Graph Homomorphisms and Universal Algebra
Course Notes

Manuel Bodirsky,
Institut für Algebra, TU Dresden,
Manuel.Bodirsky@tu-dresden.de

June 23, 2025

Disclaimer: these are course notes in draft state, and they probably contain many mistakes; please report them to *manuel.bodirsky@tu-dresden.de*. None of the material is original. I have tried to provide references to the primary sources as much as possible, but might have failed at some places, in which case I am thankful for correcting emails as well.

Contents

Prerequisites. This course is designed for students of mathematics or computer science who already had an introduction to discrete structures. Almost all notions that we use in this text will be formally introduced, with the notable exception of basic concepts from complexity theory. For example, we do not formally introduce the class of polynomial-time computable functions and NP-completeness, even though these concepts are used when we discuss computational aspects of graph homomorphisms. Here we refer to an introduction to the theory of computation as for instance the book of Papadimitriou [84]; we have also added a short appendix with the basics (Appendix B). For the basics of first-order logic (including for example the concepts of *bound* and *free* variables, *quantifier-free formulas*, and *prenex normal form*), we refer to [24] or [22].

Acknowledgements. Many thanks to Catarina Carvalho, Lukas Juhrich, Mathison Knight, Sebastian Meyer, Andrew Moorhead, Michael Pinsker, Lukas Schneider, Moritz

Schöbi, Žaneta Semanišinová, Mark Siggers, the participants of the course in the Corona spring semester 2020 and the course in spring 2023 for their suggestions and bug reports. Thanks also to Brady Zarathustra for his lecture notes [94] and answering my questions.

Exercises. The text contains 212 exercises; some of them are graded using the Mandala scale.

1 Introduction

Around the first years of this millenium, several previously separate research communities realised that many of their central questions are essentially the same: in the late 80s and 90s, the graph homomorphism community intensively studied the computational complexity of the *H-colouring problem* [62]. Independently, the theoretical artificial intelligence community studied *constraint satisfaction problems* and their computational complexity (including the important Boolean CSPs which were classified by Schaefer [88] in 1978). In the late 90s, researchers realised that universal algebra provides the right tools for this task [38, 66]. The paper by Feder and Vardi, whose conference version appeared in 1993, is probably the most influential article in the area and has inspired generations of researchers [54, 55]. It formulates for the first time the *dichotomy conjecture*, which has been solved in 2017 by Bulatov [36] and by Zhuk [96]. It also links the topic with finite model theory. For example, it identifies Datalog from database theory as an important framework that captures many of the central consistency algorithms that have been used to solve CSPs. Feder and Vardi prove that a complexity dichotomy for finite-domain CSPs implies a complexity classification for the fragment of NP called MMSNP. They also prove that every finite-domain CSP is computationally equivalent to the H-colouring problem for some finite digraph, thus further substantiating the connection between graph homomorphisms and constraint satisfaction.

This course starts very concretely, in the setting of digraphs rather than the more general setting of relational structures, because digraphs are notationally simpler than general structures. Digraphs are ideal for black-board teaching, because they are easy to draw and it is easy to come up with interesting examples. After having introduced the basics of the theory of graph homomorphisms in Section 2, we present an algorithm of outstanding importance: the so-called *arc consistency procedure*. This algorithm is theoretically very well understood. Moreover, it is practically important, because of its low time and space requirements, because it is easy to implement, and because it is widely applicable. Numerous exercises that we formulate at the end of the subsections can be easily solved if the reader properly understands the underlying principles of arc consistency.

The arc consistency procedure can be generalised to the *k-consistency procedure*, which is more powerful, and still in P, but more demanding in time and space requirements. Theoretical results in later sections of this course show that if k-consistency solves the H-colouring problem, then so does the 3-consistency algorithm. This algorithm is sometimes referred to as the *(strong) path-consistency procedure* and it is the topic of Section 4. A full description of when this procedure solves the H-colouring problem has to wait until Section 15 of the course, but we do see some sufficient conditions that can be used to answer the question for many concrete digraphs H.

At some point, the restriction to digraphs becomes unnatural; we step to general relational structures in Section 5. This will be the appropriate setting for presenting the main tools for complexity classification, which are, in increasing strength, *primitive positive definitions*, *primitive positive interpretations*, and *primitive positive constructions*. Primitive positive constructions are part of the statement of a solution to the Feder-Vardi dichotomy conjecture: if a finite structure admits a primitive positive construction of K_3 (the complete graph with three vertices), then its CSP is NP-hard (a statement that we prove in Section 5); otherwise, its CSP is in P (this is the content of the result of Bulatov [36] and of Zhuk [96]).

All three of these concepts (pp definitions, pp interpretations, and pp constructions) can also be characterised universal-algebraically, in terms of polymorphisms. For primitive positive definability, this can be found in Section 6, where we also apply it to prove Schaefer's complexity dichotomy result for CSPs of two-element structures. The universal-algebraic theory that captures primitive positive interpretations is presented in Section 8, and the universal-algebraic theory for primitive positive constructions in Section 9.

In Section 10 we show the Hell-Nešetřil dichotomy for the H-colouring problem for finite undirected graphs H. From this result we obtain a universal-algebraic formulation of the complexity dichotomy for all finite structures in terms of a Siggers polymorphism (of arity six). A much more informative formulation of the complexity dichotomy uses *cyclic polymorphisms* in Section 14, which is substantially more difficult to prove. In particular, we need the fundamental theorem of abelian algebras from Section 12, and absorption theory, developed by Barto and Kozik [12] and presented in Section 13.

Concerning algorithms for CSPs, we treat the Bulatov-Dalmau algorithm for structures with a Maltsev polymorphism (Section 7), and in Section 15 the bounded width case (i.e., the CSPs that can be solved by Datalog). A complete algorithm that solves all tractable finite-domain CSPs is outside of the scope of this course.

2 The Basics

We mostly work with *finite* graphs; results that also hold for graphs with infinitely many vertices are only treated when it comes with no extra effort.

2.1 Graphs and Digraphs

The concepts in this section are probably known to most students, and can safely be skipped; the section fixes standard terminology and conventions from graph theory and can be consulted later if needed. Almost all definitions in this section have generalisations to *relational structures*, which will be introduced in Section 5; however, we focus exclusively on graphs in this section since they allow to reach the key ideas of the underlying theory with a minimum of notation.

A *directed graph* (also *digraph*) G is a pair (V, E) of a set $V = V(G)$ of *vertices* and a binary relation $E = E(G)$ on V. Note that in general we allow that V is an infinite set. For some definitions and results, we require that V is finite, in which case we say that G is a *finite* digraph. However, since this course deals exclusively with finite digraphs, we will omit this most of the time. The elements (u, v) of E are called the *arcs* (or *directed edges*) of G. Note that we allow *loops*, i.e., arcs of the form (u, u); a digraph without loops is called *loopless*. If $(u, v) \in E(G)$ is an arc, and $w \in V(G)$ is a vertex such that $w = u$ or $w = v$, then we say that (u, v) and w are *incident*.

An *(undirected) graph* is a pair (V, E) of a set $V = V(G)$ of *vertices* and a set $E = E(G)$ of *edges*, each of which is an unordered pair of (not necessarily distinct) elements of V. In other words, we explicitly allow *loops*, which are edges that link a vertex with itself. Undirected graphs can be viewed as *symmetric* digraphs: a digraph $G = (V, E)$ is called *symmetric* if $(u, v) \in E$ if and only if $(v, u) \in E$. For a digraph G, we say that G' is the *undirected graph of* G if G' is the undirected graph with $V(G') = V(G)$ and where $\{u, v\} \in E(G')$ if $(u, v) \in E(G)$ or $(v, u) \in E(G)$. For an undirected graph G, we say that G' is an *orientation of* G if G' is a directed graph such that $V(G') = V(G)$ and $E(G')$ contains for each edge $\{u, v\} \in E(G)$ either the arc (u, v) or the arc (v, u), and no other arcs.

For some notions for digraphs G one can just use the corresponding notions for undirected graphs applied to the undirected graph of G; conversely, most notions for directed graphs, specialised to symmetric graphs, translate to notions for the respective undirected graphs.

2.1.1 Examples of graphs, and corresponding notation

- The *complete graph* on n vertices $[n] := \{1, \dots, n\}$, denoted by K_n. This is an undirected graph on n vertices in which every vertex is joined with any other distinct vertex (so K_n contains no loops).
- The *cyclic graph* on n vertices, denoted by C_n; this is the undirected graph with the vertex set $\{0, \dots, n-1\}$ and edge set

$$\big\{\{0,1\}, \dots, \{n-2, n-1\}, \{n-1, 0\}\big\} = \big\{\{u, v\} : |u - v| = 1 \mod n\big\}.$$

- The *directed cycle* on n vertices, denoted by $\vec{C}_n$; this is the digraph with the vertex set $\{0, \dots, n-1\}$ and the arcs $\big\{(0,1), \dots, (n-2, n-1), (n-1, 0)\big\}$.
- The *path* with $n+1$ vertices and n edges, denoted by P_n; this is an undirected graph with the vertex set $\{0, \dots, n\}$ and edge set $\big\{\{0,1\}, \dots, \{n-1, n\}\big\}$.
- The *directed path* with $n+1$ vertices and n edges, denoted by $\vec{P}_n$; this is a digraph with the vertex set $\{0, \dots, n\}$ and edge set $\{(0,1), \dots, (n-1, n)\}$.
- A *tournament* is a directed loopless graph G with the property that for all distinct vertices x, y either (x, y) or (y, x) is an edge of G, but not both.
- The *transitive tournament* on $n \geq 2$ vertices, denoted by T_n; this is a directed graph with the vertex set $\{1, \dots, n\}$ where (i, j) is an arc if and only if $i < j$.

Let G and H be graphs (we define the following notions both for directed and for undirected graphs). Then $G \uplus H$ denotes the *disjoint union of G and H*, which is the graph with vertex set $V(G) \cup V(H)$ (we assume that the two vertex sets are disjoint; if they are not, we take a copy of H on a disjoint set of vertices and form the disjoint union of G with the copy of H) and edge set $E(G) \cup E(H)$. A graph G' is a *subgraph* of G if $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$. A graph G' is an *induced subgraph* of G if $V' = V(G') \subseteq V(G)$ and $(u, v) \in E(G')$ if and only if $(u, v) \in E(G)$ for all $u, v \in V'$. We also say that G' is *induced by V' in G*, and write $G[V']$ for G'. We write $G - u$ for $G[V(G) \setminus \{u\}]$, i.e., for the subgraph of G where the vertex u and all incident arcs are removed.

We call $|V(G)| + |E(G)|$ the *size* of a graph G. This quantity will be important when we analyse the efficiency of algorithms on graphs.

2.1.2 Paths and Cycles

We start with definitions for directed paths; the corresponding terminology is then also used for undirected graphs as explained in the beginning of this section.

A *path* P (from u_1 to u_k in G) is a sequence $(u_1, \dots, u_k)$ of vertices of G and a sequence $(e_1, \dots, e_{k-1})$ of edges of G such that $e_i = (u_i, u_{i+1})$ or $e_i = (u_{i+1}, u_i) \in E(G)$, for every $1 \leq i < k$. The vertex u_1 is called the *start vertex* and the vertex u_k is called the *terminal* *vertex* of P, and we say that P is a path *from* u_1 *to* u_k. Edges (u_i, u_{i+1}) are called *forward edges* and edges (u_{i+1}, u_i) are called *backward edges.* If all edges are forward edges then the path is called *directed.* If $u_1, \ldots, u_k$ are pairwise distinct then the path is called *simple.* We write $|P| := k - 1$ for the *length* of P (i.e., we count the number of edges of P). The *net length* of P is the difference between the number of forward and the number of backward edges. Hence, a path is directed if and only if its length equals its net length.

A sequence $(u_0, \ldots, u_{k-1})$ of vertices and a sequence of edges $(e_0, \ldots, e_{k-1})$ is called a *cycle* (of G) if $(u_0, \ldots, u_{k-1}, u_0)$ and $(e_0, \ldots, e_{k-1})$ form a path. If all the vertices of the cycle are pairwise distinct then the cycle is called *simple.* We write $|C| := k$ for the *length* of the cycle $C = (u_0, \ldots, u_{k-1})$. The *net length* of C is the net length of the corresponding path $(u_0, \ldots, u_{k-1}, u_0)$. The cycle C is called *directed* if the corresponding path is a directed path.

A digraph G is called *(weakly) connected* if there is a path in G from any vertex to any other vertex in G. Equivalently, G is connected if and only if it cannot be written as $H_1 \uplus H_2$ for digraphs H_1, H_2 with at least one vertex each. A *connected component* of G is a maximal (with respect to inclusion of the vertex sets) connected induced subgraph of G. A digraph G is called *strongly connected* if for all vertices $x, y \in V(G)$ there is a directed path from x to y in G. Two vertices $u, v \in V(G)$ are *at distance* k in G if the shortest path from u to v in G has length k.

Some particular notions for undirected graphs G. A *(simple) cycle* of G is a sequence $(v_1, \ldots, v_k)$ of $k \geq 3$ pairwise distinct vertices of G such that $\{v_1, v_k\} \in E(G)$ and $\{v_i, v_{i+1}\} \in E(G)$ for all $1 \leq i \leq k-1$. An undirected graph is called *acyclic* if it does not contain a cycle. A sequence $u_1, \ldots, u_k \in V(G)$ is called a *(simple) path* from u_1 to u_k in G if $\{u_i, u_{i+1}\} \in E(G)$ for all $1 \leq i < k$ and if all vertices $u_1, \ldots, u_k$ are pairwise distinct. We allow the case that $k = 1$, in which case the path consists of a single vertex and no edges. Two vertices $u, v \in G$ are *at distance* k in G if the shortest path in G from u to v has length k. We say that an undirected graph G is *connected* if for all vertices $u, v \in V(G)$ there is a path from u to v. The *connected components* of G are the maximal connected induced subgraphs of G. A *forest* is an undirected acyclic graph, a *tree* is a connected forest.

A *source* in a digraph is a vertex with no incoming edges, and a *sink* is a vertex with no outgoing edges.

2.2 Graph Homomorphisms

Let G and H be directed graphs. A *homomorphism* from G to H is a mapping $h\colon V(G) \to V(H)$ which *preserves the edges*, i.e., $(h(u), h(v)) \in E(H)$ whenever $(u, v) \in E(G)$. If such a homomorphism exists between G and H we say that G *homomorphically maps* to H, and write $G \to H$. Otherwise, we write $G \not\to H$. Two directed graphs G and H are

- *homomorphically equivalent* if $G \to H$ and $H \to G$; in this case, we also write $G \leftrightarrow H$.
- *homomorphically comparable* if $G \to H$ or $H \to G$; otherwise, we say that H and G are *homomorphically incomparable.*

A homomorphism from G to H is sometimes also called an *H-colouring* of G. This terminology originates from the observation that H-colourings generalise classical colourings in the sense that a graph is n-colourable if and only if it has a K_n-colouring. Graph n-colorability is not the only natural graph property that can be described in terms of homomorphisms:

- a digraph is called *balanced* (in some articles: *layered*) if it homomorphically maps to a directed path $\vec{P}_n$;
- a digraph is called *acyclic* if it homomorphically maps to a transitive tournament T_n.

The equivalence classes of finite digraphs with respect to homomorphic equivalence will be denoted by $\mathcal{D}$. Let $\leq$ be a binary relation defined on $\mathcal{D}$ as follows: we set $C_1 \leq C_2$ if there exists a digraph $H_1 \in C_1$ and a digraph $H_2 \in C_2$ such that $H_1 \to H_2$ (note that this definition does not depend on the choice of the representatives H_1 of C_1 and H_2 of C_2). If f is a homomorphism from H_1 to H_2, and g is a homomorphism from H_2 to H_3, then the composition $f \circ g$ of these functions is a homomorphism from H_1 to H_3, and therefore the relation $\leq$ is transitive. Since every graph H homomorphically maps to H, the order $\leq$ is also reflexive. Finally, $\leq$ is antisymmetric since its elements are equivalence classes of directed graphs with respect to homomorphic equivalence. Define $C_1 < C_2$ if $C_1 \leq C_2$ and $C_1 \neq C_2$. We call $(\mathcal{D}, \leq)$ the *homomorphism order* of finite digraphs.

The homomorphism order on digraphs turns out to be a *lattice* where every two elements have a supremum (also called *join*) and an infimum (also called *meet*; see Example 8.5). In the proof of this result, we need the notion of *direct products* of graphs. This notion of graph product[1] can be seen as a special case of the notion of direct product as it is used in model theory [65]. The class of all graphs with respect to homomorphisms forms an interesting category in the sense of category theory [62] where the product introduced above is the product in the sense of category theory, which is why this product is sometimes also called the *categorical* graph product.

Definition 2.1 (direct product)**.** Let H_1 and H_2 be two graphs. Then the *(direct-, cross-, categorical-) product* $H_1 \times H_2$ of H_1 and H_2 is the graph with vertex set $V(H_1) \times V(H_2)$; the pair $((u_1, u_2), (v_1, v_2))$ is in $E(H_1 \times H_2)$ if $(u_1, v_1) \in E(H_1)$ and $(u_2, v_2) \in E(H_2)$.

Note that the product is symmetric and associative in the sense that $H_1 \times H_2$ is isomorphic to $H_2 \times H_1$ and $H_1 \times (H_2 \times H_3)$ is isomorphic to $(H_1 \times H_2) \times H_3$, and we therefore do not specify the order of multiplication when multiplying more than two graphs. The *n-th power* H^n of a graph H is inductively defined as follows. H^1 is by definition H. If H^i is already defined, then H^{i+1} is $H^i \times H$.

Proposition 2.2. *The homomorphism order $(\mathcal{D}, \leq)$ is a* lattice*; i.e., for all $A_1, A_2 \in \mathcal{D}$*

- *there exists an element $A_1 \wedge A_2 \in \mathcal{D}$, the* meet *of A_1 and A_2, such that $(A_1 \wedge A_2) \leq A_1$ and $(A_1 \wedge A_2) \leq A_2$, and such that for every $U \in \mathcal{D}$ with $U \leq A_1$ and $U \leq A_2$ we have $U \leq A_1 \wedge A_2$;*
- *there exists an element $A_1 \vee A_2 \in \mathcal{D}$, the* join *of A_1 and A_2, such that $A_1 \leq (A_1 \vee A_2)$ and $A_2 \leq (A_1 \vee A_2)$, and such that for every $U \in \mathcal{D}$ with $A_1 \leq U$ and $A_2 \leq U$ we have $A_1 \vee A_2 \leq U$.*

Proof. Let $H_1 \in A_1$ and $H_2 \in A_2$. For the meet, the equivalence class of $H_1 \times H_2$ has the desired properties. For the join, the equivalence class of the disjoint union $H_1 \uplus H_2$ has the desired properties.[2] □

[1] Warning: there are several other notions of graph products that have been studied; see e.g. [62].
[2] For this reason, $H_1 \uplus H_2$ is sometimes called the *co-product* of H_1 and H_2.

With the seemingly simple definitions of graph homomorphisms and direct products we can already formulate very difficult combinatorial questions.

Conjecture 1 (Hedetniemi)**.** *Let G and H be finite graphs, and suppose that $G \times H \to K_n$. Then $G \to K_n$ or $H \to K_n$.*

The smallest $n \in \mathbb{N}$ such that $G \to K_n$ is also called the *chromatic number* of G, and denoted by $\chi(G)$. Clearly, $\chi(G \times H) \leq \min(\chi(G), \chi(H))$. Hedetniemi's conjecture can be rephrased as

$$\chi(G \times H) = \min(\chi(G), \chi(H))\,.$$

This conjecture is easy for $n = 1$ and $n = 2$ (Exercise 4), and has been solved for $n = 3$ by El Zahar and Sauer [52]. The conjecture has been refuted in 2019 by Yaroslav Shitov [90].

Clearly, $(\mathcal{D}, \leq)$ has infinite *ascending chains*, that is, sequences $E_1, E_2, \dots$ such that $E_i < E_{i+1}$ for all $i \in \mathbb{N}$. Take for instance the equivalence class of $\vec{P}_i$ for E_i. More interestingly, $(\mathcal{D}, \leq)$ also has infinite descending chains.

Proposition 2.3. *The lattice $(\mathcal{D}, \leq)$ contains infinite descending chains $E_1 > E_2 > \cdots$.*

Proof. For this we use the following directed graphs, called *zig-zags*, which are frequently used in the theory of graph homomorphisms. We may write an orientation of a path P as a sequence of 0's and 1's, where 0 represents a forward arc and 1 represents a backward arc. For two orientations of paths P and Q with the representation $P = p_0, \dots, p_n \in \{0,1\}^*$ and $Q = q_0, \dots, q_m \in \{0,1\}^*$, respectively, the *concatenation* $P \circ Q$ of P and Q is the oriented path represented by $p_0, \dots, p_n, q_0, \dots, q_m$. For $k \geq 1$, the *zig-zag of order k*, denoted by Z_k, is the orientation of a path represented by $11(01)^{k-1}1$. We recommend the reader to draw pictures of Z_k where forward arcs point up and backward arcs point down. Now, the equivalence classes of the graphs $Z_1, Z_2, \dots$ form an infinite descending chain. □

Proposition 2.4. *The lattice $(\mathcal{D}, \leq)$ contains infinite* antichains*, that is, sets of pairwise incomparable elements of $\mathcal{D}$ with respect to $\leq$.*

Proof. Again, it suffices to work with orientations of paths. For $k, l \geq 1$, the k, l *multi zig-zag*, denoted by $Z_{k,l}$, is the orientation of a path represented by $1\big(1(01)^k\big)^l 1$. Our infinite antichain now consists of the equivalence classes containing the graphs $Z_{k,k}$ for $k \geq 1$. □

A *strong homomorphism* from a digraph G to a digraph H is a function from $V(G)$ to $V(H)$ such that $(f(u), f(v)) \in E(H)$ if and only if $(u, v) \in E(G)$ for all $u, v \in V(G)$. An *isomorphism* between two directed graphs G and H is an bijective strong homomorphism from G to H. Note that a homomorphism $h \colon G \to H$ is an isomorphism if and only if it is bijective, and h^{-1} is a homomorphism from H to G. An *automorphism* of a digraph H is an isomorphism from H to H.

Exercises.

1. How many connected components do we have in $(P_3)^3$?

2. How many weakly and strongly connected components do we have in $(\vec{C}_3)^3$?

3. Let G and H be digraphs. Prove that $G \times H$ has a directed cycle if and only if both G and H have a directed cycle.

4. Prove the Hedetniemi conjecture for $n = 1$ and $n = 2$.

5. Show that the Hedetniemi conjecture is equivalent to each of the following two statements.

   - Let n be a positive integer. If for two graphs G and H we have $G \not\to K_n$ and $H \not\to K_n$, then $G \times H \not\to K_n$.
   - Let G and H be graphs with $\chi(G) = \chi(H) = m$. Then there exists a graph K with $\chi(K) = m$ such that $K \to G$ and $K \to H$.

6. Show that Hedetniemi's conjecture is false for directed graphs.
   Hint: there are counterexamples G, H with four vertices each.

7. Show that for every $k \in \mathbb{N}$, every pair of adjacent vertices of $(K_3)^k$ has exactly one common neighbour (that is, every edge lies in a unique subgraph of $(K_3)^k$ isomorphic to K_3).

8. Show that for every $k \in \mathbb{N}$, every pair of non-adjacent vertices in $(K_3)^k$ has at least two common neighbours.

9. Show that a digraph G homomorphically maps to $\vec{P}_1 = T_2$ if and only if $\vec{P}_2$ does not homomorphically map to G.

10. Construct an orientation of a tree that is not homomorphically equivalent to an orientation of a path.

11. Construct a balanced orientation of a cycle that is not homomorphically equivalent to an orientation of a path.

12. Show that for all digraphs G we have $G \to T_3$ if and only if $\vec{P}_3 \not\to G$.

13. Show that $G \to \vec{P}_n$, for some $n \geq 1$, if and only if any two paths in G that start and end in the same vertex have the same net length.

14. Show that $G \to \vec{C}_n$, for some $n \geq 1$, if and only if any two paths in G that start and end in the same vertex have the same net length modulo n.

15. Let a be an automorphism of K_n^k. Show that there are permutations $p_1, \dots, p_k$ of $\{1, \dots, n\}$ and a permutation q of $\{1, \dots, k\}$ such that

$$a(x_1, \dots, x_k) = (p_1(x_{q(1)}), \dots, p_k(x_{q(k)})).$$

2.3 The H-colouring Problem and Variants

When does a given digraph G homomorphically map to a digraph H? For every digraph H, this question defines a computational problem, called the *H-colouring problem*. The input of this problem consists of a finite digraph G, and the question is whether there exists a homomorphism from G to H.

There are many variants of this problem. In the *precoloured H-colouring problem*, the input consists of a finite digraph G, together with a mapping f from a subset of $V(G)$ to $V(H)$. The question is whether there exists an extension of f to all of $V(G)$ which is a homomorphism from G to H. In the *list H-colouring problem*, the input consists of a finite digraph G, together with a set $S_x \subseteq V(H)$ for every vertex $x \in V(G)$. The question is whether there exists a homomorphism h from G to H such that $h(x) \in S_x$ for all $x \in V(G)$. It is clear that the H-colouring problem reduces to the precoloured H-colouring problem (it is a special case: the partial map might have an empty domain), and that the precoloured H-colouring problem reduces to the list H-colouring problem (for vertices x in the domain of f, we set $S_x := \{f(x)\}$, and for vertices x outside the domain of f, we set $S_x := V(H)$).

The *constraint satisfaction problem* is a common generalisation of all these problems, and many more. It is defined not only for digraphs H, but more generally for *relational structures*. Relational structures are the generalisation of graphs that can have many relations of arbitrary arity instead of just one binary edge relation. The constraint satisfaction problem will be introduced formally in Section 5. If H is a digraph, then the constraint satisfaction problem for H, also denoted $\mathrm{CSP}(H)$, is precisely the H-colouring problem and we use the terminology interchangeably. Note that since graphs can be seen as a special case of digraphs, H-colouring is also defined for undirected graphs H. In this case we obtain essentially the same computational problem if we only allow undirected graphs in the input; this is made precise in Exercise 18.

For every finite graph H, the H-colouring problem is obviously in NP, because for every graph G it can be verified in polynomial time whether a given mapping from $V(G)$ to $V(H)$ is a homomorphism from G to H or not. Clearly, the same holds for the precoloured and the list H-colouring problem. We have also seen that the K_n-colouring problem is the classical n-colouring problem, which is NP-complete [56] for $n \geq 3$, and therefore, no polynomial-time algorithm exists for K_n-colouring with $n \geq 3$, unless P=NP. However, for many graphs and digraphs H (see Exercise 19 and 9) the H-colouring problem can be solved in polynomial time. Since the 1990s, researchers have studied the question: for which digraphs H can the H-colouring problem be solved in polynomial time? It has been conjectured by Feder and Vardi [55] that H-colouring is for any finite digraph H either NP-complete or can be solved in polynomial time. This is the so-called *dichotomy conjecture*, and it has been confirmed in 2017, independently by Bulatov [36] and by Zhuk [96].

Theorem 2.5 (Bulatov [36], Zhuk [96])**.** *Let H be a finite digraph. Then* $\mathrm{CSP}(H)$ *is in P or NP-complete.*

It was shown by Ladner that unless P=NP there are infinitely many complexity classes between P and NP; so the conjecture states that for H-colouring these intermediate complexities do not appear. Feder and Vardi also showed that if the dichotomy conjecture holds for H-colouring problems, then also the more general class of CSPs for finite relational structures exhibits a complexity dichotomy (see Section 5.2).

The list H-colouring problem, on the other hand, is quickly NP-hard, and therefore less difficult to classify. And indeed, a complete classification has been obtained by Bulatov [33] already in 2003. Alternative proofs can be found in [6, 35]. For finite *undirected* graphs, it is known since 1990 that the dichotomy conjecture holds [60]; this text provides two fundamentally different proofs of the following.

Theorem 2.6 (of [60])**.** *Let H be a finite undirected graph. If H homomorphically maps to K_2, or contains a loop, then H-colouring can be solved in polynomial time. Otherwise, H-colouring is NP-complete.*

The case that H homomorphically maps to K_2 will be the topic of Exercise 19. The entire proof of Theorem 2.6 can be found in Section 10, and an alternative proof in Section 14.5.

Exercices.

16. Let H be a finite directed graph. Find an algorithm that decides whether there is a strong homomorphism from a given graph G to the fixed graph H. The running time of the algorithm should be polynomial in the size of G (note that we consider $|V(H)|$ to be constant).

17. Let H be a finite digraph such that $\mathrm{CSP}(H)$ can be solved in polynomial time. Find a polynomial-time algorithm that constructs for a given finite digraph G a homomorphism to H, if such a homomorphism exists.

18. Let G and H be directed graphs, and suppose that H is symmetric. Show that $f\colon V(G) \to V(H)$ is a homomorphism from G to H if and only if f is a homomorphism from the undirected graph of G to the undirected graph of H.

19. Show that for any graph H that homomorphically maps to K_2 the constraint satisfaction problem for H can be solved in polynomial time.

20. Prove that $\mathrm{CSP}(T_3)$ can be solved in polynomial time.

21. Prove that $\mathrm{CSP}(\vec{C}_3)$ can be solved in polynomial time.

22. Let $\mathcal{N}$ be the set $\{Z_1, Z_2, Z_3, \dots\}$. Show that a digraph $G \to \vec{P}_2$ if and only if no digraph in $\mathcal{N}$ homomorphically maps to G.

23. Suppose that $\mathrm{CSP}(G)$ and $\mathrm{CSP}(H)$, for two digraphs G and H, can be solved in polynomial time. Show that $\mathrm{CSP}(G \times H)$ and $\mathrm{CSP}(G \uplus H)$ can be solved in polynomial time as well.

24. Suppose that G and H are homomorphically incomparable and suppose that $\mathrm{CSP}(G) \cup \mathrm{CSP}(H)$ can be solved in polynomial time. Show that $\mathrm{CSP}(G)$ and $\mathrm{CSP}(H)$ can be solved in polynomial time as well.

25. Suppose that G and H are homomorphically incomparable and connected, and suppose that $\mathrm{CSP}(G \uplus H)$ can be solved in polynomial time. Show that $\mathrm{CSP}(G)$ and $\mathrm{CSP}(H)$ can be solved in polynomial time as well.

26. Show that the assumption in the previous exercise that G and H are connected is necessary. Specifically, find digraphs G and H such that $\mathrm{CSP}(G \uplus H)$ can be solved in polynomial time, but $\mathrm{CSP}(G)$ and $\mathrm{CSP}(H)$ are NP-hard.

27. Find digraphs G and H such that $\mathrm{CSP}(G \times H)$ can be solved in polynomial time, but $\mathrm{CSP}(G)$ and $\mathrm{CSP}(H)$ are NP-hard, or show that there are no such digraphs (unless $\mathrm{P} = \mathrm{NP}$).

2.4 Cores

An *endomorphism* of a digraph H is a homomorphism from H to H. A finite digraph H is called a *core* if every endomorphism of H is an automorphism. A graph G is called *a core of* H if H is homomorphically equivalent to G and G is a core.

Proposition 2.7. *Every finite digraph H has a core, which is unique up to isomorphism, and which is isomorphic to an induced subgraph of H.*

Proof. Any finite digraph H has a core, since we can select an endomorphism e of H such that the image of e has smallest cardinality; the subgraph of H induced by $e(V(H))$ is a core of H. Let G_1 and G_2 be cores of H, and $f_1 \colon H \to G_1$, $g_1 \colon G_1 \to H$, $f_2 \colon H \to G_2$, and $g_2 \colon G_2 \to H$ be homomorphisms. Let $e_1 := f_2 \circ g_1$ and $e_2 := f_1 \circ g_2$. See Figure 1.

We claim that e_1 is the desired isomorphism. Suppose for contradiction that e_1 is not injective, i.e., there are distinct x, y in $V(G_1)$ such that $e_1(x) = e_1(y)$. It follows that $e_2 \circ e_1$ cannot be injective, too. But $e_2 \circ e_1$ is an endomorphism of G_1, contradicting the assumption that G_1 is a core. Similarly, e_2 is an injective homomorphism from G_2 to G_1, and it follows that $|V(G_1)| = |V(G_2)|$ and both e_1 and e_2 are bijective.

Now, since $|V(G_1)|$ is finite, $e_2 \circ e_1 \circ \cdots \circ e_2 \circ e_1 = (e_2 \circ e_1)^n = \mathrm{id}$ for large enough n. Hence, $e_2 \circ e_1 \circ \cdots \circ e_2 = (e_1)^{-1}$, so the inverse of e_1 is a homomorphism, and hence an isomorphism between G_1 and G_2. □

Since a core G of a finite digraph H is unique up to isomorphism, we call G *the* core of H. We want to mention without proof that it is NP-complete to decide whether a given digraph H is not a core [61].

Cores can be characterised in many different ways; for some of them, see Exercise 30. There are examples of infinite digraphs that do not have a core in the sense defined above; see Exercise 32. Since a digraph H and its core have the same CSP, it suffices to study $\mathrm{CSP}(H)$ for core digraphs H only. Working with cores has advantages; one of them is shown in Proposition 2.9 below. In the proof of this proposition, we need a concept that we will use again in later sections.

Definition 2.8. Let H be a digraph and let $u, v \in V(H)$ be vertices of H. Then the digraph $H/\{u,v\}$ obtained from H by *contracting* u, v is defined to be the digraph with vertex set $V(H) \setminus \{u,v\} \cup \big\{\{u,v\}\big\}$ and the edge set obtained from $E(H)$ by replacing each edge in $E(H)$ of the form (x,u) or (x,v), for $x \in V(H)$, by the edge $(x, \{u,v\})$, and each edge in $E(H)$ of the form (u,x) or (v,x), for $x \in V(H)$, by the edge $(\{u,v\}, x)$.

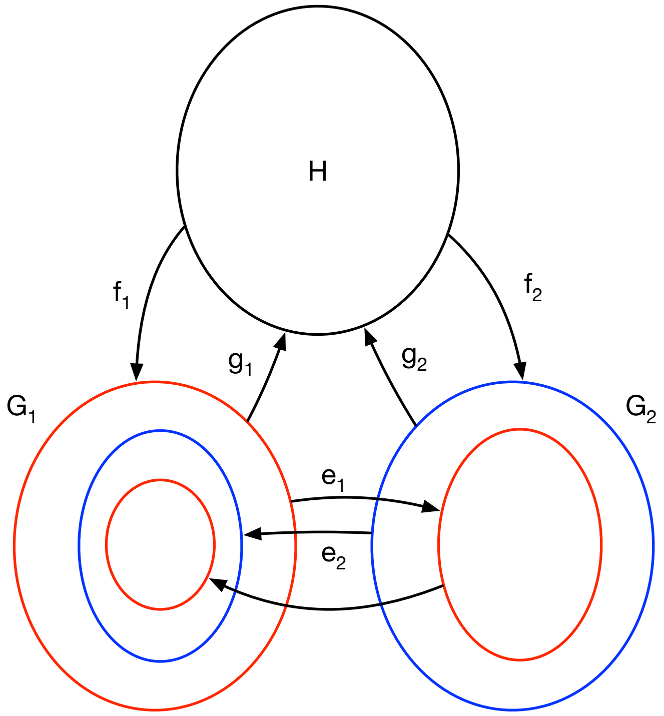

Figure 1: Illustration of the uniqueness proof for cores

.

Proposition 2.9. *Let H be a core. Then* CSP(H) *and precoloured* CSP(H) *are linear-time equivalent.*

Proof. The reduction from CSP(H) to precoloured CSP(H) is trivial, because an instance G of CSP(H) is equivalent to the instance (G, c) of precoloured CSP(H) where c is everywhere undefined.

We show the converse reduction by induction on the size of the image of the partial mapping c in instances of precoloured CSP(H). Let (G, c) be an instance of precoloured CSP(H) where c has an image of size $k \geq 1$. We show how to reduce the problem to one where the partial mapping has an image of size $k - 1$. If we compose all these reductions we finally obtain a reduction to CSP(H).

Let $x \in V(G)$ and $u \in V(H)$ be such that $c(x) = u$. We first contract all vertices y of G such that $c(y) = u$ with x. Then we create a copy of H, and attach the copy to G by contracting $x \in V(G)$ with $u \in V(H)$. Let G' be the resulting graph, and let c' be the partial map obtained from c by restricting it such that it is undefined on x, and then extending it so that $c(v) = v$ for all $v \in V(H)$, $v \neq u$, that appear in the image of c. Note that the image of c' has size $k - 1$. Note that the size of G' and the size of G only differ by a constant.

We claim that (G', c') has a solution if and only if (G, c) has a solution. If f is a homomorphism from G to H that extends c, we further extend f to the copy of H that is attached in G' by setting $f(v')$ to v if vertex v' is a copy of a vertex $v \in V(H)$. This extension of f clearly is a homomorphism from G' to H and extends c'.

Now, suppose that f' is a homomorphism from G' to H that extends c'. The restriction of f' to the vertices from the copy of H that is attached to x in G' is an endomorphism of H, and because H is a core, it is an automorphism α of H. Moreover, α fixes v for all $v \in V(H)$ in the image of c'. Let β be the inverse of α, i.e., let β be the automorphism of H such that $\beta(\alpha(v)) = v$ for all $v \in V(H)$. Let f be the mapping from $V(G)$ to $V(H)$ that maps vertices that were identified with x to $\beta(f'(x))$, and all other vertices $y \in V(G)$ to $\beta(f'(y))$. Clearly, f is a homomorphism from G to H. Moreover, f maps vertices $y \in V(G)$, $y \neq x$, where c is defined to $c(y)$, since the same is true for f' and for α. Moreover, because x in G' is identified to u in the copy of H, we have that $f(x) = \beta(f'(x)) = \beta(f'(u)) = u$, and therefore f is an extension of c. $\square$

Corollary 2.10. *If for every finite digraph H, the precoloured H-colouring problem is in P or NP-complete, then CSP(H) is in P or NP-complete for every finite digraph H as well.*

The following example shows that the assumption of Proposition 2.9 that H is a core is necessary (unless P = NP).

Example 2.11. Let H be the disjoint union of K_3 and a loop. Then CSP(H) is trivial and in P. The precoloured H-colouring problem, however, is NP-complete: we may prove this by a reduction from the NP-complete CSP(K_3) as follows. Clearly, this problem is already NP-complete if restricted to input graphs that are connected. Let G be a connected finite graph. Let c be a partial map sending one vertex of G to some element of K_3. Then G has a homomorphism to K_3 if and only if c can be extended to a homomorphism from G to H. To see this, let $f \colon G \to K_3$ be a homomorphism. Composing f with a permutation of $V(K_3)$ is also a homomorphism from G to K_3, and hence in particular to H. So we may obtain a homomorphism from G to H which extends c. Conversely, if f is a homomorphism from G to H which extends c then $f(V(G)) \subseteq V(K_3)$, since $f(x) \in V(K_3)$ and G is connected. $\triangle$

We have already seen in Exercise 17 that the computational problem to construct a homomorphism from G to H, for fixed H and given G, can be reduced in polynomial-time to the problem of deciding whether there exists a homomorphism from G to H. The intended solution of Exercise 17 requires in the worst-case $|V(G)|^2$ many executions of the decision procedure for CSP(H). Using the concept of cores and the precoloured CSP (and its equivalence to the CSP) we can give a faster method to construct homomorphisms.

Proposition 2.12. *If there is an algorithm that decides* CSP(H) *in time T, then there is an algorithm that constructs a homomorphism from a given digraph G to H (if such a homomorphism exists) which runs in time $O(|V(G)|T)$.*

Proof. Let C be the core of H; we may suppose that C is a subgraph of H. By Proposition 2.9, and since CSP(C) and CSP(H) are the same problem, there is an algorithm A for precoloured CSP(C) with a running time in $O(T)$.

To construct a homomorphism from a given finite digraph G to H, we first apply A to (G, c) for the everywhere undefined function c to decide whether there exists a homomorphism from G to C. If no, then there is also no homomorphism to H and there is nothing to be shown. If yes, we select some $x \in V(G)$, and extend c by defining $c(x) = u$ for some $u \in V(C)$. Then we use algorithm A to decide whether there is a homomorphism from G to C that extends c. If no, we try another vertex $u \in V(H)$. Clearly, for some u the algorithm must give the answer "yes". We proceed with the extension c where $c(x) = u$, and repeat the procedure with another vertex x from $V(G)$. At the end, c is defined for all vertices x of G, and c is a homomorphism from G to C. Clearly, since H and C are fixed, algorithm A is executed at most $O(|V(G)|)$ many times. $\square$

Exercises.

