb{A})$ contains p_1, we may assume that the kernels of a and b are identical. We may then also assume that a_i, a_j induce distinct sets whenever $1 \leq i, j \leq m$ and $i \neq j$, for otherwise we are done by the induction hypothesis. By the induction hypothesis, we may assume that all components of a are in E except for the second, and all components of b are in E except for the first. It is sufficient to show that there exists an increasing diagonal order type on the pair (a, b) (i.e., increasing diagonal order types for each of the pairs $(a_1, b_1), \ldots, (a_m, b_m)$ which are consistent with the kernels of a and b) such that h behaves like the first projection within the order type of (a_1, b_1) and like the second within the order type of (a_2, b_2). This, however, is obvious by our assumption. □

3.8 Bounded width

In this section, we prove a characterization of first-order expansions of $\mathbb{H}$ whose CSPs have bounded width.

Theorem 3.8.1. *Let $\mathbb{A}$ be a first-order expansion of $\mathbb{H}$. Then precisely one of the following applies.*

1. *The clone* $\mathrm{Pol}(\mathbb{A})$ *has a uniformly continuous minion homomorphism to the clone of affine maps over a finite module.*
2. *The clone* $\mathrm{Pol}(\mathbb{A})$ *has no uniformly continuous minion homomorphism to the clone of affine maps over a finite module, and* $\mathrm{CSP}(\mathbb{A})$ *has relational width* $(2\ell, \max(3\ell, b_{\mathbb{H}}))$.

We prove Theorem 3.8.1 in a similar way as Theorem 3.1.1. We prove that if $\mathbb{A}$ is a first-order expansion of $\mathbb{H}$ which is a model-complete core then $\mathrm{CSP}(\mathbb{A})$ has bounded width if, and only if, $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally non-affine. If $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally non-affine, the result follows from Corollary 3.5.2.

Let us therefore suppose that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally affine. Moreover, we may assume that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\mathrm{inj}} \curvearrowright \{E, N\}$ is equationally non-trivial as otherwise, $\mathrm{Pol}(\mathbb{A})$ has a uniformly continuous homomorphism to the clone of projections by the proof of Theorem 3.1.1. We apply the second loop lemma of smooth approximations [69, Theorem 11].

Theorem 3.8.2. *Let $k \geq 1$ and suppose that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj} \curvearrowright \{E, N\}$ is equationally non-trivial. Let $(S, \sim)$ be a minimal subfactor of $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj} \curvearrowright \{E, N\}$ with* $\mathrm{Aut}(\mathbb{H})$*-invariant $\sim$-classes. Then one of the following holds:*

- *$\sim$ is approximated by a* $\mathrm{Pol}(\mathbb{A})$*-invariant equivalence relation that is very smooth with respect to* $\mathrm{Aut}(\mathbb{H})$*;*
- *every $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},inj} \curvearrowright \{E, N\}$-invariant binary symmetric relation $R \subseteq I^2$ that contains a pair $(a, b) \in S^2$ such that $a \neq b$ and such that $a \not\sim b$ contains a pseudo-loop modulo* $\mathrm{Aut}(\mathbb{H})$*, i.e., a pair (c, c') where c, c' belong to the same orbit under* $\mathrm{Aut}(\mathbb{H})$.

In the formulation of the first item of Theorem 3.8.2, we are using [69, Lemma 8]. If the first case of the theorem applies, i.e., the equivalence relation $(S, \sim)$ on whose classes $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}}$ acts by a function from a clone $\mathscr{M}$ of affine maps over a finite module is approximated by a $\text{Pol}(\mathbb{A})$-invariant equivalence relation that is very smooth with respect to $\text{Aut}(\mathbb{H})$, Lemma 3.6.1 (with $k = \ell$, $\mathscr{C}_{\mathbb{A}}^{\mathbb{A},\text{inj}} = \mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}}$) implies that $\text{Pol}(\mathbb{A})$ has a uniformly continuous clone homomorphism to $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}} \curvearrowright (S, \sim)$ and hence to $\mathscr{M}$.

If the second case of Theorem 3.8.2 applies, we get a weakly commutative function by [69, Lemma 13], and the same argument mentioned at the end of Section 3.3 gives that $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}} \curvearrowright \{E, N\}$ contains a semilattice operation. In particular, $\mathscr{C}_{\mathbb{A}}^{\mathbb{H},\text{inj}} \curvearrowright \{E, N\}$ is equationally non-affine, which is a contradiction.

4 Bounds on the relational width of first-order expansions of structures with neoliberal automorphism group

4.1 Introduction

As we have mentioned already in Section 1.4.2, the algebraic characterization of local consistency from Theorem 1.4.4 does not translate to infinite-domain CSPs [31, 32, 55]. Moreover, no collapse of bounded width hierarchy is possible for infinite-domain CSPs with bounded width (see Example 1.4.6). This indicates that new methods have to be developed to better understand this class of CSPs.

The exploration of the exact relational width of infinite-domain CSP templates with bounded width was only started recently. First results were obtained in [82, 81] where an upper bound on the relational width of first-order expansions of certain binary structures with bounded strict width was given. In [69], algebraic conditions characterizing CSP templates with bounded width were introduced for several well-known classes of infinite-domain CSPs. Finally, in Theorem 2.1.2, we found an upper bound on the relational width of CSP templates which satisfy these conditions. Moreover, we showed in Example 2.1.4 that this bound is optimal for many templates under consideration.

Contrary to the case of bounded width, the algebraic characterization of bounded strict width can still be lifted to the ω-categorical case by Theorem 1.4.10 – this suggests that CSPs with bounded strict width form a particularly well-behaved subclass of ω-categorical CSPs with bounded width. Studying this class therefore seems to be a natural starting point in the endeavour to understand the amount of local consistency needed to solve ω-categorical CSPs. Such understanding has on the one hand algorithmic consequences – it provides us with efficient algorithms for solving the CSPs under consideration. On the other hand, it represents an imperceptible part in the constant strive of the humanity to better understand, classify, parameterize and dominate the outer world which always seems to resist our endeavour to conquer it.

4.1.1 Results

In this chapter, we build on results by Wrona [81] and prove that certain CSP templates with bounded strict width have limited expressibility in the form of *implicational simplicity*. As a corollary, we obtain a bound on the amount of local consistency needed to decide solvability of CSP instances over these templates. The relations of our templates are required to be first-order definable in an ω-categorical base structure whose automorphism group satisfies certain abstract properties. In contrast to [81], the base structure can have relations of arbitrarily large arity that is greater or equal to 3.

While for many well-known ω-categorical structures in a binary signature, their first-order reducts whose CSPs are solvable by local consistency checking are quite well-understood [69], the investigation of first-order reducts of structures whose relations have arbitrarily large arity was only started recently (see Chapter 3). The precise definition of all notions mentioned in the following theorem can be found in Sections 4.2 and 4.3.

Theorem 4.1.1. *Let $k \geq 3$, let $\mathbb{B}$ be the canonical k-ary structure of a k-neoliberal permutation group $\mathscr{G}$, and suppose that $\mathbb{B}$ has finite duality. Then any first-order expansion of $\mathbb{B}$ with bounded strict width is implicationally simple on injective instances.*

In order to prove Theorem 4.1.1, we reformulate the concept of an implication and many related concepts from [81] to our setting. Using these concepts, we will employ a similar proof strategy as in [81].

Corollary 4.1.2. *Let $k \geq 3$, let $\mathbb{B}$ be the canonical k-ary structure of a k-neoliberal permutation group $\mathscr{G}$, and suppose that $\mathbb{B}$ has finite duality. Then any first-order expansion of $\mathbb{B}$ with bounded strict width has relational width $(k, \max(k+1, b_{\mathbb{B}}))$.*

To show that Corollary 4.1.2 follows from Theorem 4.1.1, we will use results from [69] and from Chapter 3.

4.2 Implications and binary injections

In this section, we first restate some results about binary injections from [69] that will enable us to use Lemma 14 from [70] in order to reduce $\mathrm{CSP}(\mathbb{A})$ for any structure $\mathbb{A}$ under consideration to $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$. Afterwards, we introduce the notion of an implication and several related concepts that are inspired by notions from [81] and that will play a key role in the proof of Corollary 4.1.2. It is not hard to see that, unlike in the case for structures from [81], the reduction to $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$ is necessary since every structure in the scope of Theorem 4.1.1 is implicationally hard without restricting to injective instances.

4.2.1 Definitions and notation

Let $k, \ell \geq 1$, and let A be a set. For a tuple $\mathbf{t} \in A^k$, we write $\mathrm{S}(\mathbf{t})$ for the *scope* of $\mathbf{t}$, i.e., for the set of all entries of $\mathbf{t}$.

Let $\mathbb{A}$ be a relational structure and let ϕ be a first-order formula over the signature of $\mathbb{A}$. We identify the interpretation $\phi^{\mathbb{A}}$ of ϕ in $\mathbb{A}$ with the set of satisfying assignments for $\phi^{\mathbb{A}}$. Let V be the set of free variables of ϕ, and let $\mathbf{u}$ be a tuple of elements of V. We define $\mathrm{proj}_{\mathbf{u}}(\phi^{\mathbb{A}}) := \{f(\mathbf{u}) \mid f \in \phi^{\mathbb{A}}\}$.

4.2.2 Binary injections and bounded width

We will use the following results from [69]. The orbit O with the property stated in Lemma 4.2.1 is called *free* in [69].

Lemma 4.2.1 (Proposition 21 in [69]). *Let $\mathbb{B}$ be a homogeneous structure such that there exists an orbit O of pairs under $\mathrm{Aut}(\mathbb{B})$ with the property that for all $a, b \in B$, there exists $c \in B$ such that $(a, c), (b, c) \in O$, and let $\mathbb{A}$ be first-order reduct of $\mathbb{B}$. If $\mathrm{Pol}(\mathbb{A})$ contains an essential function, then it contains a binary essential function.*

Lemma 4.2.2 (Proposition 24 in [69]). *Let $\mathbb{A}$ be a first-order reduct of a transitive ω-categorical structure $\mathbb{B}$ such that the canonical binary structure of $\mathrm{Aut}(\mathbb{B})$ has finite duality. If $\mathrm{Pol}(\mathbb{A})$ contains a binary essential function preserving I_2^B, then it contains a binary injection.*

Lemmas 4.2.1 and 4.2.2 immediately yield the following proposition.

Proposition 4.2.3. *Let $k \geq 2$, let $\mathscr{G}$ be a 2-transitive oligomorphic permutation group, and let $\mathbb{B}$ be its canonical k-ary structure. Let $\mathbb{A}$ be a first-order expansion of $\mathbb{B}$, and suppose that $\mathbb{A}$ has bounded strict width. Then $\mathbb{A}$ has relational width $(k, \max(k+1, b_{\mathbb{B}}))$ if, and only if, $\mathrm{CSP}_{\mathrm{Inj}}\,\mathbb{A}$ has relational width $(k, \max(k+1, b_{\mathbb{B}}))$.*

Proof. Since $\mathbb{A}$ has bounded strict width, it has in particular bounded width, and hence $\mathrm{Pol}(\mathbb{A})$ does not have a uniformly continuous homomorphism to an affine clone by Theorem 1.4.4. In particular, it does not have a uniformly continuous clone homomorphism to the clone of projections. It is easy to see and well-known that $\mathrm{Pol}(\mathbb{A})$ then contains an essential function. Since $\mathscr{G}$ is 2-transitive, I_2^B is an orbit under $\mathscr{G}$, and it clearly satisfies the condition from Lemma 4.2.1. It follows that $\mathrm{Pol}(\mathbb{A})$ contains a binary essential operation, and since $\mathbb{A}$ is a first-order expansion of $\mathbb{B}$, every polymorphism of $\mathbb{A}$ preserves I_2^B. Lemma 4.2.2 yields that $\mathrm{Pol}(\mathbb{A})$ contains a binary injection. Now, the statement follows directly from Lemma 3.5.1. □

4.2.3 Implications

Definition 4.2.4. *Let $\mathbb{A}$ be a relational structure. Let V be a set of variables, let $\mathbf{u}, \mathbf{v}$ be injective tuples of variables in V of arity $k < |V|$ and $m < |V|$, respectively, such that $\mathrm{S}(\mathbf{u}) \cup \mathrm{S}(\mathbf{v}) = V$. Let $C \subseteq A^k$ and $D \subseteq A^m$ be pp-definable from $\mathbb{A}$ and non-empty. We say that a pp-formula ϕ over the signature of $\mathbb{A}$ with free variables from V is a $(C, \mathbf{u}, D, \mathbf{v})$-*implication in $\mathbb{A}$ *if all of the following hold:*

1. *for all distinct $x, y \in V$, $\mathrm{proj}_{(x,y)}(\phi^{\mathbb{A}}) \not\subseteq \{(a, a) \mid a \in A\}$,*
2. *$C \subsetneq \mathrm{proj}_{\mathbf{u}}(\phi^{\mathbb{A}})$,*
3. *$D \subsetneq \mathrm{proj}_{\mathbf{v}}(\phi^{\mathbb{A}})$,*
4. *for every $f \in \phi^{\mathbb{A}}$, it holds that $f(\mathbf{u}) \in C$ implies $f(\mathbf{v}) \in D$,*
5. *there exists no $D' \subsetneq D$ such that for every $f \in \phi^{\mathbb{A}}$, it holds that $f(\mathbf{u}) \in C$ implies $f(\mathbf{v}) \in D'$.*

*We say that ϕ is a $(C, \mathbf{u}, D, \mathbf{v})$-*pre-implication *if it satisfies items (2)-(5). We will call ϕ a (C, D)-*implication *if it is a $(C, \mathbf{u}, D, \mathbf{v})$-implication for some $\mathbf{u} \in I_k^V, \mathbf{v} \in I_m^V$. We say that an implication ϕ is* injective *if $\phi^{\mathbb{A}}$ contains only injective mappings.*

*Let $\mathscr{G}$ be a permutation group acting on A, and let $f \in \phi^{\mathbb{A}}$. If O, P are orbits under $\mathscr{G}$ such that $f(\mathbf{u}) \in O$, $f(\mathbf{v}) \in P$, then we say that f is an OP-*mapping.

Example 4.2.5. *Let $\mathbb{A}$ be a relational structure, let $k \geq 1$, and let O be an orbit of k-tuples under $\mathrm{Aut}(\mathbb{A})$. Suppose that $\mathbb{A}$ pp-defines the equivalence of orbits of k-tuples under $\mathrm{Aut}(\mathbb{A})$. Then the formula defining this equivalence is an (O, O)-pre-implication in $\mathbb{A}$. If $\mathbb{A}$ is such that $\mathrm{Aut}(\mathbb{A})$ does not have any fixed point in its action on A, this pre-implication is an implication. For all orbits P, Q of k-tuples under $\mathrm{Aut}(\mathbb{A})$, $\phi^{\mathbb{A}}$ contains an PQ-mapping if, and only if, $P = Q$.*

4.2.4 Implicationally simple structures

Definition 4.2.6. *Let $\mathbb{A}$ be a relational structure, and let $k \geq 1$.*

The k-ary implication graph *of $\mathbb{A}$, to be denoted by $\mathcal{G}_{\mathbb{A}}$, is a directed graph defined as follows.*

- *The set of vertices is the set of pairs (C_1, C) where $\emptyset \neq C \subsetneq C_1 \subseteq A^k$ and C, C_1 are pp-definable from $\mathbb{A}$.*
- *There is an arc from (C_1, C) to (D_1, D) if there exists a $(C, \mathbf{u}, D, \mathbf{v})$-implication ϕ in $\mathbb{A}$ such that $\mathrm{proj}_{\mathbf{u}}(\phi^{\mathbb{A}}) = C_1$, $\mathrm{proj}_{\mathbf{v}}(\phi^{\mathbb{A}}) = D_1$.*

The k-ary injective implication graph *of $\mathbb{A}$, denoted by $\mathcal{G}_{\mathbb{A}}^{\mathrm{Inj}}$, is the (non-induced) subgraph of $\mathcal{G}_{\mathbb{A}}$ which contains precisely the vertices (C_1, C) where C is injective and which contains an arc from (C_1, C) to (D_1, D) if $(C_1, C) \neq (D_1, D)$ and there exists an injective $(C, \mathbf{u}, D, \mathbf{v})$-implication ϕ in $\mathbb{A}$ with $\mathrm{proj}_{\mathbf{u}}(\phi^{\mathbb{A}}) = C_1$, $\mathrm{proj}_{\mathbf{v}}(\phi^{\mathbb{A}}) = D_1$.*

We say that $\mathbb{A}$ is implicationally simple (on injective instances) *if the (injective) implication graph $\mathcal{G}_{\mathbb{A}}$ ($\mathcal{G}_{\mathbb{A}}^{\mathrm{Inj}}$) is acyclic. Otherwise, $\mathbb{A}$ is* implicationally hard (on injective instances).

Note that by item (1) in Definition 4.2.4, the implication graph does not necessarily contain all loops – e.g., the formula over variables $\{x_1, \ldots, x_{2k}\}$ defined by $\bigwedge_{i \in [k]} (x_i = x_{i+k})$ is not an implication.

The following is essentially subsumed by Lemma 2.3.3 but we provide the reformulation to our setting as well as the proof for the convenience of the reader.

Lemma 4.2.7. *Let $k \geq 2$, let $\mathscr{G}$ be a permutation group, let $\mathbb{B}$ be its canonical k-ary structure, and suppose that $\mathbb{B}$ is finitely bounded. Let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be a non-trivial, $(k, \max(k+1, b_{\mathbb{B}}))$-minimal instance of $\mathrm{CSP}(\mathbb{B})$ such that for every $\mathbf{v} \in \mathcal{V}^k$, $\mathrm{proj}_{\mathbf{v}}(\mathcal{I})$ contains precisely one orbit under $\mathscr{G}$. Then $\mathcal{I}$ has a solution.*

Proof. Let $\sim$ be a binary relation defined on $\mathcal{V}$ such that $u \sim v$ if, and only if, $\mathcal{I}|_{\{u,v\}}$ consists of constant maps. Since $k \geq 2$, $\sim$ is an equivalence relation.

Let τ be the signature of $\mathbb{B}$, and let us define a τ-structure $\mathbb{A}$ on $\mathcal{V} / \sim$ as follows. Let $R \in \tau$; then R is of arity k. We set $([v_1]_\sim, \ldots, [v_k]_\sim) \in R$ if $\mathrm{proj}_{(v_1,\ldots,v_k)}(\mathcal{I}) = R^{\mathbb{B}}$. Note that by our assumption, for every $(v_1, \ldots, v_k) \in \mathcal{V}^k$, there is precisely one relation of $\mathbb{A}$ containing the tuple $([v_1]_\sim, \ldots, [v_k]_\sim)$.

Let us show that the definition of $\sim$ does not depend on the choice of the representatives $v_1, \ldots, v_k \in \mathcal{V}$. We will show that it does not depend on the choice of v_1, the rest can be shown similarly. Let $u_1 \sim v_1$, and let $C \in \mathcal{C}$ be such that $u_1, v_1, \ldots, v_k$ are contained in its scope. Then $C|_{\{u_1,v_1\}}$ consists of constant maps and it follows that $\text{proj}_{(u_1,v_2,\ldots,v_k)}(\mathcal{I}) = \text{proj}_{(u_1,v_2,\ldots,v_k)}(C) = \text{proj}_{(v_1,\ldots,v_k)}(C) = \text{proj}_{(v_1,\ldots,v_k)}(\mathcal{I})$.

We claim that $\mathbb{A}$ embeds into $\mathbb{B}$. Suppose for contradiction that this is not the case. Then there exists a bound $\mathbb{F} \in \mathcal{F}_\mathbb{B}$ of size b with $b \leq b_\mathbb{B}$ such that $\mathbb{F}$ embeds into $\mathbb{A}$. Let $[v_1]_\sim, \ldots, [v_b]_\sim$ be all elements in the image of this embedding. Find a constraint $C \in \mathcal{C}$ such that $v_1, \ldots, v_b$ are contained in its scope. Since C is nonempty, there exists $f \in C$. Since all relations in τ are of arity k, since $\mathcal{I}$ is k-minimal and since for every $\mathbf{v} \in \mathcal{V}^k$ such that all variables from $\mathbf{v}$ are contained in the scope of C, $\text{proj}_\mathbf{v}(C)$ contains precisely one orbit under $\mathscr{G}$, it follows that $\mathbb{F}$ embeds into the structure that is induced by the image of f in $\mathbb{B}$ which is a contradiction.

It follows that there exists en embedding $e\colon \mathbb{A} \hookrightarrow \mathbb{B}$ and it is easy to see that $f\colon \mathcal{V} \to B$ defined by $f(v) := e([v]_\sim)$ is a solution of $\mathcal{I}$. □

Proposition 4.2.8. *Let $k \geq 3$, let $\mathscr{G}$ be a $(k-1)$-transitive oligomorphic permutation group, let $\mathbb{B}$ be its canonical k-ary structure, and suppose that $\mathbb{B}$ is finitely bounded. Let $\mathbb{A}$ be a first-order expansion of $\mathbb{B}$ which is implicationally simple on injective instances and such that $\text{Pol}(\mathbb{A})$ contains a binary injection. Then $\mathbb{A}$ has relational width $(k, \max(k+1, b_\mathbb{B}))$.*

Proof. By Proposition 4.2.3, it is enough to show that $\text{CSP}_{\text{Inj}}(\mathbb{A})$ has relational width $(k, \max(k+1, b_\mathbb{B}))$. To this end, let $\mathcal{I} = (\mathcal{V}, \mathcal{C})$ be a non-trivial $(k, \max(k+1, b_\mathbb{B}))$-minimal instance of $\text{CSP}_{\text{Inj}}(\mathbb{A})$; we will show that there exists a satisfying assignment for $\mathcal{I}$.

For every $i \geq 0$, we define inductively a $(k, \max(k+1, b_\mathbb{B}))$-minimal instance $\mathcal{I}_i = (\mathcal{V}, \mathcal{C}_i)$ of $\text{CSP}_{\text{Inj}}(\mathbb{A})$ with the same variable set as $\mathcal{I}$ such that $\mathcal{I}_0 = \mathcal{I}$ and such that for every $i \geq 1$, $\mathcal{C}_i$ contains for every constraint $C_{i-1} \in \mathcal{C}_{i-1}$ a constraint C_i such that $C_i \subseteq C_{i-1}$.

Let $\mathcal{I}_0 := \mathcal{I}$. Let $i \geq 1$. We define $\mathcal{G}_i$ to be the graph that originates from $\mathcal{G}_\mathbb{A}^{\text{Inj}}$ by removing all vertices that are not of the form $(\text{proj}_\mathbf{v}(\mathcal{I}_{i-1}), F)$ for some injective $\mathbf{v} \in \mathcal{V}^k$, and some $F \subseteq A^k$. Claim 4.2.9 implies that $\mathcal{I}_{i-1}$ is k-minimal, and hence $\text{proj}_\mathbf{v}(\mathcal{I}_{i-1})$ is well-defined. If $\mathcal{G}_i$ does not contain any vertices, let $\mathcal{I}_i := \mathcal{I}_{i-1}$. Suppose now that $\mathcal{G}_i$ contains at least one vertex. In this case, since $\mathcal{G}_\mathbb{A}^{\text{Inj}}$ and hence also $\mathcal{G}_i$ is acyclic, we can find a sink $(\text{proj}_{\mathbf{v}_i}(\mathcal{I}_{i-1}), F_i)$ in $\mathcal{G}_i$ for some injective $\mathbf{v}_i \in \mathcal{V}^k$. Let us define for every $C_{i-1} \in \mathcal{C}_{i-1}$ containing all variables from $\mathbf{v}_i$ in its scope $C_i := \{f \in C_{i-1} \mid f(\mathbf{v}_i) \in F_i\}$, and let $C_i := C_{i-1}$ for every $C_{i-1} \in \mathcal{C}_{i-1}$ that does not contain all variables from $\mathbf{v}_i$ in its scope. Note that in both cases, $C_i \subseteq C_{i-1}$. Finally, we define $\mathcal{C}_i = \{C_i \mid C_{i-1} \in \mathcal{C}_{i-1}\}$.

Claim 4.2.9. *For every $i \geq 1$, $\mathcal{I}_i$ is non-trivial and $(k, \max(k+1, b_\mathbb{B}))$-minimal. Moreover, for every $\mathbf{v} \in \mathcal{V}^k \setminus \{\mathbf{v}_i\}$, $\text{proj}_\mathbf{v}(\mathcal{I}_i) = \text{proj}_\mathbf{v}(\mathcal{I}_{i-1})$ and $\text{proj}_{\mathbf{v}_i}(\mathcal{I}_i) = F_i$.*

Let $i \geq 1$ and if $i > 1$, suppose that the claim holds for $i-1$. Note that if $\mathcal{I}_i = \mathcal{I}_{i-1}$, then there is nothing to prove so we may suppose that this is not the case. Observe that for every $C_i \in \mathcal{C}_i$ containing all variables from $\mathbf{v}_i$ in its scope, $\text{proj}_{\mathbf{v}_i}(C_i) = F_i$ by the definition of C_i. We will now show that for every $\mathbf{v} \in \mathcal{V}^k \setminus \{\mathbf{v}_i\}$ and for every $C_i \in \mathcal{C}_i$ containing all variables from $\mathbf{v}$ in its scope, $\text{proj}_\mathbf{v}(C_i) = \text{proj}_\mathbf{v}(C_{i-1})$. Observe that if C_i does not contain all variables from $\mathbf{v}_i$ in its scope, then the conclusion follows immediately since $C_i = C_{i-1}$; we can therefore assume that this is not the case. Assume first that $\mathbf{v}$ is not injective, let m be the number of pairwise distinct entries of $\mathbf{v}$, and let $\mathbf{u}$ be an injective m-tuple containing all variables of $\mathbf{v}$. Hence, $\text{proj}_\mathbf{u}(C_i) = I_m^A = \text{proj}_\mathbf{u}(C_{i-1})$ by the $(k-1)$-transitivity of $\mathscr{G}$

and it follows that $\mathrm{proj}_{\mathbf{v}}(C_i) = \mathrm{proj}_{\mathbf{v}}(C_{i-1})$. Now assume that $\mathbf{v}$ is injective and, striving for a contradiction, that $\mathrm{proj}_{\mathbf{v}}(C_i) \subsetneq \mathrm{proj}_{\mathbf{v}}(C_{i-1})$. It follows that $(\mathrm{proj}_{\mathbf{v}}(C_{i-1}), \mathrm{proj}_{\mathbf{v}}(C_i))$ is a vertex in $\mathcal{G}_{\mathbb{A}}^{\mathrm{Inj}}$ and hence also in $\mathcal{G}_i$. Let $\mathbf{w} = (w_1, \ldots, w_\ell) \in \mathcal{V}^\ell$ be an enumeration of all variables of $\mathbf{v}$ and $\mathbf{v}_i$. It follows that the pp-formula $\phi(\mathbf{w})$ defining $\mathrm{proj}_{\mathbf{w}}(C_{i-1})$ is an injective $(F_i, \mathbf{v}_i, \mathrm{proj}_{\mathbf{v}}(C_i), \mathbf{v})$-implication. Hence, there is an arc from $(\mathrm{proj}_{\mathbf{v}_i}(\mathcal{I}_{i-1}), F_i)$ to $(\mathrm{proj}_{\mathbf{v}}(C_{i-1}), \mathrm{proj}_{\mathbf{v}}(C_i))$ in $\mathcal{G}_i$ and in particular, $(\mathrm{proj}_{\mathbf{v}_i}(\mathcal{I}_{i-1}), F_i)$ is not a sink in $\mathcal{G}_i$, a contradiction.

Now, it is easy to see that $\mathcal{I}_i$ is $(k, \max(k+1, b_{\mathbb{B}}))$-minimal. Indeed, since $\mathcal{I}$ is $(k, \max(k+1, b_{\mathbb{B}}))$-minimal, every subset of $\mathcal{V}$ of size at most $\max(k+1, b_{\mathbb{B}})$ is contained in the scope of some constraint of $\mathcal{I}$ and by construction also of $\mathcal{I}_i$. Moreover, by the previous paragraph, any two constraint of $\mathcal{I}_i$ agree on all k-element subsets of $\mathcal{V}$ within their scopes.