28. Prove that the core of a strongly connected digraph is strongly connected.

29. Show that $Z_{k,l}$ is a core for all $k, l \geq 2$.

30. Prove that for every finite digraph G the following is equivalent:
    - G is a core.
    - Every endomorphism of G is injective.
    - Every endomorphism of G is surjective.

31. Show that the three properties in the previous exercise are no longer equivalent if G is infinite.

32. Show that the infinite tournament $(\mathbb{Q}; <)$ has endomorphisms that are not automorphisms. Show that every digraph that is homomorphically equivalent to $(\mathbb{Q}; <)$ also has endomorphisms that are not automorphisms.

33. Prove that cores and products of digraphs without sources and sinks have no sources and sinks.

34. Let H be the core of G which we may assume to be a subgraph of G. Show that there exists a *retraction* from G to H, i.e., a homomorphism e from G to H such that $e(x) = x$ for all $x \in V(H)$.

35. A permutation group on a set V is called *transitive* if for all $a, b \in V$ there exists $g \in G$ such that $g(a) = b$. Show that if (V, E) is a graph with a transitive automorphism group, then the core of (V, E) also has a transitive automorphism group.

36. Show that the connected components of a core are cores that form an antichain in $(\mathcal{D}, \leq)$; conversely, the disjoint union of an antichain of cores is a core.

37. Prove that the core of a digraph with a transitive automorphism group is connected.

38. A permutation group G on as set X is called *primitive* if the only equivalence relations on X that are preserved by all permutations in G are the equality relation and the equivalence relation with only one equivalence class. Prove that the core of a digraph with a primitive automorphism group has a primitive automorphism group.

39. Determine the computational complexity of CSP(H) for

$$H := \big(\mathbb{Z}; \{(x, y) \ : \ |x - y| \in \{1, 2\}\}\big) \, .$$

2.5 Polymorphisms

Polymorphisms are a powerful tool for analysing the computational complexity of constraint satisfaction problems; as we will see, they are useful both for NP-hardness proofs and for proving the correctness of polynomial-time algorithms for CSPs. Polymorphisms can be seen as multi-dimensional variants of endomorphisms.

Definition 2.13. Let H be a digraph and $k \geq 1$. Then a *polymorphism of H of arity k* is a homomorphism from H^k to H.

In other words, a mapping $f\colon V(H)^k \to V(H)$ is a polymorphism of H if and only if $(f(u_1, \dots, u_k), f(v_1, \dots, v_k)) \in E(H)$ whenever $(u_1, v_1), \dots, (u_k, v_k)$ are arcs in $E(H)$. Note that any digraph H has all *projections* as polymorphisms, i.e., all mappings $\pi_i^k\colon V(H)^k \to V(H)$, for $i \leq k$ given by $\pi_i^k(x_1, \dots, x_k) = x_i$ for all $x_1, \dots, x_k \in V(H)$. The operation π_i^k is called the *i-th projection of arity k*.

Example 2.14. The operation $(x, y) \mapsto \min(x, y)$ is a polymorphism of the digraph $\vec{T}_n = (\{1, \dots, n\}; <)$. △

An operation $f\colon V(H)^k \to V(H)$ is called

- *idempotent* if $f(x, \dots, x) = x$ for all $x \in V(H)$.
- *conservative* if $f(x_1, \dots, x_k) \in \{x_1, \dots, x_k\}$ for all $x_1, \dots, x_k \in V(H)$.

A digraph H is called *projective* if every idempotent polymorphism is a projection. The following will be shown in Section 6.4.

Proposition 2.15. *For all $n \geq 3$, the graph K_n is projective.*

Exercises.

40. Show that if $f\colon H^k \to H$ is a polymorphism of a digraph H, then $\hat{f}(x) := f(x, \dots, x)$ is an endomorphism of H.
41. Show that if H is a finite core digraph with a *symmetric* binary polymorphism f, that is, $f(x, y) = f(y, x)$ for all $x, y \in V(H)$, then H also has an *idempotent* symmetric polymorphism.

3 The Arc-consistency Procedure

The *arc-consistency procedure* is one of the most fundamental and well-studied algorithms that are applied for CSPs. This procedure was first discovered for constraint satisfaction problems in artificial intelligence [80,83]; in the graph homomorphism literature, the algorithm is sometimes called the *consistency check algorithm*.

Let H be a finite digraph, and let G be an instance of CSP(H). The idea of the procedure is to maintain for each vertex of G a list of vertices of H, and each element in the list of x represents a candidate for an image of x under a homomorphism from G to H. The algorithm successively removes vertices from these lists; it only removes a vertex $u \in V(H)$ from the list for $x \in V(G)$, if there is no homomorphism from G to H that maps x to u. To detect vertices x, u such that u can be removed from the list for x, the algorithm uses two rules (in fact, one rule and a symmetric version of the same rule): if (x, y) is an edge in G, then

AC$_H(G)$
Input: a finite digraph G.
Data structure: a list $L(x) \subseteq V(H)$ for each vertex $x \in V(G)$.

Set $L(x) := V(H)$ for all $x \in V(G)$.
Do
    For each $(x, y) \in E(G)$:
        Remove u from $L(x)$ if there is no $v \in L(y)$ with $(u, v) \in E(H)$.
        Remove v from $L(y)$ if there is no $u \in L(x)$ with $(u, v) \in E(H)$.
        If $L(x)$ is empty for some vertex $x \in V(G)$ then **reject**
Loop until no list changes

Figure 2: The arc-consistency procedure for CSP(H).

- remove u from $L(x)$ if there is no $v \in L(y)$ with $(u, v) \in E(H)$;
- remove v from $L(y)$ if there is no $u \in L(x)$ with $(u, v) \in E(H)$.

If eventually we cannot remove any vertex from any list with these rules any more, the digraph G together with the lists for each vertex is called *arc-consistent*. The pseudo-code of the entire arc-consistency procedure is displayed in Figure 2.

Clearly, if the algorithm removes all vertices from one of the lists, then there is no homomorphism from G to H. It follows that if AC$_H$ rejects an instance of CSP(H), it has no solution. The converse implication does not hold in general. For instance, let H be K_2, and let G be K_3. In this case, AC$_H$ does not remove any vertex from any list, but obviously there is no homomorphism from K_3 to K_2.

However, there are digraphs H where the AC$_H$ is a complete decision procedure for CSP(H) in the sense that it rejects an instance G of CSP(H) if and only if G does not homomorphically map to H. In this case we say that AC *solves* CSP(H).

Remark 3.1. The running time of AC$_H$ is for any fixed digraph H polynomial in the size of G. Quite remarkably, it is also polynomial if H is part of the input, in which case we refer to the procedure as AC.

Implementation. In a naive implementation of the procedure, the inner loop of the algorithm would go over all edges of the digraph, in which case the running time of the algorithm is quadratic in the size of G. In the following we describe an implementation of the arc-consistency procedure, called AC-3, which is due to Mackworth [80], and has a worst-case running time that is linear in the size of G. Several other implementations of the arc-consistency procedure have been proposed in the Artificial Intelligence literature, aiming at reducing the costs of the algorithm in terms of the number of vertices of both G and H. But here we consider the size of H to be fixed, and therefore we do not follow this line of research. With AC-3, we rather present one of the simplest implementations of the arc-consistency procedure with a linear running time.

The idea of AC-3 is to maintain a *worklist*, which contains a list of arcs (x_0, x_1) of G that might help to remove a value from $L(x_0)$ or $L(x_1)$. Whenever we remove a value from a list

```
AC-3_H(G)
Input: a finite digraph G.
Data structure: a list L(x) of vertices of H for each x ∈ V(G).
            the worklist W: a list of arcs of G.

Subroutine Revise((x_0, x_1), i)
Input: an arc (x_0, x_1) ∈ E(G), an index i ∈ {0, 1}.
    change = false
    for each u_i in L(x_i)
        If there is no u_{1−i} ∈ L(x_{1−i}) such that (u_0, u_1) ∈ E(H) then
            remove u_i from L(x_i)
            change = true
        end if
    end for
    If change = true then
        If L(x_i) = ∅ then reject
    else
        For all arcs (z_0, z_1) ∈ E(G) with z_0 = x_i or z_1 = x_i add (z_0, z_1) to W
    end if

W := E(G)
Do
    remove an arc (x_0, x_1) from W
    Revise((x_0, x_1), 0)
    Revise((x_0, x_1), 1)
while W ≠ ∅
```

Figure 3: The AC-3 implementation of the arc-consistency procedure for CSP(H).

$L(x)$, we add all arcs that are in G incident to x. Note that then any arc in G might be added at most $2|V(H)|$ many times to the worklist, which is a constant in the size of G. Hence, the while loop of the implementation is iterated for at most a linear number of times. Altogether, the running time is linear in the size of G as well.

Arc-consistency for pruning search. Suppose that H is such that AC does not solve CSP(H). Even in this situation the arc-consistency procedure might be useful for *pruning the search space* in exhaustive approaches to solve CSP(H). In such an approach we might use the arc-consistency procedure as a subroutine as follows. Initially, we run AC_H on the input instance G. If it computes an empty list, we reject. Otherwise, we select some vertex $x \in V(G)$, and set $L(x)$ to $\{u\}$ for some $u \in L(x)$. Then we proceed recursively with the resulting lists. If AC_H now detects an empty list, we backtrack, but remove u from $L(x)$. Finally, if the algorithm does not detect an empty list at the first level of the recursion, we end up with singleton lists for each vertex $x \in V(G)$, which gives rise to a homomorphism from G to H.

3.1 The Power Graph

For which H does the Arc-Consistency procedure solve CSP(H)? In this section we present an elegant and effective characterisation of those finite digraphs H where AC solves CSP(H), found by Feder and Vardi [55].

Definition 3.2. For a digraph H, the *power graph* $P(H)$ is the digraph whose vertices are non-empty subsets of $V(H)$ and where two subsets U and V are joined by an arc if the following holds:

- for every vertex $u \in U$, there exists a vertex $v \in V$ such that $(u, v) \in E(H)$, and
- for every vertex $v \in V$, there exists a vertex $u \in U$ such that $(u, v) \in E(H)$.

The definition of the power graph resembles the arc-consistency algorithm, and indeed, we have the following lemma which describes the correspondence.

Lemma 3.3. AC_H *rejects* G *if and only if* $G \not\to P(H)$.

Proof. Suppose first that AC_H does *not* reject G. For $u \in V(G)$, let $L(u)$ be the list derived at the final stage of the algorithm. Then by definition of $E(P(H))$, the map $x \mapsto L(x)$ is a homomorphism from G to $P(H)$.

Conversely, suppose that $f : G \to P(H)$ is a homomorphism. We prove by induction over the execution of AC_H that for all $x \in V(G)$ the elements of $f(x)$ are never removed from $L(x)$. To see that, let $(a, b) \in E(G)$ be arbitrary. Then $(f(a), f(b)) \in E(P(H))$, and hence for every $u \in f(a)$ there exists a $v \in f(b)$ such that $(u, v) \in E(H)$. By inductive assumption, $v \in L(b)$, and hence u will not be removed from $L(a)$. This concludes the inductive step. $\square$

Theorem 3.4. *Let* H *be a finite digraph. Then* AC *solves* CSP(H) *if and only if* $P(H)$ *homomorphically maps to* H*.*

Proof. Suppose first that AC solves CSP(H). Apply AC_H to $P(H)$. Since $P(H) \to P(H)$, the previous lemma shows that AC_H does not reject $P(H)$. Hence, $P(H) \to H$ by assumption.

Conversely, suppose that $P(H) \to H$. If AC_H rejects a digraph G then $G \not\to H$. If AC_H does accept G, then the lemma asserts that $G \to P(H)$. Composing homomorphisms, we obtain that $G \to H$. $\square$

Observation 3.5. Let H be a core digraph. Note that if $P(H)$ homomorphically maps to H, then there also exists a homomorphism that maps $\{x\}$ to x for all $x \in V(H)$ (here we use the assumption that H is a core!). We claim that in this case the precoloured CSP for H can be solved by the modification of AC_H which starts with $L(x) := \{c(x)\}$ for all $x \in V(G)$ in the range of the precolouring function c, instead of $L(x) := V(H)$. This is a direct consequence of the proof of Theorem 3.4. If the modified version of AC_H solves the precoloured CSP for H, then the classical version of AC_H solves CSP(H). Hence, it follows that the following are equivalent:

- AC solves CSP(H);
- the above modification of AC_H solves the precoloured CSP for H;
- $P(H) \to H$.

Note that the condition given in Theorem 3.4 can be used to decide algorithmically whether AC solves CSP(H), because it suffices to test whether $P(H)$ homomorphically maps to H. Such problems about deciding properties of CSP(H) for given H are often called algorithmic *meta-problems*. A naive algorithm for the above test would be to first construct $P(H)$, and then to search non-deterministically for a homomorphism from $P(H)$ to H, which puts the meta-problem for solvability of CSP(H) by AC into the complexity class NExpTime (*Non-deterministic Exponential Time*). This can be improved.

Proposition 3.6. *There exists a deterministic exponential time algorithm that tests for a given finite core digraph H whether $P(H)$ homomorphically maps to H.*

Proof. We first explicitly construct $P(H)$, and then apply AC_H to $P(H)$. If AC_H rejects, then there is certainly no homomorphism from $P(H) \to H$ by the properties of AC_H, and we return 'false'. If AC_H accepts, then we cannot argue right away that $P(H)$ homomorphically maps to H, since we do not know yet whether AC_H is correct for CSP(H).

But here is the trick. What we do in this case is to pick an arbitrary $x \in V(P(H))$, and remove all but one value u from $L(x)$, and continue with the execution of AC_H. If AC_H then derives the empty list, we try the same with another value u' from $L(x)$. If we obtain failure for all values of $L(x)$, then clearly there is no homomorphism from $P(H)$ to H, and we return 'false'. Otherwise, if AC_H does not derive the empty list after removing all values but u from $L(x)$, we continue with another element y of $V(P(H))$, setting $L(y)$ to $\{v\}$ for some $v \in L(y)$. We repeat this procedure until at the end we have constructed a homomorphism from $P(H)$ to H. In this case we return 'true'.

If AC_H rejects for some $x \in V(P(H))$ when $L(x) = \{u\}$ for all possible $u \in V(H)$, then the adaptation of AC_H for the precoloured CSP would have given an incorrect answer for the previously selected variable (it said yes while it should have said no). By Observation 3.5, this means that $P(H)$ does *not* homomorphically map to H. Again, we return 'false'. □

The precise computational complexity to decide for a given digraph H whether $P(H) \to H$ is not known; we refer to [44] for related questions and results.

Question 1. *What is the computational complexity to decide for a given core digraph H whether $P(H) \to H$? Is this problem in P?*

3.2 Tree Duality

Another mathematical notion that is closely related to the arc-consistency procedure is *tree duality*. The idea of this concept is that when a digraph H has tree duality, then we can show that there is no homomorphism from a digraph G to H by exhibiting a *tree obstruction* in G. This is formalized in the following definition.

Definition 3.7. A digraph H has *tree duality* if there exists a (not necessarily finite) set $\mathcal{N}$ of orientations of finite trees such that for all digraphs G there is a homomorphism from G to H if and only if no digraph in $\mathcal{N}$ homomorphically maps to G.

We refer to the set $\mathcal{N}$ in Definition 3.7 as an *obstruction set* for CSP(H). Note that no $T \in \mathcal{N}$ homomorphically maps to H. The pair $(\mathcal{N}, H)$ is called a *duality pair*. We have already encountered such an obstruction set in Exercise 9, where $H = T_2$, and $\mathcal{N} = \{\vec{P}_2\}$. In other words, $(\{\vec{P}_2\}, T_2)$ is a duality pair. Other duality pairs are $(\{\vec{P}_3\}, T_3)$ (Exercise 12), and $(\{Z_1, Z_2, \dots\}, \vec{P}_2)$ (Exercise 22).

Theorem 3.8 is a surprising link between the completeness of the arc-consistency procedure, tree duality, and the power graph, and was discovered by Feder and Vardi [54] in the more general context of constraint satisfaction problems.

Theorem 3.8. *Let H be a finite digraph. Then the following are equivalent.*

1. *H has tree duality;*
2. *$P(H)$ homomorphically maps to H;*
3. AC *solves* CSP(H).
4. *If every orientation of a tree that homomorphically maps to G also homomorphically maps to H, then G homomorphically maps to H;*

Proof. The equivalence $2 \Leftrightarrow 3$ has been shown in the previous section. We show $3 \Rightarrow 1$, $1 \Rightarrow 4$, and $4 \Rightarrow 2$.

$3 \Rightarrow 1$: Suppose that AC solves CSP(H). We have to show that H has tree duality. Let $\mathcal{N}$ be the set of all orientations of trees that do not homomorphically map to H. We claim that if a digraph G does not homomorphically map to H, then there is $T \in \mathcal{N}$ that homomorphically maps to G.

By assumption, the arc-consistency procedure applied to G eventually derives the empty list for some vertex of G. We use the computation of the procedure to construct an orientation T of a tree, following the exposition in [73]. When deleting a vertex $u \in V(H)$ from the list of a vertex $x \in V(G)$, we define an orientation of a rooted tree $T_{x,u}$ with root $r_{x,u}$ such that

1. there is a homomorphism from $T_{x,u}$ to G mapping $r_{x,u}$ to x;
2. there is no homomorphism from $T_{x,u}$ to H mapping $r_{x,u}$ to u.

Assume that the vertex u is deleted from the list of x because we found an arc $(x, y) \in E(G)$ such that there is no arc $(u, v) \in E(H)$ with $v \in L(y)$; if it was deleted because of an arc $(y, x) \in E(H)$ the proof follows with the obvious changes.

If there is no $v \in V(H)$ such that $(u, v) \in E(H)$, then we define $T_{x,u}$ to be the tree that just contains an arc (p, q) with root $r_{x,u} = p$; clearly, $T_{x,u}$ satisfies property (1) and (2). Otherwise, for every arc $(u, v) \in E(H)$ the vertex v has already been removed from the list $L(y)$, and hence by induction $T_{y,v}$ having properties (1) and (2) is already defined. We then add a copy of $T_{y,v}$ to $T_{x,u}$, contract all the roots of all copies into one vertex q, and finally add an arc from the root vertex $r_{x,u}$ to q.

We verify that the resulting orientation of a tree $T_{x,u}$ satisfies (1) and (2). For every $v \in V(H)$ such that $(u, v) \in E(H)$, let f_v be the homomorphism from $T_{y,v}$ mapping $r_{y,v}$ to y, which exists due to (1). The common extension of all the maps f_v to $V(T_{x,u})$ that maps $r_{x,u}$ to x is a homomorphism from $T_{x,u}$ to G, and this shows that (1) holds for $T_{x,u}$. Suppose for contradiction that there exists a homomorphism h from $T_{x,u}$ to H that maps $r_{x,u}$ to u. Let $v = h(q)$; then h restricts to a homomorphism from $T_{y,v}$ to H, a contradiction. This shows that (2) holds for $T_{x,u}$. When the list $L(x)$ of some vertex $x \in V(G)$ becomes empty, we can construct an orientation of a tree T by contracting the roots of all $T_{x,u}$ into a vertex r. We then find a homomorphism from T to G by mapping r to x and extending the homomorphism independently on each $T_{x,u}$. But any homomorphism from T to H must map

r to some element $u \in V(H)$, and hence there is a homomorphism from $T_{x,u}$ to H that maps x to u, a contradiction.

$1 \Rightarrow 4$: If H has an obstruction set $\mathcal{N}$ consisting of orientations of trees, and if G does not homomorphically map to H, there exists an orientation of a tree $T \in \mathcal{N}$ that maps to G but not to H.

$4 \Rightarrow 2$: To show that $P(H)$ homomorphically maps to H, it suffices to prove that every orientation T of a tree that homomorphically maps to $P(H)$ also homomorphically maps to H. Let f be a homomorphism from T to $P(H)$, and let x be any vertex of T. We construct a sequence $f_0, \dots, f_n$, for $n = |V(T)|$, where f_i is a homomorphism from the subgraph of T induced by the vertices at distance at most i to x in T, and f_{i+1} is an extension of f_i for all $0 \leq i < n$. The mapping f_0 maps x to some vertex from $f(x)$. Suppose inductively that we have already defined f_i. Let y be a vertex at distance $i+1$ from x in T. Since T is an orientation of a tree, there is a unique $y' \in V(T)$ of distance i from x in T such that $(y, y') \in E(T)$ or $(y', y) \in E(T)$. Note that $u = f_i(y')$ is already defined. In case that $(y', y) \in E(T)$, there must be a vertex v in $f(y)$ such that $(u, v) \in E(H)$, since $(f(y'), f(y))$ must be an arc in $P(H)$, and by definition of $P(H)$. We then set $f_{i+1}(y) = v$. In case that $(y, y') \in E(T)$ we can proceed analogously. By construction, the mapping f_n is a homomorphism from T to H. □

3.3 Totally Symmetric Polymorphisms

There is also a characterisation of the power of the arc-consistency procedure which is based on polymorphisms.

Definition 3.9. A function $f \colon D^k \to D$ is called *totally symmetric* if

$$f(x_1, \dots, x_k) = f(y_1, \dots, y_k) \text{ whenever } \{x_1, \dots, x_k\} = \{y_1, \dots, y_k\}.$$

Example 3.10. The operation $(x_1, \dots, x_k) \mapsto \text{minimum}(x_1, \dots, x_k)$ is totally symmetric. △

Example 3.11. The majority operation $m \colon \{0,1\}^k \to \{0,1\}$ given by

$$m(x, x, y) = m(x, y, x) = m(y, x, x) = x$$

for all $x \in \{0,1\}$ is

- not totally symmetric because $0 = m(0,0,1) \neq m(0,1,1) = 1$;
- is *symmetric* in the sense that $m(x_1, x_2, x_3) = m(x_{\alpha(1)}, x_{\alpha(2)}, x_{\alpha(3)})$ for every permutation α of $\{1,2,3\}$. △

Theorem 3.12 (from [48])**.** *Let H be a finite digraph. Then the following are equivalent.*

1. *$P(H)$ homomorphically maps to H;*
2. *H has totally symmetric polymorphisms of all arities;*
3. *H has a totally symmetric polymorphism of arity $2|V(H)|$.*

Proof. 1. $\Rightarrow$ 2.: Suppose that g is a homomorphism from $P(H)$ to H, and let $k \in \mathbb{N}$ be arbitrary. Let f be defined by $f(x_1, \dots, x_k) = g(\{x_1, \dots, x_k\})$. If $(x_1, y_1), \dots, (x_k, y_k) \in E(H)$, then $\{x_1, \dots, x_k\}$ is adjacent to $\{y_1, \dots, y_k\}$ in $P(H)$, and hence $(f(x_1, \dots, x_k), f(y_1, \dots, y_k)) \in E(H)$. Therefore, f is a polymorphism of H, and it is clearly totally symmetric.

The implication 2. $\Rightarrow$ 3. is trivial. To prove that 3. $\Rightarrow$ 1., suppose that f is a totally symmetric polymorphism of arity $2|V(H)|$. Let $g\colon V(P(H)) \to V(H)$ be defined by

$$g(\{x_1, \dots, x_n\}) := f(x_1, \dots, x_{n-1}, x_n, x_n, \dots, x_n)$$

which is well-defined because f is totally symmetric. Let $(U, W) \in E(P(H))$, and let $x_1, \dots, x_p$ be an enumeration of the elements of U, and $y_1, \dots, y_q$ be an enumeration of the elements of W. The properties of $P(H)$ imply that there are $y'_1, \dots, y'_p \in W$ and $x'_1, \dots, x'_q \in U$ such that $(x_1, y'_1), \dots, (x_p, y'_p) \in E(H)$ and $(x'_1, y_1), \dots, (x'_q, y_q) \in E(H)$. Since f preserves E,

$$g(U) = g(\{x_1, \dots, x_p\}) = f(x_1, \dots, x_p, x'_1, \dots, x'_q, x_1, \dots, x_1)$$

is adjacent to

$$g(W) = g(\{y_1, \dots, y_q\}) = f(y'_1, \dots, y'_p, y_1, \dots, y_q, y'_1, \dots, y'_1)\ . \qquad \square$$

Given Theorem 3.12, it is natural to ask whether there exists a k so that the existence of a totally symmetric polymorphism of arity k implies totally symmetric polymorphisms of all arities. The following example shows that this is not the case.

Example 3.13. For every prime $p \geq 3$, the digraph $\vec{C}_p$ clearly does not have a totally symmetric polymorphism of arity p: if $f\colon \{0, \dots, p-1\}^p \to \{0, \dots, p-1\}$ is a totally symmetric operation, then $f(0, 1, \dots, p-1) = f(1, \dots, p-1, 0)$, and hence f does not preserve the edge relation. On the other hand, if $n < p$ then $\vec{C}_p$ has the totally symmetric polymorphism

$$f(x_1, \dots, x_n) := |S|^{-1} \sum_{x \in S} x \mod p$$

where $S = \{x_1, \dots, x_n\}$. (Note that $|S| < p$ and hence has a multiplicative inverse.) The operation is clearly totally symmetric; the verification that it preserves the edge relation of $\vec{C}_p$ is Exercise 52. $\triangle$

3.4 Semilattice Polymorphisms

Some digraphs have a single binary polymorphism that generates operations satisfying the conditions in the previous theorem. A binary operation $f\colon D^2 \to D$ is called *commutative* if it satisfies

$$f(x, y) = f(y, x) \text{ for all } x, y \in D.$$

It is called *associative* if it satisfies

$$f(x, f(y, z)) = f(f(x, y), z) \text{ for all } x, y, z \in D.$$

Definition 3.14. A binary operation is called a *semilattice operation* if it is associative, commutative, and idempotent.

Examples of semilattice operations are functions from $D^2 \to D$ defined as $(x, y) \mapsto \min(x, y)$; here the minimum is taken with respect to any fixed linear order of D.

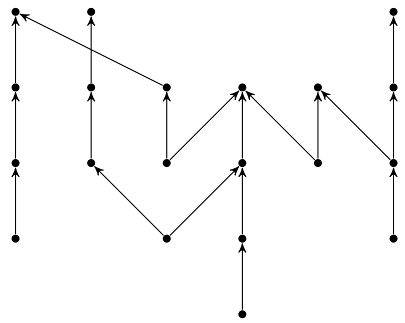

Figure 4: One of the smallest orientations of a tree H such that CSP(H) is NP-complete (assuming P $\neq$ NP; all orientations of trees with less vertices can be solved by path consistency [25]).

Theorem 3.15. *Let H be a finite digraph. Then $P(H) \to H$ if and only if H is homomorphically equivalent to a digraph with a semilattice polymorphism.*

Proof. Suppose first that $P(H) \to H$. Thus, H and $P(H)$ are homomorphically equivalent, and it suffices to show that $P(H)$ has a semilattice polymorphism. The mapping $(X, Y) \mapsto X \cup Y$ is clearly a semilattice operation; we claim that it preserves the edges of $P(H)$. Let (U, V) and (A, B) be edges in $P(H)$. Then for every $u \in U$ there is a $v \in V$ such that $(u, v) \in E(H)$, and for every $u \in A$ there is a $v \in B$ such that $(u, v) \in E(H)$. Hence, for every $u \in U \cup A$ there is a $v \in V \cup B$ such that $(u, v) \in E(H)$. Similarly, we can verify that for every $v \in V \cup B$ there is a $u \in U \cup A$ such that $(u, v) \in E(H)$. This proves the claim.

For the converse, suppose that H is homomorphically equivalent to a digraph G with a semilattice polymorphism f. Let h be the homomorphism from H to G. The operation $(x_1, \dots, x_n) \mapsto f(x_1, f(x_2, f(\dots, f(x_{n-1}, x_n) \dots)))$ is a totally symmetric polymorphism of G. Then Theorem 3.12 implies that $P(G) \to G$. The map $S \mapsto \{h(u) \mid u \in S\}$ is a homomorphism from $P(H)$ to $P(G)$. Therefore, $P(H) \to P(G) \to G \to H$, as desired. $\square$

By verifying the existence of semilattice polymorphisms for a concrete class of digraphs, we obtain the following consequence.

Corollary 3.16. AC *solves* CSP(H) *if H is an orientation of a path.*

Proof. Suppose that $1, \dots, n$ are the vertices of H such that either $(i, i+1)$ or $(i+1, i)$ is an arc in $E(H)$ for all $i < n$. It is straightforward to verify that the mapping $(x, y) \mapsto \min(x, y)$ is a polymorphism of H. The statement now follows from Theorem 3.15. $\square$

We want to remark that there are orientations of trees H with an NP-complete H-colouring problem (the smallest ones have 20 vertices [25]; see Figure 4). It can be shown (using a condition that will be presented in Section 14.3) that this digraph does not have tree-duality, without any complexity-theoretic assumptions.

Exercises.

42. Show that if G and H are homomorphically equivalent, then $P(G)$ and $P(H)$ are also homomorphically equivalent.

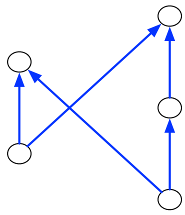

Figure 5: The graph from Exercise 48.

43. Recall that a digraph is called *balanced* if it homomorphically maps to a directed path. Let H be a finite digraph. Prove or disprove:
    - if H is balanced, then $P(H)$ is balanced;
    - if H is an orientation of a tree, then $P(H)$ is an orientation of a forest;
    - $P(H) \to H$ if and only if H is acyclic.

44. Solve the problems from the previous exercise for infinite digraphs H.

45. Show that AC solves $\mathrm{CSP}(T_n)$, for every $n \geq 1$.

46. Up to isomorphism, there is only one unbalanced cycle H on four vertices that is a core and not the directed cycle. Show that AC does not solve $\mathrm{CSP}(H)$.

47. Does the digraph $(\{0,1,2,3,4,5\}; \{(0,1),(1,2),(0,2),(3,2),(3,4),(4,5),(3,5),(0,5)\})$ have tree duality?

48. Can the CSP for the digraph depicted in Figure 5 be solved by the arc consistency procedure?

49. Let H be a finite digraph. Show that $P(H)$ contains a loop if and only if H contains a directed cycle.

50. Show that the previous statement is false for infinite digraphs H.

51. Show that an orientation of a tree homomorphically maps to H if and only if it homomorphically maps to $P(H)$.

52. Prove the final statement in Example 3.13.

53. Let H be a finite digraph. Then AC_H rejects an orientation of a tree T if and only if there is no homomorphism from T to H (in other words, AC solves $\mathrm{CSP}(H)$ if the input is restricted to orientations of trees).

54. Show that there is a linear-time algorithm that tests whether a given orientation of a tree is a core.

55. Show that the core of an orientation of a tree can be computed in polynomial time.

56. Let G and H be finite digraphs, let $x \in V(G)$, and let $L(x)$ be the list computed by the arc consistency procedure. Show that $L(x)$ is preserved by all polymorphisms of H.

57. Does Exercise 19 remain true for directed graphs?

4 The Path-consistency Procedure

The path-consistency procedure is a well-studied generalization of the arc-consistency procedure from artificial intelligence. The path-consistency procedure is also known as the pair-consistency check algorithm in the graph theory literature.

Many CSPs that can not be solved by the arc-consistency procedure can still be solved in polynomial time by the path-consistency procedure. The simplest examples are $H = K_2$ (see Exercise 19) and $H = \vec{C}_3$ (see Exercise 21). The idea is to maintain a list of pairs from $V(H)^2$ for each pair of elements from $V(G)$ (similarly to the arc-consistency procedure, where we maintained a list of vertices from $V(H)$ for each vertex in $V(G)$). We successively remove pairs from these lists when the pairs can be excluded *locally*. Some authors maintain a list only for each pair of *distinct* vertices of $V(G)$, and they refer to our (stronger) variant as the *strong path-consistency procedure*. Our procedure (where vertices need not be distinct) has the advantage that it is at least as strong as the arc-consistency procedure, because the lists $L(x, x)$ and the rules of the path-consistency procedure for $x = y$ simulate the rules of the arc-consistency procedure.

```
PC_H(G)
Input: a finite digraph G.
Data structure: for all x, y ∈ V(G) a list L(x, y) of elements of V(H)^2

For each (x, y) ∈ V(G)^2
    If (x, y) ∈ E(G) then L(x, y) := E(H),
    else L(x, y) := V(H)^2.
    If x = y then L(x, y) := L(x, y) ∩ {(u, u) | u ∈ V(H)}.
Do
    For all vertices x, y, z ∈ V(G):
        For each (u, w) ∈ L(x, z):
            If there is no v ∈ V(H) such that (u, v) ∈ L(x, y) and (v, w) ∈ L(y, z) then
                Remove (u, w) from L(x, z)
        If L(x, z) is empty then reject
Loop until no list changes
```

Figure 6: The (strong) path-consistency procedure for CSP(H).

In Subsection 4.2 we will see many examples of digraphs H where the path-consistency procedure solves the H-colouring problem, but the arc-consistency procedure does not. The greater power of the path-consistency procedure comes at the price of a bigger worst-case running time: while the arc-consistency procedure has linear-time implementations, the best known implementations of the path-consistency procedure require cubic time in the size of the input (see Exercise 58).

Remark 4.1. Similarly as for AC, the path-consistency procedure is polynomial even if H is part of the input, in which case we refer to the procedure with PC.

4.1 The k-consistency procedure

The path-consistency procedure can be generalised further to the k-consistency procedure. In fact, arc- and path-consistency procedure are just a special case of the k-consistency for $k = 2$ and $k = 3$, respectively. In other words, for digraphs H the path-consistency procedure is the 3-consistency procedure and the arc-consistency procedure is the 2-consistency procedure.

The idea of k-consistency is to maintain sets of $(k-1)$-tuples from $V(H)^{k-1}$ for each $(k-1)$-tuple from $V(G)^{k-1}$, and to successively remove tuples by local inference. It is straightforward to generalise also the details of the path-consistency procedure. For fixed H and fixed k, the running time of the k-consistency procedure is still polynomial in the size of G. But the dependency of the running time on k is clearly exponential.

However, we would like to point out that path consistency alias 3-consistency is of particular theoretical importance, due to the following recent result.

Theorem 4.2 (Barto and Kozik [11])**.** *If* CSP(H) *can be solved by k-consistency for some $k \geq 3$, then* CSP(H) *can also be solved by 3-consistency.*

Exercises

58. Show that the path-consistency procedure for CSP(H) can (for fixed H) be implemented such that the worst-case running time is cubic in the size of the input digraph. (Hint: use a worklist as in AC-3.)

59. Show that if path consistency solves CSP(H_1) and path consistency solves CSP(H_2), then path consistency solves CSP$(H_1 \uplus H_2)$.

4.2 Majority Polymorphisms

In this section, we present a powerful criterion that shows that for certain digraphs H the path-consistency procedure solves the H-colouring problem. Again, this condition was first discovered in more general form by Feder and Vardi [55]; it subsumes many criteria that were studied in artificial intelligence and in graph theory before.

Definition 4.3. Let D be a set. A function f from D^3 to D is called a *majority function* if f satisfies the following equations, for all $x, y \in D$:

$$f(x,x,y) = f(x,y,x) = f(y,x,x) = x$$

Example 4.4. As an example, let D be $\{1,\ldots,n\}$, and consider the ternary *median* operation, which is defined as follows. Let x, y be three elements from D. We define

$$\text{median}(x,x,y) = \text{median}(x,y,x) = \text{median}(y,x,x) := x.$$

If x, y, z are pairwise distinct elements of D, suppose that $\{x,y,z\} = \{a,b,c\}$, where $a < b < c$. Then median(x,y,z) is defined to be b. Note that

$$\text{median}(x,y,z) = \min(\max(x,y), \max(x,z), \max(y,z)). \qquad \triangle$$

If a digraph H has a polymorphism f that is a majority operation, then f is called a *majority polymorphism of H*.

Example 4.5. Let H be the transitive tournament on n vertices, T_n. Suppose the vertices of T_n are the first natural numbers, $\{1, \dots, n\}$, and $(u, v) \in E(T_n)$ if and only if $u < v$. Then the median operation is a polymorphism of T_n, because if $u_1 < v_1$, $u_2 < v_2$, and $u_3 < v_3$, then clearly $median(u_1, u_2, u_3) < median(v_1, v_2, v_3)$. △

Theorem 4.6 (of [55])**.** *Let H be a finite digraph. If H has a majority polymorphism, then the H-colouring problem can be solved in polynomial time (by the path-consistency procedure).*

For the proof of Theorem 4.6 we need the following lemma.

Lemma 4.7. *Let G and H be finite digraphs. Let f be a polymorphism of H of arity k and let $L := L(x, z)$ be the final list computed by the path-consistency procedure for $x, z \in V(G)$. Then f* preserves *L, i.e., if $(u_1, w_1), \dots, (u_k, w_k) \in L$, then $(f(u_1, \dots, u_k), f(w_1, \dots, w_k)) \in L$.*

Proof. Let $(u_1, w_1), \dots, (u_k, w_k) \in L$. We prove by induction over the execution of PC_H on G that at all times the pair $(u, w) := (f(u_1, \dots, u_k), f(w_1, \dots, w_k))$ is contained in L. Initially, this is true because f is a polymorphism of H. For the inductive step, let $y \in V(G)$. By definition of the procedure, for each $i \in \{1, \dots, k\}$ there exists v_i such that $(u_i, v_i) \in L(x, y)$ and $(v_i, w_i) \in L(y, z)$. By the inductive assumption, $(f(u_1, \dots, u_k), f(v_1, \dots, v_k)) \in L(x, y)$ and $(f(v_1, \dots, v_k), f(w_1, \dots, w_k)) \in L(y, z)$. Hence, $(f(u_1, \dots, u_k), f(w_1, \dots, w_k))$ will not be removed in the next step of the algorithm. □

Proof of Theorem 4.6. Let $f \colon V(H)^3 \to V(H)$ be a majority polymorphism of H. Clearly, if the path-consistency procedure derives the empty list for some pair (x, z) from $V(G)^2$, then there is no homomorphism from G to H.

Now suppose that after running the path-consistency procedure on G for all pairs (x, z) from $V(G)^2$ the list $L(x, z)$ is non-empty. We have to show that there exists a homomorphism from G to H. A function h from an induced subgraph G' of G to H is said to *preserve the lists* if $(h(x), h(z)) \in L(x, z)$ for all $x, z \in V(G')$. The proof shows by induction on i that every homomorphism from a subgraph of G with i vertices that preserves the lists can be extended to any other vertex in G such that the resulting mapping is a homomorphism to H that again preserves the lists.

For the base case of the induction, observe that for all vertices $x, z \in V(G)$ every mapping h from $\{x, z\}$ to $V(H)$ such that $(h(x), h(z)) \in L(x, z)$ can be extended to every $y \in V(G)$ such that $(h(x), h(y)) \in L(x, y)$ and $(h(y), h(z)) \in L(y, z)$ (and hence preserves the lists), because otherwise the path-consistency procedure would have removed $(h(x), h(z))$ from $L(x, z)$.

For the inductive step, let h' be any homomorphism from a subgraph G' of G on $i \geq 3$ vertices to H that preserves the lists, and let x be any vertex of G not in G'. Let x_1, x_2, and x_3 be some vertices of G', and h'_j be the restriction of h' to $V(G') \setminus \{x_j\}$, for $1 \leq j \leq 3$. By inductive assumption, h'_j can be extended to x such that the resulting mapping h_j is a homomorphism to H that preserves the lists. We claim that the extension h of h' that maps x to $f(h_1(x), h_2(x), h_3(x))$ is a homomorphism to H that preserves the lists.

For all $y \in V(G')$, we have to show that $(h(x), h(y)) \in L(x, y)$ (and that $(h(y), h(x)) \in L(y, x)$, which can be shown analogously). If $y \notin \{x_1, x_2, x_3\}$, then $h(y) = h'(y) = f(h'(y), h'(y), h'(y)) = f(h_1(y), h_2(y), h_3(y))$, by the properties of f. Since $(h_i(x), h_i(y)) \in L(x, y)$ for all $i \in \{1, 2, 3\}$, and since f preserves $L(x, y)$ by Lemma 4.7, we have $(h(x), h(y)) \in L(x, y)$, and are done in this case.