Since for every $i \geq 0$, if $\mathcal{G}_i$ is not empty, we remove at least one orbit of k-tuples under $\mathscr{G}$ from some constraint. By the oligomorphicity of $\mathscr{G}$, there exists $i_0 \geq 0$ such that $\mathcal{G}_{i_0}$ is empty. We claim that for every injective $\mathbf{v} \in \mathcal{V}^k$, $\mathrm{proj}_{\mathbf{v}}(\mathcal{I}_{i_0})$ contains precisely one orbit of k-tuples under $\mathscr{G}$: If $\mathrm{proj}_{\mathbf{v}}(\mathcal{I}_{i_0})$ contained more than one orbit, then $(\mathrm{proj}_{\mathbf{v}}(\mathcal{I}_{i_0}), O)$ would be a vertex of $\mathcal{G}_i$ for an arbitrary orbit $O \subseteq \mathrm{proj}_{\mathbf{v}}(\mathcal{I}_{i_0})$; O being a relation of $\mathbb{A}$ since $\mathbb{A}$ is a first-order expansion of $\mathbb{B}$.

It follows that $\mathcal{I}_{i_0} = (\mathcal{V}, \mathcal{C}_{i_0})$ is a non-trivial, $(k, \max(k+1, b_{\mathbb{B}}))$-minimal instance of $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{B})$ that satisfies the assumptions of Lemma 4.2.7. Hence, there exists a satisfying assignment for $\mathcal{I}_{i_0}$ and whence also for $\mathcal{I}$. ☐

4.3 Neoliberal permutation groups and bounded strict width

Definition 4.3.1. *Let $k \geq 2$, and let $\mathscr{G}$ be a permutation group acting on a set A. We say that $\mathscr{G}$ has* no k-algebraicity *if the only fixed points of any stabilizer of $\mathscr{G}$ by $k-1$ elements are these elements themselves. We say that $\mathscr{G}$ is* k-neoliberal *if it is oligomorphic, $(k-1)$-transitive, k-homogeneous, and has no k-algebraicity.*

Note that a permutation group $\mathscr{G}$ has no algebraicity if, and only if, it has no k-algebraicity, for every $k \geq 2$.

The notion of k-neoliberality is inspired by the notion of liberal binary cores introduced by Wrona in [81] – the automorphism group of every liberal binary core is 2-neoliberal. However, the opposite is not true – the automorphism group of the universal homogeneous $\mathbb{K}_3$-free graph is easily seen to be 2-neoliberal but its canonical binary structure $\mathbb{B}$ is a binary core which is not liberal. This is because a liberal binary core is supposed to be finitely bounded and the set of forbidden bounds should not contain any structure of size $3, \ldots, 6$. However, $\mathbb{K}_3$ is a 3-element graph which does not embed into $\mathbb{B}$ but all its subgraphs of size at most 2 do, and hence $\mathbb{K}_3$ has to be contained in any set of forbidden bounds for the universal homogeneous $\mathbb{K}_3$-free graph.

Example 4.3.2. *For every $k \geq 2$, the automorphism group of the universal homogeneous k-uniform hypergraph is k-neoliberal.*

Let $\mathbb{C}^2_\omega$ be the countably infinite equivalence relation where every equivalence class contains precisely 2 elements. Then $\mathrm{Aut}(\mathbb{C}^2_\omega)$ is oligomorphic, 1-transitive, and 2-homogeneous, but not 2-neoliberal. Indeed, for any element a of $\mathbb{C}^2_\omega$, the stabilizer of $\mathrm{Aut}(\mathbb{C}^2_\omega)$ by a fixes also the unique element of $\mathbb{C}^2_\omega$ which is in the same equivalence class as a.

On the other hand, the automorphism group of the countably infinite equivalence relation with equivalence classes of a fixed size $m > 2$ is easily seen to be 2-neoliberal.

Note that if $\mathscr{G}$ is a permutation group acting on a set A which is k-neoliberal for some $k \geq 2$ and which is not equal to the group of all permutations on A, then the number k is uniquely determined. Indeed, $k = \min\{i \geq 1 \mid \mathscr{G} \text{ is not } i\text{-transitive}\}$.

Note also that a permutation group $\mathscr{G}$ acting on a set A which is oligomorphic, $(k-1)$-transitive, and k-homogeneous for some $k \geq 2$ is k-neoliberal if, and only if, there exists no function $f\colon I^A_k \to A$ with the property that for every $\alpha \in \mathscr{G}$, $f(\alpha, \dots, \alpha) = \alpha f$. Indeed, if $\mathscr{G}$ is not k-neoliberal, then there exist pairwise distinct $a_1, \dots, a_k \in A$ such that a_k is a fixed point of the stabilizer of $\mathscr{G}$ by $a_1, \dots, a_{k-1}$. It is easy to see that we can obtain a function $f\colon I^A_k \to A$ with the stated property by defining for every $\mathbf{t} \in I^A_{k-1}$, $f(\mathbf{t}) := \alpha f(a_1, \dots, a_{k-1})$ where $\alpha \in \mathscr{G}$ is such that $(\alpha a_1, \dots, \alpha a_{k-1}) = \mathbf{t}$ – such α exists by the $(k-1)$-transitivity of $\mathscr{G}$. On the other hand, if such function f exists, then for every $a_1, \dots, a_{k-1}$, $f(a_1, \dots, a_{k-1})$ is a fixed point of the stabilizer of $\mathscr{G}$ by $a_1, \dots, a_{k-1}$, and $\mathscr{G}$ is not k-neoliberal.

4.3.1 Some implications with no bounded strict width

In this section, we first prove that if a structure $\mathbb{A}$ pp-defines certain implications, then it does not have bounded strict width (Lemmas 4.3.3 and 4.3.4). This will enable us to prove that if $\mathbb{A}$ has bounded strict width, and if a relation pp-definable in $\mathbb{A}$ contains a tuple with certain properties, then this relation contains an injective tuple with the same properties (Corollary 4.3.5).

Lemma 4.3.3. *Let $k \geq 3$, let $\mathscr{G}$ be a k-neoliberal permutation group, and let $\mathbb{B}$ be its canonical k-ary structure. Let $\mathbb{A}$ be a first-order expansion of $\mathbb{B}$, let $\ell \in \{2, \dots, k\}$, let $S \subseteq I^B_\ell$, and let ϕ be an $(S, =)$-implication in $\mathbb{A}$ with $\ell + 1$ variables. Then $\mathbb{A}$ does not have bounded strict width.*

Proof. Enumerate the variables of ϕ by $x_1, \dots, x_{\ell+1}$. Without loss of generality, $\mathbf{u} = (x_1, \dots, x_\ell)$ and $\mathbf{v} = (x_\ell, x_{\ell+1})$ are such that ϕ is an $(S, \mathbf{u}, =, \mathbf{v})$-implication in $\mathbb{A}$. The set $\phi^{\mathbb{A}}$ can then be viewed as an $(\ell+1)$-ary relation $R(x_1, \dots, x_{\ell+1})$.

Using the k-neoliberality of $\mathscr{G}$, we can find $a_1 \dots, a_\ell, b_1, \dots, b_\ell \in B$ with $a_\ell \neq b_\ell$ such that all of the following hold:

- $(a_1, \dots, a_\ell) \in S$,
- $(a_1, \dots, a_{\ell-1}, b_\ell) \in S$,
- $(b_1, \dots, b_{\ell-1}, a_\ell, b_\ell) \in R$.

To see this, let $(a_1, \dots, a_\ell) \in S$ be arbitrary. The fact that $\mathscr{G}$ has no k-algebraicity implies that there exists $b_\ell \in B$ which is distinct from a_ℓ but which lies in the same orbit under the stabilizer of $\mathscr{G}$ by $a_1, \ldots, a_{\ell-1}$. In particular, $(a_1, \ldots, a_\ell)$ and $(a_1, \ldots, a_{\ell-1}, b_\ell)$ lie in the same orbit under $\mathscr{G}$, and hence $(a_1, \ldots, a_{\ell-1}, b_\ell) \in S$. Finally, since $\mathscr{G}$ is 2-transitive and $\text{proj}_{(\ell,\ell+1)}(R) \not\subseteq \{(a,a) \mid a \in A\}$, we have $I_2^B \subseteq \text{proj}_{(\ell,\ell+1)}(R)$, and hence we can find $b_1, \ldots, b_{\ell-1}$ such that $(b_1, \ldots, b_{\ell-1}, a_\ell, b_\ell) \in R$.

Suppose for contradiction that $\mathbb{A}$ has bounded strict width. Then by Theorem 1.4.10, there exist $m \geq 3$ and an m-ary $f \in \text{Pol}(\mathbb{A})$ which is a local near-unanimity function on $\{a_1 \ldots, a_\ell, b_1, \ldots, b_\ell\}$. Since ϕ is an $(S,=)$-implication, it follows that $(b_1, \ldots, b_{\ell-1}, a_\ell) \notin S$ and $(a_1, \ldots, a_\ell, a_\ell), (a_1, \ldots, a_{\ell-1}, b_\ell, b_\ell) \in R$. Put

$$\mathbf{t}_1 := \begin{pmatrix} a_1 \\ \ldots \\ a_{\ell-1} \\ b_\ell \\ b_\ell \end{pmatrix}, \mathbf{t}_2 := \begin{pmatrix} b_1 \\ \ldots \\ b_{\ell-1} \\ a_\ell \\ b_\ell \end{pmatrix}, \mathbf{t}_3 = \cdots = \mathbf{t}_m := \begin{pmatrix} a_1 \\ \ldots \\ a_{\ell-1} \\ a_\ell \\ a_\ell \end{pmatrix}.$$

By the discussion above, $\mathbf{t}_i \in R$ for every $i \in [m]$. Since f preserves R, it follows that $f(\mathbf{t}_1, \ldots, \mathbf{t}_m) \in R$, i.e.,

$$\begin{pmatrix} a_1 & = & f(a_1, b_1, a_1, \ldots, a_1) \\ \ldots & & \ldots \\ a_{\ell-1} & = & f(a_{\ell-1}, b_{\ell-1}, a_{\ell-1}, \ldots, a_{\ell-1}) \\ a_\ell & = & f(b_\ell, a_\ell, \ldots, a_\ell) \\ & & f(b_\ell, b_\ell, a_\ell, \ldots, a_\ell) \end{pmatrix} \in R.$$

Since $(a_1, \ldots, a_\ell) \in S$ and ϕ is an $(S,=)$-implication, we get $f(b_\ell, b_\ell, a_\ell, \ldots, a_\ell) = a_\ell$. We can now proceed by induction to show

$$a_\ell = f(b_\ell, a_\ell, \ldots, a_\ell) = f(b_\ell, b_\ell, a_\ell, \ldots, a_\ell) = \cdots = f(b_\ell, \ldots, b_\ell) = b_\ell,$$

which is a contradiction to the choice of a_ℓ and b_ℓ. □

Lemma 4.3.4. *Let $k \geq 3$, let $\mathscr{G}$ be a k-neoliberal permutation group, and let $\mathbb{B}$ be its canonical k-ary structure. Let $\mathbb{A}$ be a first-order expansion of $\mathbb{B}$, let $\ell \in \{2, \ldots, k\}$, let $S \subseteq I_\ell^B$, and let ϕ be an $(S,=)$-implication in $\mathbb{A}$ with $\ell + 2$ variables. Then $\mathbb{A}$ does not have bounded strict width.*

Proof. Enumerate the variables of ϕ by $x_1, \ldots, x_{\ell+2}$. Without loss of generality, $\mathbf{u} = (x_1, \ldots, x_\ell)$ and $\mathbf{v} = (x_{\ell+1}, x_{\ell+2})$ are such that ϕ is an $(S, \mathbf{u}, =, \mathbf{v})$-implication in $\mathbb{A}$. The set $\phi^{\mathbb{A}}$ can then be viewed as an $(\ell+2)$-ary relation $R(x_1, \ldots, x_{\ell+2})$. Moreover, we can assume that for every $i \in \{1, \ldots, \ell\}$, and every $j \in \{\ell+1, \ell+2\}$, ϕ does not entail $S(\mathbf{u}) \Rightarrow x_i = x_j$ in $\mathbb{A}$; otherwise, the result follows immediately from Lemma 4.3.3.

Using the k-neoliberality of $\mathscr{G}$, we can find $a_1, \ldots, a_{\ell+1}, b_1, \ldots, b_{\ell+1}, c_1, d_2, \ldots, d_\ell \in B$ with $a_{\ell+1} \neq b_{\ell+1}$ such that all of the following hold:

- $(a_1, \ldots, a_\ell) \in S$,
- $(a_1, \ldots, a_\ell, a_{\ell+1}) \in I_{\ell+1}^B$,

- $(a_1, \dots, a_\ell, a_{\ell+1}, a_{\ell+1}) \in R$,
- $(c_1, a_2, \dots, a_\ell) \in S$,
- $(c_1, a_2, \dots, a_\ell, b_{\ell+1}) \in I^B_{\ell+1}$,
- $(c_1, a_2, \dots, a_\ell, b_{\ell+1}, b_{\ell+1}) \in R$,
- $(b_1, \dots, b_\ell, a_{\ell+1}, b_{\ell+1}) \in R$,
- $(b_1, d_2, \dots, d_\ell) \in S$.

To find these elements, let first $(a_1, \dots, a_\ell) \in S$ be arbitrary. By our assumption above, ϕ does not entail $S(\mathbf{u}) \Rightarrow x_i = x_j$ in $\mathbb{A}$ for any $i \in \{1, \dots, \ell\}$, $j \in \{\ell+1, \ell+2\}$, and hence there exists $a_{\ell+1} \in B$ such that $(a_1, \dots, a_\ell, a_{\ell+1}) \in I^B_{\ell+1} \cap \mathrm{proj}_{(1,\dots,\ell+1)}(R)$. Since ϕ is an $(S,=)$-implication, it follows that $(a_1, \dots, a_\ell, a_{\ell+1}, a_{\ell+1}) \in R$. By the fact that $\mathscr{G}$ has no k-algebraicity, we can find $b_{\ell+1}$ such that $a_{\ell+1} \neq b_{\ell+1}$ and both these elements lie in the same orbit under the stabilizer of $\mathscr{G}$ by $a_2, \dots, a_\ell$. In particular, it follows that $(a_2, \dots, a_\ell, b_{\ell+1})$ and $(a_2, \dots, a_{\ell+1})$ lie in the same orbit under $\mathscr{G}$, and hence there exists $c_1 \in B$ such that $(c_1, a_2, \dots, a_\ell, b_{\ell+1})$ and $(a_1, \dots, a_{\ell+1})$ lie in the same orbit under $\mathscr{G}$. In particular, $(c_1, a_2, \dots, a_\ell) \in S$ and it follows that $(c_1, a_2, \dots, a_\ell, b_{\ell+1}, b_{\ell+1}) \in R$. Since $(a_{\ell+1}, b_{\ell+1}) \in I^B_2 \subseteq \mathrm{proj}_{(\ell+1,\ell+2)}(R)$, there exist $b_1, \dots, b_\ell \in B$ such that $(b_1, \dots, b_\ell, a_{\ell+1}, b_{\ell+1}) \in R$. Finally, since $\mathscr{G}$ is 1-transitive, $\mathrm{proj}_{(1)}(S) = B$, and hence there exist $d_2, \dots, d_\ell \in B$ such that $(b_1, d_2, \dots, d_\ell) \in S$.

Suppose for contradiction that $\mathbb{A}$ has bounded strict width. Then by Theorem 1.4.10, there exist $m \geq 3$ and an m-ary $f \in \mathrm{Pol}(\mathbb{A})$ which is a local near-unanimity function on the set $\{a_1, \dots, a_{\ell+1}, b_1, \dots, b_{\ell+1}, c_1, d_2, \dots, d_\ell\}$. Put

$$\mathbf{t}_1 := \begin{pmatrix} c_1 \\ a_2 \\ \cdots \\ a_\ell \\ b_{\ell+1} \\ b_{\ell+1} \end{pmatrix}, \mathbf{t}_2 := \begin{pmatrix} b_1 \\ b_2 \\ \cdots \\ b_\ell \\ a_{\ell+1} \\ b_{\ell+1} \end{pmatrix}, \mathbf{t}_3 = \cdots = \mathbf{t}_m := \begin{pmatrix} a_1 \\ a_2 \\ \cdots \\ a_\ell \\ a_{\ell+1} \\ a_{\ell+1} \end{pmatrix}.$$

By the discussion above, $\mathbf{t}_i \in R$ for every $i \in [m]$. Since f preserves R, it follows that $\mathbf{t} := f(\mathbf{t}_1, \dots, \mathbf{t}_m) \in R$, i.e.,

$$\mathbf{t} = \begin{pmatrix} & & f(c_1, b_1, a_1, \dots, a_1) \\ a_2 & = & f(a_2, b_2, a_2, \dots, a_2) \\ \cdots & & \cdots \\ a_\ell & = & f(a_\ell, b_\ell, a_\ell, \dots, a_\ell) \\ a_{\ell+1} & = & f(b_{\ell+1}, a_{\ell+1}, \dots, a_{\ell+1}) \\ & & f(b_{\ell+1}, b_{\ell+1}, a_{\ell+1} \dots, a_{\ell+1}) \end{pmatrix} \in R.$$

Put

$$\mathbf{s}_1 := \begin{pmatrix} c_1 \\ a_2 \\ \dots \\ a_\ell \end{pmatrix}, \mathbf{s}_2 := \begin{pmatrix} b_1 \\ d_2 \\ \dots \\ d_\ell \end{pmatrix}, \mathbf{s}_3 = \dots = \mathbf{s}_m := \begin{pmatrix} a_1 \\ a_2 \\ \dots \\ a_\ell \end{pmatrix}.$$

By construction, $\mathbf{s}_i \in S$ for every $i \in [m]$. Since $f \in \mathrm{Pol}(\mathbb{A})$ and S is pp-definable from $\mathbb{A}$, f preserves S and it follows that $\mathbf{s} := f(\mathbf{s}_1, \dots, \mathbf{s}_m) \in S$, i.e.,

$$\mathbf{s} = \begin{pmatrix} & & f(c_1, b_1, a_1, \dots, a_1) \\ a_2 & = & f(a_2, d_2, a_2, \dots, a_2) \\ \dots & & \dots \\ a_\ell & = & f(a_\ell, d_\ell, a_\ell, \dots, a_\ell) \end{pmatrix} \in S.$$

Whence, $\mathbf{s}$ is precisely the tuple containing the first ℓ entries of $\mathbf{t}$. Since ϕ entails $S(\mathbf{u}) \Rightarrow x_{\ell+1} = x_{\ell+2}$ in $\mathbb{A}$, it follows that the last two entries of $\mathbf{t}$ are equal, i.e.,

$$f(b_{\ell+1}, b_{\ell+1}, a_{\ell+1}, \dots, a_{\ell+1}) = f(b_{\ell+1}, a_{\ell+1}, \dots, a_{\ell+1}) = a_{\ell+1}.$$

We can proceed by induction as follows to show that

$$a_{\ell+1} = f(b_{\ell+1}, a_{\ell+1}, \dots, a_{\ell+1}) = \dots = f(b_{\ell+1}, \dots, b_{\ell+1}) = b_{\ell+1},$$

which is a contradiction to the choice of $a_{\ell+1}$ and $b_{\ell+1}$. Suppose that for some $i \geq 2$, $i < m$ we have already shown $a_{\ell+1} = f(b_{\ell+1}, \dots, b_{\ell+1}, a_{\ell+1}, \dots, a_{\ell+1})$ where $b_{\ell+1}$ appears exactly i times. We can apply f to the rows of the matrix which has in its first i columns $(c_1, a_2, \dots, a_\ell, b_{\ell+1}, b_{\ell+1})$, in the $(i+1)$-th column $(b_1, \dots, b_\ell, a_{\ell+1}, b_{\ell+1})$, and in the remaining columns $(a_1, \dots, a_{\ell+1}, a_{\ell+1})$. Since all of its columns belong to R, we get as above that $a_{\ell+1} = f(b_{\ell+1}, \dots, b_{\ell+1}, a_{\ell+1}, \dots, a_{\ell+1})$ where $b_{\ell+1}$ appears $i+1$ times. ☐

The following corollary follows from Lemma 4.3.3 and Lemma 4.3.4.

Corollary 4.3.5. *Let $k \geq 3$, let $\mathbb{A}$ be a first-order expansion of the canonical k-ary structure $\mathbb{B}$ of a k-neoliberal permutation group $\mathscr{G}$, and suppose that $\mathbb{A}$ has bounded strict width. Let ϕ be a pp-formula over the signature of $\mathbb{A}$ with variables from a set V such that for all distinct $x, y \in V$, $\mathrm{proj}_{(x,y)}(\phi^{\mathbb{A}}) \not\subseteq \{(a,a) \mid a \in A\}$, and let $g \in \phi^{\mathbb{A}}$. Then there exists an injective $h \in \phi^{\mathbb{A}}$ with the property that for every $r \geq 1$ and for every $\mathbf{v} \in I_r^V$, if $g(\mathbf{v})$ is injective, then $g(\mathbf{v})$ and $h(\mathbf{v})$ belong to the same orbit under $\mathscr{G}$.*

Proof. Let W be the set of all tuples $\mathbf{v} \in I_r^V$ with $r \leq k$ and such that $f(\mathbf{v})$ is injective; we denote the orbit of $f(\mathbf{v})$ under $\mathscr{G}$ by $O_{\mathbf{v}}$. Note that all these orbits are pp-definable from $\mathbb{A}$ – indeed, for $r = k$ this is the fact that $\mathbb{A}$ is a first-order expansion of $\mathbb{B}$, and for $r < k$ this follows by the r-transitivity of $\mathscr{G}$. Hence, the formula

$$\psi \equiv \phi \wedge \bigwedge_{\mathbf{v} \in W} O_{\mathbf{v}}(\mathbf{v})$$

is equivalent to a pp-formula over $\mathbb{A}$. Moreover, it does not entail in $\mathbb{A}$ any equality among any two of its variables: otherwise, take a subset $W' \subseteq W$ which is maximal with respect to inclusion with the property that if we replace W by W' in the above definition, then the resulting formula ψ' does not entail in $\mathbb{A}$ any equality among any two of its variables. Then taking any $\mathbf{v} \in W \backslash W'$, ψ' entails $O_{\mathbf{v}}(\mathbf{v}) \Rightarrow x = y$ in $\mathbb{A}$ for some distinct $x, y \in V$ and by existentially quantifying all variables of ψ' except for the variables from the set $\mathrm{S}(\mathbf{v}) \cup \{x, y\}$, we obtain an $(O_{\mathbf{v}}, \mathbf{v}, =, (x, y))$-implication in $\mathbb{A}$, in contradiction with Lemmas 4.3.3 and 4.3.4.

Since $\mathscr{G}$ is $(k-1)$-transitive and $k \geq 3$, it is in particular 2-transitive, and hence the relation I_2^B is pp-definable from $\mathbb{B}$ and hence also from $\mathbb{A}$. It follows that I_m^B, where m is the number of variables of ψ, is pp-definable from $\mathbb{A}$ and hence, so is $\psi \wedge I_m$. Moreover, $\psi \wedge I_m$ is non-empty since otherwise, we would obtain an $(I_2^B, =)$-implication in contradiction with Lemmas 4.3.3 and 4.3.4 as in the previous paragraph with $I_2(v_1, v_2)$ for every $v_1 \neq v_2 \in V$ in the role of $O_{\mathbf{v}}(\mathbf{v})$.

Finally, observe that any $h \in (\psi \wedge I_m)^{\mathbb{A}}$ has the desired property by the k-homogeneity of $\mathscr{G}$. □

4.3.2 Critical relations

In this section, we adapt the notion of a critical relation from [81] to our situation and prove that no structure which satisfies the assumptions of Corollary 4.1.2 can pp-define such relation.

Definition 4.3.6. *Let $k \geq 2$, and let $\mathbb{A}$ be a relational structure. Let $C, D \subseteq I_k^A$ be disjoint and pp-definable from $\mathbb{A}$, let V be a set of $k+1$ variables, and let $\mathbf{u}, \mathbf{v} \in I_k^V$ be such that* $\mathrm{S}(\mathbf{u}) \cup \mathrm{S}(\mathbf{v}) = V$ *and such that $u_1, v_1 \notin \mathrm{S}(\mathbf{u}) \cap \mathrm{S}(\mathbf{v})$. We say that a pp-formula ϕ over the signature of $\mathbb{A}$ with variables from V is* critical *in $\mathbb{A}$* over $(C, D, \mathbf{u}, \mathbf{v})$ *if all of the following hold:*

- *ϕ is a $(C, \mathbf{u}, C, \mathbf{v})$-implication in $\mathbb{A}$,*
- *$D \subsetneq \mathrm{proj}_{\mathbf{u}}(\phi^{\mathbb{A}}) \subseteq I_k^A$,*
- *$D \subsetneq \mathrm{proj}_{\mathbf{v}}(\phi^{\mathbb{A}}) \subseteq I_k^A$,*
- *there exists no $D' \subseteq A^k$ with $D' \cap D \subsetneq D$ and such that for every $f \in \phi^{\mathbb{A}}$, it holds that $f(\mathbf{u}) \in D$ implies $f(\mathbf{v}) \in D'$.*

Lemma 4.3.7. *Let $k \geq 3$, let $\mathscr{G}$ be a k-neoliberal permutation group, and let $\mathbb{A}$ be a first-order expansion of the canonical k-ary structure of $\mathscr{G}$. Suppose that there exists a pp-formula ϕ which is critical in $\mathbb{A}$ over $(C, D, \mathbf{u}, \mathbf{v})$ for some k-ary C, D, and some $\mathbf{u}, \mathbf{v}$. Then $\mathbb{A}$ does not have bounded strict width.*

Proof. First of all, observe that the formula $\phi_{\mathrm{Inj}} := \phi \wedge I_{k+1}$ is equivalent to a pp-formula over $\mathbb{A}$ by the 2-transitivity of $\mathscr{G}$ and ϕ_{Inj} is still critical in $\mathbb{A}$ over $(C, D, \mathbf{u}, \mathbf{v})$ by Corollary 4.3.5. Indeed, all items of Definition 4.3.6 except for the first one depend only on $\mathrm{proj}_{\mathbf{u}}(\phi^{\mathbb{A}})$ and on $\mathrm{proj}_{\mathbf{v}}(\phi^{\mathbb{A}})$ and moreover, these projections are injective. Furthermore, for all distinct

$x, y \in V$, $\mathrm{proj}_{(x,y)}(\phi^{\mathbb{A}}) \not\subseteq \{(a,a) \mid a \in A\}$, and hence Corollary 4.3.5 implies that for every $g \in \phi^{\mathbb{A}}$, there exists $h \in \phi^{\mathbb{A}}_{\mathrm{Inj}}$ such that $g(\mathbf{u})$ and $h(\mathbf{u})$ belong to the same orbit under $\mathscr{G}$, and so do $g(\mathbf{v})$ and $h(\mathbf{v})$. It follows that ϕ_{Inj} satisfies also the first item of Definition 4.3.6.

Let $V = \{x_1, \dots, x_{k+1}\}$ be the set of variables of ϕ. We can assume without loss of generality that $\mathbf{u} = (x_1, \dots, x_k)$ and $S(\mathbf{v}) = \{x_2, \dots, x_{k+1}\}$. Let $\mathbf{v} = (x_{i_1}, \dots, x_{i_k})$. In the rest of the proof, for any k-tuple $(t_2, \dots, t_{k+1})$, we write $\mathrm{proj}_{\mathbf{v}}(t_2, \dots, t_{k+1})$ for the tuple $(t_{i_1}, \dots, t_{i_k})$ by abuse of notation.

Suppose for contradiction that $\mathbb{A}$ has bounded strict width – hence, by Theorem 1.4.10, there exists an oligopotent quasi near-unanimity operation $f \in \mathrm{Pol}(\mathbb{A})$ of arity $\ell \geq 2$. Let us define $\mathbf{w}^2, \dots, \mathbf{w}^{k+1} \in A^\ell$ as follows. Let $d^2, \dots, d^{k+1} \in A$ be arbitrary such that $\mathrm{proj}_{\mathbf{v}}(d^2, \dots, d^{k+1}) \in D$, and let $\mathbf{w}^j$ be constant with value d^j for all $j \in \{2, \dots, k+1\}$. Setting $J = [\ell]$ in Claim 4.3.8 below, we get that $\mathrm{proj}_{\mathbf{v}}(f(\mathbf{w}^2), \dots, f(\mathbf{w}^{k+1})) \in C$. On the other hand, $\mathrm{proj}_{\mathbf{v}}(f(\mathbf{w}^2), \dots, f(\mathbf{w}^{k+1})) \in D$ since D is pp-definable from $\mathbb{A}$, and hence preserved by f, contradicting that C and D are disjoint.