Clearly, y can be equal to at most one of $\{x_1, x_2, x_3\}$. Suppose that $y = x_1$ (the other two cases are analogous). There must be a vertex $v \in V(H)$ such that $(h_1(x), v) \in L(x, y)$

(otherwise the path-consistency procedure would have removed $(h_1(x), h_1(x_1))$ from $L(x, x_1)$). By the properties of f, we have $h(y) = h'(y) = f(v, h'(y), h'(y)) = f(v, h_2(y), h_3(y))$. Because $(h_1(x), v), (h_2(x), h_2(y)), (h_3(x), h_3(y))$ are in $L(x, y)$, Lemma 4.7 implies that $(h(x), h(y)) = (f(h_1(x), h_2(x), h_3(x)), f(v, h_2(y), h_3(y)))$ is in $L(x, y)$, and we are done.

We conclude that G has a homomorphism to H. $\square$

Corollary 4.8. *The path-consistency procedure solves the H-colouring problem for $H = T_n$.*

Another class of examples of digraphs having a majority polymorphism are *unbalanced cycles*, i.e., orientations of C_n that do not homomorphically map to a directed path [53]. We only prove a weaker result here.

Proposition 4.9. *Directed cycles have a majority polymorphism.*

Proof. Let $\vec{C}_n$ be a directed cycle. Let f be the ternary operation on the vertices of $\vec{C}_n$ that maps u, v, w to u if u, v, w are pairwise distinct, and otherwise acts as a majority operation. We claim that f is a polymorphism of $\vec{C}_n$. Let $(u, u'), (v, v'), (w, w') \in E(\vec{C}_n)$ be arcs. If u, v, w are all distinct, then u', v', w' are clearly all distinct as well, and hence $(f(u, v, w), f(u', v', w')) = (u, u') \in E(\vec{C}_n)$. Otherwise, if two elements of u, v, w are equal, say $u = v$, then u' and v' must be equal as well, and hence $(f(u, v, w), f(u', v', w')) = (u, u') \in E(\vec{C}_n)$. $\square$

Exercises.

60. Show that every orientation of a path has a majority polymorphism.

61. Show that C_4 has a majority polymorphism but C_6 does not.

62. A *quasi majority operation* is an operation from V^3 to V satisfying

$$f(x, x, y) = f(x, y, x) = f(y, x, x) = f(x, x, x)$$

for all $x, y \in V$.

- Show that every digraph with a quasi majority polymorphism is homomorphically equivalent to a digraph with a majority polymorphism.
- Use Theorem 2.6 to show that a finite undirected graph H has an H-colouring problem that can be solved in polynomial time if H has a quasi majority polymorphism, and is NP-complete otherwise.

63. There is only one unbalanced cycle H on four vertices that is a core and not the directed cycle (we have seen this digraph already in Exercise 46). Show that for this digraph H the H-colouring problem can be solved by the path-consistency procedure.

64. Determine for which $n \geq 1$ there is a linear order on the vertices of $\vec{C}_n$ such that median with respect to this linear order is a polymorphism of $\vec{C}_n$.

65. Determine for which $n \geq 1$ the operation f from Proposition 4.9 preserves $\vec{T}_n$.

66. Show that every unbalanced orientation of a cycle has a majority polymorphism.

67. Modify the path-consistency procedure such that it can deal with instances of the precoloured H-colouring problem. Show that if H has a majority polymorphism, then the modified path-consistency procedure solves the precoloured H-colouring problem.

68. Modify the path-consistency procedure such that it can deal with instances of the list H-colouring problem. Show that if H has a *conservative* majority polymorphism, then the modified path-consistency procedure solves the list H-colouring problem.

69. An *interval graph* H is an (undirected) graph $H = (V; E)$ such that there is an interval I_x of the real numbers for each $x \in V$, and $(x, y) \in E$ if and only if I_x and I_y have a non-empty intersection. Note that with this definition interval graphs are necessarily *reflexive*, i.e., $(x, x) \in E$. Show that the precoloured H-colouring problem for interval graphs H can be solved in polynomial time. Hint: use the modified path-consistency procedure in Exercise 67.

70. Show that if H is a (reflexive) interval graph, then H has a conservative majority polymorphism.

71. Let H be a reflexive graph. Show that H has a conservative majority polymorphism if and only if H is a *circular arc graph*, i.e., H can be represented by arcs on a circle so that two vertices are adjacent if and only if the corresponding arcs intersect.

72. Let H be an irreflexive graph. Then H has a conservative majority polymorphism if and only if H is bipartite and the complement of a circular arc graph.

73. Show that the digraph $(\mathbb{Z}; \{(x, y) \mid x - y = 1\})$ has a majority polymorphism and that its CSP can be solved in polynomial time.

74. Show that $(\mathbb{Z}; \neq)$ does not have a majority polymorphism, but a quasi majority polymorphism and that $\mathrm{CSP}(\mathbb{Z}; \neq)$ can be solved in polynomial time.

75. Show that the digraph $H = (\mathbb{Z}; \{(x, y) \mid x - y \in \{1, 3\}\})$ has a majority polymorphism, and give a polynomial time algorithm for its H-colouring problem.

76. Consider the digraph C_2^{++} depicted in Figure 7 (a so-called *semicomplete digraph*). Show the following statements.[3]
    - $\mathrm{CSP}(C_2^{++})$ cannot be solved by the arc-consistency procedure.
    - A finite digraph G homomorphically maps to C_2^{++} if and only if no digraph of the following form maps to G: start with any orientation of an odd cycle, and if $(u, v), (w, v)$ are edges on the cycle, append to v an outgoing directed path with two edges.
    - H does not have a majority polymorphism.
    - $\mathrm{CSP}(C_2^{++})$ can be solved by the path-consistency procedure.

[3]The author thanks Florian Starke and Sebastian Meyer for the idea for this exercise.

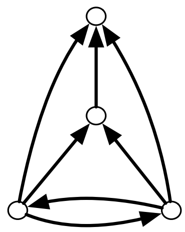

Figure 7: The graph C_2^{++} from Exercise 76.

4.3 Testing for Majority Polymorphisms

In this section we show that the question whether a given digraph has a majority polymorphism can be decided in polynomial time. The method that we present is sometimes referred to as a *self-reduction* and can be adapted for several other polymorphism conditions and several other algorithms (see Exercise 158).

```
Majority-Test(H)
Input: a finite digraph H.

Let G := H^3.
For all u, v ∈ V(H), precolour the vertices (u,u,v), (u,v,u), (v,u,u), (u,u,u) with u.
If PC_H(G) derives the empty list, reject.
For each x ∈ V(G)
    Found-Value := False.
    For each (u,u) ∈ L(x,x)
        For all y, z ∈ V(G), let L'(y,z) be a copy of L(y,z).
        L'(x,x) := {(u,u)}.
        Run PC_H(G) with the lists L'.
        If this run does not derive the empty list
            For all y, z ∈ V(G), set L(y,z) := L'(y,z).
            Found-Value := True.
    End For.
    If Found-Value = False then reject.
End For.
Accept.
```

Figure 8: A polynomial-time algorithm to find majority polymorphisms.

Theorem 4.10. *There is a polynomial-time algorithm to decide whether a given digraph H has a majority polymorphism.*

Proof. The pseudo-code of the procedure can be found in Figure 8. Given H, we construct a new digraph G as follows. We start from the third power H^3, and precolour all vertices of the form (u, u, v), (u, v, u), (v, u, u), and (u, u, u) with u. Let G be the resulting precoloured digraph. Note that there exists a homomorphism from G to H that respects the colours if and only if H has a majority polymorphism.

To decide whether G has a homomorphism to H, we run the modification of PC_H for the precoloured H-colouring problem on G (see Exercise 67). If this algorithm rejects, then we can be sure that there is no homomorphism from G to H that respects the colours, and hence H has no majority polymorphism. Otherwise, we use the same idea as in the proof of Proposition 3.6: create a copy L' of the lists L. Pick $x \in V(G)$ and remove all but one pair (u,u) from $L'(x,x)$. If PC_H derives the empty list on L' instead of L, we try the same with another pair (v,v) from $L(x,x)$.

If there exists $x \in V(G)$ such that PC_H detects an empty list for all $(u,u) \in L(x,x)$ then the adaptation of PC_H for the precoloured CSP would have given an incorrect answer for the previously selected variable: PC_H did not detect the empty list even though the input was unsatisfiable. Hence, H cannot have a majority polymorphism by Theorem 4.6.

Otherwise, if PC_H does not derive the empty list after removing all pairs but (u,u) from $L(x,x)$, we continue with another vertex $y \in V(G)$, setting $L(y,y)$ to $\{(u,u)\}$ for some $(u,u) \in L(y,y)$. We repeat this procedure; if the algorithm never rejects, then eventually all lists for pairs of the form (x,x) are singleton sets $\{(u,u)\}$; the map that sends x to u is a homomorphism from G to H that respects the colours. In this case we return '*true*'.

It is easy to see that the procedure described above has polynomial running time. □

Exercises.

77. Modify the algorithm 'Majority-Test' to obtain an algorithm that tests whether a given digraph H has a quasi majority polymorphism.

4.4 Digraphs with a Maltsev Polymorphism

If a digraph H has a *majority* polymorphism, then the path-consistency procedure solves $\mathrm{CSP}(H)$. How about digraphs H with a *minority* polymorphisms of H? It turns out that this is an even stronger restriction.

Definition 4.11. A ternary function $f\colon D^3 \to D$ is called

- a *minority operation* if it satisfies
$$\forall x,y \in D.\, f(y,x,x) = f(x,y,x) = f(x,x,y) = y$$
- and a *Maltsev operation* if it satisfies
$$\forall x,y \in D.\, f(y,x,x) = f(x,x,y) = y.$$

Example 4.12. Let $D := \{0,\dots,n-1\}$. Then the function $f\colon D^3 \to D$ given by $(x,y,z) \mapsto x - y + z \mod n$ is a Maltsev operation, since $x - x + z = z$ and $x - z + z = x$. For $n = 2$, this is even a minority operation. If $n > 2$, this function is not a minority, since then $1 - 2 + 1 = 0 \not\equiv 2 \mod n$. Note that f is a polymorphism of $\vec{C}_n$. To see this, suppose that $u_1 - v_1 \equiv 1 \mod n$, $u_2 - v_2 \equiv 1 \mod n$, and $u_3 - v_3 \equiv 1 \mod n$. Then

$$\begin{aligned} f(u_1,u_2,u_3) &\equiv u_1 - u_2 + u_3 \equiv (v_1+1) - (v_2+1) + (v_3+1) \\ &\equiv f(v_1,v_2,v_3) + 1 \mod n\,. \end{aligned}$$ △

The following result appeared in 2011.

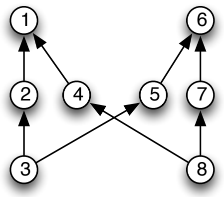

Figure 9: A totally rectangular digraph.

Theorem 4.13 (Kazda [68])**.** *If a finite digraph H has a Maltsev polymorphism then H also has a majority polymorphism.*

Hence, for finite digraphs H with a Maltsev polymorphism, the strong path-consistency procedure solves the H-colouring problem, and in fact even the precoloured H-colouring problem. Theorem 4.13 is an immediate consequence of Theorem 4.19 below; to state it, we need the following concepts.

Definition 4.14. A digraph G is called *rectangular* if $(x,y),(x',y),(x',y') \in E(G)$ implies that $(x,y') \in E(G)$.

We start with the fundamental observation: digraphs with a Maltsev polymorphism m are rectangular. This follows immediately from the definition of polymorphisms: we must have $(m(x,x',x'),m(y,y,y')) \in E(G)$, but $m(x,x',x') = x$ and $m(y,y,y') = y'$, so $(x,y') \in E(G)$. The converse does not hold, as the following example shows.

Example 4.15. The digraph $\big(\{a,b,c\};\{(a,a),(a,b),(b,c),(c,c)\}\big)$ is rectangular, but has no Maltsev polymorphism m. Indeed, such an m would have to satisfy $m(a,a,c) = c$ and $m(a,c,c) = a$. Note that

$$\begin{gathered}(m(a,a,c),m(a,b,c)) \in E(G)\\ \text{and } (m(a,b,c),m(a,c,c)) \in E(G),\end{gathered}$$

but G has no vertex x such that $(c,x) \in E(G)$ and $(x,a) \in E(G)$. △

We are therefore interested in stronger consequences of the existence of a Maltsev polymorphism.

Definition 4.16. A digraph G is called *k-rectangular* if whenever G contains directed paths of length k from x to y, from x' to y, and from x' to y', then also from x to y'. A digraph G is called *totally rectangular* if it is k-rectangular for all $k \geq 1$.

Lemma 4.17. *Every digraph with a Maltsev polymorphism m is totally rectangular.*

Proof. Let $k \geq 1$, and suppose that G is a digraph with directed paths $(x_1,\dots,x_k)$, $(y_1,\dots,y_k)$, and $(z_1,\dots,z_k)$ such that $x_k = y_k$ and $y_1 = z_1$. We have to show that there exists a directed path $(u_1,\dots,u_k)$ in G with $u_1 = x_1$ and $u_k = z_k$. It can be verified that $u_i := m(x_i,y_i,z_i)$ has the desired properties. □

An example of a totally rectangular digraph is given in Figure 9. The next lemma points out an important consequence of k-rectangularity.

Lemma 4.18. *Let G be a finite totally rectangular digraph with a cycle of net length $d > 0$. Then G contains a directed cycle of length d.*

Proof. Let $C = (u_0, \ldots, u_{k-1})$ be a cycle of G of net length d; we prove the statement by induction on k. Clearly, C can be decomposed into maximal directed paths, that is, there is a minimal set $\mathcal{D}$ of directed paths such that each pair $(u_0, u_1), (u_1, u_2), \ldots, (u_{k-1}, u_0)$ is contained in exactly one of the paths of $\mathcal{D}$. If the decomposition $\mathcal{D}$ consists of a single directed path then we have found a directed cycle and are done. Let P be the shortest directed path of $\mathcal{D}$, leading from u to v in G. Then there are directed paths Q and Q' in $\mathcal{D}$ such that Q starts in u and Q' ends in v, and $P \neq Q$ or $P \neq Q'$. By assumption, $|Q|, |Q'| \geq \ell := |P|$. By ℓ-rectangularity, there exists a directed path P' of length ℓ from the vertex s of Q' at position $|Q'| - \ell$ to the vertex t of Q at position ℓ. Now we distinguish the following cases.

- $Q = Q'$: the cycle that starts in s, follows the path Q until t, and then returns to s via the path P' is shorter than C but still has net length d.
- $Q \neq Q'$: the cycle starting in s, following Q for the final $|Q| - \ell$ vertices of Q, the cycle C until Q', the first $|Q'| - \ell$ vertices of Q' until t, and then P' back to s is a cycle which is shorter than C but still has net length d.

In both cases, the statement follows by induction. □

The following is a strengthening of Theorem 4.13; we only prove that 1 implies 2, and 2 implies 3, which suffices for the already mentioned consequence that for digraphs H with a Maltsev polymorphism, path consistency solves the H-colouring problem (cf. Exercise 59).

Theorem 4.19 (Theorem 3.3 and Corollary 4.12 in [43]). *Let G be a finite digraph. Then the following are equivalent.*

1. *G has a Maltsev polymorphism.*
2. *G is totally rectangular.*
3. *If G is acyclic, then the core of G is a directed path. Otherwise, the core of G is a disjoint union of directed cycles.*
4. *G has a minority and a majority polymorphism.*

Proof. The implication from 4 to 1 is trivial since every minority operation is in particular a Maltsev operation. The implication from 1 to 2 is Lemma 4.17. For the implication from 2 to 3, let us assume that G is connected. The general case then follows by applying the following argument to each of its connected components, and the observation that directed paths homomorphically map to longer directed paths and to directed cycles.

We first consider the case that G is acyclic, and claim that in this case G is balanced, i.e., there exists a surjective homomorphism h from G to $\vec{P}_n$ for some $n \geq 1$. Otherwise, there exist $u, v \in V(G)$ and two paths P and Q from u to v of different net lengths ℓ_1 and ℓ_2 (see Exercise 13). Put these two paths together at u and v to form an unbalanced cycle C. Then Lemma 4.18 implies that G contains a directed cycle contrary to our assumptions.

Now, choose n with $G \to \vec{P}_n$ minimal, and fix $u \in h^{-1}(0)$ and $v \in h^{-1}(n)$. Then it is easy to see from total rectangularity that there must exist a path of length n in G from u to v, and hence the core of G is $\vec{P}_n$.

Now suppose that G contains a directed cycle; let C be the shortest directed cycle of G. We prove that G homomorphically maps to C. It is easy to see that it suffices to show that for any two vertices u, v of G and for any two paths P and Q from u to v we have that their net lengths are congruent modulo $m := |C|$ (see Exercise 14). Suppose for contradiction that there are paths of net length ℓ_1 and ℓ_2 from u to v in G such that $d := \ell_1 - \ell_2 \neq 0$ modulo m; without loss of generality, $\ell_2 < \ell_1$, so $d > 0$. We can assume that u is an element of C, since otherwise we can choose a path S from a vertex of C to u by connectivity of G, and append S to both P and Q. We can also assume that $d < m$ because if not, we can append C to Q to increase the length of Q by a multiple of m, until $d = \ell_1 - \ell_2 < m$. Lemma 4.18 then implies that G contains a directed cycle of length d, a contradiction to the choice of C.

For the missing implication from 3 to 4, we refer to [43] (Corollary 4.11). $\square$

Rectangularity will be revisited from a universal algebraic perspective in Section 11.2 (Proposition 11.10).

Exercises.

78. Let H be the digraph $(\{0, 1, \ldots, 6\}; \{(0,1), (1,2), (3,2), (4,3), (4,5), (5,6)\})$. For which k is it k-rectangular?

79. Show that $G = (V, E)$ is rectangular if and only if E is a disjoint union of sets of the form $A \times B$ where $A, B \subseteq V$.

5 Logic

A *signature* is a set of relation and function symbols. The relation symbols are typically denoted by $R, S, T, \ldots$ and the function symbols are typically denoted by $f, g, h, \ldots$; each relation and function symbol is equipped with an *arity* from $\mathbb{N}$. A *τ-structure* $\mathfrak{A}$ consists of

- a set A (the *domain* or *base set*; we typically use the same letter in a different font)
- a relation $R^{\mathfrak{A}} \subseteq A^k$ for each relation symbol R of arity k from τ, and
- an operation $f^{\mathfrak{A}} \colon A^k \to A$ for each function symbol f of arity k from τ.

Function symbols of arity 0 are allowed; they are also called *constant symbols* (and the respective operations are called *constants*). In this text it causes no harm to allow structures whose domain is empty. A τ-structure $\mathfrak{A}$ is called *finite* if its domain A is finite.

A *homomorphism* h from a τ-structure $\mathfrak{A}$ to a τ-structure $\mathfrak{B}$ is a function from A to B that *preserves* each relation and each function: that is,

- if $(a_1, \ldots, a_k)$ is in $R^{\mathfrak{A}}$, then $(h(a_1), \ldots, h(a_k))$ must be in $R^{\mathfrak{B}}$;
- for all $a_1, \ldots, a_k \in A$ we have $h(f^{\mathfrak{A}}(a_1, \ldots, a_k)) = f^{\mathfrak{B}}(h(a_1), \ldots, h(a_k))$.

An *isomorphism* is a bijective homomorphism h such that the inverse mapping $h^{-1}: B \to A$ that sends $h(x)$ to x is a homomorphism, too.

A *relational* structure is a τ-structure where τ only contains relation symbols, and an *algebra* (in the sense of universal algebra) is a τ-structure where τ only contains function symbols. This section is mainly about relational structures; algebras will appear in Section 8.

Note that in a τ-structure $\mathfrak{A}$, every function symbol of arity n must be defined on all of A^n; in some settings, this requirement is not natural and we therefore also define *multi-sorted* structures.

5.1 Multisorted Structures

This section is for later reference, and can be skipped at the first reading. Multisorted structures will be used in Section 8.6 and in Section 9.

Let S be a set; the elements of S are called *sorts*. We write S^* for the set of words over S (i.e., finite sequences of elements of S) and S^+ for the set of non-empty words over S. An *S-sorted signature* τ consists of a set of *function symbols* (typically denoted by $f, g, \dots$) and a set of *relation symbols* (typically denoted by $R, S, \dots$). Each relation symbol $R \in \tau$ is equipped with a *type* $\mathrm{tp}(R) \in S^*$, and each function symbol $f \in \tau$ is equipped with a type $\mathrm{tp}(R) \in S^+$.

If τ is an S-sorted signature, then an *S-sorted τ-structure* $\mathfrak{M}$ consists of

- a set A_s for every $s \in S$;
- a relation $R^{\mathfrak{M}} \subseteq A_{s_1} \times \cdots \times A_{s_n}$ for each relation symbol $R \in \tau$ of type $\mathrm{tp}(R) = (s_1, s_2, \dots, s_n)$, $n \in \mathbb{N}$;
- a function $f^{\mathfrak{M}}: A_{s_1} \times \cdots \times A_{s_n} \to A_{s_0}$ for each function symbol $f \in \tau$ of type $\mathrm{tp}(f) = (s_0, s_1, \dots, s_n)$, for $n \in \mathbb{N}$.

Note that for the one-sorted case, i.e., if $|S| = 1$, we recover the notion of a structure as introduced earlier. Vector spaces or, more generally, modules may be viewed naturally as two-sorted structures; see Example 8.3.

The syntax and semantics of first-order logic over an S-sorted signature τ are defined as follows. Let V be a set; the elements of V are called *variables*. Each variable $x \in V$ is equipped with a *type* $\mathrm{tp}(x) \in S$; we require that for every $s \in S$ there are infinitely many variables of type s.

If $x_1, \dots, x_n$ are variables, then a *τ-term t of type* $\mathrm{tp}(t)$ *over the variables* $x_1, \dots, x_n$ is defined inductively as follows.

- each variable from $x_1, \dots, x_n$ is a τ-term of type $\mathrm{tp}(t) = (\mathrm{tp}(x_i), \mathrm{tp}(x_1), \dots, \mathrm{tp}(x_n))$.
- if $f \in \tau$ is a function symbol of type $(s_0, s_1, \dots, s_n)$, for $n \in \mathbb{N}$, and for each $i \in \{1, \dots, n\}$ we have a τ-term t_i of type $(s_i, \mathrm{tp}(x_1), \dots, \mathrm{tp}(x_n))$ over the variables $x_1, \dots, x_n$, then (the syntactic object) $f(t_1, \dots, t_n)$ is a τ-term over $x_1, \dots, x_n$ of type $(s_0, \mathrm{tp}(x_1), \dots, \mathrm{tp}(x_n))$. Note that the case $n = 0$ is another base case of the induction, which covers terms without any occurrence of variables.

All τ-terms t over the variables $x_1, \dots, x_n$ are built in this way; we often write $t(x_1, \dots, x_n)$ to indicate that t is a term over $x_1, \dots, x_n$.

If $\mathfrak{M}$ is an S-sorted τ-structure, $x_1, \dots, x_n \in V$, and $t(x_1, \dots, x_n)$ is a τ-term of type $\mathrm{tp}(t) = (s_0, \mathrm{tp}(x_1), \dots, \mathrm{tp}(x_n))$, then the *term function* $t^{\mathfrak{M}}$ (in the one-sorted case called *term operation*) is the function of type $\mathrm{tp}(t)$ defined inductively as follows:

- if t is of the form x_i, for $i \in \{1, \dots, n\}$, then $t^{\mathfrak{M}}$ is the function $(a_1, \dots, a_n) \mapsto a_i$.
- if t is of the form $f(t_1, \dots, t_k)$, for f of type $(s_0, \mathrm{tp}(t_1), \dots, \mathrm{tp}(t_k))$, then $t^{\mathfrak{M}}$ is the function $(a_1, \dots, a_n) \mapsto f^{\mathfrak{M}}(t_1^{\mathfrak{M}}(a_1, \dots, a_n), \dots, t_k^{\mathfrak{M}}(a_1, \dots, a_n))$.

An *atomic S-sorted τ-formula over variables* $x_1, \dots, x_n$ *of type* $(\mathrm{tp}(x_1), \dots, \mathrm{tp}(x_n))$ is

- an expression of the form $t_1 = t_2$, for S-sorted τ-terms $t_1(x_1, \dots, x_n)$ and $t_2(x_1, \dots, x_n)$ of type $\mathrm{tp}(t_1) = (s_0, \mathrm{tp}(x_1), \dots, \mathrm{tp}(x_n))$ and $\mathrm{tp}(t_2) = (s_0, \mathrm{tp}(x_1), \dots, \mathrm{tp}(x_n))$,
- an expression of the form $R(t_1, \dots, t_n)$, for S-sorted τ-terms $t_1, \dots, t_n$ and a relation symbol $R \in \tau$ of type $(\mathrm{tp}(t_1), \dots, \mathrm{tp}(t_n))$.

An S-sorted first-order τ-formula over the variables $x_1, \dots, x_n \in V$ of type $(\mathrm{tp}(x_1), \dots, \mathrm{tp}(x_n))$ is defined inductively as one of the following expressions:

- an atomic S-sorted first-order τ-formula over $x_1, \dots, x_n$;
- $\phi_1 \wedge \phi_2$ for S-sorted first-order τ-formula ϕ_1, ϕ_2 over the variables $x_1, \dots, x_n$;
- $\neg\phi$ for an S-sorted first-order τ-formula ϕ over the variables $x_1, \dots, x_n$;
- $\exists x_0.\phi$ where $x_0 \in V$ and ϕ is an S-sorted first-order τ-formula over the variables $x_0, x_1, \dots, x_n$.

If $\mathfrak{M}$ is an S-sorted τ-structure, $x_1, \dots, x_n \in V$, and $\phi(x_1, \dots, x_n)$ is an τ-formula, then $\phi^{\mathfrak{M}}$ is the relation defined as follows.

- If ϕ is atomic and of the form $t_1 = t_2$, then $\phi^{\mathfrak{M}}$ consists of all tuples $(a_1, \dots, a_n) \in A_{\mathrm{tp}(x_1)} \times \cdots \times A_{\mathrm{tp}(x_n)}$ such that $t_1^{\mathfrak{M}}(a_1, \dots, a_n) = t_2^{\mathfrak{M}}(a_1, \dots, a_n)$.
- If ϕ is atomic and of the form $R(t_1, \dots, t_k)$ for S-sorted τ-terms $t_1, \dots, t_k$ and $R \in \tau$ of type $(\mathrm{tp}(t_1), \dots, \mathrm{tp}(t_k))$, then $\phi^{\mathfrak{M}}$ consists of all tuples $(a_1, \dots, a_n) \in A_{\mathrm{tp}(x_1)} \times \cdots \times A_{\mathrm{tp}(x_n)}$ such that
$$\left(t_1^{\mathfrak{M}}(a_1, \dots, a_n), \dots, t_k^{\mathfrak{M}}(a_1, \dots, a_n)\right) \in R^{\mathfrak{M}}.$$
- If ϕ is of the form $\phi_1 \wedge \phi_2$ for S-sorted τ-formulas ϕ_1 and ϕ_2 over $x_1, \dots, x_n$, then $\phi^{\mathfrak{M}} := \phi_1^{\mathfrak{M}} \cap \phi_2^{\mathfrak{M}}$.
- If ϕ is of the form $\neg\psi$ for some S-sorted τ-formula ψ over $x_1, \dots, x_n$, then $\phi^{\mathfrak{M}} := (A_{\mathrm{tp}(x_1)} \times \cdots \times A_{\mathrm{tp}(x_n)}) \setminus \psi^{\mathfrak{M}}$;
- If ϕ is of the form $\exists x_0.\psi$ for some $x_0 \in V$ and some S-sorted τ-formula ψ over $x_0, x_1, \dots, x_n$, then $\phi^{\mathfrak{M}}$ consists of all tuples $(a_1, \dots, a_n) \in A_{\mathrm{tp}(x_1)} \times \cdots \times A_{\mathrm{tp}(x_n)}$ such that there exists $a_0 \in A_{\mathrm{tp}(x_0)}$ such that $(a_0, a_1, \dots, a_n) \in \psi^{\mathfrak{M}}$.

We recover the syntax and semantics of usual first-order logic as the special case of the one-sorted case.

Exercises.

80. Show that one can decide in polynomial time whether a given string is an S-sorted τ-term over variables $x_1, \dots, x_n$.

81. Generalise the notion of a homomorphism between τ-structures to S-sorted τ-structures.

5.2 Primitive Positive Formulas

A first-order τ-formula $\phi(x_1, \dots, x_n)$ is called *primitive positive* (in database theory also *conjunctive query*) if it is of the form

$$\exists x_{n+1}, \dots, x_\ell (\psi_1 \wedge \dots \wedge \psi_m)$$

where $\psi_1, \dots, \psi_m$ are *atomic τ-formulas*, i.e., formulas of the form $R(y_1, \dots, y_k)$ with $R \in \tau$ and $y_i \in \{x_1, \dots, x_\ell\}$, of the form $y = y'$ for $y, y' \in \{x_1, \dots, x_\ell\}$, or $\top$ and $\bot$ (for *true* and *false*). As usual, formulas without free variables are called *sentences*. If $\mathfrak{A}$ is a τ-structure and ϕ a τ-sentence, then we write $\mathfrak{A} \models \phi$ if $\mathfrak{A}$ satisfies ϕ (i.e., ϕ holds in $\mathfrak{A}$).

Note that if we would require that all our structures have a non-empty domain, we would not need the symbol $\top$ since we can use the primitive positive sentence $\exists x.\, x = x$ to express it. It is possible to rephrase the H-colouring problem and its variants using primitive positive sentences.

Definition 5.1. Let $\mathfrak{B}$ be a structure with a finite relational signature τ. Then $\mathrm{CSP}(\mathfrak{B})$ is the computational problem of deciding whether a given primitive positive τ-sentence ϕ is true in $\mathfrak{B}$.

The given primitive positive τ-sentence ϕ is also called an *instance* of $\mathrm{CSP}(\mathfrak{B})$. The conjuncts of an instance ϕ are called the *constraints* of ϕ. A mapping from the variables of ϕ to the elements of B that is a satisfying assignment for the quantifier-free part of ϕ is also called a *solution* to ϕ.

Example 5.2 (Disequality constraints)**.** Consider the problem $\mathrm{CSP}(\{1, 2, \dots, n\}; \neq)$. An instance of this problem can be viewed as an (existentially quantified) set of variables, some linked by disequality[4] constraints. Such an instance holds in $(\{1, 2, \dots, n\}; \neq)$ if and only if the graph whose vertices are the variables, and whose edges are the disequality constraints, has a homomorphism to K_n. $\triangle$

The *dichotomy conjecture* of Feder and Vardi was that $\mathrm{CSP}(\mathfrak{B})$ is always in P or NP-complete, for every finite structure $\mathfrak{B}$ with finite relational signature; this conjecture was proved by Bulatov [36] and by Zhuk [96]. For a first more informative formulation of their result, see Theorem 5.19; many more reformulations can be found later in the text. Feder and Vardi showed that their conjecture is equivalent to the special case of their conjecture for finite digraphs (see Theorem 2.5), because for every relational structure $\mathfrak{B}$ there exists a finite digraph $\mathfrak{H}$ such that $\mathrm{CSP}(\mathfrak{B})$ and $\mathrm{CSP}(\mathfrak{H})$ are polynomial time equivalent; this result has later been refined in [41].

[4] We deliberately use the word *disequality* instead of *inequality*, since we reserve the word *inequality* for the relation $x \leq y$.

Exercises.

82. Show that $\mathrm{CSP}\big(\{0,1\}; R_{0,1}, R_{1,1}, R_{0,0}\big)$, where the relation $R_{i,j}$ equals $\{0,1\}^2 \setminus \{(i,j)\}$, can be solved in polynomial time.

83. Generalise the notion of *direct products* from digraphs (Definition 2.1) to general relational τ-structures.

84. Generalise the arc-consistency procedure and the power graph to general relational structures, and prove a generalisation of Theorem 3.4.

85. Generalise the concept of tree duality to general relational structures, and prove a generalisation of Theorem 3.8.

86. Does the arc-consistency procedure (see Exercise 84) solve $\mathrm{CSP}(\mathfrak{B})$ where $\mathfrak{B}$ has domain $B = \{0,1,2,3\}$, the unary relation $U_i^{\mathfrak{B}}$ for every $i \in B$, and the binary relations $B^4 \setminus \{(0,0)\}$ and $\{(1,2),(2,3),(3,1),(0,0)\}$?

87. Generalise the k-consistency procedure from digraphs to general relational structures.

88. Verify that the structure $\mathfrak{B}$ from Exercise 86 has the binary idempotent commutative polymorphism $*$ defined as $1*2 = 2$, $2*3 = 3$, $3*1 = 1$, and $0*b = b$ for all $b \in \{1,2,3\}$. Verify that $*$ satisfies 'restricted associativity', i.e., $x*(x*y) = (x*x)*y$ for all $x, y \in B$ (and since it is additionally idempotent and commutative it is called a *2-semilattice*).

89. Does the structure $\mathfrak{B}$ from Exercise 86 have a majority polymorphism?

90. Does the path-consistency procedure solve $\mathrm{CSP}(\mathfrak{B})$ for the structure $\mathfrak{B}$ from Exercise 86?

91. Let $\mathfrak{B}$ be the structure with domain $B := \{-1, 0, +1\}$ and the ternary relations

$$R^{\mathfrak{B}} := \{(x,y,z) \in B^3 \mid x+y+z \geq 1\}$$
$$S^{\mathfrak{B}} := \{(x,y,z) \in B^3 \mid x+y+z \leq -1\}.$$

- Prove that for every $k \in \mathbb{N}$, the k-ary operation s defined by

$$(x_1,\dots,x_k) \mapsto \begin{cases} +1 & (x_1+\dots+x_k)/k \geq 1/3 \\ -1 & (x_1+\dots+x_k)/k \leq 1/3 \\ 0 & \text{otherwise} \end{cases}$$

  is a polymorphism of $\mathfrak{B}$.
- Show that $\mathrm{CSP}(\mathfrak{B})$ cannot be solved by the arc-consistency procedure (see Exercise 84).
- Show that the k-consistency procedure solves $\mathrm{CSP}(\mathfrak{B})$, for a sufficiently large k (see Exercise 87).

5.3 From Structures to Formulas

To every finite relational τ-structure $\mathfrak{A}$ we can associate a τ-sentence, called the *canonical conjunctive query* of $\mathfrak{A}$, and denoted by $\phi(\mathfrak{A})$. The variables of this sentence are the elements of $\mathfrak{A}$, all of which are existentially quantified in the quantifier prefix of the formula, which is followed by the conjunction of all formulas of the form $R(a_1, \dots, a_k)$ for $R \in \tau$ and tuples $(a_1, \dots, a_k) \in R^{\mathfrak{A}}$.

For example, the canonical conjunctive query $\phi(K_3)$ of the complete graph on three vertices K_3 is the formula

$$\exists u, v, w \, \big(E(u,v) \wedge E(v,u) \wedge E(v,w) \wedge E(w,v) \wedge E(u,w) \wedge E(w,u)\big) \, .$$

The proof of the following proposition is straightforward.

Proposition 5.3. *Let $\mathfrak{B}$ be a structure with finite relational signature τ, and let $\mathfrak{A}$ be a finite τ-structure. Then there is a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$ if and only if $\mathfrak{B} \models \phi(\mathfrak{A})$.*

5.4 From Formulas to Structures

To present a converse of Proposition 5.3, we define the *canonical structure* $\mathfrak{S}(\phi)$ (in database theory this structure is called the *canonical database*) of a primitive positive τ-sentence, which is a relational τ-structure defined as follows. We require that ϕ does not contain $\bot$. If ϕ contains an atomic formula of the form $x = y$, we remove it from ϕ, and replace all occurrences of x in ϕ by y. Repeating this step if necessary, we may assume that ϕ does not contain atomic formulas of the form $x = y$.

Then the domain of $\mathfrak{S}(\phi)$ is the set of variables that occur in ϕ. There is a tuple $(v_1, \dots, v_k)$ in a relation R of $\mathfrak{S}(\phi)$ if and only if ϕ contains the conjunct $R(v_1, \dots, v_k)$. The following is similarly straightforward as Proposition 5.3.

Proposition 5.4. *Let $\mathfrak{B}$ be a relational τ-structure and let ϕ be a primitive positive τ-sentence that does not contain $\bot$. Then $\mathfrak{B} \models \phi$ if and only if $\mathfrak{S}(\phi)$ homomorphically maps to $\mathfrak{B}$.*

Due to Proposition 5.4 and Proposition 5.3, we may freely switch between the homomorphism and the logic perspective whenever this is convenient. In particular, instances of $\mathrm{CSP}(\mathfrak{B})$ can from now on be either finite structures $\mathfrak{A}$ or primitive positive sentences ϕ.

Note that the H-colouring problem, the precoloured H-colouring problem, and the list H-colouring problem can be viewed as constraint satisfaction problems for appropriately chosen relational structures.

5.5 Primitive Positive Definability

Let $\mathfrak{A}$ be a τ-structure, and let $\mathfrak{A}'$ be a τ'-structure with $\tau \subseteq \tau'$. If $\mathfrak{A}$ and $\mathfrak{A}'$ have the same domain and $R^{\mathfrak{A}} = R^{\mathfrak{A}'}$ for all $R \in \tau$, then $\mathfrak{A}$ is called the *τ-reduct* (or simply *reduct*) of $\mathfrak{A}'$, and $\mathfrak{A}'$ is called a *τ'-expansion* (or simply *expansion*) of $\mathfrak{A}$. If $\mathfrak{A}$ is a structure, and R is a relation over the domain of $\mathfrak{A}$, then we denote the expansion of $\mathfrak{A}$ by R by $(\mathfrak{A}, R)$.

If $\mathfrak{A}$ is a τ-structure, and $\phi(x_1, \dots, x_k)$ is a formula with k free variables $x_1, \dots, x_k$, then the *relation defined by ϕ* is the relation

$$\{(a_1, \dots, a_k) \mid \mathfrak{A} \models \phi(a_1, \dots, a_k)\} \, .$$

If the formula is primitive positive, then this relation is called *primitive positive definable.*

Example 5.5. The relation $\{(a, b) \in \{0, 1, 2, 3, 4\}^2 \mid a \neq b\}$ is primitive positive definable in C_5: the primitive positive definition is

$$\exists p_1, p_2 \left(E(x_1, p_1) \land E(p_1, p_2) \land E(p_2, x_2)\right). \qquad \triangle$$

Example 5.6. The non-negative integers are primitively positively definable in $(\mathbb{Z}; 0, 1, +, *)$, namely by the following formula $\phi(x)$ which states that x is the sum of four squares.

$$\exists x_1, x_2, x_3, x_4 (x = x_1^2 + x_2^2 + x_3^2 + x_4^2)$$

Clearly, every integer that satisfies $\phi(x)$ is non-negative; the converse is the famous four-square theorem of Lagrange [59]. $\triangle$

Definition 5.7 (Relational product)**.** For binary relations $R_1, R_2 \subset B^2$, define $R_1 \circ R_2$ to be the binary relation

$$\left\{(x, y) \mid \exists z (R_1(x, z) \land R_2(z, y))\right\}. \tag{1}$$

For $R \subseteq B^2$ and $k \geq 1$, define $R^1 := R$ and $R^{k+1} := R^k \circ R$. Note that R^k is primitively positively definable in $(B; R)$.

The following lemma says that we can expand structures by primitive positive definable relations without changing the complexity of the corresponding CSP. Hence, primitive positive definitions are an important tool to prove NP-hardness: to show that $\mathrm{CSP}(\mathfrak{B})$ is NP-hard, it suffices to show that there is a primitive positive definition of a relation R such that $\mathrm{CSP}(\mathfrak{B}, R)$ is already known to be NP-hard. Stronger tools to prove NP-hardness of CSPs will be introduced in Sections 5.7 and 5.9.