Claim 4.3.8. *For every $J \subseteq [\ell]$, the following holds. Let $\mathbf{w}^2, \dots, \mathbf{w}^{k+1} \in A^\ell$ be such that $\mathbf{w}^2, \dots, \mathbf{w}^k$ are constant tuples, and such that $\mathrm{proj}_{\mathbf{v}}(w_i^2, \dots, w_i^{k+1}) \in C$ for all $i \in [\ell] \backslash J$ and $\mathrm{proj}_{\mathbf{v}}(w_i^2, \dots, w_i^{k+1}) \in D$ for all $i \in J$. Then $\mathrm{proj}_{\mathbf{v}}(f(\mathbf{w}^2), \dots, f(\mathbf{w}^{k+1})) \in C$.*

We will prove the claim by induction on $n := |J|$. For $n = 0$ the claim follows by the assumption that C is pp-definable from $\mathbb{A}$, and hence it is preserved by f.

Let now $n > 0$, and suppose that Claim 4.3.8 holds for $n-1$. The set $\phi^{\mathbb{A}}_{\mathrm{Inj}}$ can be viewed as a $(k+1)$-ary relation $R(x_1, \dots, x_{k+1})$. Let $m \in J$ be arbitrary, and set $J' := J \backslash \{m\}$. Using the k-neoliberality of $\mathscr{G}$, we will find $\mathbf{w}^1 \in A^\ell$ such that all of the following hold:

- $(w_i^1, \dots, w_i^{k+1}) \in R$ for all $i \in [\ell] \backslash \{m\}$,
- $(w_i^1, \dots, w_i^k) \in C$ for all $i \in [\ell] \backslash J$ and $(w_i^1, \dots, w_i^k) \in D$ for all $i \in J'$,
- $(w_m^1, \dots, w_m^k) \in C$ and $w_m^1 \neq w_m^{k+1}$.

To find $\mathbf{w}^1$, let first $i \in [\ell] \backslash J$. Note that by the definition of a $(C, \mathbf{u}, C, \mathbf{v})$-implication in $\mathbb{A}$, there exists no $C' \subsetneq C$ such that for every $g \in \phi^{\mathbb{A}}_{\mathrm{Inj}}$, it holds that $g(\mathbf{u}) \in C$ implies $g(\mathbf{v}) \in C'$. Since $\mathrm{proj}_{\mathbf{v}}(w_i^2, \dots, w_i^{k+1}) \in C$, it follows that there exists $h \in \phi^{\mathbb{A}}$ such that $h(\mathbf{v}) = \mathrm{proj}_{\mathbf{v}}(w_i^2, \dots, w_i^{k+1})$ and $h(\mathbf{u}) \in C$. Set $w_i^1 := h(x_1)$. Let now $i \in J'$. Using the definition of a critical formula over $(C, D, \mathbf{u}, \mathbf{v})$, we can find $h \in \phi^{\mathbb{A}}$ such that $h(\mathbf{v}) = \mathrm{proj}_{\mathbf{v}}(w_i^2, \dots, w_i^{k+1})$ and $h(\mathbf{u}) \in D$ similarly as above, and set $w_i^1 := h(x_1)$. It remains to find w_m^1 satisfying the last item. By the $(k-1)$-transitivity of $\mathscr{G}$ and the fact that it has no k-algebraicity, we can extend any tuple $(a^2, \dots, a^k) \in I^A_{k-1}$ to injective tuples $(a^1, \dots, a^k), (b^1, a^2, \dots, a^k) \in O$ with $a^1 \neq b^1$, for an arbitrary injective orbit O of k-tuples under $\mathscr{G}$. Applying this fact to the tuple $(w_m^2, \dots, w_m^k)$, we get w_m^1 such that $(w_m^1, \dots, w_m^k) \in C$ and $w_m^1 \neq w_m^{k+1}$ as desired.

Note that $\mathbf{w}^1, \dots, \mathbf{w}^k$ satisfy the assumptions of Claim 4.3.8 for J' in the role of J up to permuting the order of the tuples. Indeed, $\mathbf{w}^2, \dots, \mathbf{w}^k$ are constant, and it holds that $(w_i^1, w_i^2, \dots, w_i^k) \in C$ for all $i \in [\ell] \backslash J'$ and $(w_i^1, w_i^2, \dots, w_i^k) \in D$ for all $i \in J'$. Since $|J'| = n-1$, the induction hypothesis yields that $(f(\mathbf{w}^1), \dots, f(\mathbf{w}^k)) \in C$.

Since $(w_m^1, w_m^{k+1}) \in I_2^A = \mathrm{proj}_{(1,k+1)}(R)$, there exist $a^2, \dots, a^k \in A$ such that

$$(w_m^1, a^2, \dots, a^k, w_m^{k+1}) \in R.$$

For all $j \in \{2, \ldots, k\}$, let $\mathbf{q}^j$ be the tuple obtained by replacing the m-th coordinate of $\mathbf{w}^j$ by a^j. It follows that $(w_i^1, q_i^2, \ldots, q_i^k, w_i^{k+1}) \in R$ for all $i \in [\ell]$, and since R is preserved by f, it holds that $(f(\mathbf{w}^1), f(\mathbf{q}^2), \ldots, f(\mathbf{q}^k), f(\mathbf{w}^{k+1})) \in R$.

Since $\mathbf{w}^i$ is constant, and since f is a quasi near-unanimity operation, we have $f(\mathbf{w}^j) = f(\mathbf{q}^j)$ for all $j \in \{2, \ldots, k\}$. This implies that $(f(\mathbf{w}^1), f(\mathbf{q}^2), \ldots, f(\mathbf{q}^k)) \in C$. Since ϕ entails $C(\mathbf{u}) \Rightarrow C(\mathbf{v})$ in $\mathbb{A}$, we get $\operatorname{proj}_{\mathbf{v}}(f(\mathbf{q}^2), \ldots, f(\mathbf{q}^k), f(\mathbf{w}^{k+1})) \in C$, and hence $\operatorname{proj}_{\mathbf{v}}(f(\mathbf{w}^2), \ldots, f(\mathbf{w}^{k+1})) \in C$. □

4.3.3 Composition of implications

In this section, we introduce the notion of composition of implications which will play an important role in the rest of Section 4.3.

Definition 4.3.9. *Let $\mathbb{A}$ be a relational structure, let $k \geq 1$, let $C, D, E \subseteq A^k$ be non-empty, let ϕ_1 be a $(C, \mathbf{u}^1, D, \mathbf{v}^1)$-implication in $\mathbb{A}$, and let ϕ_2 be a $(D, \mathbf{u}^2, E, \mathbf{v}^2)$-implication in $\mathbb{A}$. Let us rename the variables of ϕ_2 so that $\mathbf{v}^1 = \mathbf{u}^2$ and so that ϕ_1 and ϕ_2 do not share any other variables. We define $\phi_1 \circ \phi_2$ to be the pp-formula arising from the formula $\phi_1 \wedge \phi_2$ by existentially quantifying all variables that are not contained in $\mathrm{S}(\mathbf{u}^1) \cup \mathrm{S}(\mathbf{v}^2)$.*

Let ψ be a $(C, \mathbf{u}, C, \mathbf{v})$-implication. For $n \geq 2$, we write $\psi^{\circ n}$ for the pp-formula $\psi \circ \cdots \circ \psi$ where ψ appears exactly n times.

Lemma 4.3.10. *Let $k \geq 3$, let $\mathbb{A}$ be a first-order expansion of the canonical k-ary structure of a permutation group $\mathscr{G}$. Let ϕ_1, ϕ_2 be as in Definition 4.3.9, and suppose that $\operatorname{proj}_{\mathbf{v}^1}(\phi_1) = \operatorname{proj}_{\mathbf{u}^2}(\phi_2)$. Then $\phi := \phi_1 \circ \phi_2$ is a $(C, \mathbf{u}^1, E, \mathbf{v}^2)$-pre-implication in $\mathbb{A}$. Moreover, for all orbits $O_1 \subseteq \operatorname{proj}_{\mathbf{u}^1}(\phi_1^{\mathbb{A}})$, $O_3 \subseteq \operatorname{proj}_{\mathbf{v}^2}(\phi_2^{\mathbb{A}})$ under $\mathscr{G}$, $\phi^{\mathbb{A}}$ contains an O_1O_3-mapping if, and only if, there exists an orbit O_2 under $\mathscr{G}$ such that $\phi_1^{\mathbb{A}}$ contains an O_1O_2-mapping and $\phi_2^{\mathbb{A}}$ contains an O_2O_3-mapping.*

Suppose moreover that $\mathscr{G}$ is k-neoliberal, that $\mathbb{A}$ has bounded strict width and that ϕ_1 and ϕ_2 are injective implications. Then ϕ is a $(C, \mathbf{u}^1, E, \mathbf{v}^2)$-implication in $\mathbb{A}$. Restricting $\phi^{\mathbb{A}}$ to injective mappings, one moreover obtains an injective $(C, \mathbf{u}^1, E, \mathbf{v}^2)$-implication, which for all injective orbits $O_1 \subseteq C, O_3 \subseteq E$ under $\mathscr{G}$ contains an O_1O_3-mapping if, and only if, $\phi^{\mathbb{A}}$ contains such mapping.

Proof. Let us assume as in Definition 4.3.9 that $\mathbf{v}^1 = \mathbf{u}^2$ and that ϕ_1, ϕ_2 do not share any further variables. Let V_1 be the set of variables of ϕ_1, let V_2 be the set of variables of ϕ_2, and let V be the set of variables of ϕ. We will first prove the last sentence of the first part of Lemma 4.3.10 about O_1O_3-mappings. To this end, let $O_1 \subseteq \operatorname{proj}_{\mathbf{u}^1}(\phi_1^{\mathbb{A}}), O_2 \subseteq \operatorname{proj}_{\mathbf{v}^1}(\phi_1^{\mathbb{A}}), O_3 \subseteq \operatorname{proj}_{\mathbf{v}^2}(\phi_2^{\mathbb{A}})$ be orbits under $\mathscr{G}$, and suppose that $\phi_1^{\mathbb{A}}$ contains an O_1O_2-mapping f and $\phi_2^{\mathbb{A}}$ contains an O_2O_3-mapping g. Using that $f|_{V_1 \cap V_2}$ is contained in the same orbit under $\mathscr{G}$ as $g|_{V_1 \cap V_2}$, find a mapping $h \colon V_1 \cup V_2 \to A$ such that $h|_{V_1}$ is contained in the same orbit under $\mathscr{G}$ as f and $h|_{V_2}$ is contained in the same orbit as g. It follows that $h \in (\phi_1 \wedge \phi_2)^{\mathbb{A}}$, and hence $h|_V \in \phi^{\mathbb{A}}$ is an O_1O_3 mapping. On the other hand, if $\phi^{\mathbb{A}}$ contains an O_1O_3-mapping h', we can extend it to a mapping $h \in (\phi_1 \wedge \phi_2)^{\mathbb{A}}$ such that $f := h|_{V_1} \in \phi_1^{\mathbb{A}}$ and $g := h|_{V_2} \in \phi_2^{\mathbb{A}}$ by the definition of ϕ. Setting O_2 to be the orbit of $h(\mathbf{v}^1)$, we get that $f \in \phi_1^{\mathbb{A}}$ is an O_1O_2-mapping and $g \in \phi_2^{\mathbb{A}}$ is an O_2O_3-mapping as desired.

Observe now that the fact that ϕ is a $(C, \mathbf{u}^1, E, \mathbf{v}^2)$-pre-implication in $\mathbb{A}$ follows from the previous paragraph. Indeed, it follows immediately that ϕ satisfies items (4) and (5) of Definition 4.2.4. To see that items (2) and (3) are satisfied as well, take any $g \in \phi_2^{\mathbb{A}}$ with $g(\mathbf{v}^2) \notin E$, let O_2 be the orbit of $g(\mathbf{u}^2)$ under $\mathscr{G}$, and let O_3 be the orbit of $g(\mathbf{v}^2)$. Since $\mathrm{proj}_{\mathbf{v}^1}(\phi_1) = \mathrm{proj}_{\mathbf{u}^2}(\phi_2)$, we can find an O_1O_2-mapping in $\phi_1^{\mathbb{A}}$, and $\phi^{\mathbb{A}}$ contains an O_1O_3-mapping witnessing that $C \subsetneq \mathrm{proj}_{\mathbf{u}^1}(\phi^{\mathbb{A}}), C \subsetneq \mathrm{proj}_{\mathbf{v}^2}(\phi^{\mathbb{A}})$.

To prove the second part of the lemma, we will prove that for all orbits $O_1 \subseteq C, O_2 \subseteq D, O_3 \subseteq E$ under $\mathscr{G}$ such that $\phi_1^{\mathbb{A}}$ contains an O_1O_2-mapping f and $\phi_2^{\mathbb{A}}$ contains an O_2O_3-mapping g, $\phi^{\mathbb{A}}$ contains an injective O_1O_3-mapping h. Note that as in the previous paragraph, it is enough to find an injective mapping $h \in (\phi_1 \wedge \phi_2)^{\mathbb{A}}$ such that $h|_{V_1}$ is contained in the same orbit under $\mathscr{G}$ as f and $h|_{V_2}$ is contained in the same orbit as g. Let U_1 be the set of all injective tuples of variables from V_1 of length at most k, and for every $\mathbf{v} \in U_1$, let us denote the orbit of $f(\mathbf{v})$ by $O^f_{\mathbf{v}}$. Similarly, let U_2 be the set of all injective tuples of variables from V_2 of length at most k, and for every tuple $\mathbf{v} \in U_2$, let $O^g_{\mathbf{v}}$ be the orbit of $g(\mathbf{v})$. Let us define a formula ψ with variables from $V_1 \cup V_2$ by

$$\psi \equiv \bigwedge_{\mathbf{v} \in U_1} O^f_{\mathbf{v}}(\mathbf{v}) \wedge \bigwedge_{\mathbf{v} \in U_2} O^g_{\mathbf{v}}(\mathbf{v}).$$

Note that since $\mathbb{A}$ is a first-order expansion of the canonical k-ary structure of $\mathscr{G}$, all orbits $O^f_{\mathbf{v}}, O^g_{\mathbf{v}}$ are pp-definable from $\mathbb{A}$, and hence ψ is equivalent to a pp-formula over $\mathbb{A}$. Since $\mathscr{G}$ is k-neoliberal, and since $\mathbb{A}$ has bounded strict width, we can proceed as in the proof of Corollary 4.3.5 and use Lemmas 4.3.3 and 4.3.4 to show that $\psi^{\mathbb{A}}$ contains an injective mapping h. By the construction and by the k-homogeneity of $\mathscr{G}$, this mapping satisfies our assumptions. □

The following observation states a few properties of implications and their compositions which will be used later.