Lemma 5.8 (Jeavons, Cohen, Gyssens [66])**.** *Let $\mathfrak{B}$ be a structure with finite relational signature, and let R be a relation that has a primitive positive definition in $\mathfrak{B}$. Then* $\mathrm{CSP}(\mathfrak{B})$ *and* $\mathrm{CSP}(\mathfrak{B}, R)$ *are linear-time equivalent.*

Proof. It is clear that $\mathrm{CSP}(\mathfrak{B})$ reduces to the new problem. So suppose that ϕ is an instance of $\mathrm{CSP}((\mathfrak{B}, R))$. Replace each conjunct $R(x_1, \ldots, x_l)$ of ϕ by its primitive positive definition $\psi(x_1, \ldots, x_l)$. Move all quantifiers to the front, such that the resulting formula is in *prenex normal form* and hence primitive positive. Finally, equalities can be eliminated one by one: for equality $x = y$, remove y from the quantifier prefix, and replace all remaining occurrences of y by x. Let ϕ' be the formula obtained in this way.

It is straightforward to verify that ϕ is true in $(\mathfrak{B}, R)$ if and only if ϕ' is true in $\mathfrak{B}$, and it is also clear that ϕ' can be constructed in linear time in the representation size of ϕ. $\square$

Recall from Section 2.3 that $\mathrm{CSP}(K_5)$ is NP-hard. Since the edge relation of K_5 is primitively positively definable in C_5 (Example 5.5), Lemma 5.8 implies that $\mathrm{CSP}(C_5)$ is NP-hard, too.

Exercices.

92. Let $f\colon A^k \to A$ be an operation. The *graph of* f is the relation

$$G_f := \{(a_1, \dots, a_k, a_0) \mid a_0 = f(a_1, \dots, a_k)\}.$$

Show that a relation is primitively positively definable in the structure $(A; f)$ if and only if it is primitively positive definable in $(A; G_f)$.

93. Show that if $E \subseteq B^2$, then $E^n = E^2$ for some $n \in \mathbb{N}$ if and only if the digraph $(B; E)$ is strongly connected.

94. For a binary relation $R \subseteq A \times B$, define $R^{-1} := \{(b, a) \mid (a, b) \in R\}$. For $n \in \mathbb{N}$, define $R^{-n} := (R^{-1})^n$ (see Definition 5.7). Show that $(R^n)^{-1} = R^{-n}$.

95. Show that $(R \cup R^{-1})^n = B^2$, for some $n \in \mathbb{N}$, if and only if the graph $(B; E)$ is weakly connected.

96. Show that the relation $R := \{(a, b, c) \in \{1, 2, 3\}^3 \mid a = b \text{ or } b = c \text{ or } a = c\}$ has a primitive positive definition over K_3.

97. Show that the relation $\neq$ on $\{1, 2, 3\}$ has a primitive positive definition in the structure $(\{1, 2, 3\}; R, \{1\}, \{2\}, \{3\})$ where R is the relation from the previous exercise.

98. Let R_+, R_* be the relations defined as follows.

$$R_+ := \{(x, y, z) \in \mathbb{Q}^3 \mid x + y = z\}$$
$$R_* := \{(x, y, z) \in \mathbb{Q}^3 \mid x * y = z\}.$$

Show that R_* is primitive positive definable in the structure $(\mathbb{Q}; R_+, \{(x, y) \mid y = x^2\})$.

99. Let B be any set, and for $n \in \mathbb{N}$ define the relation P_B^{2n} of arity $2n$ as follows.

$$P_B^{2n} := \{(x_1, y_1, x_2, y_2, \dots, x_n, y_n) \in B^{2n} \mid \bigvee_{i \in \{1,\dots,n\}} x_i = y_i\}$$

Show that for every n the relation P_B^{2n} has a primitive positive definition in $(B; P_B^4)$.

100. Let $n \geq 4$. Is there a primitive positive definition of $\neq$ over the structure

$$M_n := (\{1, \dots, n\}; R, \{1\}, \{2\}, \dots, \{n\})$$
$$\text{where } R := \{(1, \dots, 1), (2, \dots, 2), \dots, (n, \dots, n), (1, 2, \dots, n)\}?$$

5.6 Cores and Constants

An *automorphism* of a structure $\mathfrak{B}$ with domain B is an isomorphism between $\mathfrak{B}$ and itself. The set of all automorphisms α of $\mathfrak{B}$ is denoted by $\mathrm{Aut}(\mathfrak{B})$, and forms a permutation group. If G is a permutation group on a set B, and $b \in B$, then a set of the form

$$S = \{\alpha(b) \mid \alpha \in G\}$$

is called an *orbit* of G (the *orbit of* b). Let $(b_1, \dots, b_k)$ be a k-tuple of elements of $\mathfrak{B}$. A set of the form

$$S = \{(\alpha b_1, \dots, \alpha b_k) \mid \alpha \in \mathrm{Aut}(\mathfrak{B})\}$$

is called an *orbit of k-tuples* of $\mathfrak{B}$; it is an orbit of the componentwise action of G on the set B^k of k-tuples from B.

Lemma 5.9. *Let $\mathfrak{B}$ be a structure with a finite relational signature and domain B, and let $R = \{(b_1, \dots, b_k)\}$ be a k-ary relation that only contains one tuple $(b_1, \dots, b_k) \in B^k$. If the orbit of $(b_1, \dots, b_k)$ in $\mathfrak{B}$ is primitive positive definable, then there is a polynomial-time reduction from* $\mathrm{CSP}(\mathfrak{B}, R)$ *to* $\mathrm{CSP}(\mathfrak{B})$.

Proof. Let ϕ be an instance of $\mathrm{CSP}(\mathfrak{B}, R)$ with variable set V. If ϕ contains two constraints $R(x_1, \dots, x_k)$ and $R(y_1, \dots, y_k)$, then replace each occurrence of y_1 by x_1, then each occurrence of y_2 by x_2, and so on, and finally each occurrence of y_k by x_k. We repeat this step until all constrains that involve R are imposed on the same tuple of variables $(x_1, \dots, x_k)$. Replace $R(x_1, \dots, x_k)$ by the primitive positive definition θ of its orbit in $\mathfrak{B}$. Finally, move all quantifiers to the front, such that the resulting formula ψ is in prenex normal form and thus an instance of $\mathrm{CSP}(\mathfrak{B})$. Clearly, ψ can be computed from ϕ in polynomial time. We claim that ϕ is true in $(\mathfrak{B}, R)$ if and only if ψ is true in $\mathfrak{B}$.

Suppose ϕ has a solution $s \colon V \to B$. Since $(b_1, \dots, b_k)$ satisfies θ, we can extend s to the existentially quantified variables of θ to obtain a solution for ψ. In the opposite direction, suppose that s' is a solution to ψ over $\mathfrak{B}$. Let s be the restriction of s' to V. Since $(s(x_1), \dots, s(x_k))$ satisfies θ, it lies in the same orbit as $(b_1, \dots, b_k)$. Thus, there exists an automorphism α of $\mathfrak{B}$ that maps $(s(x_1), \dots, s(x_k))$ to $(b_1, \dots, b_k)$. Then the extension of the map $x \mapsto \alpha s(x)$ that maps variables y_i of ϕ that have been replaced by x_i in ψ to the value b_i is a solution to ϕ over $(\mathfrak{B}, R)$. $\square$

The definition of cores can be extended from finite digraphs to finite structures: as in the case of finite digraphs, we require that every endomorphism be an automorphism. All results we proved for cores of digraphs remain valid for cores of structures. In particular, every finite structure $\mathfrak{C}$ is homomorphically equivalent to a core structure $\mathfrak{B}$, which is unique up to isomorphism (see Section 2.4). The following proposition can be shown as in the proof of Proposition 2.9.

Proposition 5.10. *Let $\mathfrak{B}$ be a finite core structure. Then the orbits of k-tuples of $\mathfrak{B}$ are primitive positive definable.*

Proposition 5.10 and Lemma 5.9 have the following consequence.

Corollary 5.11. *Let $\mathfrak{B}$ be a finite core with a finite relational signature. Let $b_1, \dots, b_n \in B$. Then* $\mathrm{CSP}(\mathfrak{B})$ *and* $\mathrm{CSP}(\mathfrak{B}, \{b_1\}, \dots, \{b_n\})$ *are polynomial time equivalent.*

Exercises.

101. Show that if m is the number of orbits of k-tuples of a finite structure $\mathfrak{A}$, and $\mathfrak{C}$ is the core of $\mathfrak{A}$, then $\mathfrak{C}$ has at most m orbits of k-tuples.

102. Show that if $\mathfrak{A}$ is a finite structure, and $\mathfrak{C}$ its core, and if $\mathfrak{A}$ and $\mathfrak{C}$ have the same number of orbits of pairs, then $\mathfrak{A}$ and $\mathfrak{C}$ are isomorphic.

5.7 Primitive Positive Interpretations

Primitive positive interpretations are a powerful generalisation of primitive positive definability that can be used to also relate structures with *different* domains. They are a special case of *(first-order) interpretations* that play an important role in model theory (see, e.g., [64]).

If C and D are sets and $g \colon C \to D$ is a map, then the *kernel* of g is the equivalence relation E on C where $(c, c') \in E$ if $g(c) = g(c')$. For $c \in C$, we denote by c/E the equivalence class of c in E, and by C/E the set of all equivalence classes of elements of C. The *index* of E is defined to be $|C/E|$.

Definition 5.12. Let σ and τ be relational signatures, let $\mathfrak{A}$ be a τ-structure, and let $\mathfrak{B}$ be a σ-structure. A *primitive positive interpretation* I of $\mathfrak{B}$ in $\mathfrak{A}$ consists of

- a natural number d, called the *dimension* of I,
- a primitive positive τ-formula $\delta_I(x_1, \dots, x_d)$, called the *domain formula*,
- for each atomic σ-formula $\phi(y_1, \dots, y_k)$ a primitive positive τ-formula $\phi_I(\overline{x}_1, \dots, \overline{x}_k)$, called the *defining formulas*, and
- the *coordinate map*: a surjective map $h \colon D \to B$ where

$$D := \{(a_1, \dots, a_d) \in A^d \mid \mathfrak{A} \models \delta_I(a_1, \dots, a_d)\}$$

such that for all atomic σ-formulas ϕ and all tuples $\overline{a}_i \in D$

$$\mathfrak{B} \models \phi(h(\overline{a}_1), \dots, h(\overline{a}_k)) \;\Leftrightarrow\; \mathfrak{A} \models \phi_I(\overline{a}_1, \dots, \overline{a}_k)\,.$$

Sometimes, the same symbol is used for the interpretation I and the coordinate map. Note that the dimension d, the set D, and the coordinate map h determine the defining formulas up to logical equivalence; hence, we sometimes denote an interpretation by $I = (d, D, h)$. Note that the kernel of h coincides with the relation defined by $(y_1 = y_2)_I$, for which we also write $=_I$, the defining formula for equality. Also note that the structures $\mathfrak{A}$ and $\mathfrak{B}$ and the coordinate map determine the defining formulas of the interpretation up to logical equivalence.

Example 5.13. Let G be a digraph and let F be an equivalence relation on $V(G)$. Then G/F is the digraph whose vertices are the equivalence classes of F, and where S and T are adjacent if there are $s \in S$ and $t \in T$ such that $\{s, t\} \in E(G)$. If F has a primitive positive definition in G, then G/F has a primitive positive interpretation in G. $\triangle$

Example 5.14. The field of rational numbers $(\mathbb{Q}; 0, 1, +, *)$ has a primitive positive 2-dimensional interpretation I in $(\mathbb{Z}; 0, 1, +, *)$. Example 5.6 presented a primitive positive definition $\phi(x)$ of the set of non-negative integers. The interpretation is now given as follows.

- The domain formula $\delta_I(x, y)$ is $y \geq 1$ (using $\phi(x)$, it is straightforward to express this with a primitive positive formula);
- The formula $=_I (x_1, y_1, x_2, y_2)$ is $x_1 y_2 = x_2 y_1$;
- The formula $0_I(x, y)$ is $x = 0$, the formula $1_I(x, y)$ is $x = y$;
- The formula $+_I(x_1, y_1, x_2, y_2, x_3, y_3)$ is $y_3 * (x_1 * y_2 + x_2 * y_1) = x_3 * y_1 * y_2$;
- The formula $*_I(x_1, y_1, x_2, y_2, x_3, y_3)$ is $x_1 * x_2 * y_3 = x_3 * y_1 * y_2$. △

Theorem 5.15. *Let $\mathfrak{B}$ and $\mathfrak{C}$ be structures with finite relational signatures. If there is a primitive positive interpretation of $\mathfrak{B}$ in $\mathfrak{C}$, then there is a polynomial-time reduction from* CSP$(\mathfrak{B})$ *to* CSP$(\mathfrak{C})$*.*

Proof. Let d be the dimension of the primitive positive interpretation I of the τ-structure $\mathfrak{B}$ in the σ-structure $\mathfrak{C}$, let $\delta_I(x_1, \dots, x_d)$ be the domain formula, and let $h \colon \delta_I(\mathfrak{C}^d) \to B$ be the coordinate map. Let ϕ be an instance of CSP$(\mathfrak{B})$ with variable set $U = \{x_1, \dots, x_n\}$. We construct an instance ψ of CSP$(\mathfrak{C})$ as follows. For distinct variables $V := \{y_1^1, \dots, y_n^d\}$, we set ψ_1 to be the formula

$$\bigwedge_{1 \leq i \leq n} \delta_I(y_i^1, \dots, y_i^d) \; .$$

Let ψ_2 be the conjunction of the formulas $\theta_I(y_{i_1}^1, \dots, y_{i_1}^d, \dots, y_{i_k}^1, \dots, y_{i_k}^d)$ over all conjuncts $\theta = R(x_{i_1}, \dots, x_{i_k})$ of ϕ. By moving existential quantifiers to the front, the sentence

$$\exists y_1^1, \dots, y_n^d \; (\psi_1 \wedge \psi_2)$$

can be re-written to a primitive positive σ-sentence ψ, and clearly ψ can be constructed in polynomial time in the size of ϕ.

We claim that ϕ is true in $\mathfrak{B}$ if and only if ψ is true in $\mathfrak{C}$. Suppose that $f \colon V \to C$ satisfies all conjuncts of ψ in $\mathfrak{C}$. Hence, by construction of ψ, if ϕ has a conjunct $\theta = R(x_{i_1}, \dots, x_{i_k})$, then

$$\mathfrak{C} \models \theta_I\big((f(y_{i_1}^1), \dots, f(y_{i_1}^d)), \dots, (f(y_{i_k}^1), \dots, f(y_{i_k}^d))\big) \; .$$

By the definition of interpretations, this implies that

$$\mathfrak{B} \models R\big(h(f(y_{i_1}^1), \dots, f(y_{i_1}^d)), \dots, h(f(y_{i_k}^1), \dots, f(y_{i_k}^d))\big) \; .$$

Hence, the mapping $g \colon U \to B$ that sends x_i to $h(f(y_i^1), \dots, f(y_i^d))$ satisfies all conjuncts of ϕ in $\mathfrak{B}$.

Now, suppose that $f \colon U \to B$ satisfies all conjuncts of ϕ over $\mathfrak{B}$. Since h is a surjective mapping from $\delta_I(\mathfrak{C}^d)$ to B, there are elements $c_i^1, \dots, c_i^d$ in $\mathfrak{C}$ such that $h(c_i^1, \dots, c_i^d) = f(x_i)$, for all $i \in \{1, \dots, n\}$. We claim that the mapping $g \colon V \to C$ that maps y_i^j to c_i^j satisfies ψ in $\mathfrak{C}$. By construction, any constraint in ψ either comes from ψ_1 or from ψ_2. If it comes from ψ_1 then it must be of the form $\delta_I(y_i^1, \dots, y_i^d)$, and is satisfied since the pre-image of h is $\delta_I(\mathfrak{C}^d)$. If the constraint comes from ψ_2, then it must be a conjunct of a formula

$\theta_I(y^1_{i_1}, \dots, y^d_{i_1}, \dots, y^1_{i_k}, \dots, y^d_{i_k})$ that was introduced for a constraint $\theta = R(x_{i_1}, \dots, x_{i_k})$ in ϕ. It therefore suffices to show that

$$\mathfrak{C} \models \theta_I\big(g(y^1_{i_1}), \dots, g(y^d_{i_1}), \dots, g(y^1_{i_k}), \dots, g(y^d_{i_k})\big)\,.$$

By assumption, $R(f(x_{i_1}), \dots, f(x_{i_k}))$ holds in $\mathfrak{B}$. By the choice of $c^1_1, \dots, c^d_n$, this shows that $R(h(c^1_{i_1}, \dots, c^d_{i_1}), \dots, h(c^1_{i_k}, \dots, c^d_{i_k}))$ holds in $\mathfrak{C}$. By the definition of interpretations, this is the case if and only if $\theta_I(c^1_{i_1}, \dots, c^d_{i_1},\ \dots, c^1_{i_k}, \dots, c^d_{i_k})$ holds in $\mathfrak{C}$, which is what we had to show. $\square$

In many hardness proofs we use Theorem 5.15 in the following way.

Corollary 5.16. *Let $\mathfrak{B}$ be a finite relational structure. If there is a primitive positive interpretation of K_3 in $\mathfrak{B}$, then* $\mathrm{CSP}(\mathfrak{B})$ *is NP-hard.*

Proof. This is a direct consequence of Theorem 5.15 and the fact that $\mathrm{CSP}(K_3)$ is NP-hard (see, e.g., [56]). $\square$

Indeed, K_3 is one of the most expressive finite structures, in the following sense.

Theorem 5.17. *If $n \geq 3$ then every finite structure has a primitive positive interpretation in K_n.*

Proof. Let $\mathfrak{A}$ be a finite τ-structure with the domain $A = \{1, \dots, k\}$. Our interpretation I of $\mathfrak{A}$ in K_n is $2k$-dimensional. The domain formula $\delta_I(x_1, \dots, x_k, x'_1, \dots, x'_k)$ expresses that for exactly one $i \leq k$ we have $x_i = x'_i$. Note that this formula is preserved by all permutations of $\{1, \dots, k\}$. We will see in Proposition 6.19 that every such formula is equivalent to a primitive positive formula over K_n. Equality is interpreted by the formula

$$=_I (x_1, \dots, x_k, x'_1, \dots, x'_k, y_1, \dots, y_k, y'_1, \dots, y'_k) := \bigwedge_{i=1}^{k} \big((x_i = x'_i) \Leftrightarrow (y_i = y'_i)\big)$$

Note that $=_I$ defines an equivalence relation on the set of all $2k$-tuples $(u_1, \dots, u_k, u'_1, \dots, u'_k)$ that satisfy δ_I. The coordinate map sends this tuple to i if and only if $u_i = u'_i$. When $R \in \tau$ is m-ary, then the formula $R(x_1, \dots, x_m)_I$ is any primitive positive formula which is equivalent to the following disjunction of conjunctions with $2mk$ variables $x_{1,1}, \dots, x_{m,k}, x'_{1,1}, \dots, x'_{m,k}$: for each tuple $(t_1, \dots, t_m) \in R^{\mathfrak{A}}$ the disjunction contains the conjunct $\bigwedge_{i \leq m} x_{i,t_i} = x'_{i,t_i}$; again, Proposition 6.19 implies that such a primitive positive formula exists. $\square$

Exercises.

103. Show that the digraph

$$(\{a, b, c, d, e\}; \{(a, b), (b, c), (c, d), (d, e), (b, d), (a, d), (d, e)\})$$

has a pp-interpretation in $(\{0, 1\}; \{0, 1\}^3 \setminus \{(1, 1, 0)\}, \{0\}, \{1\})$, and vice versa.

Hints. Exercise 120. There is a 1-dimensional pp-interpretation of the second structure in the digraph.

5.7.1 Composing Interpretations

Primitive positive interpretations can be composed: if

- $\mathfrak{C}_1$ has a d_1-dimensional pp-interpretation I_1 in $\mathfrak{C}_2$, and
- $\mathfrak{C}_2$ has an d_2-dimensional pp-interpretation I_2 in $\mathfrak{C}_3$,

then $\mathfrak{C}_1$ has a natural $(d_1 d_2)$-dimensional pp-interpretation in $\mathfrak{C}_3$, which we denote by $I_1 \circ I_2$. To formally describe $I_1 \circ I_2$, suppose that the signature of $\mathfrak{C}_i$ is τ_i for $i = 1, 2, 3$, and that $I_1 = (d_1, S_1, h_1)$ and $I_2 = (d_2, S_2, h_2)$. When ϕ is a primitive positive τ_2-formula, let ϕ_{I_2} denote the τ_3-formula obtained from ϕ by replacing each atomic τ_2-formula ψ in ϕ by the τ_3-formula ψ_{I_2}. Note that ϕ_{I_2} is again primitive positive. The coordinate map of $I_1 \circ I_2$ is defined by

$$(a_1^1, \dots, a_{d_2}^1, \dots, a_1^{d_1}, \dots, a_{d_2}^{d_1}) \;\mapsto\; h_1\big(h_2(a_1^1, \dots, a_{d_2}^1), \dots, h_2(a_1^{d_1}, \dots, a_{d_2}^{d_1})\big) \; .$$

Two pp-interpretations I_1 and I_2 of $\mathfrak{B}$ in $\mathfrak{A}$ are called *homotopic*[5] if the relation

$$\{(\bar{x}, \bar{y}) \mid I_1(\bar{x}) = I_2(\bar{y})\}$$

of arity $d_1 + d_2$ is pp-definable in $\mathfrak{A}$. Note that id_C is a pp-interpretation of $\mathfrak{C}$ in $\mathfrak{C}$, called the *identity interpretation* of $\mathfrak{C}$ (in $\mathfrak{C}$).

Definition 5.18. Two structures $\mathfrak{A}$ and $\mathfrak{B}$ with an interpretation I of $\mathfrak{B}$ in $\mathfrak{A}$ and an interpretation J of $\mathfrak{A}$ in $\mathfrak{B}$ are called *mutually pp-interpretable.* If both $I \circ J$ and $J \circ I$ are homotopic to the identity interpretation (of $\mathfrak{A}$ and of $\mathfrak{B}$, respectively), then we say that $\mathfrak{A}$ and $\mathfrak{B}$ are *primitively positively bi-interpretable (via I and J).*

We close this section with a more informative version of Theorem 2.5. It has been conjectured (in slightly different, but equivalent form) by Bulatov, Jeavons, and Krokhin in [39], which is known under the name *tractability conjecture.*

Theorem 5.19 (Tractability Theorem, Version 1)**.** *Let $\mathfrak{B}$ be a relational structure with finite domain and finite signature, and let $\mathfrak{C}$ be the expansion of the core of $\mathfrak{B}$ by all singleton unary relations. If K_3 has a primitive positive interpretation in $\mathfrak{C}$, then* $\mathrm{CSP}(\mathfrak{B})$ *is NP-complete. Otherwise,* $\mathrm{CSP}(\mathfrak{B})$ *is in P.*

Proof. The first part of the theorem easily follows from the results that we have already shown: $\mathfrak{B}$ and its core have the same CSP, and $\mathfrak{C}$ has the same complexity by Lemma 5.9. The first statement then follows from Corollary 5.16. The second statement was shown by Bulatov [36] and by Zhuk [96]. □

A reformulation of this result can be found in Section 5.9.

[5] We follow the terminology from [2].

5.8 Reduction to Binary Signatures

In this section we prove that every structure $\mathfrak{C}$ with a relational signature of maximal arity $m \in \mathbb{N}$ is primitively positively bi-interpretable with a *binary structure* $\mathfrak{B}$, i.e., a relational structure where every relation symbol has arity at most two. Moreover, if $\mathfrak{C}$ has a finite signature, then $\mathfrak{B}$ can be chosen to have a finite signature, too. It follows from Theorem 5.15 that every CSP is polynomial-time equivalent to a binary CSP. This transformation is known under the name *dual encoding* [45, 49]. We want to stress that the transformation works for relational structures with domains of arbitrary cardinality.

A d-dimensional primitive positive interpretation I of $\mathfrak{B}$ in $\mathfrak{A}$ is called *full* if for every $R \subseteq B^k$ we have that R is primitively positively definable in $\mathfrak{B}$ if and only if the relation $I^{-1}(R)$ of arity kd is primitively positively definable in $\mathfrak{A}$. Note that every structure with a primitive positive interpretation in $\mathfrak{A}$ is a reduct of a structure with a full primitive positive interpretation in $\mathfrak{A}$.

Definition 5.20. Let $\mathfrak{C}$ be a structure and $d \in \mathbb{N}$. Then a *d-th full power of* $\mathfrak{C}$ is a structure $\mathfrak{D}$ with domain C^d such that the identity map on C^d is a full d-dimensional primitive positive interpretation of $\mathfrak{D}$ in $\mathfrak{C}$.

In particular, for all $i, j \in \{1, \dots, d\}$ the relation

$$E_{i,j} := \big\{((x_1, \dots, x_d), (y_1, \dots, y_d)) \mid x_1, \dots, x_d, y_1, \dots, y_d \in C \text{ and } x_i = y_j\big\}$$

is primitively positively definable in $\mathfrak{D}$.

Proposition 5.21. *Let $\mathfrak{C}$ be a structure and $\mathfrak{D}$ a d-th full power of $\mathfrak{C}$ for $d \geq 1$. Then $\mathfrak{C}$ and $\mathfrak{D}$ are primitively positively bi-interpretable.*

Proof. Let I be the identity map on C^d which is a full interpretation of $\mathfrak{D}$ in $\mathfrak{C}$. Our interpretation J of $\mathfrak{C}$ in $\mathfrak{D}$ is one-dimensional and the coordinate map is the first projection. The domain formula is *true* and the pre-image of the equality relation in $\mathfrak{C}$ under the coordinate map has the primitive positive definition $E_{1,1}(x, y)$. To define the pre-image of a k-ary relation R of $\mathfrak{C}$ under the coordinate map it suffices to observe that the k-ary relation

$$S := \big\{((a_{1,1}, \dots, a_{1,d}), \dots, (a_{k,1}, \dots, a_{k,d})) \mid (a_{1,1}, \dots, a_{k,1}) \in R\big\}$$

is primitively positively definable in $\mathfrak{D}$ and $J(S) = R$.

To show that $\mathfrak{C}$ and $\mathfrak{D}$ are primitively positive bi-interpretable we prove that $I \circ J$ and $J \circ I$ are pp-homotopic to the identity interpretation. The relation

$$\big\{(u_0, u_1, \dots, u_k) \mid u_0 = I(J(u_1), \dots, J(u_k)), u_1, \dots, u_k \in C^{k+1}\big\}$$

has the primitive positive definition $\bigwedge_{i \in \{1,\dots,k\}} E_{i,1}(u_0, u_i)$ and the relation

$$\big\{(v_0, v_1, \dots, v_k) \mid v_0 = J(I(v_1, \dots, v_k)), v_1, \dots, v_k \in D^{k+1}\big\}$$

has the primitive positive definition $v_0 = v_1$. $\square$

Note that for every relation R of arity $k \leq d$ of $\mathfrak{C}$, in a d-th full power $\mathfrak{D}$ of $\mathfrak{C}$ the unary relation

$$R' := \{(a_1, \dots, a_d) \mid (a_1, \dots, a_k) \in R\}$$

must be primitively positively definable. We now define a particular full power.

Definition 5.22. Let $\mathfrak{C}$ be a relational structure with maximal arity m and let $d \geq m$. Then the structure $\mathfrak{B} := \mathfrak{C}^{[d]}$ with domain C^d is defined as follows:

- for every relation $R \subseteq C^k$ of $\mathfrak{C}$ the structure $\mathfrak{B}$ has the unary relation $R' \subseteq B = C^d$ defined above, and
- for all $i, j \in \{1, \ldots, d\}$ the structure $\mathfrak{B}$ has the binary relation symbol $E_{i,j}$.

It is clear that the signature of $\mathfrak{B}$ is finite if the signature of $\mathfrak{C}$ is finite. Also note that the signature of $\mathfrak{C}^{[d]}$ is always binary.

Lemma 5.23. *Let $\mathfrak{C}$ be a relational structure with maximal arity m and let $d \geq m$. Then the binary structure $\mathfrak{C}^{[d]}$ is a full power of $\mathfrak{C}$.*

Proof. The identity map is a d-dimensional primitive positive interpretation I of $\mathfrak{B} := \mathfrak{C}^{[d]}$ in $\mathfrak{C}$. Our interpretation J of $\mathfrak{C}$ in $\mathfrak{B}$ is one-dimensional and the coordinate map is the first projection. The domain formula is *true* and the pre-image of the equality relation in $\mathfrak{C}$ under the coordinate map has the primitive positive definition $E_{1,1}(x, y)$. The pre-image of the relation R of $\mathfrak{C}$ under the coordinate map is defined by the primitive positive formula

$$\exists y \big(\bigwedge_{i \in \{1,\ldots,k\}} E_{1,i}(x_i, y) \wedge R'(y) \big).$$

The proof that $I \circ J$ and $J \circ I$ are pp-homotopic to the identity interpretation is as in the proof of Proposition 5.21. □

Corollary 5.24. *For every structure $\mathfrak{C}$ with maximal arity m there exists a structure $\mathfrak{B}$ with maximal arity 2 such that $\mathfrak{B}$ and $\mathfrak{C}$ are primitively positively bi-interpretable. If the signature of $\mathfrak{C}$ is finite, then the signature of $\mathfrak{B}$ can be chosen to be finite, too.*

Proof. An immediate consequence of Lemma 5.23 and Proposition 5.21. □

We will revisit primitive positive interpretations in Section 8 where we study them from a universal-algebraic perspective.

5.9 Primitive Positive Constructions

In the previous three sections we have seen several conditions on $\mathfrak{A}$ and $\mathfrak{B}$ that imply that CSP($\mathfrak{A}$) reduces to CSP($\mathfrak{B}$); in this section we compare them. Let $\mathcal{C}$ be a class of structures. We write

1. H($\mathcal{C}$) for the class of structures homomorphically equivalent to structures in $\mathcal{C}$.
2. C($\mathcal{C}$) for the class of all structures obtained by expanding a core structure in $\mathcal{C}$ by singleton relations $\{a\}$. In the setting of relational structures, they play the role of constants (which formally are operation symbols of arity 0).
3. I($\mathcal{C}$) for the class of all structures with a primitive positive interpretation in a structure from $\mathcal{C}$.

Let $\mathcal{D}$ be the smallest class containing $\mathcal{C}$ and closed under H, C, and I. Barto, Opršal, and Pinsker [17] showed that $\mathcal{D} = \mathrm{HI}(\mathcal{C}) := \mathrm{H}(\mathrm{I}(\mathcal{C}))$. In other words, if there is a chain of applications of the three operators H, C, and I to derive $\mathfrak{A}$ from $\mathfrak{B}$, then there is also a two-step chain to derive $\mathfrak{A}$ from $\mathfrak{B}$, namely by interpreting a structure $\mathfrak{B}'$ that is homomorphically equivalent to $\mathfrak{A}$. This insight is conceptually important for the CSP since it leads to a better understanding of the power of the available tools. If $\mathfrak{A} \in \mathrm{HI}(\mathfrak{B})$, then we also say that $\mathfrak{A}$ has a *primitive positive (pp) construction* in $\mathfrak{B}$, following [17].

Proposition 5.25 (from [17]). *Suppose that $\mathfrak{B}$ is a core, and that $\mathfrak{C}$ is the expansion of $\mathfrak{B}$ by a relation of the form $\{c\}$ for $c \in B$. Then $\mathfrak{B}$ pp-constructs $\mathfrak{C}$. In symbols,*

$$\mathrm{C}(\mathcal{C}) \subseteq \mathrm{HI}(\mathcal{C}) .$$

Proof. By Proposition 5.10, the orbit O of c has a primitive positive definition $\phi(x)$ in $\mathfrak{B}$. We give a 2-dimensional primitive positive interpretation in $\mathfrak{B}$ of a structure $\mathfrak{A}$ with the same signature τ as $\mathfrak{C}$. The domain formula $\delta_I(x_1, x_2)$ for $\mathfrak{A}$ is $\phi(x_2)$. Let $R \in \tau$. If R is from the signature of $\mathfrak{B}$ and has arity k then

$$R^{\mathfrak{A}} := \{((a_1, b_1), \ldots, (a_k, b_k)) \in A^k \mid (a_1, \ldots, a_k) \in R^{\mathfrak{B}} \text{ and } b_1 = \cdots = b_k \in O\}.$$

Otherwise, $R^{\mathfrak{C}}$ is of the form $\{c\}$ and we define $R^{\mathfrak{A}} := \{(a, a) \mid a \in O\}$. It is clear that $\mathfrak{A}$ has a primitive positive interpretation in $\mathfrak{B}$.

We claim that $\mathfrak{A}$ and $\mathfrak{C}$ are homomorphically equivalent. The homomorphism from $\mathfrak{C}$ to $\mathfrak{A}$ is given by $a \mapsto (a, c)$:

- if $(a_1, \ldots, a_k) \in R^{\mathfrak{C}} = R^{\mathfrak{B}}$ then $((a_1, c), \ldots, (a_k, c)) \in R^{\mathfrak{A}}$;
- the relation $R^{\mathfrak{C}} = \{c\}$ is preserved since $(c, c) \in R^{\mathfrak{A}}$.

To define a homomorphism h from $\mathfrak{A}$ to $\mathfrak{C}$ we pick for each $a \in O$ an automorphism $\alpha_a \in \mathrm{Aut}(\mathfrak{B})$ such that $\alpha_a(a) = c$. Note that $b \in O$ since $\mathfrak{B} \models \delta(a, b)$, and we define $h(a, b) := \alpha_b(a)$. To check that this is indeed a homomorphism, let $R \in \tau$ be k-ary, and let $t = ((a_1, b_1), \ldots, (a_k, b_k)) \in R^{\mathfrak{A}}$. Then $b_1 = \cdots = b_k =: b \in O$ and we have that $h(t) = (\alpha_b(a_1), \ldots, \alpha_b(a_k))$ is in $R^{\mathfrak{C}}$ since $(a_1, \ldots, a_k) \in R^{\mathfrak{B}} = R^{\mathfrak{C}}$ and α_b preserves $R^{\mathfrak{B}} = R^{\mathfrak{C}}$. If $R^{\mathfrak{A}} = \{(a, a) \mid a \in O\}$, then R is preserved as well, because

$$h((a, a)) = \alpha_a(a) = c \in \{c\} = R^{\mathfrak{C}}.$$

Hence, $\mathfrak{C} \in \mathrm{H}(\mathfrak{A}) \subseteq \mathrm{HI}(\mathfrak{B})$. □

Theorem 5.26 (from [17]). *Suppose that $\mathfrak{A}$ can be obtained from $\mathcal{C}$ by repeatedly applying* H, C, *and* I. *Then $\mathfrak{A} \in \mathrm{HI}(\mathcal{C})$, that is, $\mathfrak{C}$ pp-constructs $\mathfrak{A}$.*

Proof. We have to show that $\mathrm{HI}(\mathcal{C})) =$ is closed under H, C, and I. Homomorphic equivalence is transitive so $\mathrm{H}(\mathrm{H}(\mathcal{C})) \subseteq \mathrm{H}(\mathcal{C})$.

We show that if $\mathfrak{A}$ and $\mathfrak{B}$ are homomorphically equivalent, and $\mathfrak{C}$ has a d-dimensional primitive positive interpretation I_1 in $\mathfrak{B}$, then $\mathfrak{C}$ is homomorphically equivalent to a structure $\mathfrak{D}$ with a d-dimensional primitive positive interpretation I_2 in $\mathfrak{A}$. Let $h_1 \colon A \to B$ be the homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$, and h_2 the homomorphism from $\mathfrak{B}$ to $\mathfrak{A}$. The interpreting formulas of I_2 are the same as the interpreting formulas of I_1; this describes the structure $\mathfrak{D}$

up to isomorphism. We claim that the map $g_1(I_2(a_1,\dots,a_d)) := I_1(h_1(a_1),\dots,h_1(a_d))$ is a homomorphism from $\mathfrak{D}$ to $\mathfrak{C}$. Indeed, for a k-ary relation symbol from the signature of $\mathfrak{C}$ and $\mathfrak{D}$, let $((a^1_1,\dots,a^1_d),\dots,(a^k_1,\dots,a^k_d)) \in R^{\mathfrak{D}}$; hence, the dk-tuple $(a^1_1,\dots,a^1_d,\dots,a^k_1,\dots,a^k_d)$ satisfies the primitive positive defining formula for $R(x^1_1,\dots,x^k_d)$, and

$$(h_1(a^1_1),\dots,h_1(a^1_d),\dots,h_1(a^k_1),\dots,h_1(a^k_d))$$

satisfies this formula, too. This in turn implies that

$$(I_1(h_1(a^1_1),\dots,h_1(a^1_d)),\dots,I_1(h_1(a^k_1),\dots,h_1(a^k_d))) \in R^{\mathfrak{C}}.$$

Similarly, $g_2(I_1(b_1,\dots,b_d)) := I_2(h_2(b_1),\dots,h_2(b_d))$ is a homomorphism from $\mathfrak{C}$ to $\mathfrak{D}$. So we conclude that

$$\mathrm{I}(\mathrm{HI}(\mathcal{C})) \subseteq \mathrm{H}(\mathrm{I}(\mathrm{I}(\mathcal{C}))) \subseteq \mathrm{HI}(\mathcal{C})$$

because primitive positive interpretability is transitive, too. Finally, Proposition 5.25 shows that

$$\mathrm{C}(\mathrm{HI}(\mathcal{C}))) \subseteq \mathrm{HI}(\mathrm{HI}(\mathcal{C})) \subseteq \mathrm{HI}(\mathcal{C})$$

where the last inclusion again follows from the observations above. □

The following example shows that there are finite structures $\mathfrak{B}$ all of whose polymorphisms are idempotent such that $\mathrm{HI}(\mathfrak{B})$ is strictly larger than $\mathrm{I}(\mathfrak{B})$.

Example 5.27. Let $\mathfrak{B}$ be the structure with domain $(\mathbb{Z}_2)^2$ and signature $\{R_{a,b} \mid a,b \in \mathbb{Z}_2\}$ such that

$$R^{\mathfrak{B}}_{a,b} := \{(x,y,z) \in ((\mathbb{Z}_2)^2)^3 \mid x+y+z = (a,b)\}.$$

Let $\mathfrak{B}'$ be the reduct of $\mathfrak{B}$ with the signature $\tau := \{R_{0,0}, R_{1,0}\}$. Let $\mathfrak{A}$ be the τ-structure with domain $\mathbb{Z}_2$ such that for $a = 0$ and $a = 1$

$$R^{\mathfrak{A}}_{a,0} := \{(x,y,z) \in (\mathbb{Z}_2)^3 \mid x+y+z = a\}\,.$$

Now observe that

- $(x_1,x_2) \mapsto x_1$ is a homomorphism from $\mathfrak{B}'$ to $\mathfrak{A}$, and $x \mapsto (x,0)$ is a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}'$. Therefore $\mathfrak{A} \in \mathrm{H}(\mathfrak{B}')$.
- Trivially, $\mathfrak{B}' \in \mathrm{I}(\mathfrak{B})$ and consequently $\mathfrak{A} \in \mathrm{HI}(\mathfrak{B})$.
- All polymorphisms of $\mathfrak{B}$ are idempotent.

We finally show that $\mathfrak{A} \notin \mathrm{I}(\mathfrak{B})$. Suppose for contradiction that there is a pp-interpretation of $\mathfrak{A}$ in $\mathfrak{B}$ with coordinate map $c\colon C \to A$ where $C \subseteq B^n$ is primitive positive definable in $\mathfrak{B}$. The kernel K of c has a primitive positive definition ϕ in $\mathfrak{B}$. The two equivalence classes of K are pp-definable relations over $\mathfrak{B}$, too: the formula $\exists x(\phi(x,y) \wedge R_{a,b}(x))$ defines the equivalence class of (a,b). But the relations with a primitive positive definition in $\mathfrak{B}$ are precisely affine linear subspaces of the vector space $(\mathbb{Z}_2)^2$, so their cardinality must be a power of 4. And two powers of 4 cannot add up to a power of 4. △

Using the operator HI, we reformulate the tractability theorem (Theorem 5.19) as follows.

Theorem 5.28 (Tractability Theorem, Version 2)**.** *Let $\mathfrak{B}$ be a relational structure with finite domain and finite signature. If $K_3 \in \mathrm{HI}(\mathfrak{B})$, then $\mathrm{CSP}(\mathfrak{B})$ is NP-complete. Otherwise, $\mathrm{CSP}(\mathfrak{B})$ is in P.*

Proof. If $K_3 \in \mathrm{HI}(\mathfrak{B})$, then the NP-hardness of $\mathrm{CSP}(\mathfrak{B})$ follows from Corollary 5.16. Otherwise, let $\mathfrak{C}$ be the expansion of $\mathfrak{B}$ by all singleton relations. Note that $\mathfrak{C} \in \mathrm{HI}(\mathfrak{B})$ by Proposition 5.25. Hence, if $K_3 \in \mathrm{HI}(\mathfrak{C})$, then $K_3 \in \mathfrak{B}$ by Theorem 5.26, a contradiction. Hence, Theorem 2.5 implies that $\mathrm{CSP}(\mathfrak{C})$ is in P, and therefore $\mathrm{CSP}(\mathfrak{B})$ is in P. $\square$

Assuming that $\mathrm{P} \neq \mathrm{NP}$, it follows that $K_3 \in \mathrm{HI}(\mathfrak{B})$ if and only if K_3 has a primitive positive interpretation in the expansion of the core of $\mathfrak{B}$ by all singleton unary relations; we will see a proof of this fact without complexity-theoretic assumptions (Corollary 9.18).

We will revisit primitive positive constructions in Section 9 where we study them from a universal-algebraic perspective; in particular, then next reformulation of the tractability conjecture can be found in Section 9.5.

Exercises.

103. Prove that $\vec{C}_6$ pp-constructs $\vec{C}_3$.

104. Prove that $\vec{C}_2 \uplus \vec{C}_3$ pp-constructs $\vec{C}_6$.

105. Prove that $\vec{C}_3$ pp-constructs $\vec{C}_9$.

6 Relations and Operations

In this section we introduce *operation clones*. Most of our results concern operation clones on a *finite* domain; however, some results can naturally be proved for arbitrary domains without extra effort and we of course then state the general results.

6.1 Operation Clones

For $n \geq 1$ and a set D (the *domain*), denote by $\mathscr{O}_D^{(n)}$ the set $D^{D^n} := (D^n \to D)$ of n-ary functions on D. The elements of $\mathscr{O}_D^{(n)}$ will typically be called the *operations* of arity n on D, and D will be called the *domain*. The set of all operations on D of finite arity will be denoted by $\mathscr{O}_D := \bigcup_{n \geq 1} \mathscr{O}_D^{(n)}$. An *operation clone* (over D) is a subset $\mathscr{C}$ of $\mathscr{O}_D$ satisfying the following two properties:

- $\mathscr{C}$ contains all projections, that is, for all $1 \leq k \leq n$ it contains the operation $\pi_k^n \in \mathscr{O}_D^{(n)}$ defined by $\pi_k^n(x_1, \ldots, x_n) = x_k$, and
- $\mathscr{C}$ is *closed under composition*, that is, for all $f \in \mathscr{C} \cap \mathscr{O}_D^{(n)}$ and $g_1, \ldots, g_n \in \mathscr{C} \cap \mathscr{O}_D^{(m)}$ it contains the operation $f(g_1, \ldots, g_n) \in \mathscr{O}_D^{(m)}$ defined by

$$(x_1, \ldots, x_m) \mapsto f(g_1(x_1, \ldots, x_m), \ldots, g_n(x_1, \ldots, x_m)) .$$

A *clone* is an abstraction of an operation clone that will be introduced later in the course. In the literature, operation clones are often called clones, or *concrete clones*; we prefer to use the terms 'operation clone' and 'clone' in analogy to 'permutation group' and 'group'.

If $\mathscr{C}$ is an operation clone, then $\mathscr{C}'$ is called a *subclone* of $\mathscr{C}$ if $\mathscr{C}'$ is an operation clone and $\mathscr{C}' \subseteq \mathscr{C}$. If $\mathscr{F}$ is a set of functions, we write $\langle \mathscr{F} \rangle$ for the smallest operation clone $\mathscr{C}$ which contains $\mathscr{F}$, and call $\mathscr{C}$ the clone *generated* by $\mathscr{F}$; similarly, we also say that the elements of $\mathscr{C}$ are generated by $\mathscr{F}$. Note that the set of all clones over a set B forms a lattice: the meet of two operation clones $\mathscr{C}$ and $\mathscr{D}$ is their intersection $\mathscr{C} \cap \mathscr{D}$ (which is again a clone!); the join of $\mathscr{C}$ and $\mathscr{D}$ is the clone generated by their union, $\langle \mathscr{C} \cup \mathscr{D} \rangle$.

Remark 6.1. Clones on a two-element set have been classified by Post [86]; the set of such clones is countably infinite. In contrast, there are 2^ω many clones over the set $\{0, 1, 2\}$ [93].

6.2 Inv-Pol

The most important source of operation clones in this text are *polymorphism clones* of digraphs and, more generally, structures. For simplicity, we only discuss relational structures; the step to structures that also involve function symbols is straightforward.

Let f be from $\mathscr{O}_B^{(n)}$, and let $R \subseteq B^m$ be a relation. Then we say that f *preserves* R (and that R is *invariant under* f) if $f(r_1, \dots, r_n) \in R$ whenever $r_1, \dots, r_n \in R$, where $f(r_1, \dots, r_n)$ is calculated componentwise. If $\mathfrak{B}$ is a relational structure with domain B then $\mathrm{Pol}(\mathfrak{B})$ contains precisely those operations that preserve $\mathfrak{B}$.

Observation 6.2. $\mathrm{Pol}(\mathfrak{B})$ is an operation clone.

Conversely, if $\mathscr{F}$ is a set of operations on B, then we write $\mathrm{Inv}(\mathscr{F})$ for the set of all relations on B that are invariant under all functions in $\mathscr{F}$. It will be convenient to define the operator Pol also for sets $\mathcal{R}$ of relations over B, writing $\mathrm{Pol}(\mathcal{R})$ for the set of operations of $\mathscr{O}_B$ that preserve all relations from $\mathcal{R}$.

Proposition 6.3. *Let $\mathscr{F}$ be a set of operations on a set B. Then $\langle \mathscr{F} \rangle \subseteq \mathrm{Pol}(\mathrm{Inv}(\mathscr{F}))$.*

Proposition 6.4. *Let $\mathscr{F}$ be a set of operations on a finite set B. Then $\mathrm{Pol}(\mathrm{Inv}(\mathscr{F})) = \langle \mathscr{F} \rangle$.*

Proof. Exercise 112. □

Proposition 6.5. *Let $\mathfrak{B}$ be any relational structure. Then $\mathrm{Inv}(\mathrm{Pol}(\mathfrak{B}))$ contains the set of all relations that have a primitive positive definition in $\mathfrak{B}$.*

Proof. Suppose that R is k-ary, has a primitive positive definition $\psi(x_1, \dots, x_k)$, and let f be an l-ary polymorphism of $\mathfrak{B}$. To show that f preserves R, let $t_1, \dots, t_l$ be k-tuples from R. Let $x_{k+1}, \dots, x_n$ be the existentially quantified variables of ψ. Write s_i for the n-tuple which extends the k-tuple t_i such that s_i satisfies the quantifier-free part $\psi'(x_1, \dots, x_k, x_{k+1}, \dots, x_n)$ of ψ. Then the tuple $f(s_1, \dots, s_l)$ satisfies ψ' since f is a polymorphism. This shows that $\mathfrak{B} \models \psi(f(t_1, \dots, t_l))$ which is what we had to show. □

Note that Proposition 6.5 also holds (and is useful!) for structures with an infinite domain; see, e.g., Exercise 111.

Theorem 6.6 (of [32,57])**.** *Let $\mathfrak{B}$ be a finite relational structure. A relation R has a primitive positive definition in $\mathfrak{B}$ if and only if R is preserved by all polymorphisms of $\mathfrak{B}$.*

Proof. One direction has been shown in Proposition 6.5. For the other direction, let $a^1, \dots, a^w$ be an enumeration of R. Let $b_1 = (b_1^1, \dots, b_1^w), b_2 = (b_2^1, \dots, b_2^w), \dots, b_\ell = (b_\ell^1, \dots, b_\ell^w)$ be an enumeration of B^w. Let ϕ be the quantifier-free part of the canonical query of $\mathfrak{B}^w$ (see Exercise 83 for the definition of $\mathfrak{B}^w$ and Section 5.3 for the definition of canonical queries). Note that for every $i \in [k]$ there exists $j_i \in [w]$ such that $(a_i^1, \dots, a_i^w) = b_{j_i}$.

We claim that

$$\psi(x_1, \dots, x_k) := \exists b_1, \dots, b_\ell (\phi \wedge \bigwedge_{i \in [k]} x_i = b_{j_i})$$

is a primitive positive definition of R.

We first show that a^j satisfies ψ for every $j \in [w]$. The elements $b_1^j, \dots, b_\ell^j \in B$ provide witnesses for the existentially quantified variables showing that $a^j = (b_{j_1}^j, \dots, b_{j_k}^j)$ satisfies ψ.

Conversely, suppose that $(t_1, \dots, t_k)$ satisfies ψ. The witnesses for the existentially quantified variables $b_1, \dots, b_\ell$ define a homomorphism f from $\mathfrak{B}^w$ to $\mathfrak{B}$. Since ψ contains the conjuncts $x_i = b_{j_i}$, for $i \in [k]$, we have that $t_i = f(b^1, \dots, b^w)_{j_i}$. Note that f is a polymorphism of $\mathfrak{B}$ and by assumption preserves R. Since the tuples $(a^1, \dots, a^w)$ are from R and

$$f(a^1, \dots, a^w) = (f(b^1, \dots, b^w)_{j_1}, \dots, f(b^1, \dots, b^w)_{j_k}) = (t_1, \dots, t_k)$$

we obtain that $(t_1, \dots, t_k) \in R$. □

Corollary 6.7. *The complexity of* CSP($\mathfrak{B}$) *only depends on* Pol($\mathfrak{B}$). *If* $\mathfrak{C}$ *is such that* Pol($\mathfrak{B}$) $\subseteq$ Pol($\mathfrak{C}$), *then* CSP($\mathfrak{C}$) *reduces in linear time to* CSP($\mathfrak{B}$).

Proof. Direct consequence of Theorem 6.6 and Lemma 5.8. □

Remark 6.8. One direction in Theorem 6.6 is false in general for infinite structures; there are e.g. infinite digraphs $\mathfrak{B}$ that are rigid cores and projective, so Pol($\mathfrak{B}$) has uncountably many invariant relations; in particular, many of these relations do not have a primitive positive definition in $\mathfrak{B}$ because there are only countably many primitive positive formulas over the signature of graphs. However, there is a modified version of the theorem, where primitive positive definitions are replaced by formulas that additionally allow to form unions of chains of relations and infinite intersections; see [21]. Theorem 6.6 is true without modification if the structure $\mathfrak{B}$ is countably infinite and *ω-categorical* [26]. There are also many other infinite structures where Theorem 6.6 remains true; we given an example below.

Example 6.9. Let $\mathbf{F}$ be a field. Let R_+ be the graph $\{(x, y, z) \in F^3 \mid x + y = z\}$ of the addition in $\mathfrak{F}$, and for $\alpha \in F$ let S_α be the binary relation $\{(x, \alpha x) \mid x \in F\}$. We write $\mathfrak{F}$ for the structure $(F; R_+, (S_\lambda)_{\lambda \in F})$ (so this structure can be viewed as a relational version of the reduct of $\mathbf{F}$ in the signature of modules; see Section 8.3).

The following statement is true for arbitrary fields, but already interesting for finite fields, and will be used in later sections.

$$
\begin{aligned}
\langle \mathfrak{F} \rangle &= \operatorname{Inv}(\operatorname{Pol}(\mathfrak{F})) && (2)\\
&= \operatorname{Inv}(\{(x,y) \mapsto x+y\} \cup \{x \mapsto \alpha x \mid \alpha \in F\}) && (3)\\
&= \big\{R \mid n \in \mathbb{N}, R \text{ linear subspace of } \mathbf{F}^n\big\} \cup \{\emptyset\} && (4)\\
&= \big\{\{x \in F^n \mid Ax = 0\} \mid n,m \in \mathbb{N}, A \in F^{m\times n}\big\} \cup \{\emptyset\} && (5)\\
\operatorname{Pol}(\mathfrak{F}) &= \langle \{(x,y) \mapsto x+y\} \cup \{x \mapsto \alpha x \mid \alpha \in F\}\rangle && (6)\\
&= \Big\{(x_1,\dots,x_k) \mapsto \sum_{i=1}^{k} \alpha_i x_i \mid k \geq 1, \alpha_1,\dots,\alpha_k \in F\Big\} && (7)
\end{aligned}
$$

We prove (2), (3), (4), and (5) by showing inclusions in cyclic order. The inclusion $\langle \mathfrak{F} \rangle \subseteq \operatorname{Inv}(\operatorname{Pol}(\mathfrak{F}))$ follows from Proposition 6.5.

Clearly, the relation R_+ is preserved by $(x,y) \mapsto x+y$ and preserved by $x \mapsto \alpha x$ for every $\alpha \in F$. Similarly, for every $\lambda \in F$ the relation S_λ is preserved by these operations. Hence,

$$
\{(x,y) \mapsto x+y\} \cup \{x \mapsto \alpha x \mid \alpha \in F\} \subseteq \operatorname{Pol}(\mathfrak{F}). \tag{8}
$$

Since the Galois-connection Inv-Pol is antitone, we obtain the left-to-right inclusion for (3). The equality (4) is by definition of linear subspaces.

For the left-to-right inclusion of (5), suppose that U is a linear subspace of $\mathbf{F}^n$ and that $(u_1,\dots,u_m)$ is a basis of U. By Steinitz' theorem there exists a basis B of $\mathbf{F}^n$ of the form $(u_1,\dots,u_m,u_{m+1},\dots,u_n)$. Let T be the basis change matrix which maps u_i to $e_i := (0,\dots,0,\overset{i}{1},0,\dots,0)$, which can be written as $T = \begin{pmatrix} X \\ R \end{pmatrix}$ where $X \in \mathbb{K}^{(m,n)}$ and $R \in \mathbb{K}^{(n-m,n)}$. In the following, if $S \subseteq F^n$, we write $\langle S \rangle := \big\{\sum_{i \in \{1,\dots,m\}} \alpha_i u_i \mid u_1,\dots,u_m \in S, \alpha_1,\dots,\alpha_m \in F\big\}$ for the linear hull of S (i.e., for the smallest subalgebra of $\mathbf{F}^n$ that contains S, in the signature of modules). Then

$$
\begin{aligned}
v \in U &\Leftrightarrow v \in \langle u_1,\dots,u_m\rangle\\
&\Leftrightarrow Tv \in \langle Tu_1,\dots,Tu_m\rangle = \langle e_1,\dots,e_m\rangle\\
&\Leftrightarrow (Tv)_{m+1} = \dots = (Tv)_n = 0\\
&\Leftrightarrow Rv = 0\\
&\Leftrightarrow v \in \{x \mid Rx = 0\}.
\end{aligned}
$$

Finally, it is a good exercise to write primitive positive definitions of the solution sets of homogeneous linear equations systems over the structure $\mathfrak{F}$, which closes the chain of implications (Exercise 106).

We prove (6) and (7) by showing inclusions in cyclic order. The left-to-right inclusion in (6) follows from (8) and Observation 6.2. Clearly, every operation that can be composed from the operations in $\{(x,y) \mapsto x+y\} \cup \{x \mapsto \alpha x \mid \alpha \in F\}$ can be written in the form $(x_1,\dots,x_k) \mapsto \sum_{i=1}^k \alpha_i x_i$ for some $k \geq 1$ and some elements $\alpha_1,\dots,\alpha_k \in F$, and this shows the left-to-right inclusion in (7). Finally, these operations preserve addition and multiplication with scalars, which closes the chain of inclusions. △

Example 6.10. Using the same notation as in the previous example, let $\mathfrak{G}$ be the expansion of $\mathfrak{F}$ by all unary relations of the form $\{\alpha\}$ for $\alpha \in F$.

$$
\begin{aligned}
\langle \mathfrak{G} \rangle &= \operatorname{Inv}\big(\operatorname{Pol}(\mathfrak{G})\big) && (9)\\
&= \operatorname{Inv}(\{(x,y,z) \mapsto x - y + z\} \cup \{(x,y) \mapsto \alpha_1 x + \alpha_2 y \mid \alpha_1 + \alpha_2 = 1\}) && (10)\\
&= \big\{R \mid n \in \mathbb{N}, R \text{ affine subspace of } \mathbf{F}^n\big\} \cup \{\emptyset\} && (11)\\
&= \big\{\{x \in F^n \mid Ax = b\} \mid n, m \in \mathbb{N}, A \in F^{m \times n}, b \in F^m\big\} && (12)\\
\operatorname{Pol}(\mathfrak{G}) &= \langle \{(x,y,z) \mapsto x - y + z\} \cup \{(x,y) \mapsto \alpha_1 x + \alpha_2 y \mid \alpha_1 + \alpha_2 = 1\} \rangle && (13)\\
&= \big\{(x_1, \dots, x_n) \mapsto \sum_{i=1}^n \alpha_i x_i \mid n \geq 1, \alpha_1 + \dots + \alpha_n = 1\big\} && (14)
\end{aligned}
$$

We show the inclusions in cyclic order. The inclusion $\langle \mathfrak{G} \rangle \subseteq \operatorname{Inv}(\operatorname{Pol}(\mathfrak{G}))$ follows from Proposition 6.5.

Clearly, the operation $m \colon F^3 \to F$ given by $(x,y,z) \mapsto x - y + z$ clearly preserves not only the relations from $\mathfrak{F}$, but also all unary relations of the form $\{\alpha\}$ for $\alpha \in F$, so $m \in \operatorname{Pol}(\mathfrak{G})$. Similarly, we verify that $(x,y) \mapsto \alpha_1 x + \alpha_2 y \in \operatorname{Pol}(\mathfrak{G})$ whenever $\alpha_1 + \alpha_2 = 1$. Since the Galois-connection Inv-Pol is antitone, we obtain the left-to-right inclusion for (10).

For the left-to-right inclusion in (11), let $R \in \operatorname{Inv}(\{m\})$. If $R = \emptyset$ there is nothing to be shown. Otherwise, let $o \in R$. We have to show that $S := \{v - o \mid o \in R\}$ is a linear subspace. Let $u_1, u_2 \in S$. Then

$$u_1 + u_2 = v_1 - o + v_2 - o = (v_1 - o + v_2) - o = \underbrace{m(v_1, o, v_2)}_{\in R} - o \in S.$$

To see that S is invariant under scalar multiplication, let $u \in S$ and $\alpha \in F$. By definition of S there exists $v \in R$ such that $u = v - o$. Then

$$\alpha u = \alpha(v - o) = \underbrace{(\alpha v + (1 - \alpha) o)}_{\in R} - o \in S.$$

For the left-to-right inclusion of (12), let R be an affine subspace of $\mathbf{F}^n$. Then $R = \{w + u \mid u \in U\}$ for a linear subspace U of V. In Example 6.9 we have seen that there exists $A \in F^{m \times n}$ with $U = \{x \mid Ax = 0\}$. Then

$$R = \{w + u \mid Au = 0\} = \{w' \mid A(w' - w) = 0\} = \{x \mid Ax = Aw\},$$

so R is the solution set of a system of linear equations. Since $\emptyset$ is the solution space of the unsatisfiable equation system $\{0 = 1\}$, this completes the proof.

Similarly as in Example 6.9, it is a good exercise to find a primitive positive definition of the solution space of a system of linear equations $Ax = b$ in $\mathfrak{G}$, which closes the chain of inclusions.

We prove (13) and (14) by showing inclusions in cyclic order. The left-to-right inclusion in (14) follows from the fact that all relations of $\mathfrak{G}$ are preserved by m and by $(x,y) \mapsto \alpha_1 x + \alpha_2 y$ whenever $\alpha_1, \alpha_2 \in F$ are such that $\alpha_1 + \alpha_2 = 1$, and Observation 6.2. Clearly, every operation that can be composed from these operations can be written in the form $(x_1, \dots, x_n) \mapsto \sum_{i=1}^n \alpha_i x_i$ for some $k \geq 1$ and some elements $\alpha_1, \dots, \alpha_n \in F$ such that $\alpha_1 + \dots + \alpha_n = 1$, and this shows the left-to-right inclusion in (7). Finally, these operations preserve all relations of $\mathfrak{G}$, which closes the chain of inclusions. △

Exercises.

106. Let $n, m \in \mathbb{N}$, $A \in F^{m \times n}$, $b \in F^m$. Give a primitive positive definition of $\{x \in F^n \mid Ax = b\}$ in the structure $\mathfrak{F}$ from Example 6.9.

107. Show that the relation $\neq$ is not primitively positively definable in the graph C_6 (the undirected cycle with 6 vertices).

108. For an operation $f\colon A^k \to A$ and a relation R on A, we write $\langle R \rangle_f$ for the smallest relation that contains R and is preserved by f. Similarly, if $\mathscr{F}$ is a set of operations, we write $\langle R \rangle_{\mathscr{F}}$ for the smallest relation that contains R and is preserved by all operations of $\mathscr{F}$. Show that if $\mathfrak{A}$ is a structure with a finite domain, then $\langle R \rangle_{\mathrm{Pol}(\mathfrak{A})}$ equals the smallest relation that contains R and has a primitive positive definition over $\mathfrak{A}$.

109. Show that (6), (7), (4), and (14) also hold if $\mathbf{F}$ is a ring rather than a field.

110. Show that (5) fails in general if $\mathbf{F}$ is a ring (Example 8.2) rather than a field.

111. Let R_+ and R_* be the relations as defined in Exercise 98. Show that R_* is not primitively positively definable in the structure $(\mathbb{Q}; R_+, \{(x, y) \mid y \geq x^2\})$.

112. Prove Proposition 6.4.

113. Find a digraph with the properties described in Remark 6.8.

6.3 Essentially Unary Clones

An operation $f\colon B^k \to B$ is called *essentially unary* if there is an $i \in \{1, \dots, k\}$ and a unary operation f_0 such that $f(x_1, \dots, x_k) = f_0(x_i)$ for all $x_1, \dots, x_k \in B$. Operations that are not essentially unary are called *essential.*[6] We say that f *depends on argument* i if there are $r, s \in B^k$ such that $f(r) \neq f(s)$ and $r_j = s_j$ for all $j \in \{1, \dots, k\} \setminus \{i\}$.

Lemma 6.11. *Let $f \in \mathscr{O}_B$ be an operation. Then the following are equivalent.*

1. *f is essentially unary.*
2. *f preserves $P_B^3 := \big\{(a, b, c) \in B^3 \mid a = b$ or $b = c\big\}$.*
3. *f preserves $P_B^4 := \big\{(a, b, c, d) \in B^4 \mid a = b$ or $c = d\big\}$.*
4. *f depends on at most one argument.*

Proof. Let k be the arity of f. The implication from (1) to (2) is obvious, since unary operations clearly preserve P_B^3.

To show the implication from (2) to (3), we show the contrapositive, and assume that f violates P_B^4. By permuting arguments of f, we can assume that there are 4-tuples $a^1, \dots, a^k \in P_B^4$ with $f(a^1, \dots, a^k) \notin P_B^4$ and $l \leq k$ such that in $a^1, \dots, a^l$ the first two coordinates are equal,

[6] This is standard in clone theory, and it makes sense also when studying the complexity of CSPs, since the essential operations are those that are essential for complexity classification.

and in $a^{l+1}, \ldots, a^k$ the last two coordinates are equal. Let $c := (a_1^1, \ldots, a_1^l, a_4^{l+1}, \ldots, a_4^k)$. Since $f(a^1, \ldots, a^k) \notin P_B^4$ we have $f(a_1^1, \ldots, a_1^k) \neq f(a_2^1, \ldots, a_2^k)$, and therefore $f(c) \neq f(a_1^1, \ldots, a_1^k)$ or $f(c) \neq f(a_2^1, \ldots, a_2^k)$. Let $d = (a_1^1, \ldots, a_1^k)$ in the first case, and $d = (a_2^1, \ldots, a_2^k)$ in the second case. Likewise, we have $f(c) \neq f(a_3^1, \ldots, a_3^k)$ or $f(c) \neq f(a_4^1, \ldots, a_4^k)$, and let $e = (a_3^1, \ldots, a_3^k)$ in the first, and $e = (a_4^1, \ldots, a_4^k)$ in the second case. Then for each $i \leq k$, the tuple (d_i, c_i, e_i) is from P_B^3, but $(f(d), f(c), f(e)) \notin P_B^3$.

The proof of the implication from (3) to (4) is again by contraposition. Suppose f depends on the i-th and j-th argument, $1 \leq i \neq j \leq k$. Hence there exist tuples $a_1, b_1, a_2, b_2 \in B^k$ such that a_1, b_1 and a_2, b_2 only differ at the entries i and j, respectively, and such that $f(a_1) \neq f(b_1)$ and $f(a_2) \neq f(b_2)$. Then $(a_1(l), b_1(l), a_2(l), b_2(l)) \in P_B^4$ for all $l \leq k$, but $(f(a_1), f(b_1), f(a_2), f(b_2)) \notin P_B^4$, which shows that f violates P_B^4.

For the implication from (4) to (1), suppose that f depends only on the first argument. Let $i \leq k$ be minimal such that there is an operation g with $f(x_1, \ldots, x_k) = g(x_1, \ldots, x_i)$. If $i = 1$ then f is essentially unary and we are done. Otherwise, observe that since f does not depend on the i-th argument, neither does g, and so there is an $(i-1)$-ary operation g' such that for all $x_1, \ldots, x_n \in B$ we have $f(x_1, \ldots, x_n) = g(x_1, \ldots, x_i) = g'(x_1, \ldots, x_{i-1})$, contradicting the choice of i. $\square$

6.4 Minimal Clones

A *trivial* clone is a clone all of whose operations are projections. Note that it follows from Lemma 6.11 that for any set $B = \{b_1, \ldots, b_n\}$ the clone $\mathrm{Pol}(B; P_B^4, \{b_1\}, \ldots, \{b_n\})$ is trivial.

Definition 6.12. A clone $\mathscr{C}$ is *minimal* if it is non-trivial, and for every non-trivial $\mathscr{E} \subseteq \mathscr{C}$ we have $\mathscr{E} = \mathscr{C}$.

Recall that $\langle \mathscr{F} \rangle$ denotes the smallest clone that contains $\mathscr{F}$. If $g \in \langle \{f\} \rangle$, then we say that f *generates* g.

Definition 6.13. An operation $f \in \mathscr{O}_B$ is *minimal* if f is not a projection and of minimal arity such that every g generated by f is either a projection or generates f.

The following is straightforward from the definitions.

Proposition 6.14. *Every minimal f generates a minimal clone, and every minimal clone is generated by a minimal operation.*

Theorem 6.15. *Every non-trivial operation clone $\mathscr{C} \subseteq \mathscr{O}_B$ over a finite set B contains a minimal operation.*

Proof. Consider the set of all clones contained in $\mathscr{C}$, partially ordered by inclusion. From this poset we remove the trivial clone; the resulting poset will be denoted by P. We use Zorn's lemma to show that P contains a minimal element. Observe that in P, all chains $(\mathscr{C}_i)_{i \in \kappa}$ that are *descending*, i.e., $\mathscr{C}_i \supseteq \mathscr{C}_j$ for $i < j$, are *bounded*, i.e., for all such chains there exists a $\mathscr{D} \in P$ such that $\mathscr{C}_i \supseteq \mathscr{D}$ for all $i \in \kappa$. To see this, observe that the set $\bigcup_{i \in \kappa} \mathrm{Inv}(\mathscr{C}_i)$ is *closed under primitive positive definability* in the sense that it is the set of relations that is primitively positively definable over some relational structure $\mathfrak{B}$ (since only a finite number of relations can be mentioned in a formula, and since $\mathrm{Inv}(\mathscr{C}_i)$ is closed under primitive positive definability, for each $i \in \kappa$). Moreover, one of the relations $P_B^4, \{b_1\}, \ldots, \{b_n\}$, for $B = \{b_1, \ldots, b_n\}$, is not contained in $\bigcup_{i \in \kappa} \mathrm{Inv}(\mathscr{C}_i)$; otherwise, there would be a $j \in \kappa$ such that $\mathrm{Inv}(\mathscr{C}_j)$ contains all these relations, and hence $\mathscr{C}_j$ is the trivial clone contrary to our assumptions. Hence, $\mathrm{Pol}(\mathfrak{B})$ is a non-trivial lower bound of the descending chain $(\mathscr{C}_i)_{i\in\kappa}$. By Zorn's lemma, P contains a minimal element, and this element contains a minimal operation in $\mathscr{C}$. $\square$

Remark 6.16. Note that the statement above would be false if B is infinite: take for example the clone over the domain $B := \mathbb{N}$ of the integers generated by the operation $x \mapsto x+1$. Every operation in this clone is essentially unary, and every unary operation in this clone is of the form $x \mapsto x+c$ for $c \in \mathbb{N}$. Note that for $c > 0$, the operation $x \mapsto x+c$ generates $x \mapsto x+2c$, but not vice versa, so the clone does not contain a minimal operation.

In the remainder of this section, we show that a minimal operation has one out of the following five types, due to Rosenberg [87]. An n-ary operation f is called a *semiprojection* if there exists an $i \leq n$ such that $f(x_1,\dots,x_n) = x_i$ whenever $|\{x_1,\dots,x_n\}| < n$. For the purpose of proving the next lemma, we call an n-ary operation f a *weak semiprojection* if for all distinct $i,j \in \{1,\dots,n\}$ there exists an index $s(i,j)$ such that

$$\forall x_1,\dots,x_n \colon f(x_1,\dots,x_n) = x_{s(i,j)}$$

holds whenever x_i and x_j are the same variable. Note that in this case, f is a semiprojection if and only if $s(i,j)$ is constant.

In the proof of the following lemma the following notation for weak semiprojections will be practical. Let f be a weak semiprojection, let $S \subseteq \{1,\dots,n\}$ be of cardinality at least two, and let $(x_1,\dots,x_n)$ be a tuple of variables such that $x_i = x_j$ for all $i,j \in S$. Then for some $k \in \{1,\dots,n\}$ it holds that $f(x_1,\dots,x_n) = x_k$. If $k \in S$ define $E(S) := S$. Otherwise, define $E(S) := \{k\}$. Note that if $S \subseteq T \subseteq \{1,\dots,n\}$, then $E(S) \subseteq E(T)$. Also note that if there exists a $k \in \{1,\dots,n\}$ such that $k \in E(S)$ for every $S \subseteq \{1,\dots,n\}$ with at least two elements, then f is a semiprojection.

Lemma 6.17 (Świerczkowski)**.** *Let f be a weak semiprojection of arity at least $n \geq 4$. Then f is a semiprojection.*

Proof. We first show that $E(\{1,2\}) \cap E(\{3,4\}) \neq \emptyset$. If $E(\{1,2,3,4\}) = \{\ell\}$ for some $\ell \notin \{1,2,3,4\}$, then $E(\{1,2\}) = \{\ell\} = E(\{3,4\})$ and we are done. So we assume that $E(\{1,2,3,4\}) = \{1,2,3,4\}$. First consider the case that $E(\{1,2\}) = \{i\} \subseteq \{3,4\}$ so that $f(x,x,y,y,x_5,\dots,x_n) = y$. If $E(\{3,4\}) = \{j\} \subseteq \{1,2\}$ then and $f(x,x,y,y,x_5,\dots,x_n) = x$ for $i \neq j$, which is a contradiction since $|B| \geq 2$. Hence, $E(\{3,4\}) = \{3,4\}$ and we have found $i \in E(\{1,2\}) \cap E(\{3,4\})$. Similarly we can treat the case that $E(\{3,4\}) = \{i\} \subseteq \{1,2\}$. If $E(\{1,2\}) = \{1,2\}$ and $E(\{3,4\}) = \{3,4\}$ then $f(x,x,y,y,x_5,\dots,x_n) = x$ because of $E(\{1,2\}) \subseteq \{1,2\}$ and $f(x,x,y,y,x_5,\dots,x_n) = y$ because of $E(\{3,4\}) \subseteq \{3,4\}$, a contradiction.

Let $i \in E(\{1,2\}) \cap E(\{3,4\})$. Note that if $i \notin \{1,2\}$, then $E(\{1,2\}) = \{i\}$. Similarly, if $i \notin \{3,4\}$ then $E(\{3,4\}) = \{i\}$. We therefore have a set $S \subseteq \{1,\dots,n\}$ of size two with $E(S) = \{i\}$. Let $T \subseteq \{1,\dots,n\}$ be of cardinality at least two. We will show that $i \in E(T)$. Observe that if $T \subseteq \{1,\dots,n\} \setminus \{i\}$, then $E(T) = E(\{1,\dots,n\} \setminus \{i\}) = E(S) = \{i\}$. Now suppose that $T = \{i,j\}$ for some $j \in \{1,\dots,n\} \setminus \{i\}$. Then $\{1,\dots,n\} \setminus T$ has at least two elements (since $n \geq 4$). We can therefore apply the argument from the first paragraph, up to renaming argument, to conclude that $E(\{i,j\}) \cap E(\{1,\dots,n\} \setminus \{i,j\})$ contains an element k. If $k \notin \{i,j\}$, then $E(\{1,\dots,n\} \setminus \{i,j\}) = \{1,\dots,n\} \setminus \{i,j\}$, which is in contradiction to $E(\{1,\dots,n\} \setminus \{i\}) = \{i\}$. Hence, $E(\{i,j\}) = \{i,j\}$. This implies that $E(T) = T$ for all

$T \subseteq \{1, \ldots, n\}$ of cardinality at least 2 containing i. We conclude that $i \in E(T)$ for every $T \subseteq \{1, \ldots, n\}$ with at least two elements, so f is a semiprojection. $\square$

In other words,

Theorem 6.18 (Rosenberg's five types theorem)**.** *Let f be a minimal operation. Then f has one of the following types:*

1. *a unary operation. If f is an operation on a finite set, then it is either a permutation such that $f^p(x) = x$, for some prime p, or satisfies $f(f(x)) = f(x)$ for all x;*
2. *a binary idempotent operation;*
3. *a majority operation;*
4. *a minority operation;*
5. *a k-ary semiprojection, for $k \geq 3$, which is not a projection.*

Proof. The statement is easy to prove if f is unary (see Exercises 116 and 117). If f is at least binary, then $\hat{f}$ (see Exercise 40) must be the identity by the minimality of f, and hence f is idempotent. In particular, we are done if f is binary. If f is ternary, we have to show that f is majority, Maltsev, or a semiprojection. By the minimality of f, the binary operation $f_1(x, y) := f(y, x, x)$ is a projection, that is, $f_1(x, y) = x$ or $f_1(x, y) = y$. Note that in particular $f(x, x, x) = x$. Similarly, the other operations $f_2(x, y) := f(x, y, x)$, and $f_3(x, y) := f(x, x, y)$ obtained by identifications of two variables must be projections. We therefore distinguish eight cases.

1. $f(y, x, x) = x, f(x, y, x) = x, f(x, x, y) = x$.
   In this case, f is a majority.
2. $f(y, x, x) = x, f(x, y, x) = x, f(x, x, y) = y$.
   In this case, f is a semiprojection.
3. $f(y, x, x) = x, f(x, y, x) = y, f(x, x, y) = x$.
   In this case, f is a semiprojection.
4. $f(y, x, x) = x, f(x, y, x) = y, f(x, x, y) = y$.
   The operation $g(x, y, z) := f(y, x, z)$ is a Maltsev operation.
5. $f(y, x, x) = y, f(x, y, x) = x, f(x, x, y) = x$.
   In this case, f is a semiprojection.
6. $f(y, x, x) = y, f(x, y, x) = x, f(x, x, y) = y$.
   In this case, f is a Maltsev operation.
7. $f(y, x, x) = y, f(x, y, x) = y, f(x, x, y) = x$.
   The operation $g(x, y, z) := f(x, z, y)$ is a Maltsev operation.
8. $f(y, x, x) = y, f(x, y, x) = y, f(x, x, y) = y$.
   In this case, f is a Maltsev operation.

We claim that if f is a Maltsev operation, then either it is a minority operation (and we are done) or it generates a Majority operation. Indeed, if f is not a minority then minimality of f implies that $f(x,y,x) = x$. Now consider the function g defined by $g(x,y,z) = f(x, f(x,y,z), z)$. We have

$$\begin{aligned} g(x,x,y) &= f(x, f(x,x,y), y) = f(x,y,y) = x \\ g(x,y,x) &= f(x, f(x,y,x), x) = f(x,x,x) = x \\ g(y,x,x) &= f(y, f(y,x,x), x) = f(y,y,x) = x\,. \end{aligned}$$

Note that every ternary function generated by a majority is again a majority. Also note that a function cannot be a majority and a minority at the same time unless the domain has only one element, so we obtain in this case a contradiction to the minimality of f.

Finally, let f be k-ary, where $k \geq 4$. By minimality of f, the operations obtained from f by identifications of arguments of g must be projections. The lemma of Świerczkowski (Lemma 6.17) implies that f is a semiprojection. □

Proposition 6.19. *For all $n \geq 3$, the graph K_n is projective (i.e., all idempotent polymorphisms of K_n are projections). All relations that are preserved by* $\mathrm{Sym}(\{0,\dots,n-1\})$ *are primitive positive definable in K_n.*

This provides for example a solution to Exercise 96.

Proof. By Theorem 6.15, it suffices to show that the clone of idempotent polymorphisms of K_n does not contain a minimal operation. Hence, by Theorem 6.18, we have to verify that $\mathrm{Pol}(K_n)$ does not contain a binary idempotent, a Maltsev, a majority, or a k-ary semiprojection for $k \geq 3$.

1. Let f be a binary idempotent polymorphism of K_n.

   Observation 1. $f(u,v) \in \{u,v\}$: otherwise, $i := f(u,v)$ is adjacent to both u and v, but $f(i,i) = i$ is not adjacent to i, in contradiction to f being a polymorphism.

   Observation 2. If $f(u,v) = u$, then $f(v,u) = v$: this is clear if $u = v$, and if $u \neq v$ it follows from f being a polymorphism.

   By Observation 1, it suffices to show that there cannot be distinct u, v and distinct u', v' such that $f(u,v) = u$ and $f(u',v') = v'$. Suppose for contradiction that there are such u, v, u', v'.

   Case 1. $u = u'$. Since $f(u,v') = f(u',v') = v'$, we have $f(v',u) = u$ by Observation 2. This is in contradiction to $f(u,v) = u$ since $u = u'$ is adjacent to v', and $E(v,u)$.

   Case 2. $u \neq u'$.
   Case 2.1. $f(u',u) = u$: this is impossible because $f(u,v) = u$, $E(u,u')$, and $E(u,v)$.
   Case 2.2. $f(u',u) = u'$: this is impossible because $f(v',u') = u'$, $E(u',v')$, and $E(u',u)$.

2. Since $(1,0),(1,2),(0,2) \in E(K_n)$, but $(0,0) \notin E(K_n)$, the graph K_n has no Maltsev polymorphism (it is not rectangular; see Section 4.4).

3. If f is a majority, note that $f(0,1,2) = f(x_0,x_1,x_2)$ where x_i is some element distinct from i if $f(0,1,2) = i$, and $x_i := f(0,1,2)$ otherwise. But $(i,x_i) \in E(K_n)$, so f is not a polymorphism of K_n.

4. Finally, let f be a k-ary semiprojection for $k \geq 3$ which is not a projection. Suppose without loss of generality that $f(x_1, \dots, x_k) = x_1$ whenever $|\{x_1, \dots, x_k\}| < k$ (otherwise, permute the arguments of f). Since f is not a projection, there exist pairwise distinct $a_1, \dots, a_k \in V(K_n)$ such that $c := f(a_1, \dots, a_k) \neq a_1$. Let $b_1, \dots, b_k$ be such that b_i is any element of $V(K_n) \setminus \{c\}$ if $c = a_i$, and $b_i := c$ otherwise. Note that $b_1 = a_1$ since $c \neq a_1$, and that $f(b_1, \dots, b_k) = b_1 = a_1$ because f is a semiprojection. But $(a_i, b_i) \in E(K_n)$ for all $i \leq k$, so f is not a polymorphism of K_n.

The second part of the statement follows from Theorem 6.6. □

The presentation of the proof of the following result is inspired by (but not identical to[7]) a presentation of Csákány [46].