Observation 4.3.11. *Let $\mathbb{A}$ be a relational structure, let $k \geq 2$, let $C \subseteq A^k$, let ϕ_1 be a $(C, \mathbf{u}^1, C, \mathbf{v}^1)$-implication in $\mathbb{A}$, let ϕ_2 be a $(C, \mathbf{u}^2, C, \mathbf{v}^2)$-implication in $\mathbb{A}$. Let p_1 be the number of variables of ϕ_1, let p_2 be the number of variables of ϕ_2, and let ℓ be the number of variables of $\phi := \phi_1 \circ \phi_2$. Then both of the following hold.*

1. $\ell \geq \max(p_1, p_2)$.
2. $\ell = p_1 = p_2$ *if, and only if,* $\mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{v}^2) = \mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{v}^1) = \mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{u}^2) \cap \mathrm{S}(\mathbf{v}^2)$.
3. *Suppose that $\phi_1 = \phi_2$ and $\ell = p_1$. Then for every $i \in [k]$, it holds that if v^1_i is contained in the intersection in (2), then so is u^1_i.*

Proof. Let us rename the variables of ϕ_2 as in Definition 4.3.9 so that $\mathbf{v}^1 = \mathbf{u}^2$, and ϕ_1 and ϕ_2 do not share any further variables.

For (1), observe that the number of variables of a $(C, \mathbf{u}, C, \mathbf{v})$-implication is equal to $2k - |\,\mathrm{S}(\mathbf{u}) \cap \mathrm{S}(\mathbf{v})|$. Hence, for ϕ we get that $\ell = 2k - |\,\mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{v}^2)|$, and since $\mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{v}^2)$ is contained both in $\mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{v}^1)$ and in $\mathrm{S}(\mathbf{u}^2) \cap \mathrm{S}(\mathbf{v}^2)$, (1) follows by applying the same reasoning to p_1 and p_2.

For (2), observe that by the previous paragraph, $\ell = p_1$ if, and only if, $\mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{v}^2) = \mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{v}^1)$. Similarly, $\ell = p_2$ if, and only if, $\mathrm{S}(\mathbf{u}^1) \cap \mathrm{S}(\mathbf{v}^2) = \mathrm{S}(\mathbf{u}^2) \cap \mathrm{S}(\mathbf{v}^2)$, and (2) follows by the fact that $\mathbf{v}^1 = \mathbf{u}^2$.

For (3), suppose that $v_i^1 \in \mathrm{S}(\mathbf{u}^1)$. It follows by (2) that $v_i^1 = u_i^2 \in \mathrm{S}(\mathbf{v}^2)$, and since $\phi_1 = \phi_2$ and $\mathbf{u}^2$ was obtained by renaming $\mathbf{v}^1$, it follows that $u_i^1 \in \mathrm{S}(\mathbf{v}^1)$. □

4.3.4 Digraphs of implications

In this section, we reformulate the notion of digraph of implications from [81] and prove a few auxiliary statements about these digraphs.

Definition 4.3.12. *Let $k \geq 3$, let $\mathbb{A}$ be a first-order expansion of the canonical k-ary structure of a k-neoliberal permutation group $\mathscr{G}$. Let $\emptyset \neq C \subseteq A^k$, and let ϕ be $(C, \mathbf{u}, C, \mathbf{v})$-implication in $\mathbb{A}$ such that $\mathrm{proj}_{\mathbf{u}}(\phi^{\mathbb{A}}) = \mathrm{proj}_{\mathbf{v}}(\phi^{\mathbb{A}}) =: E$. Let $\mathbf{Vert}(E)$ be the set of all orbits under $\mathscr{G}$ contained in E. Let $\mathcal{B}_\phi \subseteq \mathbf{Vert}(E) \times \mathbf{Vert}(E)$ be the directed graph such that $\mathcal{B}_\phi$ contains an arc $(O, P) \in \mathbf{Vert}(E) \times \mathbf{Vert}(E)$ if $\phi^{\mathbb{A}}$ contains an OP-mapping.*

We say that $S \subseteq \mathbf{Vert}(E)$ is a strongly connected component *if it is a maximal set with respect to inclusion such that for all (not necessary distinct) vertices $O, P \in S$, there exists a path in $\mathcal{B}_\phi$ connecting O and P. We say that S is a* sink *in $\mathcal{B}_\phi$ if every arc originating in S ends in S; S is a* source *in $\mathcal{B}_\phi$ if every arc finishing in S originates in S.*

Note that the digraph $\mathcal{B}_\phi$ can be defined also for relational structures which do not satisfy the assumptions on $\mathbb{A}$ from Definition 4.3.12; however, these assumptions are needed in the proof of Lemma 4.3.15 so we chose to include them already in Definition 4.3.12. Note also that $\mathcal{B}_\phi$ can contain vertices which are not contained in any strongly connected component.

Observation 4.3.13. *Let ϕ be as in Definition 4.3.12. Then there exist strongly connected components $S_1 \subseteq \mathbf{Vert}(C), S_2 \subseteq \mathbf{Vert}(E \backslash C)$ in $\mathcal{B}_\phi$ such that S_1 is a sink in $\mathcal{B}_\phi$, and S_2 is a source in $\mathcal{B}_\phi$. Moreover, since $\mathrm{proj}_{\mathbf{u}}(\phi^{\mathbb{A}}) = \mathrm{proj}_{\mathbf{v}}(\phi^{\mathbb{A}}) = E$, any vertex $O \in \mathbf{Vert}(E)$ has an outgoing and an incoming arc, i.e., the digraph $\mathcal{B}_\phi$ is* smooth.

Proof. The second part of the lemma is immediate. To prove the first part, observe that since ϕ is a (C, C)-implication in $\mathbb{A}$, it follows that $\mathbf{Vert}(C)$ is a sink in $\mathcal{B}_\phi$. Using the oligomorphicity of $\mathscr{G}$ and the definition of a (C, C)-implication, we get that the induced subgraph of $\mathcal{B}_\phi$ on $\mathbf{Vert}(C)$ is finite and smooth. Hence, there exists a strongly connected component $S_1 \subseteq \mathbf{Vert}(C)$ in $\mathcal{B}_\phi$ which is a sink in the induced subgraph, and hence also in $\mathcal{B}_\phi$. Similarly, one sees that $\mathbf{Vert}(E \backslash C)$ is a source in $\mathcal{B}_\phi$, and it contains a strongly connected component S_2 which is itself a source in $\mathcal{B}_\phi$. □

Let ϕ be as in Definition 4.3.12, and set $I_\phi := \{i \in [k] \mid u_i \in \mathrm{S}(\mathbf{v})\}$. If the number of variables of ϕ is equal to the number of variables of $\phi \circ \phi$, then item (3) in Observation 4.3.11 yields that $I_\phi = \{i \in [k] \mid v_i \in \mathrm{S}(\mathbf{u})\}$.

Definition 4.3.14. *Let ϕ be as in Definition 4.3.12. We say that ϕ is* complete *if the number of variables of $\phi \circ \phi$ is equal to the number of variables of ϕ, $u_i = v_i$ for every $i \in I_\phi$, and each strongly connected component of $\mathcal{B}_\phi$ contains all possible arcs including loops.*

The following is a modification of Lemma 36 in [81]:

Lemma 4.3.15. *Let ϕ be as in Definition 4.3.12. Then there exists a complete injective (C, C)-implication in $\mathbb{A}$.*

Proof. We will construct the desired complete implication as a conjunction of a power of ϕ and I_ℓ, where ℓ is the number of variables of the power of ϕ. Note that for every $n \geq 1$, the number of variables of $\phi^{\circ n}$ is at most $2k$ by Definition 4.3.9 and this number never decreases with increasing n by item (1) in Observation 4.3.11. Hence, there is $n \geq 1$ such that the number of variables of $\phi^{\circ n}$ is the biggest among all choices of n. Let us denote the number of variables of $\phi^{\circ n}$ by ℓ; it follows that the number of variables of $(\phi^{\circ n})^{\circ m}$ is equal to ℓ for every $m \geq 1$. Let us replace ϕ by $\phi^{\circ n}$. It follows from Lemma 4.3.10 that for every $m \geq 1$, $\phi^{\circ m} \wedge I_\ell$ is an injective (C, C)-implication. Now, it follows by item (3) in Observation 4.3.11 that there exists a unique bijection $\sigma\colon I_\phi \to I_\phi$ such that $u_i = v_{\sigma(i)}$ for every $i \in I_\phi$. Replacing ϕ with a power of ϕ again, we can assume that σ is the identity.

Now, we can find $m \geq 1$ such that the number of strongly connected components of $\phi^{\circ m}$ is maximal among all possible choices of m. It follows that composing $\phi^{\circ m}$ with itself arbitrarily many times does not disconnect any vertices from $\mathbf{Vert}(C)$ which are contained in the same strongly connected component of $\mathcal{B}_{\phi^{\circ m}}$; we replace ϕ by $\phi^{\circ m}$. Taking another power of ϕ and replacing ϕ again, we can assume that every strongly connected component of $\mathcal{B}_\phi$ contains all loops. Now, setting p to be the number of vertices of $\mathcal{B}_\phi$, we have that, replacing ϕ with $\phi^{\circ p}$, every strongly connected component of $\mathcal{B}_\phi$ contains all arcs, whence $\phi \wedge I_\ell$ is a complete injective (C, C)-implication. □

4.3.5 Proof of the main result

Theorem 4.1.1. *Let $k \geq 3$, let $\mathbb{B}$ be the canonical k-ary structure of a k-neoliberal permutation group $\mathscr{G}$, and suppose that $\mathbb{B}$ has finite duality. Then any first-order expansion of $\mathbb{B}$ with bounded strict width is implicationally simple on injective instances.*

Proof. Let $\mathbb{A}$ be a first-order expansion of $\mathbb{B}$ with bounded strict width. Striving for contradiction, suppose that $\mathbb{A}$ is implicationally hard on injective instances. Then the injective implication graph $\mathcal{G}_{\mathbb{A}}^{\mathrm{Inj}}$ contains a directed cycle $(D_1, C_1), \ldots, (D_{n-1}, C_{n-1}), (D_n, C_n) = (D_1, C_1)$. This means that for all $i \in [n-1]$, there exists an injective $(C_i, \mathbf{u}^i, C_{i+1}, \mathbf{u}^{i+1})$-implication ϕ_i in $\mathbb{A}$ with $\mathrm{proj}_{\mathbf{u}^i}(\phi_i^{\mathbb{A}}) = D_i$, and $\mathrm{proj}_{\mathbf{u}^{i+1}}(\phi_i^{\mathbb{A}}) = D_{i+1}$.

Let us define $\phi := ((\phi_1 \circ \phi_2) \circ \ldots \circ \phi_{n-1})$. Restricting $\phi^{\mathbb{A}}$ to injective mappings, we obtain an injective $(C_1, \mathbf{u}^1, C_1, \mathbf{u}^n)$-implication by Lemma 4.3.10. Lemma 4.3.15 asserts us that there exists a complete injective (C_1, C_1)-implication ψ in $\mathbb{A}$.

Observation 4.3.13 yields that there exist $C \subseteq C_1$, and $D \subseteq D_1 \backslash C_1$ such that $\mathbf{Vert}(C)$ is a strongly connected component which is a sink in $\mathcal{B}_\psi$, and $\mathbf{Vert}(D)$ is a strongly connected component which is a source in $\mathcal{B}_\psi$. Observe that since $\mathbb{A}$ is a first-order expansion of $\mathbb{B}$ and since ψ is complete, C is pp-definable from $\mathbb{A}$. Indeed, for any fixed orbit $O \subseteq C$ under $\mathscr{G}$, C is equal to the set of all orbits $P \subseteq C_1$ such that $\psi^{\mathbb{A}}$ contains an OP-mapping. We can observe in a similar way that D is pp-definable from $\mathbb{A}$ as well. Moreover, ψ is easily seen to be a complete (C, C)-implication in $\mathbb{A}$.

Since $\mathbb{B}$ has finite duality, there exists a number $d \geq 3$ such that for every finite structure $\mathbb{X}$ in the signature of $\mathbb{B}$, it holds that if every substructure of $\mathbb{X}$ of size at most $d-2$ maps homomorphically to $\mathbb{B}$, then so does $\mathbb{X}$. Set $\rho := \psi^{\circ d}$. Let V be the set of variables of ρ. It follows from Lemma 4.3.10 that ρ is a $(C, \mathbf{u}, C, \mathbf{v})$-implication in $\mathbb{A}$ for some $\mathbf{u}, \mathbf{v}$. We are going to prove the following claim using the finite duality of $\mathbb{B}$ and by the completeness of ψ.

Claim 4.3.16. *Every $f \in A^V$ with $f(\mathbf{u}) \in C$ and $f(\mathbf{v}) \in C$ is an element of $\rho^{\mathbb{A}}$. The same holds for any $f \in A^V$ with $f(\mathbf{u}) \in D$ and $f(\mathbf{v}) \in D$.*

To prove Claim 4.3.16, let $f \in A^V$ be as in the statement of the claim. Up to renaming variables, we can assume that ψ is a $(C, \mathbf{u}, C, \mathbf{v})$-implication. Completeness of ψ implies that $I_\psi = \{i \in [k] \mid u_i \in \mathrm{S}(\mathbf{v})\} = \{i \in [k] \mid v_i \in \mathrm{S}(\mathbf{u})\}$, and that $u_i = v_i$ for every $i \in I_\psi$. Let $\mathbf{w}^0, \dots, \mathbf{w}^d$ be k-tuples of variables such that $w_i^j = u_i = v_i$ for all $i \in I_\psi$ and all $j \in \{0, \dots, d\}$, and disjoint otherwise. We can moreover assume that $\mathbf{w}^0 = \mathbf{u}$, $\mathbf{w}^d = \mathbf{v}$. For every $j \in \{0, \dots, d-1\}$, let ψ_{j+1} be the $(C, \mathbf{w}^j, C, \mathbf{w}^{j+1})$-implication obtained from ψ by renaming $\mathbf{w}^j$ by $\mathbf{v}^1$, and $\mathbf{w}^{j+1}$ by $\mathbf{v}^2$. It follows that ρ is equivalent to the formula obtained from $\psi_1 \wedge \cdots \wedge \psi_d$ by existentially quantifying all variables that are not contained in $\mathrm{S}(\mathbf{u}) \cup \mathrm{S}(\mathbf{v})$. In order to proof Claim 4.3.16, it is therefore enough to show that f can be extended to a mapping $h \in (\psi_1 \wedge \cdots \wedge \psi_d)^{\mathbb{A}}$.

Now, for all $p, q \in [k]$, we identify the variables $w_q^0 = u_q$ and $w_p^d = v_p$ if $f(v_p) = f(u_q)$. Observe that $f|_{\mathrm{S}(\mathbf{u}) \cap \mathrm{S}(\mathbf{v})}$ is injective since $f(\mathbf{u}) \in C \subseteq I_k^A$, and hence this identification does not force any variables from $\mathrm{S}(\mathbf{u}) \cap \mathrm{S}(\mathbf{v})$ to be equal. Moreover, since $d \geq 2$, this identification does not identify any variables from the tuples $\mathbf{w}^{d-1}$ and $\mathbf{w}^d$. Let us define $f_0 := f|_{\mathrm{S}(\mathbf{w}^0) \cup \mathrm{S}(\mathbf{w}^d)}$. It is enough to show that f_0 can be extended to a mapping $h \in (\psi_1 \wedge \cdots \wedge \psi_d)^{\mathbb{A}}$.

Let O be the orbit of $f(\mathbf{u})$ under $\mathscr{G}$, and let P be the orbit of $f(\mathbf{v})$. Let $f_1 \in \psi_1^{\mathbb{A}}$ be an injective OP-mapping, and for every $j \in \{2, \dots, d\}$, let $f_j \in \psi_j^{\mathbb{A}}$ be an injective PP-mapping. Note that such f_j exists for every $j \in [d]$ since ψ_j is complete and since $\mathbf{Vert}(C)$ and $\mathbf{Vert}(D)$ are strongly connected components in $\mathcal{B}_{\psi_j}$.

Let τ be the signature of $\mathbb{B}$. Let $X := \mathrm{S}(\mathbf{w}^0) \cup \cdots \cup \mathrm{S}(\mathbf{w}^d)$, and let us define a τ-structure $\mathbb{X}$ on X as follows. Recall that the relations from τ correspond to the orbits of injective k-tuples under $\mathscr{G}$. For every relation $O \in \tau$, we define $O^{\mathbb{X}}$ to be the set of all tuples $\mathbf{w} \in I_k^X$ such that there exists $j \in \{0, \dots, d\}$ such that $\mathrm{S}(\mathbf{w}) \subseteq \mathrm{S}(\mathbf{w}^j) \cup \mathrm{S}(\mathbf{w}^{j+1})$ and $f_j(\mathbf{w}) \in O$; here and in the following, the addition $+$ on indices is understood modulo $d+1$. We will show that $\mathbb{X}$ has a homomorphism h to $\mathbb{B}$. If this is the case, it follows by the construction and by the k-homogeneity of $\mathscr{G}$ that $h \in (\psi_1 \wedge \cdots \wedge \psi_d)^{\mathbb{A}}$. Moreover, we can assume that $h|_{\mathrm{S}(\mathbf{w}^0) \cup \mathrm{S}(\mathbf{w}^d)} = f_0$ as desired.

Let now $Y \subseteq X$ be of size at most $d-2$. Then for cardinality reasons there exists $j \in \{1, \dots, d-1\}$ such that $Y \subseteq \mathrm{S}(\mathbf{w}^0) \cup \cdots \cup \mathrm{S}(\mathbf{w}^{j-1}) \cup \mathrm{S}(\mathbf{w}^{j+1}) \cup \cdots \cup \mathrm{S}(\mathbf{w}^d)$. Observe that f_{j+1} is a homomorphism from the induced substructure of $\mathbb{X}$ on $\mathrm{S}(\mathbf{w}^{j+1}) \cup \mathrm{S}(\mathbf{w}^{j+2})$ to $\mathbb{A}$. Since the orbits of $f_{j+1}(\mathbf{w}^{j+2})$ and $f_{j+2}(\mathbf{w}^{j+2})$ agree by definition, by composing f_{j+2} with an element of $\mathscr{G}$ we can assume that $f_{j+1}(\mathbf{w}^{j+2}) = f_{j+2}(\mathbf{w}^{j+2})$. We can proceed inductively and extend f_{j+1} to $\mathrm{S}(\mathbf{w}^0) \cup \cdots \cup \mathrm{S}(\mathbf{w}^{j-1}) \cup \mathrm{S}(\mathbf{w}^{j+1}) \cup \cdots \cup \mathrm{S}(\mathbf{w}^d)$ such that it is a homomorphism from the induced substructure of $\mathbb{X}$ on this set to $\mathbb{A}$. It follows that the substructure of $\mathbb{X}$ induced on Y maps homomorphically to $\mathbb{A}$. Finite duality of $\mathbb{B}$ yields that $\mathbb{X}$ has a homomorphism to $\mathbb{A}$ as desired and Claim 4.3.16 follows.

Assume without loss of generality that $1 \notin I_\rho$, and identify u_i with v_i for every $i \neq 1$. Note that this is possible by item (3) in Observation 4.3.11, and since ρ is easily seen to be complete. Set $W := \mathrm{S}(\mathbf{u}) \cup \mathrm{S}(\mathbf{v})$, and let ρ' be the formula arising from ρ by this identification. We will argue that ρ' is critical in $\mathbb{A}$ over $(C, D, \mathbf{u}, \mathbf{v})$. To this end, let us first show that ρ' is a $(C, \mathbf{u}, C, \mathbf{v})$-implication in $\mathbb{A}$. Observe that for every orbit $O \subseteq C$ under $\mathscr{G}$, $(\rho')^{\mathbb{A}}$ contains an injective OO-mapping. Indeed, there exists an injective $g \in A^W$ such that $g(\mathbf{u}) \in O$ and $g(\mathbf{v}) \in O$; this easily follows by the $(k-1)$-transitivity of $\mathscr{G}$ and by the fact that it has no k-algebraicity. Forgetting the identification of variables, we can understand g as an element of A^V, and Claim 4.3.16 yields that $g \in \rho^{\mathbb{A}}$, whence $g \in (\rho')^{\mathbb{A}}$. Now, it immediately follows that ρ' satisfies the items (1)-(3) and (5) from Definition 4.2.4. Moreover, the satisfaction of item (4) follows immediately from the fact that ρ is a $(C, \mathbf{u}, C, \mathbf{v})$-implication in $\mathbb{A}$.

It remains to verify the last three items of Definition 4.3.6. Observe similarly as above that for any orbit $O \subseteq D$ under $\mathscr{G}$, $(\rho')^{\mathbb{A}}$ contains an OO-mapping, which immediately yields that D is contained both in $\mathrm{proj}_{\mathbf{u}}((\rho')^{\mathbb{A}})$ and in $\mathrm{proj}_{\mathbf{v}}((\rho')^{\mathbb{A}})$, it also yields that there exists no $D' \subseteq A^k$ with $D' \cap D \subsetneq D$ and such that for every $f \in (\rho')^{\mathbb{A}}$, it holds that $f(\mathbf{u}) \in D$ implies $f(\mathbf{v}) \in D'$. Hence, ρ' is indeed critical in $\mathbb{A}$ over $(C, D, \mathbf{u}, \mathbf{v})$, contradicting Lemma 4.3.7. □

Corollary 4.1.2. *Let $k \geq 3$, let $\mathbb{B}$ be the canonical k-ary structure of a k-neoliberal permutation group $\mathscr{G}$, and suppose that $\mathbb{B}$ has finite duality. Then any first-order expansion of $\mathbb{B}$ with bounded strict width has relational width $(k, \max(k+1, b_{\mathbb{B}}))$.*

Proof. Let $\mathbb{A}$ be a first-order expansion of $\mathbb{B}$ with bounded strict width. By Proposition 4.2.3, it is enough to prove that $\mathrm{CSP}_{\mathrm{Inj}}(\mathbb{A})$ has relational width $(k, \max(k+1, b_{\mathbb{B}}))$. Theorem 4.1.1 yields that $\mathbb{A}$ is implicationally simple on injective instances and the result follows from Proposition 4.2.8. □

Bibliography

[1] Samson Abramsky, Anuj Dawar, and Pengming Wang. The pebbling comonad in finite model theory. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12. IEEE Computer Society, 2017. doi:10.1109/LICS.2017.8005129.

[2] Reza Akhtar and Alistair H. Lachlan. On countable homogeneous 3-hypergraphs. *Archive for Mathematical Logic*, 34:331–344, 1995. doi:10.1007/BF01387512.

[3] Albert Atserias, Andrei A. Bulatov, and Víctor Dalmau. On the power of k -consistency. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland, July 9-13, 2007, Proceedings*, volume 4596 of *Lecture Notes in Computer Science*, pages 279–290. Springer, 2007. doi:10.1007/978-3-540-73420-8_26.

[4] Albert Atserias, Andrei A. Bulatov, and Anuj Dawar. Affine systems of equations and counting infinitary logic. *Theor. Comput. Sci.*, 410(18):1666–1683, 2009. doi:10.1016/j.tcs.2008.12.049.

[5] László Babai. Graph isomorphism in quasipolynomial time [extended abstract]. In Daniel Wichs and Yishay Mansour, editors, *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 684–697. ACM, 2016. doi:10.1145/2897518.2897542.

[6] Libor Barto. The collapse of the bounded width hierarchy. *J. Log. Comput.*, 26(3):923–943, 2016. doi:10.1093/logcom/exu070.

[7] Libor Barto, Michael Kompatscher, Miroslav Olšák, Trung Van Pham, and Michael Pinsker. Equations in oligomorphic clones and the constraint satisfaction problem for ω-categorical structures. *Journal of Mathematical Logic*, 19(02):1950010, 2019. doi:10.1142/S0219061319500107.

[8] Libor Barto, Michael Kompatscher, Miroslav Olšák, Trung Van Pham, and Michael Pinsker. The equivalence of two dichotomy conjectures for infinite domain constraint satisfaction problems. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science – LICS'17*, 2017. doi:10.1109/lics.2017.8005128.

[9] Libor Barto and Marcin Kozik. Congruence distributivity implies bounded width. *SIAM J. Comput.*, 39(4):1531–1542, 2009. doi:10.1137/080743238.

[10] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem. *Log. Methods Comput. Sci.*, 8(1), 2012. doi:10.2168/LMCS-8(1:7)2012.

[11] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *J. ACM*, 61(1):3:1–3:19, 2014. doi:10.1145/2556646.

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

[13] Libor Barto and Michael Pinsker. The algebraic dichotomy conjecture for infinite domain constraint satisfaction problems. In *Proceedings of the 31th Annual IEEE Symposium on Logic in Computer Science – LICS'16*, pages 615–622, 2016. Preprint arXiv:1602.04353. doi:10.1145/2933575.2934544.

[14] Libor Barto and Michael Pinsker. Topology is irrelevant. *SIAM Journal on Computing*, 49(2):365–393, 2020. doi:10.1137/18M1216213.

[15] Meghyn Bienvenu, Balder ten Cate, Carsten Lutz, and Frank Wolter. Ontology-based data access: A study through disjunctive datalog, CSP, and MMSNP. *ACM Trans. Database Syst.*, 39(4):33:1–33:44, 2014. doi:10.1145/2661643.

[16] Manuel Bodirsky. Cores of countably categorical structures. *Logical Methods in Computer Science*, 3(1):1–16, 2007. doi:10.2168/LMCS-3(1:2)2007.

[17] Manuel Bodirsky. Ramsey classes: Examples and constructions. In *Surveys in Combinatorics. London Mathematical Society Lecture Note Series 424*. Cambridge University Press, 2015. Invited survey article for the British Combinatorial Conference; preprint arXiv:1502.05146. doi:10.1017/CBO9781316106853.002.

[18] Manuel Bodirsky and Bertalan Bodor. Canonical polymorphisms of Ramsey structures and the unique interpolation property. In *Proceedings of the 36th Annual ACM/IEEE Symposium on Logic in Computer Science – LICS'21*, 2021. doi:10.1109/LICS52264.2021.9470683.

[19] Manuel Bodirsky and Víctor Dalmau. Datalog and constraint satisfaction with infinite templates. *J. Comput. Syst. Sci.*, 79(1):79–100, 2013. A conference version appeared in the Proceedings of the 34th Symposium on Theoretical Aspects of Computer Science (STACS 2006), pages 646–659. doi:10.1016/j.jcss.2012.05.012.

[20] Manuel Bodirsky and Martin Grohe. Non-dichotomies in constraint satisfaction complexity. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *Proceedings of the 35th International Colloquium on Automata, Languages and Programming, ICALP 2008*, pages 184–196, Berlin, Heidelberg, 2008. Springer-Verlag. doi:10.1007/978-3-540-70583-3_16.

[21] Manuel Bodirsky, Peter Jonsson, and Trung Van Pham. The complexity of phylogeny constraint satisfaction problems. *ACM Transactions on Computational Logic*, 18(3),

2017. An extended abstract appeared in the conference STACS 2016. doi:https://doi.org/10.1145/3105907.

[22] Manuel Bodirsky and Jan Kára. The complexity of equality constraint languages. *Theory of Computing Systems*, 3(2):136–158, 2008. A conference version appeared in the proceedings of Computer Science Russia (CSR'06). doi:10.1007/s00224-007-9083-9.