Theorem 6.20 (Płonka [85]). *Let $\mathfrak{G}$ the structure obtained from a finite field $\mathbf{F}$ of prime order p as in Example 6.10. Then the clone*

$$\operatorname{Pol}(\mathfrak{G}) = \big\{(x_1, \dots, x_n) \mapsto \sum_{i=1}^{n} \alpha_i x_i \mid n \geq 1, \alpha_1, \dots, \alpha_n \in F, \alpha_1 + \cdots + \alpha_n = 1\big\}$$

is minimal.

Proof. The statement is easy to prove for $p = 2$, and it also follows from Theorem 6.22 that we prove later. For $p > 2$, let $f \in \operatorname{Pol}(\mathfrak{G})$ be non-trivial. We have to show that the clone $\mathscr{C} := \langle f \rangle$ equals $\operatorname{Pol}(\mathfrak{G})$. We already know that f can be written as $\sum_{i=1}^{n} \alpha_i x_i$ for $\alpha_1, \dots, \alpha_n \in F$ such that $\alpha_1 + \cdots + \alpha_n = 1$ (Example 6.10).

We first show that f generates a non-trivial binary operation. Since f is non-trivial, there are distinct $p, q \in \{1, \dots, n\}$ such that $\alpha_p, \alpha_q \neq 0$. If $(1 - \alpha_p) = \alpha_1 + \cdots + \alpha_n - \alpha_p \neq 0$, then equating all arguments of f except the p-th argument with the q-th argument yields the non-trivial binary operation $\alpha_p x_p + (1 - \alpha_p) x_q$. So we may assume that $\alpha_p = 1$. Similar reasoning applies to all $j \in \{1, \dots, n\}$ with $\alpha_j \neq 0$ instead of p. We therefore may suppose without loss of generality that $\alpha_1 = \cdots = \alpha_k = 1$ and $\alpha_{k+1} = \cdots = \alpha_n = 0$, for some $k \in \{2, \dots, n\}$. We know that $k \equiv 1 \mod p$, because $\alpha_1 + \cdots + \alpha_k = 1$. Since $p > 2$, if we identify the first two arguments, and identify the remaining arguments, we again obtain a non-trivial binary operation.

Let $s \in \mathscr{C}$ be the resulting non-trivial binary operation; s is of the form $\beta x_1 + (1 - \beta) x_2$ for some $\beta \in F \setminus \{0, 1\}$. Note that $\gamma^{p-1} = 1$ for every $\gamma \in F \setminus \{0\}$ by Fermat's lemma. Let l and r be the binary operations defined as follows.

$$l(x, y) := \underbrace{s(s(\dots s(x, y), \dots, y), y)}_{p-2 \text{ occurrences of } s} = \beta^{p-2} x + \beta^{p-3}(1 - \beta) y + \cdots + (1 - \beta) y$$

$$r(y, z) := \underbrace{s(y, s(y, \dots, s(y, z) \dots))}_{p-2 \text{ occurrences of } s} = \beta y + (1 - \beta)\beta y + \cdots + (1 - \beta)^{p-3}\beta y + (1 - \beta)^{p-2} x$$

[7] I thank Andrew Moorhead for a hint!

Note that for all $x, y, z \in F$ we have

$$\begin{aligned}
m(x,y,z) &:= s(l(x,y), r(y,z)) \\
&= \underbrace{\beta^{p-1}}_{=1} x + (1-\beta)(\beta^{p-2} + \beta^{p-3} + \dots + \beta^2 + \beta)y \\
&\quad + \beta((1-\beta)^{p-2} + \dots + (1-\beta)^2 + (1-\beta))y + \underbrace{(1-\beta)^{p-1}}_{=1} z \\
&= x + (\beta - 1)y + ((1-\beta) - 1)y + z \\
&= x - y + z.
\end{aligned}$$

We conclude that $m \in \mathscr{C}$. Next, we show that $\mathscr{C}$ contains all binary operations $\alpha_1 x_1 + \alpha_2 x_2$ with $\alpha_1 + \alpha_2 = 1$. Indeed, for all $x, y \in F$ we have

$$\underbrace{m(x,y,m(x,y,\dots(m(x,y,x))\dots))}_{\alpha_1 \text{ occurrences of } m} = \alpha_1 x - \alpha_2 y.$$

Finally, m and all operations of the form $(x,y) \mapsto \alpha_1 x + \alpha_2 y$ with $\alpha_1 + \alpha_2 = 1$ generate $\mathrm{Pol}(\mathfrak{G})$ (Example 6.10). Hence, $\mathscr{C} = \mathrm{Pol}(\mathfrak{G})$, which concludes the proof. □

Corollary 6.21. *Let $\mathfrak{G}$ be the structure obtained from a finite field $\mathbf{F}$ of prime order p as in Example 6.10. Then for every $n \in \mathbb{N}$ and $R \subseteq F^n$ which is not primitively positively definable in $\mathfrak{G}$, we have that $\mathrm{CSP}(\mathfrak{G}, R)$ is NP-complete.*

Proof. Since $\mathrm{Pol}(\mathfrak{G})$ is minimal, $\mathrm{Pol}(\mathfrak{G}, R)$ only contains the projections, and therefore every relation $R \subseteq F^k$ is primitively positively definable in $(\mathfrak{G}, R)$. The statement then follows via Lemma 5.8 from the existence of an NP-hard CSP with a domain of size p (e.g., $\mathrm{CSP}(K_p)$ if $p > 2$; we will also see such examples for $p = 2$, see Theorem 6.28). □

Exercises.

114. Show that every semilattice operation (Definition 3.14) generates a minimal clone.

6.5 Schaefer's Theorem

Schaefer's theorem states that every CSP for a 2-element structure is either in P or NP-hard. By the general results in Section 6.2, most of the classification arguments in Schaefer's article follow from earlier work of Post [86] (also see [77]), who classified all clones on a two-element domain. We present a short proof of Schaefer's theorem here.

Note that on Boolean domains, there is precisely one minority operation, and precisely one majority operation.

Theorem 6.22 (Post [86])**.** *Every minimal operation on $\{0,1\}$ is among one of the following:*

- *a unary constant function.*
- *the unary function $x \mapsto 1 - x$.*
- *the binary function $(x,y) \mapsto \min(x,y)$.*
- *the binary function $(x,y) \mapsto \max(x,y)$.*

- *the Boolean minority operation.*
- *the Boolean majority operation.*

Proof. If f is unary the statement is trivial, so let f be a minimal at least binary idempotent function above $\mathscr{C}$. There are only four binary idempotent operations on $\{0,1\}$, two of which are projections and therefore cannot be minimal. The other two operations are min and max. Next, note that a semiprojection of arity at least three on a Boolean domain must be a projection. Thus, Theorem 6.18 implies that f is the majority or a minority operation. $\square$

Definition 6.23. A Boolean relation $R \subseteq \{0,1\}^n$ is called *affine* if it is the solution space of a system of linear equalities modulo 2 (see Example 6.10).

Lemma 6.24. *A Boolean relation is affine if and only if it is preserved by the Boolean minority operation.*

Proof. Let R be n-ary. We view R as a subset of the Boolean vector space $\{0,1\}^n$. We have seen in Example 6.10 that affine spaces are precisely those that are closed under *affine combinations*, i.e., linear combinations of the form $\alpha_1 x_1 + \cdots + \alpha_k x_k$ such that $\alpha_1 + \cdots + \alpha_k = 1$. In particular, if R is affine then it is preserved by $(x_1, x_2, x_3) \mapsto x_1 + x_2 + x_3$ which is the minority operation. Conversely, if R is preserved by the minority operation, then $x_1 + \cdots + x_k$, for odd k, can be written as

$$\text{minority}(x_1, x_2, \text{minority}(x_3, x_4, \ldots \text{minority}(x_{n-2}, x_{k-1}, x_k) \ldots))$$

and hence R is preserved by all affine combinations, and thus affine. $\square$

It is well-known and easy to see (see, for example, [23]) that for every relation $R \subseteq \{0,1\}^n$ there exists a propositional formula $\phi(x_1, \ldots, x_n)$ that defines R, and that ϕ can even be chosen to be in *conjunctive normal form (CNF)*. That is, there is a conjunction of disjunctions of variables or negated variables from $x_1, \ldots, x_n$ such that a tuple $(t_1, \ldots, t_n) \in \{0,1\}^n$ is in R if and only if the formula ϕ evaluates to true after replacing x_i by t_i, for $i \in \{1, \ldots, n\}$. The following definition is useful for proving that certain Boolean relations R can be defined in syntactically restricted propositional logic.

Definition 6.25. If ϕ is a propositional formula in CNF that defines a Boolean relation R, we say that ϕ is *reduced* if the following holds: whenever we remove a literal from a clause in ϕ, then the resulting formula no longer defines R.

Clearly, every Boolean relation has a reduced definition: simply remove literals from any definition in CNF until the formula becomes reduced. A propositional formula in CNF is called *Horn* if every clause contains at most one positive literal.

Lemma 6.26. *A Boolean relation has a Horn definition if and only if it is preserved by* min.

Proof. It is easy to see that min preserves every relation defined by clauses that contains at most one positive literal, and hence every relation with a Horn definition. Conversely, let R be a Boolean relation preserved by min. Let ϕ be a reduced propositional formula in CNF that defines R. Now suppose for contradiction that ϕ contains a clause C with two positive literals u and v. Since ϕ is reduced, there is an assignment s_1 that satisfies ϕ such that $s_1(u) = 1$, and such that all other literals of C evaluate to 0. Similarly, there is a satisfying assignment s_2 for ϕ such that $s_2(v) = 1$ and all other literals of C evaluate to 0. Then $s_0 \colon x \mapsto \min(s_1(x), s_2(y))$ does not satisfy C, and does not satisfy ϕ, in contradiction to the assumption that min preserves R. $\square$

A binary relation is called *bijunctive* if it can be defined by a propositional formula in CNF where each disjunction has at most two disjuncts.

Lemma 6.27. *A Boolean relation R is bijunctive if and only if it is preserved by the Boolean majority operation.*

Proof. It is easy to see that the majority operation preserves every Boolean relation of arity two, and hence every bijunctive Boolean relation. We present the proof that if R is preserved by majority, and ϕ is a reduced definition of R, then all clauses C have at most two literals. Suppose for contradiction that C has three literals l_1, l_2, l_3. Since ϕ is reduced, there must be satisfying assignments s_1, s_2, s_3 to ϕ such that under s_i all literals of C evaluate to 0 except for l_i. Then the mapping $s_0 \colon x \mapsto \text{majority}(s_1(x), s_2(x), s_3(x))$ does not satisfy C and therefore does not satisfy ϕ, in contradiction to the assumption that majority preserves R. $\square$

The following relation is called the *(Boolean) not-all-equal relation.*

$$\text{NAE} := \{(0,0,1), (0,1,0), (1,0,0), (1,1,0), (1,0,1), (0,1,1)\} \tag{15}$$

Theorem 6.28 (Schaefer [88])**.** *Let $\mathfrak{B}$ be a structure over the two-element universe $\{0,1\}$. Then either $(\{0,1\}; \text{NAE})$ has a primitive positive definition in $\mathfrak{B}$, and $\text{CSP}(\mathfrak{B})$ is NP-complete, or*

1. *$\mathfrak{B}$ is preserved by a constant operation.*
2. *$\mathfrak{B}$ is preserved by* min*. Equivalently, every relation of $\mathfrak{B}$ has a definition by a propositional Horn formula.*
3. *$\mathfrak{B}$ is preserved by* max*. Equivalently, every relation of $\mathfrak{B}$ has a definition by a* dual-Horn *formula, that is, by a propositional formula in CNF where every clause contains at most one negative literal.*
4. *$\mathfrak{B}$ is preserved by the majority operation. Equivalently, every relation of $\mathfrak{B}$ is bijunctive.*
5. *$\mathfrak{B}$ is preserved by the minority operation. Equivalently, every relation of $\mathfrak{B}$ can be defined by a conjunction of linear equations modulo 2.*

In case (1) *to case* (5)*, then for every finite-signature reduct $\mathfrak{B}'$ of $\mathfrak{B}$ the problem* $\text{CSP}(\mathfrak{B}')$ *can be solved in polynomial time.*

Proof. If $\text{Pol}(\mathfrak{B})$ contains a constant operation, then we are in case one; so suppose in the following that this is not the case. If NAE is primitive positive definable in $\mathfrak{B}$, then $\text{CSP}(\mathfrak{B})$ is NP-hard by reduction from positive not-all-equal-3SAT [56]. Otherwise, by Theorem 6.6 there is an operation $f \in \text{Pol}(\mathfrak{B})$ that violates NAE. If $\hat{f}$ defined as $x \mapsto f(x, \dots, x)$ equals the identity then f is idempotent. Otherwise, $\hat{f}$ equals $\neg$. But then $\neg f \in \text{Pol}(\mathfrak{B})$ is idempotent and also violates NAE. So let us assume in the following that f is idempotent. Then f generates an at least binary minimal operation $g \in \text{Pol}(\mathfrak{B})$.

By Theorem 6.22, the operation g equals min, max, the Boolean minority, or the Boolean majority function.

- $g = \min$ or $g = \max$. By Lemma 6.26, the relations of $\mathfrak{B}$ are preserved by min if and only if they can be defined by propositional Horn formulas. It is well-known that positive unit-resolution is a polynomial-time decision procedure for the satisfiability problem of propositional Horn-clauses [89]. The case that $g = \max$ is dual to this case.
- $g =$ majority. By Lemma 6.27, the relations of $\mathfrak{B}$ are preserved by majority if and only if they are bijunctive. Hence, in this case the instances of CSP($\mathfrak{B}$) can be viewed as instances of the 2SAT problem, and can be solved in linear time [4].
- $g =$ minority. By Lemma 6.24 every relation of $\mathfrak{B}$ has a definition by a conjunction of linear equalities modulo 2. Then CSP($\mathfrak{B}$) can be solved in polynomial time by Gaussian elimination.

This concludes the proof of the statement. $\square$

Exercises.

115. Show that if A is a finite set and $f \colon A \to A$, then $g := f^{|A|!}$ satisfies $g(g(x)) = g(x)$ for all $x \in A$.
116. Show that if f is a permutation on a finite set A, then either f is the identity of f generates a permutation g which is not the identity and additionally satisfies $g^p(x) = x$ for some prime p.
117. Show that if A is a finite set and $f \colon A \to A$ is not the identity, then f generates a non-identity operation g which additionally satisfies $g(g(x)) = g(x)$.
118. The Rosenberg theorem is only a *preclassification* in the sense that not every operation which has one of the five types is minimal. For each of the following five questions, either present a proof or give a counterexample.
    (a) Which unary operations which are a permutation such that $f^p(x) = x$ for some prime p, or which satisfy $f(f(x)) = f(x)$, are minimal?
    (b) Is every binary idempotent operation minimal?
    (c) Is every majority operation minimal?
    (d) Is every minority operation minimal?
    (e) Is every k-ary semiprojection, for $k \geq 3$, which is not a projection, minimal?
119. Determine the complexity of the following CSPs.
$$\begin{aligned}
&\mathrm{CSP}(\{0,1\}; \{(0,0,1,1),(1,1,0,0)\})\\
&\mathrm{CSP}(\{0,1\}; \{(0,0,1),(0,1,0),(1,0,0),(1,1,1)\}, \{(0,1),(1,0)\})\\
&\mathrm{CSP}(\{0,1\}; \{0,1\}^3 \setminus \{(1,1,0)\}, \{(0,1),(1,0)\}).
\end{aligned}$$
120. Show that a Boolean relation $R \subseteq \{0,1\}^k$ can be defined by a propositional Horn formula if and only if it is primitively positively definable in $(\{0,1\}; \{0,1\}^3 \setminus \{(1,1,0)\}, \{0\}, \{1\})$.
121. Show that all polymorphisms of $(\{0,1\}; \mathrm{NAE})$ are essentially unary. Hint: one way to prove this is to use Theorem 6.18.

122. Show that all polymorphisms of $(\{0,1\};\{(0,0,1),(0,1,0),(1,0,0)\})$ are projections. Hint: one way to prove this is to use Theorem 6.18.

123. Show that

$$\begin{aligned}&\mathrm{Pol}(\{0,1\};\{0\},\{1\},\{0,1\}^3\setminus\{1,1,0\})\\&=\langle\min\rangle\\&=\big\{(x_1,\dots,x_k)\mapsto\min(x_{i_1},\dots,x_{i_l})\mid l\leq k, i_1,\dots,i_l\in\{1,\dots,n\}\big\}.\end{aligned}$$

124. Show that the operation f from case 6 in the proof of Rosenberg's five types theorem (Theorem 6.18) not only generates a majority operation, but also a minority operation.

6.6 Near Unanimity Polymorphisms

An operation f of arity at least 3 is a *quasi near-unanimity operation* if it satisfies the identities

$$f(x,\dots,x,y)\approx f(x,\dots,x,y,x)\approx\cdots\approx f(y,x,\dots,x)\approx f(x,\dots,x).$$

If f is additionally idempotent, then it is called a *near-unanimity operation.* Note that majority operations are exactly the ternary near-unanimity operations.

Example 6.29. If D has two elements, say $D=\{0,1\}$, then there is a near unanimity of arity $k\geq 3$ which returns 1 if at least two of its arguments are 1, and returns 0 otherwise. An example of a relation that is preserved by f_{k+1}, but not by f_k, is the relation

$$B_k:=\{0,1\}^k\setminus\{(0,\dots,0)\}.\qquad\triangle$$

We later often need a more flexible notation concerning projections.

Definition 6.30. For $I=\{i_1,\dots,i_k\}\in\binom{[n]}{k}$, with $i_1<\cdots<i_k$, we write π_I^n for the function from $A^n\to A^k$ defined by $\pi_I^n(t):=(t_{i_1},\dots,t_{i_k})$. Sometimes, it will also be convenient to define $\pi_s^n(t)$ for $t\in A^n$ and $s\in[n]^k$, as follows: $\pi_s^n(t):=(t_{s_1},\dots,t_{s_k})\in A^k$. If $s=(s_1,\dots,s_k)$, we may also omit the brackets in the subscript and write $\pi^n_{s_1,\dots,s_k}(t)$ instead of $\pi_s^n(t)$.

Note that Definition 6.30 is compatible with our earlier definition of the projection operations π_i^n. If n is clear from the context, the superscript n may also be omitted. We use compact notation for applying functions pointwise or setwise. In particular, if $R\subseteq A^n$, we write $\pi_{i_1,\dots,i_k}(R)$ for the relation $\{\pi_{i_1,\dots,i_k}(t)\mid t\in R\}$. We also apply our notation for projections for relations $R\subseteq A_1\times\cdots\times A_k$ instead of $R\subseteq A^k$.

Theorem 6.31. *A finite structure $\mathfrak{B}$ has a $k+1$-ary near unanimity polymorphism if and only if every relation R with a primitive positive definition in $\mathfrak{B}$ satisfies*

$$R=\bigcap_{S\in\binom{[m]}{k}}\pi_S(R).\tag{16}$$

The next example illustrates that there are some clones $\mathscr{C}$ on a finite set B such that *every* set $\mathcal{R}$ of relations over B such that $\mathscr{C}=\mathrm{Pol}(\mathcal{R})$ has to be infinite. In fact, we already know that such clones on a three-element set B must exist, because otherwise there would be only countably many such clones, which is false (Remark 6.1). There are even concrete examples of clones on a *two-element* set which have this property.

Example 6.32. Let $\mathfrak{D}$ be the structure with domain $\{0,1\}$ which contains

- the unary relation $\{0\}$,
- the binary relation $\leq := \{(0,0),(0,1),(1,1)\}$, and
- for every $n \in \mathbb{N}$ the relation $B_n := \{0,1\}^n \setminus \{(0,\dots,0)\}$.

Note that all of these relations are preserved by the operation $p\colon \{0,1\}^3 \to \{0,1\}$ given by $(x,y,z) \mapsto x \wedge (y \vee z)$. We claim that $\text{Pol}(\mathfrak{D}) = \langle p \rangle$. We already know that $\langle p \rangle \subseteq \text{Pol}(\mathfrak{D})$. To show the reverse inclusion, we have to show that every relation that is preserved by p belongs to the set $\mathcal{R}$ of relations that are primitively positively definable in $\mathfrak{D}$. Indeed,

$$\begin{aligned}
\langle \{p\} \rangle &= \text{Pol}(\text{Inv}(\{p\})) && \text{(Proposition 6.4)}\\
&\supseteq \text{Pol}(\mathcal{R}) && \text{(by assumption)}\\
&\supseteq \text{Pol}(\text{Inv}(\text{Pol}(\mathfrak{D}))) && \text{(Proposition 6.5)}\\
&= \text{Pol}(\mathfrak{D}).
\end{aligned}$$

Since every relation R that is preserved by p is also preserved by $p(x,y,y) = \min(x,y)$, it is Horn (Lemma 6.26). Let ϕ be a propositional Horn formula in CNF that defines R; we may assume that ϕ is reduced (Definition 6.25). Suppose for contradiction that ϕ contains a clause ψ with one positive literal u and two negative literals $\neg v$ and $\neg w$. Since ϕ is reduced, this means that R has satisfying assignments s_1, s_2, s_3 such that u is the only literal in ψ satisfied by s_1, $\neg v$ is the only literal of ψ satisfied by s_2, and $\neg w$ is the only literal of ψ satisfied by s_3. Define $s := p(s_1, s_2, s_3)$. Then $s(u) = p(1,0,0) = 0$, $s(v) = p(1,0,1) = 1$, and $s(w) = p(1,1,0) = 1$. Hence, s satisfies none of the literals $u, \neg v, \neg w$; moreover, the other literals of ψ aren't satisfied as well, a contradiction to the assumption that R is preserved by p. Therefore, all clauses of ϕ either consist of a single positive literal, or of one positive and one negative literal, or only of negative literals. This shows that R can even be defined by a conjunction of relations of $\mathfrak{D}$ (no existential quantification is needed).

Suppose that $\mathfrak{C}$ is a reduct of $\mathfrak{D}$ with finite signature; let k be the maximal arity of the relations in $\mathfrak{C}$. Then $\mathfrak{C}$ is preserved by the $k+1$-ary near unanimity polymorphism f from Example 6.29. We have already mentioned that this function does not preserve the relation $\{0,1\}^{k+1} \setminus \{(0,\dots,0)\}$. It follows that $\text{Pol}(\mathfrak{D})$ is a proper subclone of $\text{Pol}(\mathfrak{C})$. Note that this shows that there is no structure $\mathfrak{D}'$ over the domain $\{0,1\}$ with finitely many relations such that $\text{Pol}(\mathfrak{D}') = \text{Pol}(\mathfrak{D})$, because the relations of $\mathfrak{D}$ would have a primitive positive definition in a reduct $\mathfrak{C}$ of $\mathfrak{D}$ with finite relational signature, as we have seen above. We have also seen that $\text{Pol}(\mathfrak{D})$ is a proper subclone of $\text{Pol}(\mathfrak{C})$, and hence $\text{Pol}(\mathfrak{D}') \subseteq \text{Pol}(\mathfrak{C})$ is a proper subclone as well. $\triangle$

Exercises.

125. Prove Theorem 6.31.

126. Show that if H is a digraph with a $k+1$-ary near unanimity polymorphism, then the k-consistency procedure (see Section 15.1) solves CSP(H).

127. Show that the digraph C_2^{++} from Exercise 76 does not have near unanimity polymorphisms.

128. Let H be an irreflexive graph. Then H has a conservative near unanimity polymorphism if and only if is has a conservative majority polymorphism.

7 Maltsev Polymorphisms

Recall from Section 4.4 the definition of a Maltsev operation: a ternary operation $f\colon D^3 \to D$ satisfying

$$\forall x, y \in D.\, f(y, x, x) = f(x, x, y) = y\,.$$

As we have seen in Theorem 4.19, every digraph with a Maltsev polymorphism can be solved by the path-consistency procedure. However, when considering arbitrary relational structures then there are many examples with a Maltsev polymorphism that cannot be solved by the path-consistency procedure [55] (see Theorem 7.2 below). In this section, we present the algorithm of Bulatov and Dalmau for CSP($\mathfrak{A}$) when $\mathfrak{A}$ is preserved by a Maltsev polymorphism [37].

Theorem 7.1. *Let $\mathfrak{A}$ be a finite structure with finite relational signature and a Maltsev polymorphism. Then* CSP($\mathfrak{A}$) *can be solved in polynomial time.*

7.1 Affine Maltsev Operations

The most prominent class of structures $\mathfrak{A}$ with a Maltsev polymorphism comes from groups. For any group $\mathbf{G}$ (see Example 8.1), the operation m given by $(x, y, z) \mapsto x - y + z$ is obviously Maltsev. If $\mathbf{G}$ is abelian, then m is called an *affine Maltsev operation.* Structures with an affine Maltsev polymorphism are also called affine Maltsev. Note that if the group $\mathbf{G}$ is $F = (\mathbb{Z}_p; +, -, 0)$, for some prime number p, then the k-ary relations preserved by m are precisely the affine subspaces of F^k (Example 6.10). In this case one can use Gaussian elimination to solve CSP($\mathfrak{A}$).

Theorem 7.2 (from [55])**.** *Let $\mathbf{G}$ be an abelian group (see Example 8.1) with at least two elements. For $c \in G$ and $k \in \mathbb{N}$, define*

$$R^k_c := \{(x_1, \dots, x_k) \in G^k \mid x_1 + \cdots + x_k = c\}.$$

For some $a \in G \setminus \{0\}$, let $\mathfrak{B}$ be the structure $(G; R^3_0, R^2_0, R^3_a)$. Then for any $k \in \mathbb{N}$, the problem CSP($\mathfrak{B}$) *cannot be solved by k-consistency.*

Proof. We construct an unsatisfiable instance $\mathfrak{A}$ of CSP($\mathfrak{B}$) as follows. In the proof of this theorem, we work with structures of large *girth.* The girth of a graph G is the length of the shortest cycle in G. It is known that there are finite graphs of arbitrarily large girth that are *cubic*, i.e., all vertices have degree three (much stronger results are known; see, e.g., [18]). Let $(V; E)$ be a finite cubic graph of girth at least $4k + 1$. Orient the edges E arbitrarily.

The domain of $\mathfrak{A}$ is $V \times E$. For each $v \in V$ we add $\big((v, e_1), (v, e_2), (v, e_3)\big)$ to $(R^3_0)^{\mathfrak{A}}$. For each $e = (v, w) \in E$ we add $\big((v, e), (w, e)\big)$ to $(R^2_0)^{\mathfrak{A}}$. Finally, we move exactly one of the tuples from $(R^3_0)^{\mathfrak{A}}$ to $(R^3_a)^{\mathfrak{A}}$. Suppose for contradiction that $s\colon A \to G$ is a solution for $\mathfrak{A}$. Sum over all constraints. Since each element of A appears once in a two-variable constraint and once in a three-variable constraint, we obtain $2 \sum_{e \in E, u \in e} s(u, e)$ on the left-hand side. Since s satisfies $s(u, e) + s(v, e) = 0$ for every edge $e = \{u, v\} \in E$, the left-hand side can be rewritten as $2 \sum_{e \in E} \sum_{u \in e} s(u, e) = \sum_{\{u,v\} \in E} \big(s(u, e) + s(v, e)\big) = 0$. On the right-hand side we obtain a since we have precisely one tuple in S^i_a in $\mathfrak{A}$. Hence, s cannot be a homomorphism. Using high girth, it can be shown that the k-consistency procedure does not derive *false* on $\mathfrak{A}$; for the details of this last part, see Theorem 8.6.11 in [21]. □

7.2 Further Examples

For general finite groups $\mathbf{G}$, and if all relations of $\mathfrak{B}$ are cosets $gH := \{gh \mid h \in H\}$ of subgroups H of $\mathbf{G}^k$, then Feder and Vardi [55] showed how to solve CSP($\mathfrak{B}$) in polynomial time using a previously known algorithm to find small generating sets for a permutation group. We will not discuss this approach, but rather present the more general algorithm of Bulatov and Dalmau which works for all finite structures preserved by a Maltsev polymorphism. The following proposition shows that this is indeed more general.

Proposition 7.3. *Let $\mathbf{G}$ be a finite group and let $m\colon G^3 \to G$ be the Maltsev operation defined by $m(x, y, z) := xy^{-1}z$. Then m preserves a relation $R \subseteq G^k$ if and only if R is a coset of a subgroup of $\mathbf{G}^k$, for all $k \in \mathbb{N}$.*

Proof. Let $k \in \mathbb{N}$ and let H be a subgroup of $\mathbf{G}^k$. Let $a \in G^3$ and $h_1, h_2, h_3 \in H$. As usual, we may apply m to elements of G^3 componentwise; then

$$m(ah_1, ah_2, ah_3) = ah_1(ah_2)^{-1}ah_3 = ah_1h_2h_3 \in aH$$

so m indeed preserves all cosets of H.

Conversely, suppose that $R \subseteq G^k$ is preserved by m. Choose $y \in R$ arbitrarily. We claim that $y^{-1}R$ is a subgroup H of $\mathbf{G}^k$. This will show that $R = yH$ is a coset of a subgroup of $\mathbf{G}^k$. Arbitrarily choose $a, b \in y^{-1}R$. Then $x := ya \in R$ and $z := yb \in R$. Hence, $m(x, y, z) = yay^{-1}yb = yab \in R$, so $ab \in y^{-1}R$ and $y^{-1}R$ is closed under the group operation. Moreover, $m(y, ya, y) = y(ya)^{-1}y = ya^{-1} \in R$, so $a^{-1} \in y^{-1}R$ and $y^{-1}R$ is also closed under taking inverses. □

We now present examples of Maltsev operations that do not come from groups in the way described above.

Example 7.4. On the domain $\{0, 1, 2\}$, let m be the minority defined by

$$m(x, y, z) = x \text{ whenever } |\{x, y, z\}| = 3.$$

Note that m preserves all unary relations, for every permutation α of $\{0, 1, 2\}$ the relation $\{(x, y) \mid \alpha(x) = y\}$, and for $i \in \{0, 1\}$ the relation

$$R_n := \{(x_1, \ldots, x_n) \in \{0, 1\}^n \mid x_1 + x_2 + \cdots + x_n = i \text{ mod } 2\}.$$

This example will be revisited in Exercise 161, Example 13.12, Exercise 194, and Exercise 195. △

Example 7.5. On the domain $\{0, 1, 2\}$, let m be the minority operation defined by

$$m(x, y, z) := 2 \text{ whenever } |\{x, y, z\}| = 3.$$

Let $\mathbf{A}$ be the algebra $(\{0, 1, 2\}, m)$. Note that m preserves all unary relations, the graph H of the endomorphism of $\mathbf{A}$ which maps $0, 1$ to 1 and which maps 2 to 0, and

$$L = \big\{(x, y, z) \mid x = y = z = 2 \text{ or } x, y, z \in \{0, 1\} \text{ and } x + y + z = 0 \mod 2\big\}.$$

The following relations are primitively positively definable with these relations, and hence belong to $\text{Inv}(m)$ as well:

- The graph T of the transposition (12) has the primitive positive definition
$$\exists z\big(L(x,y,z) \wedge z \in \{1,2\}\big).$$
- The relation $C := \big\{(2,0),(0,1),(2,1),(1,2)\big\}$, has the primitive positive definition
$$\exists u,v\big(H(x,u) \wedge T(u,v) \wedge H(y,v)\big).$$
- The equivalence relation E with the equivalence classes $\{2\}$ and $\{0,1\}$ has the primitive positive definition $\exists z\big(C(x,z) \wedge C(z,y)\big)$.
- For every $n \in \mathbb{N}$ the relation
$$L_n = \big\{(2,\ldots,2)\big\} \cup \big\{(x_1,\ldots,x_n) \in \{0,1\}^n \mid x_1 + x_2 + \cdots + x_n = 0 \bmod 2\big\}.$$
- The relation
$$\big\{(x_1,\ldots,x_n) \in \{0,1,2\}^n \mid \text{ an even number of the } x_i\text{'s is from } \{0,1\}\big\}$$
has the primitive positive definition
$$\exists y_1,\ldots,y_n\big(L_n(y_1,\ldots,y_n) \wedge H(x_1,y_1) \wedge \cdots \wedge H(x_n,y_n)\big).$$

This example will be revisited in Exercise 196. △

Remark 7.6. It is unclear whether Maltsev operations on finite sets can be classified completely. However, it is known that for every $n \in \mathbb{N}$ there are only countably many clones on $\{1,\ldots,n\}$ that contain a Maltsev operation [3].

Exercises.

129. Check the claims made in Example 7.4.

130. Check the claims made in Example 7.5.

131. Let $\mathfrak{A}$ be the structure $(\{0,1,2\};\{1,2\},H,L)$ with the relations H and L as defined in Example 7.5, and let $\mathfrak{B}$ be its substructure with domain $\{0,1\}$. Show that $\mathfrak{B}$ has a primitive positive construction in $\mathfrak{A}$ (easy), and that $\mathfrak{A}$ has a primitive positive construction in $\mathfrak{B}$ (not so easy).

132. Show that if a Maltsev operation $m\colon D^n \to D$ preserves the graph $\{(x,y,z,u) \mid m'(x,y,z) = u\}$ of a Maltsev operation $m'\colon D^n \to D$, then $m = m'$.

7.3 Compact Representations of Relations

Our presentation of the proof closely follows that of Bulatov and Dalmau [37].

Definition 7.7 (Forks and Representations)**.** Let $R \subseteq A^n$ be a relation.

- A *fork of* R is a triple (i, a, b) such that there exist $s, t \in R$ with $(s_1, \dots, s_{i-1}) = (t_1, \dots, t_{i-1})$, $s_i = a$, and $t_i = b$. We say that s and t *witness* (i, a, b).
- $R' \subseteq R$ is called a *representation of* R if every fork of R is also a fork of R'.
- A representation R' of R is called *compact* if its cardinality is at most twice the number of forks of R.

Clearly, every relation has a compact representation. Recall Exercise 108 for the relevance of the following lemma.

Lemma 7.8. *Let A be a finite set and let $m \colon A^3 \to A$ be a Maltsev operation. Let $R \subseteq A^k$ be a relation preserved by m, and let R' be a representation of R. Then $R = \langle R' \rangle_m$.*

Proof. We show by induction on $i \in \{1, \dots, n\}$ that $\pi_{1,\dots,i}(\langle R' \rangle_m) = \pi_{1,\dots,i}(R)$. Clearly, $\pi_{1,\dots,i}(\langle R' \rangle_m) \subseteq \pi_{1,\dots,i}(R)$, so we only need to show the converse inclusion. The case $i = 1$ follows from that fact that R has for every $t \in R$ the fork $(1, t_1, t_1)$, and since R' must also have this fork it must contain a tuple t' such that $t'_1 = t_1$.

So let us assume that the statement holds for $i < n$. We have to show that for every $t \in R$ we have $(t_1, \dots, t_{i+1}) \in \pi_{1,\dots,i+1}(\langle R' \rangle_m)$. By induction hypothesis there exists a tuple $s \in \langle R' \rangle_m$ such that $(s_1, \dots, s_i) = (t_1, \dots, t_i)$. Then $(i+1, s_{i+1}, t_{i+1})$ is a fork of R, so there exist tuples $s', s'' \in R'$ witnessing it. Then the tuple $t' := m(s, s', s'') \in \langle R' \rangle_m$ is such that

$$\begin{aligned}(t'_1, \dots, t'_i, t'_{i+1}) &= (m(t_1, s'_1, s''_1), \dots, m(t_i, s'_i, s''_i), m(s_{i+1}, s_{i+1}, t_{i+1})) \\ &= (t_1, \dots, t_i, t_{i+1}) \qquad \text{(since } s'_i = s''_i\text{)}.\end{aligned}$$

Hence, $(t_1, \dots, t_i, t_{i+1})$ is a tuple from $\pi_{1,\dots,i+1}(\langle R' \rangle_m)$, as required. □

Exercises.

133. Let A be a finite set. How many forks does the n-ary relation $R := A^n$ have? Explicitly construct a compact representation for R.

134. Let R be the relation $\{(x, y, z, u) \in \{0, 1\}^4 \mid x + y + z = 1 \mod 2\}$. Find a smallest possible representation R' for R. Explicitly compute $\langle R' \rangle_m$ where m is the Boolean minority.

7.4 The Bulatov-Dalmau Algorithm

Let $\exists x_1, \dots, x_n(\phi_1 \wedge \cdots \wedge \phi_m)$ be an instance of CSP($\mathfrak{A}$). For $\ell \leq n$, we write R_ℓ for the relation

$$\{(s_1, \dots, s_n) \in A^n \mid \mathfrak{A} \models (\phi_1 \wedge \cdots \wedge \phi_\ell)(s_1, \dots, s_n)\}.$$

The idea of the algorithm is to inductively construct a compact representation R'_ℓ of R_ℓ, adding constraints one by one. Initially, for $\ell = 0$, we have $R_\ell = A^n$, and it is easy to come up with a compact representation for this relation. Note that when we manage to compute the compact representation R'_n for R_n, we can decide satisfiability of the instance: it is unsatisfiable if and only if R'_n is empty. For the inductive step, we need a procedure called *Next* which is more involved; we first introduce two auxiliary procedures.

Procedure *Nonempty*$(R', i_1, \ldots, i_k, S)$.

Set $U := R'$.
While $\exists r, s, t \in U$ such that $\pi_{i_1,\ldots,i_k}(m(r,s,t)) \notin \pi_{i_1,\ldots,i_k}(U)$:
  Set $U := U \cup \{m(r,s,t)\}$
If $\exists t \in U$ such that $(t_{i_1}, \ldots, t_{i_k}) \in S$ then return t
else return 'No'.

Figure 10: The procedure *Nonempty*.

The procedure *Nonempty*

The procedure *Nonempty* receives as input

- a compact representation R' of a relation R,
- a sequence $i_1, \ldots, i_k$ of elements in $[n]$ where n is the arity of R, and
- a k-ary relation S which is also preserved by m.

The output of the procedure is either a tuple $t \in R$ such that $(t_{i_1}, \ldots, t_{i_k}) \in S$, or 'No' if no such tuple exists. The procedure can be found in Figure 10.

Correctness. For the correctness of *Nonempty* we note the following:

- $R' \subseteq U \subseteq R$: initially we start from $U := R' \subseteq R$, and only add tuples to U obtained by applying m to tuples in U, so the added tuples are again in R.
- It follows that if *Nonempty* returns a tuple $(t_{i_1}, \ldots, t_{i_k})$, then this tuple is indeed from $\pi_{i_1,\ldots,i_k}(R)$ and the output of the algorithm is correct.
- When the algorithm exits the while loop then $\pi_{i_1,\ldots,i_k}(\langle U\rangle_m) = \pi_{i_1,\ldots,i_k}(U)$. Since $R' \subseteq U$ we have that $\langle U\rangle_m = R$. Hence, every tuple $t \in \pi_{i_1,\ldots,i_k}(R) = \pi_{i_1,\ldots,i_k}(\langle U\rangle_m)$ is contained in $\pi_{i_1,\ldots,i_k}(U)$, and so the answer of the algorithm is also correct when it returns 'No'.

We mention that this procedure does not use the particular properties of a Maltsev polymorphism, but works for any explicitly given polymorphism.

Running time. The number of iterations of the while loop can be bounded by the size $|U|$ of the set U at the end of the execution of the procedure. Hence, when we want to use this procedure to obtain a polynomial-time running time, we have to make sure that the size of U remains polynomial in the input size. The way this is done in the Bulatov-Dalmau algorithm is to guarantee that at each call of *Nonempty* the size L of $\pi_{i_1,\ldots,i_k}(R)$ is polynomial in the input size. Then $|U|$ is bounded by $L + |R'|$ which is also polynomial.

We have to test all tuples $r, s, t \in U$; this can be implemented so that $|U|^3$ steps suffice. In each step we have to compute $m(r,s,t)$ and test whether $\pi_{i_1,\ldots,i_k}(m(r,s,t)) \in \pi_{i_1,\ldots,i_k}(U)$, which can be done in $O(kL)$. In the important case that L is bounded by a constant in the size of the input N, the running time of *Nonempty* is in $O(N^4)$.

Procedure *Fix-values*$(R', c_1, \dots, c_k)$.

Set $j := 0$; $U_j := R'$.
While $j < k$ do:
  Set $U_{j+1} := \emptyset$.
  For each $(i, a, b) \in [n] \times A^2$:
    If $\exists s, t \in U_j$ witnessing (i, a, b) (assuming $s = t$ if $a = b$):
      If $r :=$ Nonempty$(U_j, j + 1, i, \{(c_{j+1}, a)\}) \neq$ 'No'
        If $(i > j + 1)$ or $(a = b = c_i)$:
          Set $U_{j+1} := U_{j+1} \cup \{r, m(r, s, t)\}$
  Set $j := j + 1$.
Return U_k.

Figure 11: The procedure *Fix-values*.

The procedure *Fix-values*

The procedure *Fix-values* receives as input

- a compact representation R' of an n-ary relation R preserved by m, and
- a sequence $c_1, \dots, c_k \in A$ for $k \leq n$.

The output of *Fix-values* is a compact representation of the relation

$$R \cap (\{c_1\} \times \cdots \times \{c_k\} \times A \times \cdots \times A).$$

The procedure can be found in Figure 11. The algorithm computes inductively a compact representation U_j of the relation

$$R_j = R \cap (\{c_1\} \times \cdots \times \{c_j\} \times A \times \cdots \times A)$$

This is immediate for $U_0 = R'$, and the set U_k is the relation that we have to compute.

For its correctness, suppose inductively that U_j is a compact representation of R_j. We have to show that the set U_{j+1} computed by the procedure is a compact representation of R_{j+1}:

1. $U_{j+1} \subseteq R_{j+1}$. Suppose that the procedure adds $\{r, m(r, s, t)\}$ to U_{j+1}, where r and s witness the fork (i, a, b) of U_j processed in the for-loop of the procedure. Note that $r \in R_{j+1}$ since $r \in \langle U_j \rangle_m \subseteq R_j$ and $r_{j+1} = c_{j+1}$. Since m preserves R and is idempotent, it also preserves R_j, and since $r, s, t \in R_j$ it follows that $m(r, s, t) \in R_j$. To show that $m(r, s, t) \in R_{j+1}$ it suffices to show that $s_{j+1} = t_{j+1}$ because then $m(r, s, t)_{j+1} = r_{j+1} = c_{j+1}$ since m is Maltsev. If $i > j + 1$ then we have that $s_{j+1} = t_{j+1}$ since s, t witness (i, a, b). Otherwise, we must have $a = b = c_i$ because of the innermost if-clause of the procedure. But then $s = t$ by the stipulation of the algorithm on the choice of s and t.
2. All forks (i, a, b) of R_{j+1} are forks of U_{j+1}. If R_{j+1} has the fork (i, a, b), then by inductive assumption U_j must contain witnesses s, t for (i, a, b). Therefore, the first if-clause of the procedure is positive. Moreover, $s_{j+1} = c_{j+1}$ and $s_i = a$, so $r :=$

Procedure *Next*$(R', i_1, \ldots, i_k, S)$.

Set $U := \emptyset$.
For each $(i, a, b) \in [n] \times A^2$:
  If *Nonempty*$(R', i_1, \ldots, i_k, i, S \times \{a\}) =: t \neq$ 'No':
    If *Nonempty*$($*Fix-values*$(R', t_1, \ldots, t_{i-1}), i_1, \ldots, i_k, i, S \times \{b\}) =: t' \neq$ 'No':
      Set $U := U \cup \{t, t'\}$.
Return U_k.

Figure 12: The procedure *Next*.

Nonempty$(U_j, j+1, i, \{(c_{j+1}, a)\}) \neq$ 'No'. Also note that if $i \leq j+1$, then $a = s_i = c_i = t_i = b$. So all the if-clauses of the procedure are positive, and the procedure adds r and $m(r, s, t)$ to U_{j+1}. The tuples r and $m(r, s, t)$ witness (i, a, b). Since s, t witness (i, a, b) we have that $(s_1, \ldots, s_{i-1}) = (t_1, \ldots, t_{i-1})$. Hence, $\pi_{1,\ldots,i-1}(m(r, s, t)) = (r_1, \ldots, r_{i-1})$. Furthermore, we have that $\pi_i(m(r, s, t)) = m(a, a, b) = b$.

3. The representation U_{j+1} of R_{j+1} is compact since at most two tuples are added to U_{j+1} for each fork of R_{j+1}.

Running time. The while loop is performed $k \leq n$ times; the inner for-loop is executed for each $(i, a, b) \in [n] \times A^2$, which is linear for fixed $\mathfrak{A}$. The cost of each iteration is dominated by the cost of calling the procedure *Nonempty*. Note that when calling *Nonempty*, the size of $\pi_{j+1,i}(U_j)$ is polynomial in the input size (even constant size when $\mathfrak{A}$ is fixed), so the cost of *Nonempty* is in $O(N^4)$ where N is the size of the input. Therefore, the total time complexity of the procedure *Fix-values* is polynomial in the input size (for fixed $\mathfrak{A}$ it is in $O(N^5)$).

The procedure *Next*

Now comes the heart of the algorithm, which is the procedure *Next* that updates a compact representation of the solution space when constraints are added one by one. The input of the procedure is

- a compact representation R' of a relation $R \subseteq A^n$ that is preserved by m,
- a sequence $i_1, \ldots, i_k$ of elements from $[n]$,
- a k-ary relation S which is also preserved by m.

The output of the procedure is a compact representation of the relation

$$R^* := \{t \in R \mid (t_{i_1}, \ldots, t_{i_k}) \in S\}.$$

The procedure *Next* can be found in Figure 12. Observe that

- the condition *Nonempty*$(R', i_1, \ldots, i_k, i, S \times \{a\}) \neq$ 'No' from the first if-clause is satisfied if and only if there exists a tuple $t \in R$ such that $(t_{i_1}, \ldots, t_{i_k}) \in S$ and $t_i = a$. Hence, if such a tuple does not exist, then (i, a, b) cannot be a fork of R^*, and nothing needs to be done.

- the condition *Nonempty*(*Fix-values*$(R', t_1, \ldots, t_{i-1}), i_1, \ldots, i_k, i, S \times \{b\}) \neq$ 'No' from the second if-clause is satisfied if and only if there exists a tuple $t' \in R$ such that
  - $(t'_1, \ldots, t'_{i-1}) = (t_1, \ldots, t_{i-1})$,
  - $(t'_{i_1}, \ldots, t'_{i_k}) \in S$, and
  - $t'_i = b$.

  If this condition holds, and since $t_i = a$, we have that t and t' witness (i, a, b). It only remains to show that if (i, a, b) is a fork of R^*, then such a tuple t' must exist. So let r and s be witnesses for (i, a, b) in R^*. Then the tuple $t' := m(t, r, s)$ has the desired properties:
  - for $j < i$ we have that $t'_j = m(t_j, r_j, s_j) = t_j$;
  - $t' \in S$ because $(r_{i_1}, \ldots, r_{i_k}), (s_{i_1}, \ldots, s_{i_k}), (t_{i_1}, \ldots, t_{i_k}) \in S$ and m preserves S.
  - $t'_i = m(t_i, r_i, s_i) = m(a, a, b) = b$.
- The cardinality of U is bounded by twice the number of forks of R^*, so the representation computed by the algorithm is compact.

Running time. The for-loop of the procedure *Next* is performed $n|A|^2$ times and the cost of each iteration is polynomial in the cost of *Nonempty* and *Fix-values*. Also note that k is bounded by the maximal arity of the relations in $\mathfrak{A}$, so constant for fixed $\mathfrak{A}$. It follows that $\pi_{i_1,\ldots,i_k,i}(R)$ is polynomial, so the running time of the calls to *Nonempty* are polynomial. For fixed $\mathfrak{A}$, the global running time of the procedure *Next* is in $O(N^6)$ where N is the size of the input.

Proof of Theorem 7.1. Starting from an empty list of constraints, we add constraints on the variables $x_1, \ldots, x_n$ one by one, and maintain a compact representation of the n-ary relation defined by the constraints considered so far. Initially, we start with a compact representation of the full relation A^n. In later steps, we use the procedure *Next* to compute a compact representation when a constraint is added, in $O(N^6)$ for fixed $\mathfrak{A}$ and N the size of the input. The instance is unsatisfiable if and only if at the final stage we end up with an empty representation. The entire running time of the algorithm is in $O(N^7)$. □

Exercises.

135. Let $\mathfrak{A}$ be the structure $(\{0,1\}; L_0, L_1)$ where $L_i := \{(x, y, z) \mid x + y + z = i \mod 2\}$, which has the Boolean minority m as polymorphism. Consider the instance

$$\exists x_1, \ldots, x_5 \big(L_1(x_1, x_2, x_3) \wedge L_1(x_2, x_3, x_4) \wedge L_1(x_3, x_4, x_5) \wedge L_0(x_1, x_3, x_5)\big)$$

Compute compact representations R'_ℓ of R_ℓ, for $\ell \in \{1, 2, 3, 4\}$.

136. Let $\mathfrak{B}$ be a structure with a Maltsev polymorphism f and an *infinite* relational signature. Note that we have defined CSP($\mathfrak{B}$) only if $\mathfrak{B}$ has a finite signature. If we want to define CSP($\mathfrak{B}$) also for structures $\mathfrak{B}$ with an infinite signature, it is important to discuss how the relation symbols in the signature of $\mathfrak{B}$ are represented in the input. We choose to represent a relation symbol R from $\mathfrak{B}$ by listing the tuples in $R^{\mathfrak{B}}$. Adapt the Dalmau algorithm such that it can solve CSP($\mathfrak{B}$) in polynomial time for this choice of representing the relations in $\mathfrak{B}$.

137. The *graph isomorphism problem (GI)* is a famous computational problem that is neither known to be solvable in polynomial time, nor expected to be NP-hard. An instance of GI consists of two graphs G and H, and the question is to decide whether G and H are isomorphic. Consider the variant of the graph-isomorphism problem where the vertices are coloured, each color appears at most k times for some constant k, and the isomorphism between H and G that we are looking for is required to additionally preserve the colours. Show that this problem can be solved in polynomial time using Dalmau's algorithm (use the previous exercise).

8 Universal Algebra

We have seen in Section 6 that for finite relational structures $\mathfrak{B}$ with finite relational signature, the computational complexity of CSP($\mathfrak{B}$) only depends on the polymorphisms of $\mathfrak{B}$. For more advanced results that use this perspective, it will be useful to view the set of all polymorphisms of $\mathfrak{B}$ as an algebra, since we may then use ideas and results from universal algebra.

8.1 Algebras and Clones

In universal algebra, an *algebra* is simply a structure with a purely functional signature. We will typically use bold font letters, like $\mathbf{A}$, to denote algebras, and the corresponding capital roman letters, like A, to denote their domain.

Example 8.1 (Group)**.** A *group* is an algebra with a binary function symbol $\circ$ for composition, a unary function symbol $^{-1}$ for taking the inverse, and a constant denoted by e, satisfying

- $\forall x, y, z.\, x \circ (y \circ z) = (x \circ y) \circ z$,
- $\forall x.\, x \circ x^{-1} = e$,
- $\forall x.\, e \circ x = x$, and $\forall x.\, x \circ e = x$.

Note that all axioms are *universal* in the sense that all the variables are universally quantified (more on that comes later). A group is called *abelian* if it additionally satisfies

$$\forall x, y.\, x \circ y = y \circ x.$$

For abelian groups we sometimes use the signature $\{+, -, 0\}$ instead of $\{\circ, ^{-1}, e\}$. △

Example 8.2 (Ring)**.** A *(unital) ring* is an algebra $\mathbf{A}$ with the signature $\{\cdot, +, -, 0, 1\}$ where $\cdot, +$ are binary, $-$ is unary, and $0, 1$ are constants, such that $(A; +, -, 0)$ is an abelian group and additionally

$$\begin{aligned} \forall x, y, z.\ (xy)z &= x(yz) && \text{(associativity)} \\ \forall x.\ 1 \cdot x &= x && \text{(multiplicative unit)} \\ \forall x, y, z.\ x(y + z) &= xy + xz && \text{(distributivity)} \end{aligned}$$

A ring is called *commutative* if it additionally satisfies

$$\forall x, y.\ xy = yx \qquad \text{(commutativity).}$$ △

The next example generalises vector spaces.

Example 8.3 (Module)**.** Let $\mathbf{R}$ be a ring. An $\mathbf{R}$-*module* is an algebra $\mathbf{M}$ with the signature $\{+, -, 0\} \cup \{f_r \mid r \in R\}$ such that $(M; +, -, 0)$ is an abelian group and for all $r, s \in R$ it holds that

$$\forall x, y.\ f_r(x + y) = f_r(x) + f_r(y) \tag{17}$$

$$\forall x.\ f_{r+s}(x) = f_r(x) + f_s(x) \tag{18}$$

$$\forall x.\ f_r(f_s(x)) = f_{rs}(x). \tag{19}$$

An $\mathbf{R}$-module is called *unitary* if it additionally satisfies $\forall x. f_1(x) = x$. We usually write rx instead of $f_r(x)$.

An alternative formalisation of modules is to view them as structures with two sorts, one sort for R and one for M above (see Section 5.1). The details of this perspective are the content of Exercise 139 and omitted because we do not need it in this text. △

Example 8.4 (Semilattice)**.** A *meet-semilattice* $\mathfrak{S}$ is a $\{\leq\}$-structure with domain S such that $\leq^{\mathfrak{S}}$ denotes a partial order where any two $u, v \in S$ have a (unique) *greatest lower bound* $u \wedge v$, i.e., an element w such that $w \leq u$, $w \leq v$, and for all w' with $w' \leq u$ and $w \ \leq v$ we have $w' \leq w$. Dually, a *join-semilattice* is a partial order with least upper bounds, denoted by $u \vee v$. A *semilattice* is a meet-semilattice or a join-semilattice where the distinction between meet and join is either not essential or clear from the context.

Semilattices can also be characterised as $\{\wedge\}$-algebras where $\wedge$ is a binary operation that must satisfy the following axioms

$$\begin{aligned} \forall x, y, z\colon\ x \wedge (y \wedge z) &= (x \wedge y) \wedge z && \textit{(associativity)} \\ \forall x, y\colon\ x \wedge y &= y \wedge x && \textit{(commutativity)} \\ \forall x\colon x \wedge x &= x && \textit{(idempotency, or idempotence).} \end{aligned}$$

Clearly, the operation $\wedge^{\mathfrak{S}}$, defined as above in a semilattice $\mathfrak{S}$ viewed as a poset, satisfies these axioms. Conversely, if $(S; \wedge)$ is a semilattice, then the formula $x \wedge y = x$ defines a partial order on S which is a meet-semilattice (and $x \wedge y = y$ defines a partial order on S which is a join-semilattice).

Note that the two ways of formalising semilattices differ when it comes to the notion of a substructure; a *subsemilattice* is referring to the substructure of a semilattice when formalised as an algebraic structure. △

Example 8.5 (Lattice)**.** A *lattice* $\mathfrak{L}$ is a $\{\leq\}$-structure with domain L such that $\leq^{\mathfrak{L}}$ denotes a partial order such that any two $u, v \in L$ have a largest lower bound $u \wedge v$ and a least upper bound, denoted by $u \vee v$. Lattices can also be characterised as $\{\wedge, \vee\}$-algebras where $\wedge$ and $\vee$ are semilattice operations (Example 8.4) that additionally satisfy

$$\forall x, y\colon\ x \wedge (x \vee y) = x \text{ and } x \vee (x \wedge y) = x \qquad \text{(absorption).}$$

If $\mathfrak{L}$ is a lattice and the operations $\wedge$ and $\vee$ are defined as above for semilattices, then these two operations also satisfy the absorption axiom. Conversely, if we are given an algebra $(S; \wedge, \vee)$ satisfying the mentioned axioms, then the formula $x \wedge y = x$ (equivalently, the formula $x \vee y = y$) defines a partial order on S which is a lattice. Of course, there is potential danger of confusion of the symbols for lattice operations $\wedge$ and $\vee$ with the propositional connectives $\wedge$ for conjunction and $\vee$ for disjunction (which can be seen as lattice operations on the set $\{0,1\}$) which luckily should not cause trouble here. A lattice $\mathfrak{L} = (L; \wedge, \vee)$ is called *distributive* if it satisfies

$$\forall x, y \colon\ x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z) \qquad \text{(distributivity).}$$

△

Exercises.

138. Let $\mathbf{A} = (A; +, -, 0)$ be an abelian group. Let R be the set of all endomorphisms of $\mathbf{A}$. Prove the following operations define a ring $\mathbf{R}$ on R: addition is defined pointwise, and multiplication is defined as function composition. The constant $0^{\mathbf{R}}$ denotes the endomoprphism which is constant 0, and the constant $1^{\mathbf{R}}$ denotes the identity.

139. Formalise modules (Example 8.3) as two-sorted structures as introduced in Section 5.1.

The clone of an algebra. If $\mathbf{A}$ is an algebra with the signature τ, then a τ-term $t(x_1, \dots, x_n)$ gives rise to a *term operation* $t^{\mathbf{A}} \colon A^n \to A$; the value of $t^{\mathbf{A}}$ at $a_1, \dots, a_n \in A$ can be obtained by replacing the variables $x_1, \dots, x_n$ by $a_1, \dots, a_n$ and evaluating in $\mathbf{A}$.

Example 8.6. If $\mathbf{A}$ is a group, then the term operation for the term $(x \circ y^{-1}) \circ z$ is a Maltsev operation on A. △

Example 8.7. If $t(x_1, x_2)$ is the term that just consists of the variable x_1, then $t^{\mathbf{A}}$ equals the projection π^2_1. △

Algebras give rise to clones in the following way. We denote by $\mathrm{Clo}(\mathbf{A})$ the set of all term operations of $\mathbf{A}$ of arity at least one. Clearly, $\mathrm{Clo}(\mathbf{A})$ is an operation clone since it is closed under compositions, and contains the projections.

Polymorphism algebras. In the context of complexity classification of CSPs, algebras arise as follows.

Definition 8.8. Let $\mathfrak{B}$ be a relational structure with domain B. An algebra $\mathbf{B}$ with domain B such that $\mathrm{Clo}(\mathbf{B}) = \mathrm{Pol}(\mathfrak{B})$ is called a *polymorphism algebra of* $\mathfrak{B}$.

Note that a structure $\mathfrak{B}$ has many different polymorphism algebras, since Definition 8.8 does not prescribe how to assign function symbols to the polymorphisms of $\mathfrak{B}$.

Any clone $\mathcal{C}$ on a set D can be viewed as an algebra $\mathbf{A}$ with domain D whose signature consists of the operations of $\mathcal{C}$ themselves; that is, if $f \in \mathcal{C}$, then $f^{\mathbf{A}} := f$. We will therefore use concepts defined for algebras also for clones. In particular, the polymorphism clone $\mathrm{Pol}(\mathfrak{B})$ of a structure $\mathfrak{B}$ might be viewed as an algebra, which we refer to as *the* polymorphism algebra of $\mathfrak{B}$. Note that the signature of the polymorphism algebra is always infinite, since we have polymorphisms of arbitrary finite arity.

8.2 Subalgebras, Products, Homomorphic Images

In this section we recall some basic universal-algebraic facts that will be used in the following subsections.

Subalgebras. Let $\mathbf{A}$ be a τ-algebra with domain A. A τ-algebra $\mathbf{B}$ with domain $B \subseteq A$ is called a *subalgebra* of $\mathbf{A}$ if for each $f \in \tau$ of arity k we have $f^{\mathbf{B}}(b_1, \ldots, b_k) = f^{\mathbf{A}}(b_1, \ldots, b_k)$ for all $b_1, \ldots, b_k \in B$; in this case, we write $\mathbf{B} \leq \mathbf{A}$. A *subuniverse* of $\mathbf{A}$ is the domain of some subalgebra of $\mathbf{A}$. Note that as for structures, we do not exclude algebras whose domain is empty (which is of course only possible if the signature does not contain any constant symbols). A subalgebra $\mathbf{B}$ of $\mathbf{A}$ is called *proper* if $\emptyset \neq B \neq A$. The smallest subuniverse of $\mathbf{A}$ that contains a given set $S \subseteq A$ is called the *subuniverse of* $\mathbf{A}$ *generated by* S, and the corresponding subalgebra is called the *subalgebra of* $\mathbf{A}$ *generated by* S, and denoted by $\langle S \rangle_{\mathbf{A}}$.

Products. Let $\mathbf{A}, \mathbf{B}$ be τ-algebras with domain A and B, respectively. Then the *product* $\mathbf{A} \times \mathbf{B}$ is the τ-algebra with domain $A \times B$ such that for each $f \in \tau$ of arity k we have $f^{\mathbf{A} \times \mathbf{B}}\big((a_1, b_1), \ldots, (a_k, b_k)\big) = \big(f^{\mathbf{A}}(a_1, \ldots, a_k), f^{\mathbf{B}}(b_1, \ldots, b_k)\big)$ for all $a_1, \ldots, a_k \in A$ and $b_1, \ldots, b_k \in B$. More generally, when $(\mathbf{A}_i)_{i \in I}$ is a sequence of τ-algebras, indexed by some set I, then $\prod_{i \in I} \mathbf{A}_i$ is the τ-algebra $\mathbf{A}$ with domain $\prod_{i \in I} A_i$ such that for $a_i^1, \ldots, a_i^k \in A_i$

$$f^{\mathbf{A}}\big((a_i^1)_{i \in I}, \ldots, (a_i^k)_{i \in I}\big) := \big(f^{\mathbf{A}_i}(a_i^1, \ldots, a_i^k)\big)_{i \in I}.$$

Lemma 8.9. *Let* $\mathbf{A}$ *be the polymorphism algebra of a finite structure* $\mathfrak{A}$*. Then the (domains of the) subalgebras of* $\mathbf{A}^k$ *are precisely the relations that have a primitive positive definition in* $\mathfrak{A}$*.*

Proof. A relation $R \subseteq A^k$ is a subalgebra of $\mathbf{A}^k$ if and only if for all m-ary f in the signature of $\mathbf{A}$ and $t^1, \ldots, t^m \in R$, we have $\big(f(t_1^1, \ldots, t_1^m), \ldots, f(t_k^1, \ldots, t_k^m)\big) \in R$, which is the case if and only if R is preserved by all polymorphisms of $\mathfrak{A}$, which is the case if and only if R is primitive positive definable in $\mathfrak{A}$ by Theorem 6.6. $\square$

Homomorphic Images. Let $\mathbf{A}$ and $\mathbf{B}$ be τ-algebras. Then a *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$ is a mapping $h \colon A \to B$ such that for all k-ary $f \in \tau$ and $a_1, \ldots, a_k \in A$ we have

$$h\big(f^{\mathbf{A}}(a_1, \ldots, a_k)\big) = f^{\mathbf{B}}\big(h(a_1), \ldots, h(a_k)\big).$$

Note that if h is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ then the image of h is the domain of a subalgebra of $\mathbf{B}$, which is called a *homomorphic image* of $\mathbf{A}$.

Definition 8.10. A *congruence* of an algebra $\mathbf{A}$ is an equivalence relation that is preserved by all operations in $\mathbf{A}$.

Lemma 8.11. *Let* $\mathfrak{B}$ *be a finite structure, and* $\mathbf{B}$ *be a polymorphism algebra of* $\mathfrak{B}$*. Then the congruences of* $\mathbf{B}$ *are exactly the primitively positively definable equivalence relations over* $\mathfrak{B}$*.*

Proof. A direct consequence of Theorem 6.6. $\square$

Proposition 8.12 (see [42])**.** *Let* $\mathbf{A}$ *be an algebra. Then* E *is a congruence of* $\mathbf{A}$ *if and only if* E *is the kernel of a homomorphism from* $\mathbf{A}$ *to some other algebra* $\mathbf{B}$*.*

Example 8.13. Let $G = (V, E)$ be the undirected graph with $V = \{a_1, \ldots, a_4, b_1, \ldots, b_4\}$ such that $a_1, \ldots, a_4$ and $b_1, \ldots, b_4$ induce a clique, for each $i \in \{1, \ldots, 4\}$ there is an edge between a_i and b_i, and otherwise there are no edges in G. Let $\mathbf{A}$ be a polymorphism algebra of G. Then $\mathbf{A}$ homomorphically maps to a two-element algebra $\mathbf{B}$. By Proposition 8.12, it suffices to show that $\mathbf{A}$ has a congruence with two equivalence classes. By Lemma 8.11, it suffices to show that an equivalence relation of index two is primitive positive definable. Here is the primitive positive definition:

$$\exists u, v \big(E(x,u) \wedge E(y,u) \wedge E(x,v) \wedge E(y,v) \wedge E(u,v)\big)$$

The equivalence classes of this relation are precisely $\{a_1, \dots, a_4\}$ and $\{b_1, \dots, b_4\}$. $\triangle$

Example 8.14. Let $\mathbf{A}$ be the algebra with domain

$$A := S_3 = \big\{ \mathrm{id}, (231), (312), (12), (23), (13) \big\}$$

(the symmetric group on three elements), and a single binary operation, the composition function of permutations. Note that $\mathbf{A}$ has the subalgebra induced by $\big\{ \mathrm{id}, (123), (321) \big\}$. Also note that $\mathbf{A}$ homomorphically maps to $(\{0,1\}, +)$ where $+$ is addition modulo 2: the preimage of 0 is $\{\mathrm{id}, (123), (321)\}$ and the preimage of 1 is $\{(12), (23), (13)\}$. $\triangle$

When $\mathbf{A}$ is a τ-algebra, and $h \colon A \to B$ is a mapping such that the kernel of h is a congruence of $\mathbf{A}$, we define the *quotient algebra* $\mathbf{A}/h$ *of* $\mathbf{A}$ *under* h to be the algebra with domain $h(A)$ where

$$f^{\mathbf{A}/h}(h(a_1), \dots, h(a_k)) = h(f^{\mathbf{A}}(a_1, \dots, a_k))$$

where $a_1, \dots, a_k \in A$ and $f \in \tau$ is k-ary. This is well-defined since the kernel of h is preserved by all operations of $\mathbf{A}$. Note that h is a surjective homomorphism from $\mathbf{A}$ to $\mathbf{A}/h$. The following is well known (see e.g. Theorem 6.3 in [42]).

Lemma 8.15. *Let* $\mathbf{A}$ *and* $\mathbf{B}$ *be algebras with the same signature, and let* $h \colon \mathbf{A} \to \mathbf{B}$ *be a homomorphism. Then the image of any subalgebra* $\mathbf{A}'$ *of* $\mathbf{A}$ *under* h *is a subalgebra of* $\mathbf{B}$*, and the preimage of any subalgebra* $\mathbf{B}'$ *of* $\mathbf{B}$ *under* h *is a subalgebra of* $\mathbf{A}$*.*

Proof. Let $f \in \tau$ be k-ary. Then for all $a_1, \dots, a_k \in A'$,

$$f^{\mathbf{B}}(h(a_1), \dots, h(a_k)) = h(f^{\mathbf{A}}(a_1, \dots, a_k)) \in h(A') ,$$

so $h(A')$ is a subalgebra of $\mathbf{B}$. Now suppose that $h(a_1), \dots, h(a_k)$ are elements of B'; then $f^{\mathbf{B}}(h(a_1), \dots, h(a_k)) \in B'$ and hence $h(f^{\mathbf{A}}(a_1, \dots, a_k)) \in B'$. So, $f^{\mathbf{A}}(a_1, \dots, a_k) \in h^{-1}(B')$ which shows that $h^{-1}(B')$ induces a subalgebra of $\mathbf{A}$. $\square$

Exercices.

140. Show that for all τ-algebras $\mathbf{A}$ and $\mathbf{B}$ with $\mathrm{Clo}(\mathbf{A}) = \mathrm{Clo}(\mathbf{B})$ we have $\mathrm{Clo}(\mathbf{A}^2) = \mathrm{Clo}(\mathbf{B}^2)$.

141. Find a relational signature τ and τ-structures $\mathbf{A}$, $\mathbf{B}$ such that $\mathrm{Clo}(\mathbf{A}) = \mathrm{Clo}(\mathbf{B})$ but $\mathrm{Clo}(\mathbf{A} \times \mathbf{B}) \neq \mathrm{Clo}(\mathbf{A} \times \mathbf{A})$.

142. Prove Proposition 8.12.

143. Consider the algebra $\mathbf{A}_n := (\{0, \dots, n-1\}; m)$ where $m(x,y,z) := x - y + z$. Then for every $k \geq 1$ the clone $\mathrm{Clo}(\mathbf{A}_n)^{(k)}$ consists of precisely the operations defined as

$$g(x_0, \dots, x_{k-1}) := \sum_i a_i x_i$$

where $a_0, \dots, a_{k-1} \in \mathbb{Z}$ with $\sum_i a_i = 1$.

144. Let p be a positive prime. Show that the only proper subalgebras of $\mathbf{A}_p$ from the previous exercise are of the form $\{a\}$ for some $a \in \{0, \dots, p-1\}$.

Hint. Use Exercise 143 and Example 6.9.

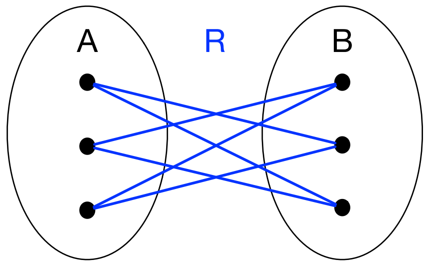

Figure 13: Illustration of $E(K_3) \leq \mathbf{A}^2$, where $\mathrm{Clo}(\mathbf{A}) = \mathrm{Pol}(K_3)$, as a bipartite graph.

8.3 Algebras and CSPs

Let $\mathbf{A}$ and $\mathbf{B}$ be algebras with the same signature, and let $R \leq \mathbf{A} \times \mathbf{B}$ be a subalgebra. The relation R can be viewed as the edge relation of a bipartite graph with colour classes A and B. Note that if $\mathbf{A} = \mathbf{B}$ and $\mathrm{Clo}(\mathbf{A}) = \mathrm{Pol}(\mathfrak{A})$, then the relations R that arise in this way are precisely the binary relations on A that are primitively positively definable in $\mathfrak{A}$. For example, if $\mathrm{Clo}(\mathbf{A}) = \mathrm{Pol}(K_3)$, then $E(K_3) \leq \mathbf{A}^2$ and the corresponding bipartite graph is drawn in Figure 13.

The importance of the set-up of $R \leq \mathbf{A} \times \mathbf{B}$ for CSPs is that we may imagine A as the possible values for a variable x in an instance of the CSP, and B as the possible values for a variable y in the CSP, and R represents a binary constraint between x and y. The advantage of this perspective is that many important definitions are very intuitively phrased in the language of bipartite graphs.

We start with the following fundamental definition from universal algebra which is highly relevant for the universal-algebraic approach to CSPs, in particular in Section 13.

Definition 8.16. Let $k \geq 1$ and $\mathbf{A}_1, \dots, \mathbf{A}_k$ be τ-algebras. Then $\mathbf{R} \leq \mathbf{A}_1 \times \cdots \times \mathbf{A}_k$ is called *subdirect* if $\pi_i(R) = A_i$ for every $i \in \{1, \dots, k\}$.

If $\mathbf{A}$ is the polymorphism algebra of a finite digraph H, then there is a link between the notion of subdirect subalgebras of $\mathbf{A}^2$ and arc consistency. Let G be a finite digraph and let $L(x) \subseteq V(H)$ be the list for $x \in V(G)$ at the final stage of the evaluation of $\mathrm{AC}_H(G)$. Note that for every $x \in V(G)$, the set $L(x)$ is a subuniverse of $\mathbf{A}$. Also note that every $(x, y) \in E(G)$ we have that $E(H) \cap (L(x) \times L(y))$ is subdirect in $L(x) \times L(y)$.

Also note that $E(H)$ is subdirect in $\mathbf{A}^2$ if and only if H has no sources and no sinks. Digraphs without sources and sinks are also called *smooth*.

Let $x, y \in V(G)$ and let $L(x)$ and $L(y)$ be the lists computed by the arc consistency procedure. Recall from Exercise 56 that $L(x)$ and $L(y)$ are subuniverses of the polymorphism algebra $\mathbf{A}$ of H. Note that $E(G) \cap (L(x) \times L(y)) \leq \mathbf{A}^2$ is subdirect!

Exercices.

145. Show that a digraph $G = (V, E)$ is rectangular if and only if E, when regarded as a bipartite graph with color classes A and B as described in this section, is a disjoint union of *bicliques*, i.e., if $a \in A$ has a path to $b \in B$, then $(a, b) \in E$.

8.4 Pseudovarieties and Varieties

Varieties are a fascinatingly powerful concept to study classes of algebras. The fundamental result about varieties is Birkhoff's theorem, which links varieties with equational theories (Section 8.5). By Birkhoff's theorem, there is also a close relationship between varieties and the concept of an *abstract clone* (Section 8.6).

If $\mathcal{K}$ is a class of algebras of the same signature, then

- $\mathrm{P}(\mathcal{K})$ denotes the class of all products of algebras from $\mathcal{K}$.
- $\mathrm{P}^{\mathrm{fin}}(\mathcal{K})$ denotes the class of all finite products of algebras from $\mathcal{K}$.
- $\mathrm{S}(\mathcal{K})$ denotes the class of all subalgebras of algebras from $\mathcal{K}$.
- $\mathrm{H}(\mathcal{K})$ denotes the class of all homomorphic images of algebras from $\mathcal{K}$.

Note that closure under homomorphic images implies in particular closure under isomorphism. For the operators P, $\mathrm{P}^{\mathrm{fin}}$, S and H we often omit the brackets when applying them to single singleton classes that just contain one algebra, i.e., we write $\mathrm{H}(\mathbf{A})$ instead of $\mathrm{H}(\{\mathbf{A}\})$. The elements of $\mathrm{HS}(\mathbf{A})$ are also called the *factors* of $\mathbf{A}$.

A class $\mathcal{V}$ of algebras with the same signature τ is called a *pseudovariety* if $\mathcal{V}$ contains all homomorphic images, subalgebras, and direct products of algebras in $\mathcal{V}$, i.e., $\mathrm{H}(\mathcal{V}) = \mathrm{S}(\mathcal{V}) = \mathrm{P}^{\mathrm{fin}}(\mathcal{V}) = \mathcal{V}$. The class $\mathcal{V}$ is called a *variety* if $\mathcal{V}$ also contains all (finite and infinite) products of algebras in $\mathcal{V}$. So the only difference between pseudovarieties and varieties is that pseudovarieties need not be closed under direct products of infinite cardinality. The smallest pseudovariety (variety) that contains an algebra $\mathbf{A}$ is called the pseudovariety (variety) *generated* by $\mathbf{A}$.

Lemma 8.17 (HSP lemma)**.** *Let* $\mathbf{A}$ *be an algebra.*

- *The pseudovariety generated by* $\mathbf{A}$ *equals* $\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})$.
- *The variety generated by* $\mathbf{A}$ *equals* $\mathrm{HSP}(\mathbf{A})$.

Proof. Clearly, $\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})$ is contained in the pseudovariety generated by $\mathbf{A}$, and $\mathrm{HSP}(\mathbf{A})$ is contained in the variety generated by $\mathbf{A}$. For the converse inclusion, it suffices to verify that $\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})$ is closed under H, S, and $\mathrm{P}^{\mathrm{fin}}$. It is clear that $\mathrm{H}(\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})) = \mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})$. The second part of Lemma 8.15 implies that $\mathrm{S}(\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})) \subseteq \mathrm{HS}(\mathrm{SP}^{\mathrm{fin}}(\mathbf{A})) = \mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})$. Finally,

$$\mathrm{P}^{\mathrm{fin}}(\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})) \subseteq \mathrm{H}\,\mathrm{P}^{\mathrm{fin}}\,\mathrm{S}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{A}) \subseteq \mathrm{HSP}^{\mathrm{fin}}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{A}) = \mathrm{HSP}^{\mathrm{fin}}(\mathbf{A}) \,.$$

The proof that $\mathrm{HSP}(\mathbf{A})$ is closed under H, S, and P is analogous. □

Pseudo-varieties are linked to primitive positive interpretability from Section 5.7.

Theorem 8.18. *Let* $\mathfrak{C}$ *be a finite structure with polymorphism algebra* $\mathbf{C}$. *Then* $\mathfrak{B} \in \mathrm{I}(\mathfrak{C})$ *if and only if there exists* $\mathbf{B} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{C})$ *such that* $\mathrm{Clo}(\mathbf{B}) \subseteq \mathrm{Pol}(\mathfrak{B})$.

Proof. We only prove the 'if' part of the statement here; the proof of the 'only if' part is similarly easy. There exists a finite number $d \geq 1$, a subalgebra $\mathbf{D}$ of $\mathbf{C}^d$, and a surjective homomorphism h from $\mathbf{D}$ to $\mathbf{B}$. We claim that $\mathfrak{B}$ has a primitive positive interpretation I of dimension d in $\mathfrak{C}$. All operations of $\mathbf{C}$ preserve D (viewed as a d-ary relation over $\mathfrak{C}$), since $\mathbf{D}$

is a subalgebra of $\mathbf{C}^d$. By Theorem 6.6, this implies that D has a primitive positive definition $\delta(x_1, \dots, x_d)$ in $\mathfrak{C}$, which becomes the domain formula δ_I of I. As coordinate map we choose the mapping h. Since h is an algebra homomorphism, the kernel K of h is a congruence of $\mathbf{D}$. It follows that K, viewed as a $2d$-ary relation over C, is preserved by all operations from $\mathbf{C}$. Theorem 6.6 implies that K has a primitive positive definition in $\mathfrak{C}$. This definition becomes the formula $=_I$. Finally, let R be a relation of $\mathfrak{B}$ and let f be a function symbol from the signature of $\mathbf{B}$. By assumption, $f^{\mathbf{B}}$ preserves R. It is easy to verify that then $f^{\mathbf{C}}$ preserves $h^{-1}(R)$. Hence, all polymorphisms of $\mathfrak{C}$ preserve $h^{-1}(R)$, and the relation $h^{-1}(R)$ has a primitive positive definition in $\mathfrak{C}$ (Theorem 6.6), which becomes the defining formula for the atomic formula $R(x_1, \dots, x_k)$ in I. This concludes our construction of the primitive positive interpretation I of $\mathfrak{B}$ in $\mathfrak{C}$. □

Primitive positive bi-interpretability can also be characterised with the varieties and pseudo-varieties generated by polymorphism algebras. The following is a special case of Proposition 25 in [27] (where it is proved for a must larger class of countable structures).

Proposition 8.19. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures with finite domains. Then the following are equivalent.*

- *there are polymorphism algebras $\mathbf{A}$ of $\mathfrak{B}$ and $\mathbf{B}$ of $\mathfrak{B}$ such that* $\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A}) = \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$*;*
- *$\mathfrak{A}$ and $\mathfrak{B}$ are primitively positively bi-interpretable.*

Proof. For the forward implication, we assume that there is a $d_1 \geq 1$, a subalgebra $\mathbf{S}_1$ of $\mathbf{A}^{d_1}$, and a surjective homomorphism h_1 from $\mathbf{S}_1$ to $\mathbf{B}$. Moreover, we assume that there is a $d_2 \geq 1$, a subalgebra $\mathbf{S}_2$ of $\mathbf{B}^{d_2}$, and a surjective homomorphisms h_2 from $\mathbf{S}_2$ to $\mathbf{A}$. The proof of Theorem 8.18 shows that $I_1 := (d_1, S_1, h_1)$ is an interpretation of $\mathfrak{B}$ in $\mathfrak{A}$, and $I_2 := (d_2, S_2, h_2)$ is an interpretation of $\mathfrak{A}$ in $\mathfrak{B}$. Because the statement is symmetric it suffices to show that the (graph of the) function $h_1 \circ h_2 \colon (S_2)^{d_1} \to B$ defined by

$$(y_{1,1}, \dots, y_{1,d_2}, \dots, y_{d_1,1}, \dots, y_{d_1,d_2}) \mapsto h_1(h_2(y_{1,1}, \dots, y_{1,d_2}), \dots, h_2(y_{d_1,1}, \dots, y_{d_1,d_2}))$$

is primitively positively definable in $\mathfrak{B}$. Theorem 6.6 asserts that this is equivalent to showing that $h_1 \circ h_2$ is preserved by all operations $f^{\mathbf{B}}$ of $\mathbf{B}$. So let k be the arity of $f^{\mathbf{B}}$ and let $b^i = (b^i_1, \dots, b^i_{d_1})$ be elements of $(S_2)^{d_1}$, for $1 \leq i \leq k$. Then indeed

$$\begin{aligned}
& f^{\mathbf{B}}((h_1 \circ h_2)(b^1), \dots, (h_1 \circ h_2)(b^k)) \\
&= h_1\big(f^{\mathbf{A}}(h_2(b^1_1), \dots, h_2(b^k_1)), \dots, f^{\mathbf{A}}(h_2(b^1_{d_1}), \dots, h_2(b^k_{d_1}))\big) \\
&= (h_1 \circ h_2)(f^{\mathbf{B}}(b^1, \dots, b^k)) \,.
\end{aligned}$$

For the backwards implication, suppose that $\mathfrak{A}$ and $\mathfrak{B}$ are primitive positive bi-interpretable via an interpretation $I_1 = (d_1, S_1, h_1)$ of $\mathfrak{B}$ in $\mathfrak{A}$ and an interpretation $I_2 = (d_2, S_2, h_2)$ of $\mathfrak{A}$ in $\mathfrak{B}$. Let $\mathbf{A}$ be a polymorphism algebra of $\mathfrak{A}$. The proof of Theorem 8.18 shows that S_1 induces an algebra $\mathbf{S}_1$ in $\mathbf{A}^{d_1}$ and h_1 is a surjective homomorphism from $\mathbf{S}_1$ to an algebra $\mathbf{B}$ satisfying $\mathrm{Clo}(\mathbf{B}) \subseteq \mathrm{Pol}(\mathfrak{B})$. Similarly, S_2 is the domain of a subalgebra $\mathbf{S}_2$ of $\mathbf{B}^{d_2}$ and h_2 is a homomorphism from $\mathbf{S}_2$ onto an algebra $\mathbf{A}'$ such that $\mathrm{Clo}(\mathbf{A}') \subseteq \mathrm{Pol}(\mathfrak{A})$.