[23] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2), 2010. An extended abstract appeared in the Proceedings of the Symposium on Theory of Computing (STOC). doi:10.1145/1667053.1667058.

[24] Manuel Bodirsky and Simon Knäuer. Network satisfaction for symmetric relation algebras with a flexible atom. *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, 35(7):6218–6226, 2021. doi:10.1609/aaai.v35i7.16773.

[25] Manuel Bodirsky, Florent Madelaine, and Antoine Mottet. A proof of the algebraic tractability conjecture for monotone monadic SNP. *SIAM Journal on Computing*, 50(4):1359–1409, 2021. arXiv:https://doi.org/10.1137/19M128466X, doi:10.1137/19M128466X.

[26] Manuel Bodirsky, Florent R. Madelaine, and Antoine Mottet. A universal-algebraic proof of the complexity dichotomy for monotone monadic SNP. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 105–114. ACM, 2018. doi:10.1145/3209108.3209156.

[27] Manuel Bodirsky, Barnaby Martin, and Antoine Mottet. Discrete temporal constraint satisfaction problems. *J. ACM*, 65(2):9:1–9:41, 2018. doi:10.1145/3154832.

[28] Manuel Bodirsky, Barnaby Martin, Michael Pinsker, and András Pongrácz. Constraint satisfaction problems for reducts of homogeneous graphs. *SIAM J. Comput.*, 48(4):1224–1264, 2019. A conference version appeared in the Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, pages 119:1-119:14. doi:10.1137/16M1082974.

[29] Manuel Bodirsky and Antoine Mottet. Reducts of finitely bounded homogeneous structures, and lifting tractability from finite-domain constraint satisfaction. In *Proceedings of the 31th Annual ACM/IEEE Symposium on Logic in Computer Science – LICS'16*, 2016. doi:10.1145/2933575.2934515.

[30] Manuel Bodirsky and Antoine Mottet. A dichotomy for first-order reducts of unary structures. *Logical Methods in Computer Science*, 14(2), 2018.

[31] Manuel Bodirsky, Antoine Mottet, Miroslav Olšák, Jakub Opršal, Michael Pinsker, and Ross Willard. Topology is relevant (in a dichotomy conjecture for infinite-domain constraint satisfaction problems). In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, pages 1–12. IEEE, 2019. doi:10.1109/LICS.2019.8785883.

[32] Manuel Bodirsky, Antoine Mottet, Miroslav Olšák, Jakub Opršal, Michael Pinsker, and Ross Willard. ω-categorical structures avoiding height 1 identities. *Transactions of the AMS*, 374:327–350, 2021. doi:10.1090/tran/8179.

[33] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006. A conference version appeared in the proceedings of Computer Science Logic (CSL 2003). doi:10.1093/logcom/exi083.

[34] Manuel Bodirsky, Wied Pakusa, and Jakub Rydval. Temporal constraint satisfaction problems in fixed-point logic. In Holger Hermanns, Lijun Zhang, Naoki Kobayashi, and Dale Miller, editors, *LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Saarbrücken, Germany, July 8-11, 2020*, pages 237–251. ACM, 2020. doi:10.1145/3373718.3394750.

[35] Manuel Bodirsky and Michael Pinsker. Schaefer's theorem for graphs. *J. ACM*, 62(3):19:1–19:52, 2015. A conference version appeared in the Proceedings of STOC 2011, pages 655-664. doi:10.1145/2764899.

[36] Manuel Bodirsky and Michael Pinsker. Topological Birkhoff. *Transactions of the American Mathematical Society*, 367:2527–2549, 2015. doi:10.1016/j.jpaa.2017.06.016.

[37] Manuel Bodirsky and Michael Pinsker. Canonical functions: a proof via topological dynamics. *Homogeneous Structures, A Workshop in Honour of Norbert Sauer's 70th Birthday, Contributions to Discrete Mathematics*, 16(2):36–45, 2021.

[38] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Projective clone homomorphisms. *Journal of Symbolic Logic*, 86(1):148–161, 2021. doi:10.1017/jsl.2019.23.

[39] Manuel Bodirsky, Michael Pinsker, and Todor Tsankov. Decidability of definability. *Journal of Symbolic Logic*, 78(4):1036–1054, 2013. A conference version appeared in the Proceedings of the 26th Annual ACM/IEEE Symposium on Logic in Computer Science – LICS'11.

[40] Manuel Bodirsky and Michał Wrona. Equivalence constraint satisfaction problems. In *Computer Science Logic (CSL'12) - 26th International Workshop/21st Annual Conference of the EACSL, CSL 2012, September 3-6, 2012, Fontainebleau, France*, pages 122–136, 2012. doi:10.4230/LIPIcs.CSL.2012.122.

[41] Bertalan Bodor. CSP dichotomy for ω-categorical monadically stable structures. PhD dissertation, Institute of Algebra, Technische Universität Dresden, 2021.

[42] Pierre Bourhis and Carsten Lutz. Containment in monadic disjunctive datalog, mmsnp, and expressive description logics. In Chitta Baral, James P. Delgrande, and Frank Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016*, pages 207–216. AAAI Press, 2016. URL: http://www.aaai.org/ocs/index.php/KR/KR16/paper/view/12847.

[43] Hans-Jürgen Bückert and Bernhard Nebel. Reasoning about temporal relations: a maximal tractable subclass of allen's interval algebra. *J. ACM*, 42(1):43–66, 1995. doi:10.1145/200836.200848.

[44] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005. doi:10.1137/S0097539700376676.

[45] Andrei A. Bulatov. Bounded relational width, 2009. URL: https://www2.cs.sfu.ca/~abulatov/papers/relwidth.pdf.

[46] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 319–330. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.37.

[47] Andrei A. Bulatov, Andrei A. Krokhin, and Benoît Larose. Dualities for constraint satisfaction problems. In Nadia Creignou, Phokion G. Kolaitis, and Heribert Vollmer, editors, *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar]*, volume 5250 of *Lecture Notes in Computer Science*, pages 93–124. Springer, 2008. doi:10.1007/978-3-540-92800-3_5.

[48] Hubie Chen and Benoît Larose. Asking the metaquestions in constraint tractability. *ACM Trans. Comput. Theory*, 9(3):11:1–11:27, 2017. doi:10.1145/3134757.

[49] Jan Hubička and Jaroslav Nešetřil. All those ramsey classes (ramsey classes with closures and forbidden homomorphisms). *Advances in Mathematics*, 356:106791, 2019. doi:https://doi.org/10.1016/j.aim.2019.106791.

[50] Víctor Dalmau. There are no pure relational width 2 constraint satisfaction problems. *Information Processing Letters*, 109(4):213–218, 2009. URL: https://www.sciencedirect.com/science/article/pii/S0020019008003074, doi:https://doi.org/10.1016/j.ipl.2008.10.005.

[51] Víctor Dalmau and Justin Pearson. Closure functions and width 1 problems. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*, pages 159–173, 1999. doi:10.1007/978-3-540-48085-3_12.

[52] Roland Fraïssé. Une hypothèse sur l'extension des relations finies et sa vérification dans certaines classes particulières (deuxième partie). *Synthese*, 16(1):34–46, 1966. URL: http://www.jstor.org/stable/20114493.

[53] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998. doi:10.1137/S0097539794266766.

[54] Cristina Feier, Antti Kuusisto, and Carsten Lutz. Rewritability in monadic disjunctive datalog, mmsnp, and expressive description logics. *Log. Methods Comput. Sci.*, 15(2),

2019. A conference version appeared in the Proceedings of the 20th International Conference on Database Theory (ICDT17) 2017, pages 1:1–1:17. doi:10.23638/LMCS-15(2:15)2019.

[55] Pierre Gillibert, Julius Jonušas, Michael Kompatscher, Antoine Mottet, and Michael Pinsker. Hrushovski's encoding and ω-categorical CSP monsters. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 131:1–131:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICALP.2020.131.

[56] Pierre Gillibert, Julius Jonušas, Michael Kompatscher, Antoine Mottet, and Michael Pinsker. When symmetries are not enough: a hierarchy of hard constraint satisfaction problems. *SIAM Journal on Computing*, 51(2):175–213, 2022.

[57] Martin Grohe. Fixed-point definability and polynomial time on graphs with excluded minors. *J. ACM*, 59(5):27:1–27:64, 2012. doi:10.1145/2371656.2371662.

[58] Wilfrid Hodges. *Model Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993. doi:10.1017/CBO9780511551574.

[59] Peter Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200(1):185–204, 1998. URL: https://www.sciencedirect.com/science/article/pii/S0304397597002302, doi:https://doi.org/10.1016/S0304-3975(97)00230-2.

[60] Alexander Kechris, Vladimir Pestov, and Stevo Todorčević. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geometric and Functional Analysis*, 15(1):106–189, 2005. doi:10.1007/s00039-005-0503-1.

[61] Karim Khanaki. Dividing lines in unstable theories and subclasses of baire 1 functions. *Archive for Mathematical Logic*, 61(7-8):977–993, 2022. doi:10.1007/s00153-022-00816-8.

[62] Phokion G. Kolaitis and Moshe Y. Vardi. Conjunctive-query containment and constraint satisfaction. *J. Comput. Syst. Sci.*, 61(2):302–332, 2000. A conference version appeared in the Proceedings of Symposium on Principles of Database Systems (PODS) 1998, pages 205-213. doi:10.1006/jcss.2000.1713.

[63] Michael Kompatscher and Trung Van Pham. A complexity dichotomy for poset constraint satisfaction. *IfCoLog Journal of Logics and their Applications (FLAP)*, 5(8):1663–1696, 2018. A conference version appeared in the Proceedings of the 34th Symposium on Theoretical Aspects of Computer Science (STACS 2017), pages 47:1–47:12. URL: https://www.collegepublications.co.uk/downloads/ifcolog00028.pdf, doi:10.4230/LIPIcs.STACS.2017.47.

[64] Marcin Kozik, Andrei Krokhin, Matt Valeriote, and Ross Willard. Characterizations of several Maltsev conditions. *Algebra universalis*, 73(3-4):205–224, 2015.

[65] Richard Emil Ladner. On the structure of polynomial time reducibility. *J. ACM*, 22(1):155–171, jan 1975. doi:10.1145/321864.321877.

[66] Benoit Larose and László Zádori. Bounded width problems and algebras. *Algebra Universalis*, 56(3-4):439–466, 2007. doi:10.1007/s00012-007-2012-6.

[67] Miklós Maróti and Ralph McKenzie. Existence theorems for weakly symmetric operations. *Algebra Universalis*, 59(3):463–489, 2008. doi:10.1007/s00012-008-2122-9.

[68] Antoine Mottet and Michael Pinsker. Cores over Ramsey structures. *The Journal of Symbolic Logic*, 86(1):352–361, 2021. doi:10.1017/jsl.2021.6.

[69] Antoine Mottet and Michael Pinsker. Smooth approximations and CSPs over finitely bounded homogeneous structures. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science – LICS'22*, 2022. doi:10.1145/3531130.3533353.

[70] Antoine Mottet, Michael Pinsker, and Tomáš Nagy. An order out of nowhere: a new algorithm for infinite-domain CSPs. arXiv.2301.12977, 2023. arXiv:2301.12977, doi:10.48550/arXiv.2301.12977.

[71] Antoine Mottet, Tomáš Nagy, Michael Pinsker, and Michał Wrona. Smooth approximations and relational width collapses. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 138:1–138:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.138.

[72] Antoine Mottet, Tomáš Nagy, Michael Pinsker, and Michał Wrona. When symmetries are enough: collapsing the bounded width hierarchy for infinite-domain CSPs. arxiv:2102.07531, 2022. URL: https://doi.org/10.48550/arXiv.2102.07531, doi:10.48550/ARXIV.2102.07531.

[73] Michael Pinsker. Current challenges in infinite-domain constraint satisfaction: Dilemmas of the infinite sheep. In *2022 IEEE 52nd International Symposium on Multiple-Valued Logic (ISMVL)*, pages 80–87, Los Alamitos, CA, USA, 2022. IEEE Computer Society. URL: https://doi.ieeecomputersociety.org/10.1109/ISMVL52857.2022.00019, doi:10.1109/ISMVL52857.2022.00019.

[74] Michael Pinsker, Pierre Gillibert, and Julius Jonušas. Pseudo-loop conditions. *Bulletin of the London Mathematical Society*, 51(5):917–936, 2019. doi:10.1112/blms.12286.

[75] Emil L. Post. The two-valued iterative systems of mathematical logic. *Annals of Mathematics Studies*, 5, 1941.

[76] Jaroslav Nešetřil and Vojtěch Rödl. Ramsey classes of set systems. *Journal of Combinatorial Theory, Series A*, 34(2):183–201, 1983.

[77] Mark H. Siggers. A strong mal'cev condition for locally finite varieties omitting the unary type. *Algebra universalis*, 64:15–20, 2010. doi:10.1007/s00012-010-0082-3.

[78] Ágnes Szendrei. Clones in universal algebra. In *Séminaire de Mathématiques Supérieures*. Les Presses de l'Université de Montréal, 1986.

[79] Simon Thomas. Reducts of random hypergraphs. *Annals of Pure and Applied Logic*, 80(2):165–193, 1996. doi:10.1016/0168-0072(95)00061-5.

[80] Matthew A. Valeriote. A subalgebra intersection property for congruence distributive varieties. *Canadian Journal of Mathematics*, 61(2):451–464, 2009. doi:10.4153/CJM-2009-023-2.

[81] Michał Wrona. On the relational width of first-order expansions of finitely bounded homogeneous binary cores with bounded strict width. In Holger Hermanns, Lijun Zhang, Naoki Kobayashi, and Dale Miller, editors, *LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Saarbrücken, Germany, July 8-11, 2020*, pages 958–971. ACM, 2020. doi:10.1145/3373718.3394781.

[82] Michał Wrona. Relational width of first-order expansions of homogeneous graphs with bounded strict width. In Christophe Paul and Markus Bläser, editors, *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France*, volume 154 of *LIPIcs*, pages 39:1–39:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.STACS.2020.39.

[83] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 331–342. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.38.

[84] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *Journal of the ACM*, 67(5):30:1–30:78, 2020. doi:10.1145/3402029.

[85] Dmitriy Zhuk. Strong subalgebras and the constraint satisfaction problem. *Journal of Multiple-Valued Logic and Soft Computing*, 36(4-5):455–504, 2021.