We claim that $\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A}) = \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$. The inclusion '$\supseteq$' is clear since $\mathbf{B} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})$. For the reverse inclusion it suffices to show that $\mathbf{A} = \mathbf{A}'$ since $\mathbf{A}' \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$. Let $f \in \tau$ be

k-ary; we show that $f^{\mathbf{A}} = f^{\mathbf{A}'}$. Let $a_1, \dots, a_k \in A$. Since $h_2 \circ h_1$ is surjective onto A, there are $c^i = (c^i_{1,1}, \dots, c^i_{d_1,d_2}) \in A^{d_1 d_2}$ such that $a_i = h_2 \circ h_1(c^i)$. Then

$$\begin{aligned}
f^{\mathbf{A}'}(a_1, \dots, a_k) &= f^{\mathbf{A}'}(h_2 \circ h_1(c^1), \dots, h_2 \circ h_1(c^k)) \\
&= h_2\big(f^{\mathbf{B}}(h_1(c^1_{1,1}, \dots, c^1_{d_1,1}), \dots, h_1(c^k_{1,1}, \dots, c^k_{d_1,1})), \dots, \\
&\qquad f^{\mathbf{B}}(h_1(c^1_{1,d_2}, \dots, c^1_{d_1,d_2}), \dots, h_1(c^k_{1,d_2}, \dots, c^k_{d_1,d_2}))\big) \\
&= h_2 \circ h_1(f^{\mathbf{A}}(c^1, \dots, c^k)) \\
&= f^{\mathbf{A}}(h_2 \circ h_1(c^1), \dots, h_2 \circ h_1(c^k)) \\
&= f^{\mathbf{A}}(a_1, \dots, a_k)
\end{aligned}$$

where the second and third equations hold since h_2 and h_1 are algebra homomorphisms, and the fourth equation holds because $f^{\mathbf{A}}$ preserves $h_2 \circ h_1$, because $I_2 \circ I_1$ is pp-homotopic to the identity. $\square$

Exercices.

146. Show that an algebra has the empty algebra as a subalgebra if and only if the signature does not contain constants (i.e., function symbols of arity 0).

147. Let B be a subuniverse of an algebra $\mathbf{A}$ generated by $S \subseteq A$. Show that an element $a \in A$ belongs to B if and only if there exists a term $t(x_1, \dots, x_k)$ and elements $s_1, \dots, s_k \in S$ such that $a = t^{\mathbf{A}}(s_1, \dots, s_k)$.

148. Show that the operators HS and SH are distinct.

149. Show that the operators SP and PS are distinct.

8.5 Birkhoff's Theorem

Birkhoff's theorem provides a characterisation of varieties in terms of sets of *identities*. A sentence in a functional signature τ is called a (τ-) *identity* if it is of the form

$$\forall x_1, \dots, x_n \colon s = t$$

where s and t are τ-terms over the variables $x_1, \dots, x_n$ (such sentences are also called *universally conjunctive*). We follow the usual notation in universal algebra and sometimes write such sentences as

$$s \approx t.$$

If $\mathcal{K}$ is a class of τ-algebras, then we say that $\mathcal{K}$ *satisfies* $s \approx t$ (or: $s \approx t$ *holds in* $\mathcal{K}$), in symbols $\mathcal{K} \models s \approx t$, if every algebra in $\mathcal{K}$ satisfies $s \approx t$.

Theorem 8.20 (Birkhoff [19]; see e.g. [64] or [42])**.** *Let τ be a functional signature, let $\mathcal{K}$ be a class of τ-algebras, and let $\mathbf{A}$ be a τ-algebra. Then the following are equivalent.*

1. *All identities that hold in $\mathcal{K}$ also hold in $\mathbf{A}$;*

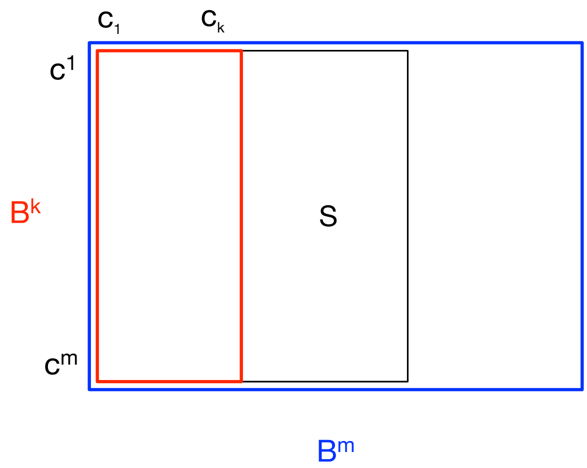

Figure 14: Illustration for the proof of Birkhoff's theorem

.

2. $\mathbf{A} \in \mathrm{HSP}(\mathcal{K})$*.*

If $\mathbf{A}$ *has a finite domain, and* $\mathcal{K} = \{\mathbf{B}\}$ *for some algebra* $\mathbf{B}$ *with a finite domain, then this is also equivalent to*

3. $\mathbf{A} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$*.*

Proof. To show that 2. implies 1., let $s(x_1, \ldots, x_n) \approx t(x_1, \ldots, x_n)$ be an identity that holds in $\mathcal{K}$. Then $s \approx t$ is preserved in products $\mathbf{A} = \prod_{i \in I} \mathbf{B}_i$ of algebras $\mathbf{B}_i \in \mathcal{K}$. To see this, let $a_1, \ldots, a_n \in A$ be arbitrary. Since $\mathbf{B} \models \phi$ we have $s^{\mathbf{B}_i}(a_1[i], \ldots, a_n[i]) = t^{\mathbf{B}_i}(a_1[i], \ldots, a_n[i])$ for all $j \in I$, and thus $s^{\mathbf{A}}(a_1, \ldots, a_n) = t^{\mathbf{A}}(a_1, \ldots, a_n)$ by the definition of products. Since $a_1, \ldots, a_n$ were chosen arbitrarily, we have $\mathbf{A} \models \phi$. Moreover, universal sentences are preserved by taking subalgebras. Finally, suppose that $\mathbf{B}$ is an algebra that satisfies $s \approx t$, and μ is a surjective homomorphism from $\mathbf{B}$ to some algebra $\mathbf{A}$. Let $a_1, \ldots, a_n \in A$. By the surjectivity of μ we can choose $b_1, \ldots, b_n$ such that $\mu(b_i) = a_i$ for all $i \leq n$. Then

$$\begin{aligned} s^{\mathbf{B}}(b_1, \ldots, b_n) = t^{\mathbf{B}}(b_1, \ldots, b_n) \Rightarrow \quad & \mu(s^{\mathbf{B}}(b_1, \ldots, b_n)) = \mu(t^{\mathbf{B}}(b_1, \ldots, b_n)) \\ \Rightarrow \quad & t^{\mathbf{A}}(\mu(b_1), \ldots, \mu(b_n)) = s^{\mathbf{A}}(\mu(b_1), \ldots, \mu(b_n)) \\ \Rightarrow \quad & t^{\mathbf{A}}(a_1, \ldots, a_n) = s^{\mathbf{A}}(a_1, \ldots, a_n) \; . \end{aligned}$$

We only show the implication from 1. to 3. (and hence to 2.) if $\mathbf{A}$ and $\mathbf{B}$ have finite domains and $\mathcal{K} = \{\mathbf{B}\}$; the proof of the general case is similar (see Exercise 151). Let $a_1, \ldots, a_k$ be the elements of $\mathbf{A}$, define $m := |B|^k$ and $C := B^k$. Let $c^1, \ldots, c^m$ be the elements of C; write c_i for $(c_i^1, \ldots, c_i^m)$. Let $\mathbf{S}$ be the smallest subalgebra of $\mathbf{B}^m$ that contains $c_1, \ldots, c_k$; so the elements of $\mathbf{S}$ are precisely those of the form $t^{\mathbf{B}^m}(c_1, \ldots, c_k)$, for a k-ary τ-term t. See Figure 14.

Define $\mu \colon S \to A$ by

$$\mu(t^{\mathbf{B}^m}(c_1, \ldots, c_k)) := t^{\mathbf{A}}(a_1, \ldots, a_k).$$

Claim 1: μ is well-defined. Suppose that $t^{\mathbf{B}^m}(c_1,\ldots,c_k) = s^{\mathbf{B}^m}(c_1,\ldots,c_k)$; then $t^{\mathbf{B}} = s^{\mathbf{B}}$ by the choice of S, and by assumption we have $t^{\mathbf{A}}(a_1,\ldots,a_k) = s^{\mathbf{A}}(a_1,\ldots,a_k)$.

Claim 2: μ is surjective. For all $i \leq k$, the element c_i is mapped to a_i.

Claim 3: μ is a homomorphism from $\mathbf{S}$ to $\mathbf{A}$. Let $f \in \tau$ be of arity n and let $s_1,\ldots,s_n \in S$. For $i \leq n$, write $s_i = t_i^{\mathbf{S}}(c_1,\ldots,c_k)$ for some τ-term t_i (see Exercise 147). Then

$$\begin{aligned}
\mu\big(f^{\mathbf{S}}(s_1,\ldots,s_n)\big) &= \mu\big(f^{\mathbf{S}}(t_1^{\mathbf{S}}(c_1,\ldots,c_k),\ldots,t_n^{\mathbf{S}}(c_1,\ldots,c_k))\big) \\
&= \mu\big(f^{\mathbf{S}}(t_1^{\mathbf{S}},\ldots,t_n^{\mathbf{S}})(c_1,\ldots,c_k)\big) \\
&= \mu\big((f(t_1,\ldots,t_n))^{\mathbf{S}}(c_1,\ldots,c_k)\big) \\
&= \big(f(t_1,\ldots,t_n)\big)^{\mathbf{A}}(a_1,\ldots,a_k) \\
&= f^{\mathbf{A}}\big(t_1^{\mathbf{A}}(a_1,\ldots,a_k),\ldots,t_n^{\mathbf{A}}(a_1,\ldots,a_k)\big) \\
&= f^{\mathbf{A}}(\mu(s_1),\ldots,\mu(s_n)).
\end{aligned}$$

Therefore, $\mathbf{A}$ is the homomorphic image of the subalgebra $\mathbf{S}$ of $\mathbf{B}^m$, and so $\mathbf{A} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$. $\square$

Theorem 8.20 is important for analysing the constraint satisfaction problem for a structure $\mathfrak{B}$, since it can be used to transform the 'negative' statement of not interpreting certain finite structures, which is equivalent to not having a certain finite algebra in the pseudo-variety generated by a polymorphism algebra of $\mathfrak{B}$, into a 'positive' statement of having polymorphisms satisfying non-trivial identities. We will learn several concrete identities that must be satisfied in later sections, e.g., in Section 9.5, Section 10.2, Section 14.2, and Section 14.3.

8.5.1 The Free Algebra

In the following we extract an important idea from the proof of Birkhoff's theorem and present it in different words which will be useful later. Fix a functional signature τ and a class of τ-algebras $\mathcal{K}$.

Definition 8.21. Let $\mathbf{F}$ be a τ-algebra generated by $X \subseteq F$. We say that $\mathbf{F}$ has the *universal mapping property* for $\mathcal{K}$ over X if for every $\mathbf{A} \in \mathcal{K}$ and $f \colon X \to A$ there exists a (unique) extension of f to a homomorphism from $\mathbf{F}$ to $\mathbf{A}$.

Proposition 8.22 (Uniqueness)**.** *Suppose that $\mathbf{F}_1, \mathbf{F}_2 \in \mathcal{K}$ have the universal mapping property for $\mathcal{K}$ over X_i, for $i \in \{1,2\}$. If $|X_1| = |X_2|$ then $\mathbf{F}_1$ and $\mathbf{F}_2$ are isomorphic.*

Proof. Fix any bijection between X_1 and X_2; the bijection has a unique extension to an isomorphism between $\mathbf{F}_1$ and $\mathbf{F}_2$. $\square$

Proposition 8.23 (Existence)**.** *For every class $\mathcal{K}$ of τ-algebras and and for every set X there exists a τ-algebra $\mathbf{F} \in \mathrm{SP}(\mathcal{K})$ which has the universal mapping property for* $\mathrm{HSP}(\mathcal{K})$ *over X.*

By Proposition 8.22, the algebra $\mathbf{F} \in \mathrm{SP}(\mathcal{K})$ is unique up to isomorphism and called the *free algebra for $\mathcal{K}$ over X*, and will be denoted by $\mathbf{F}_{\mathcal{K}}(X)$.

Lemma 8.24. *Let $\mathbf{F}$ be free for $\mathcal{K}$ over $\{x_1,\ldots,x_n\}$ and let $s(y_1,\ldots,y_n)$, $t(y_1,\ldots,y_n)$ be τ-terms. Then the following are equivalent.*

1. $\mathcal{K} \models s(y_1,\ldots,y_n) \approx t(y_1,\ldots,y_n)$

2. $s(y_1,\dots,y_n) \approx t(y_1,\dots,y_n)$ *holds in* $\mathbf{F}$;

3. $s^{\mathbf{F}}(x_1,\dots,x_n) = t^{\mathbf{F}}(x_1,\dots,x_n)$.

Proof. 1. $\Rightarrow$ 2. : If $s \approx t$ holds in every algebra of $\mathcal{K}$, then it also holds in products and subalgebras of algebras in $\mathcal{K}$, and hence also in $\mathbf{F}$.

2. $\Rightarrow$ 3. holds trivially.

3. $\Rightarrow$ 1. and Let $\mathbf{A} \in \mathcal{K}$. If $a_1,\dots,a_n \in A$, then the map that sends x_i to a_i for all $i \in \{1,\dots,n\}$ can be extended to a homomorphism from $\mathbf{F}$ to $\mathbf{A}$, and since $s^{\mathbf{F}}(x_1,\dots,x_n) = t^{\mathbf{F}}(x_1,\dots,x_n)$ we have $s^{\mathbf{A}}(a_1,\dots,a_n) = t^{\mathbf{A}}(a_1,\dots,a_n)$. Since $a_1,\dots,a_n \in A$ were chosen arbitrarily, this shows that $\mathbf{A} \models s(y_1,\dots,y_n) \approx t(y_1,\dots,y_n)$. $\square$

Note that if $\mathcal{K} := \mathrm{HSP}(\mathbf{B})$, and B and X are finite, then $\mathbf{F}_{\mathcal{K}}(X) \leq \mathbf{B}^{B^X}$ is finite as well.

8.5.2 Equational Theories

Birkhoff's theorem provides for every class $\mathcal{K}$ of τ-algebras a characterisation of the class of all τ-algebras that satisfies all identities satisfied by $\mathcal{K}$. Conversely, there is for every set of τ-identities Σ a syntactic characterisation of the set of all τ-identities that are satisfied in all algebras that satisfy Σ (i.e., we have a proof-theoretic characterisation of the set of all universal conjunctive consequences of Σ; the this can be seen as a special case of the completeness theorem of first-order logic for universally conjunctive sentences in an algebraic signature τ).

Theorem 8.25. *A τ-identity $f \approx g$ is implied by Σ (i.e., holds in all algebras that satisfy Σ) if and only if (f,g) is contained in the smallest equivalence relation E which*

- *contains (r,s) for every $(r \approx s) \in \Sigma$,*
- *is compatible: if $f \in \tau$ is a function symbol of arity n and $(r_1,s_1),\dots,(r_n,s_n) \in E$, then $(f(r_1,\dots,r_n), f(s_1,\dots,s_n)) \in E$,*
- *is fully invariant: if $(r(x_1,\dots,x_n), s(x_1,\dots,x_n)) \in E$ and $t_1,\dots,t_n$ are τ-terms, then $(r(t_1,\dots,t_n), s(t_1,\dots,t_n)) \in E$.*

Exercices.

150. Prove Proposition 8.23.

151. Show the implication from 1. to 2. in Birkhoff's theorem in full generality.

8.6 Abstract Clones

Clones (in the literature often *abstract clones*) relate to operation clones in the same way as (abstract) groups relate to permutation groups: the elements of a clone correspond to the functions of an operation clone, and the signature contains composition symbols to code how functions compose. Since an operation clone contains functions of various arities, a clone will be formalized as a multi-sorted structure, with a sort for each arity.

Definition 8.26. An (abstract) *clone* $\mathbf{C}$ is a multi-sorted structure with sorts $\{C^{(i)} \mid i \in \mathbb{N}_{\geq 1}\}$ and the signature $\{\pi_i^k \mid 1 \leq i \leq k\} \cup \{\mathrm{comp}_l^k \mid k, l \geq 1\}$. The elements of the sort $C^{(k)}$ will be called the *k-ary operations* of $\mathbf{C}$. We denote a clone by

$$\mathbf{C} = (C^{(1)}, C^{(2)}, \dots; (\pi_i^k)_{1 \leq i \leq k}, (\mathrm{comp}_l^k)_{k,l \geq 1})$$

and require that π_i^k is a constant in $C^{(k)}$, and that $\mathrm{comp}_l^k \colon C^{(k)} \times (C^{(l)})^k \to C^{(l)}$ is an operation of arity $k+1$. Moreover, it holds that

$$\mathrm{comp}_k^k(f, \pi_1^k, \dots, \pi_k^k) = f \tag{20}$$

$$\mathrm{comp}_l^k(\pi_i^k, f_1, \dots, f_k) = f_i \tag{21}$$

$$\begin{aligned}\mathrm{comp}_l^k\left(f, \mathrm{comp}_l^m(g_1, h_1, \dots, h_m), \dots, \mathrm{comp}_l^m(g_k, h_1, \dots, h_m)\right) = \\ \mathrm{comp}_l^m\left(\mathrm{comp}_m^k(f, g_1, \dots, g_k), h_1, \dots, h_m\right).\end{aligned} \tag{22}$$

The final equation generalises associativity in groups and monoids, and we therefore refer to it by *associativity*. We also write $f(g_1, \dots, g_k)$ instead of $\mathrm{comp}_l^k(f, g_1, \dots, g_k)$ when l is clear from the context. So associativity might be more readable as

$$f(g_1(h_1, \dots, h_m), \dots, g_k(h_1, \dots, h_m)) = f(g_1, \dots, g_k)(h_1, \dots, h_m).$$

Every operation clone $\mathscr{C}$ gives rise to an abstract clone $\mathbf{C}$ in the obvious way: $\pi_i^k \in C^{(k)}$ denotes the k-ary i-th projection in $\mathscr{C}$, and $\mathrm{comp}_l^k(f, g_1, \dots, g_k) \in C^{(l)}$ denotes the composed function $(x_1, \dots, x_l) \mapsto f(g_1(x_1, \dots, x_l), \dots, g_k(x_1, \dots, x_l)) \in \mathscr{C}$. Conversely, every abstract clone arises from an operation clone - this will follow from Proposition 8.34.

Example 8.27. An algebra $\mathbf{A}$ satisfies $f(x_1, x_2) \approx f(x_2, x_1)$ if and only if

$$\mathrm{Clo}(\mathbf{A}) \models \mathrm{comp}_2^2(f^{\mathbf{A}}, \pi_1^2, \pi_2^2) = \mathrm{comp}_2^2(f^{\mathbf{A}}, \pi_2^2, \pi_1^2). \qquad \triangle$$

In the following, we will also use the term ‘abstract clone’ in situations where we want to stress that we are working with a clone and *not* with an operation clone. The notion of a *homomorphism* between clones is just the usual notion of homomorphisms for algebras, adapted to the multi-sorted case. Since we didn’t formally introduce homomorphisms for multi-sorted structures, we spell out the definition in the special case of clones.

Definition 8.28. Let $\mathbf{C}$ and $\mathbf{D}$ be clones. A function $\xi \colon C \to D$ is called a *(clone) homomorphism* if

1. ξ preserves arities of functions, i.e., $\xi(C^{(i)}) \subseteq D^{(i)}$ for all $i \in \mathbb{N}$;
2. $\xi((\pi_i^k)^{\mathbf{C}}) = (\pi_i^k)^{\mathbf{D}}$ for all $1 \leq i \leq k$;
3. $\xi(f(g_1, \dots, g_n)) = \xi(f)(\xi(g_1), \dots, \xi(g_n))$ for all $n, m \geq 1$, $f \in C^{(n)}$, $g_1, \dots, g_n \in C^{(m)}$.

We say that ξ is a *(clone) isomorphism* if ξ is bijective and both ξ and ξ^{-1} is a homomorphism.

Example 8.29. We write **Proj** for the abstract clone of an algebra with at least two elements all of whose operations are projections; note that any such algebra has the same abstract clone (up to isomorphism), and that **Proj** has a homomorphism into any other clone. $\triangle$

Example 8.30. All abstract clones of an algebra on a one-element set are isomorphic, too, but of course not isomorphic to **Proj**. Any clone homomorphically maps to this trivial clone. △

Example 8.31. Using Proposition 6.19, it is easy to see that there exists a clone homomorphism from $\mathrm{Pol}(K_3)$ to **Proj**. △

The following definition plays an important role throughout the later sections in this text.

Definition 8.32 (Star composition)**.** Let $l, m \in \mathbb{N}$ and $n = lm$. We write $f * g$ as a shortcut for

$$\mathrm{comp}^l_n(f, \mathrm{comp}^m_n(g, \pi^n_1, \dots, \pi^n_m), \dots, \mathrm{comp}^m_n(g, \pi^n_{(l-1)m+1}, \dots, \pi^n_n))\,.$$

Note that if $f\colon A^l \to A$ and $g\colon A^m \to A$, then $f * g$ denotes the operation from $A^{lm} \to A$ given by

$$(x_{1,1}, \dots, x_{l,m}) \mapsto f\big(g(x_{1,1}, \dots, x_{1,m}), \dots, g(x_{l,1}, \dots, x_{l,m})\big).$$

8.7 Clone Formulation of Birkhoff's Theorem

One can translate back and forth between varieties and abstract clones.

Definition 8.33 ($\mathrm{Var}(\mathbf{C})$)**.** For any abstract clone $\mathbf{C}$, the variety $\mathrm{Var}(\mathbf{C})$ is defined as follows. We use the elements of C as a functional signature τ, where the elements of $C^{(n)}$ are n-ary function symbols. We consider the set Σ of τ-identities defined as follows. If $f \in C^{(k)}$ and $g_0, g_1, \dots, g_k \in C^{(m)}$ are such that $\mathbf{C} \models (g_0 = \mathrm{comp}^k_m(f, g_1, \dots, g_k))$, then we add the identity $g_0(y_1, \dots, y_m) \approx f(g_1(y_1, \dots, y_m), \dots, g_k(y_1, \dots, y_m))$ to Σ. Moreover, we add the identities $\pi^n_i(y_1, \dots, y_n) \approx y_i$ to Σ. Then $\mathrm{Var}(\mathbf{C})$ denotes the class of τ-algebras that satisfy Σ.

Conversely, to every variety $\mathcal{V}$ we may associate the clone $\mathrm{Clo}(\mathcal{V}) := \mathrm{Clo}(\mathbf{F}_{\mathcal{V}}(\{x_1, x_2, \dots\}))$ of the algebra $\mathbf{F}_{\mathcal{V}}(\{x_1, x_2, \dots\})$ which is free for $\mathcal{V}$ over countably many generators.

Proposition 8.34. *Let* $\mathbf{C}$ *be an abstract clone. Then* $\mathrm{Clo}(\mathrm{Var}(\mathbf{C}))$ *is isomorphic to* $\mathbf{C}$.

Proof. Let $\mathcal{V} := \mathrm{Var}(\mathbf{C})$, and let Σ be the set of τ-identities that defines $\mathcal{V}$. Let $\mathbf{F} := \mathbf{F}_{\mathcal{V}}(\{x_1, x_2, \dots\})$ and let $\mathbf{D} := \mathrm{Clo}(\mathbf{F})$.

Claim 1. The map ξ that sends $f \in C^{(n)}$ to $f^{\mathbf{F}} \in D^{(n)}$ is a clone homomorphism $\mathbf{C} \to \mathbf{D}$:

- It clearly preserves arities.
- If $i \in [n]$ then Σ contains $\pi^n_i(y_1, \dots, y_n) \approx y_i$. Since $\mathbf{F} \models \Sigma$ we have $(\pi^n_i)^{\mathbf{F}}(a_1, \dots, a_n) = a_i$ for all $a_1, \dots, a_n \in F$. Hence, $\xi((\pi^n_i)^{\mathbf{C}}) = (\pi^n_i)^{\mathbf{F}} = (\pi^n_i)^{\mathbf{D}}$.
- If $n, m \geq 1$, $f \in C^{(n)}$, $g_1, \dots, g_n \in C^{(m)}$, and $g_0 = f(g_1, \dots, g_n)$, then Σ contains $g_0(y_1, \dots, y_m) \approx f(g_1(y_1, \dots, y_m), \dots, g_n(y_1, \dots, y_m))$, and since $\mathbf{F} \models \Sigma$ it follows that $\xi(g_0) = \xi(f)(\xi(g_1), \dots, \xi(g_n))$.

Claim 2. ξ is surjective. Every element of D is of the form $t^{\mathbf{F}}$, for some τ-term t. It can be shown by induction over the term structure that there exists $s \in \tau$ such that $s^{\mathbf{F}} = t^{\mathbf{F}}$, and hence $\xi(s) = t^{\mathbf{F}}$.

Claim 3. ξ is injective. Suppose that $\xi(f) = \xi(g)$ for some $f, g \in C^{(n)}$. Then $f^{\mathbf{F}} = g^{\mathbf{F}}$ and hence $\mathbf{F} \models f(y_1, \dots, y_n) \approx g(y_1, \dots, y_n)$. By Lemma 8.24, $f(y_1, \dots, y_n) \approx g(y_1, \dots, y_n)$

holds in all algebras of $\mathcal{V}$. By Theorem 8.25 we have that $(f(y_1,\dots,y_n),g(y_1,\dots,y_n))$ is in the smallest compatible fully invariant equivalence relation E that contains Σ. We show by induction on the structure of E that if there are $h \in \tau$ and τ-terms $t_1,\dots,t_k,s$ with variables from $y_1,\dots,y_m$ such that E contains

$$(s, h(t_1,\dots,t_k))$$

and $\xi(q) = s, \xi(p_1) = t_1^{\mathbf{F}},\dots,\xi(p_k) = t_k^{\mathbf{F}}$, then

$$q = \mathrm{comp}_m^k(h, p_1,\dots,p_k).$$

In particular, if $t_1 = y_1,\dots,t_k = y_k$, and $s = h'(y_1,\dots,y_k)$, so that E in fact contains $(h'(y_1,\dots,y_k),h(y_1,\dots,y_k))$, then $h' = \mathrm{comp}_m^k(h,\pi_1^k,\dots,\pi_k^k) = h$, because $\xi(\pi_i^n) = (y_i)^{\mathbf{F}}$.

- If s is of the form $g_0(y_1,\dots,y_m)$ and t_i is of the form $g_i(y_1,\dots,y_m)$ for each $i \in \{1,\dots,k\}$, then the statement is true by the definition of Σ.
- If s is of the form $h(s_1,\dots,s_k)$ for τ-terms $s_1,\dots,s_k$, and $(s_1,t_1),\dots,(s_k,t_k) \in E$, then choose $q_1,\dots,q_k$ such that $\xi(q_i) = s_i^{\mathbf{F}}$, which exist by the surjectivity of ξ. We further distinguish for each $i \in \{1,\dots,k\}$ whether t_i or s_i is a variable of the form $h'(r_1,\dots,r_u)$. If one is of the form $h'(r_1,\dots,r_u)$, then $(s_i,t_i) \in E$ and the inductive assumption implies that $p_i = q_i$. If both s_i and t_i are a variable, then it must be the same variable y_i, and since $\xi(\pi_i^m) = (y_i)^{\mathbf{F}}$ we have that $p_i = \pi_i^m = q_i$ in this case as well. Hence, $q = \mathrm{comp}_m^k(h,q_1,\dots,q_k) = \mathrm{comp}_m^k(h,p_1,\dots,p_k)$.
- If s is of the form $h'(p_1,\dots,p_k)$ and $(h'(x_1,\dots,x_k),h(x_1,\dots,x_k)) \in E$, then $h' = h$ by the inductive assumption. Hence, $q = \mathrm{comp}_m^k(h',p_1,\dots,p_k) = \mathrm{comp}_m^k(h,p_1,\dots,p_k)$.

Since $(f(y_1,\dots,y_n),g(y_1,\dots,y_n)) \in E$, we therefore have $f = g$. This proves the injectivity of ξ. □

Proposition 8.34 in particular shows the following, which can be seen as an analog of Cayley's theorem for clones.

Corollary 8.35. *Every abstract clone is isomorphic to an operation clone.*

Proposition 8.34 has a converse; to state it, we need the following definition.

Definition 8.36. Let $\mathcal{V}$, $\mathcal{W}$ be varieties with signatures σ and ρ, respectively. An *interpretation* of $\mathcal{V}$ in $\mathcal{W}$ is a map I from σ to ρ-terms such that $\mathcal{V}$ contains $\{I(\mathbf{A}) \mid \mathbf{A} \in \mathcal{W}\}$ where $I(\mathbf{A})$ is the σ-algebra with domain A and the operation $I(f)^{\mathbf{A}}$ for $f \in \sigma$.

The following lemma is straightforward from the definitions.

Lemma 8.37. *Let $\mathcal{V}$ and $\mathcal{W}$ be varieties. Then there is an interpretation of $\mathcal{V}$ in $\mathcal{W}$ if and only if there exists a clone homomorphism from* $\mathrm{Clo}(\mathcal{V})$ *to* $\mathrm{Clo}(\mathcal{W})$.

Proposition 8.38. *Let $\mathcal{V}$ be a variety. Then* $\mathrm{Var}(\mathrm{Clo}(\mathcal{V}))$ *and $\mathcal{V}$ mutually interpret each other.*

Proof. Let σ be the signature of $\mathcal{V}$ and let ρ be the signature of $\mathcal{W} := \mathrm{Var}(\mathrm{Clo}(\mathcal{V}))$. Let $\mathbf{F} := \mathbf{F}_{\mathcal{V}}(\{x_1, x_2, \dots\})$. The identities that hold in every algebra of $\mathcal{V}$ are precisely those that hold in $\mathbf{F}$ by Lemma 8.24. Then the map that sends $f \in \sigma$ of arity k to $f^{\mathbf{F}}(x_1, \dots, x_k)$, viewed as a ρ-term, is an interpretation of $\mathcal{V}$ in $\mathcal{W}$.

Conversely, every $f \in \rho$ has been introduced for an element of $\mathbf{F}$ which equals $t^{\mathbf{F}}(x_{i_1}, \dots, x_{i_n})$ for some $i_1, \dots, i_n \in \mathbb{N}$ and some σ-term $t(y_1, \dots, y_n)$. The map J that sends f to $t(y_1, \dots, y_n)$ is an interpretation of $\mathcal{W}$ in $\mathcal{V}$. □

The following proposition links the existence of clone homomorphisms with the language of algebras, and in particular identities and (pseudo-) varieties.

Proposition 8.39. *Let $\mathscr{C}$ and $\mathscr{D}$ be operation clones on finite sets. Then the following are equivalent.*

1. *There is a surjective clone homomorphism from $\mathscr{C}$ to $\mathscr{D}$;*
2. *there are algebras $\mathbf{A}$ and $\mathbf{B}$ with the same signature τ such that $\mathrm{Clo}(\mathbf{A}) = \mathscr{D}$, $\mathrm{Clo}(\mathbf{B}) = \mathscr{C}$, and all universal conjunctive τ-sentences that hold in $\mathbf{B}$ also hold in $\mathbf{A}$;*
3. *there are algebras $\mathbf{A}$ and $\mathbf{B}$ with the same signature such that $\mathrm{Clo}(\mathbf{A}) = \mathscr{D}$, $\mathrm{Clo}(\mathbf{B}) = \mathscr{C}$, and $\mathbf{A} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$ (equivalently, $\mathbf{A} \in \mathrm{HSP}(\mathbf{B})$).*

Moreover, the following are equivalent.

- *There is a clone isomorphism between $\mathscr{C}$ and $\mathscr{D}$.*
- *there are algebras $\mathbf{A}$ and $\mathbf{B}$ with the same signature such that $\mathrm{Clo}(\mathbf{A}) = \mathscr{D}$, $\mathrm{Clo}(\mathbf{B}) = \mathscr{C}$, and $\mathrm{HSP}^{\mathrm{fin}}(\mathbf{A}) = \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$ (equivalently: $\mathrm{HSP}(\mathbf{A}) = \mathrm{HSP}(\mathbf{B})$).*

In the study of the complexity of CSPs, the equivalence between (1) and (3) in the above is the most relevant, since (3) is related to our most important tool to prove NP-hardness of CSPs (because of the link between pseudovarieties and primitive positive interpretations from Theorem 8.18), and since (1) is the universal-algebraic property that will be used in the following (see e.g. Theorem 9.15 below). The following lemma is central for our applications of abstract clones when studying the complexity of CSPs; it applies to all operation clones $\mathbf{F}$ on a finite set.

Lemma 8.40. *Let $\mathbf{C}$ be a clone and let $\mathbf{F}$ be the clone that has finitely many elements of each sort such that there is no clone homomorphism from $\mathbf{C}$ to $\mathbf{F}$. Then there is a primitive positive sentence in the language τ of (abstract) clones that holds in $\mathbf{C}$ but not in $\mathbf{F}$.*

Proof. Let $\mathbf{E}$ be the expansion of $\mathbf{C}$ by constant symbols such that every element e of $\mathbf{E}$ is named by a constant c_e. Let V be the set of atomic sentences that hold in $\mathbf{E}$. Let U be the first-order theory of $\mathbf{F}$. Suppose that $U \cup V$ has a model $\mathbf{M}$. There might be elements in $\mathbf{M}$ outside of $\bigcup_i M^{(i)}$. But the τ-reduct of the restriction of $\mathbf{M}$ to $\bigcup_i M^{(i)}$ must be isomorphic to $\mathbf{F}$, since each of the $M^{(i)}$ is finite; we identify it with $\mathbf{F}$. Note that for all constants c_e we have that $c_e^{\mathbf{M}} \in \mathbf{F}$. Since $\mathbf{M}$ satisfies all atomic formulas that hold in $\mathbf{E}$, we have that the mapping $e \mapsto c_e^{\mathbf{M}}$, for e an element of $\mathbf{E}$, is a homomorphism from $\mathbf{C}$ to $\mathbf{F}$, in contradiction to our assumptions.

So $U \cup V$ is unsatisfiable, and by compactness of first-order logic there exists a finite subset V' of V such that $V' \cup U$ is unsatisfiable. Replace each of the new constant symbols in V' by an existentially quantified variable; then the conjunction of the resulting sentences from V is a primitive positive sentence, and it must be false in $\mathbf{F}$. □

A set of identities Σ is called *trivial* if there exists an algebra $\mathbf{A}$ that satisfies Σ and $\mathrm{Clo}(\mathbf{A})$ is isomorphic to $\mathbf{Proj}$.

Corollary 8.41. *Let* $\mathbf{A}$ *be an algebra. If there is no clone homomorphism from* $\mathrm{Clo}(\mathbf{A})$ *to* $\mathbf{Proj}$*, then there exists a non-trivial finite set of identities that holds in* $\mathbf{A}$*.*

Remark 8.42. Recall from Remark 6.1 that there are uncountably many clones on a three-element set. In fact, there are uncountably many even when considered up to homomorphic equivalence [31].

8.8 Clone Homomorphisms and Primitive Positive Interpretations

Clone homomorphisms can be linked to pseudovarieties of algebras, and pseudo-varieties of polymorphism algebras can be linked to primitive positive interpretations; in this section, we present shortcuts that directly link the existence of clone homomorphisms of polymorphism clones with primitive positive interpretations. The proofs will be merely combinations of previous results, but the combinations are often easier to cite and this will be convenient later in the text.

Corollary 8.43. *A finite structure* $\mathfrak{A}$ *has a primitive positive interpretation in a finite structure* $\mathfrak{B}$ *if and only if there exists a clone homomorphism from* $\mathrm{Pol}(\mathfrak{B})$ *to* $\mathrm{Pol}(\mathfrak{A})$*.*

Proof. The proof is a straightforward combination of Theorem 8.18 with Proposition 8.39. Let $\mathbf{B}$ be a polymorphism algebra of $\mathfrak{B}$. If $\mathfrak{A}$ has a primitive positive interpretation in $\mathfrak{B}$ then by Theorem 8.18 there exists $\mathbf{A} \in \mathrm{HSP}^{\mathrm{fin}}(\mathfrak{B})$ such that $\mathrm{Clo}(\mathbf{A}) \subseteq \mathrm{Pol}(\mathfrak{A})$. Then Proposition 8.39 implies that there exists a surjective clone homomorphism from $\mathrm{Clo}(\mathbf{B})$ to $\mathrm{Clo}(\mathbf{A})$, which is a clone homomorphism from $\mathrm{Pol}(\mathfrak{B})$ to $\mathrm{Pol}(\mathfrak{A})$. Conversely, suppose that there exists a clone homomorphism from $\mathrm{Pol}(\mathfrak{B})$ to $\mathrm{Pol}(\mathfrak{A})$. Let $\mathcal{C} \subseteq \mathrm{Pol}(\mathfrak{A})$ be the image of this clone homomorphism. Then by Proposition 8.39 there are algebras $\mathbf{A}$ and $\mathbf{B}$ with the same signature such that $\mathrm{Clo}(\mathbf{A}) = \mathcal{C}$, $\mathrm{Clo}(\mathbf{B}) = \mathrm{Pol}(\mathfrak{B})$, and $\mathbf{A} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$. This in turn means that $\mathfrak{A}$ has a primitive positive interpretation in $\mathfrak{B}$ by Theorem 8.18. $\square$

Corollary 8.44. *Two finite structures* $\mathfrak{A}$ *and* $\mathfrak{B}$ *are primitively positively bi-interpretable if and only if* $\mathrm{Pol}(\mathfrak{A})$ *and* $\mathrm{Pol}(\mathfrak{B})$ *are isomorphic as abstract clones.*

Proof. Combine Proposition 8.19 and the second part of Proposition 8.39. $\square$

8.9 Hardness from Factors

An algebra is called *idempotent* if all of its operations are idempotent. For idempotent algebras $\mathbf{A}$ there is another characterisation for the existence of a clone homomorphism to $\mathbf{Proj}$ by Bulatov and Jeavons [38, Proposition 4.14] (Corollary 8.46 below). We present a slightly strengthened version of their result by Zhuk [95, Lemma 4.2].

Theorem 8.45. *Let* $\mathbf{B}$ *be an idempotent algebra and suppose that* $\mathbf{A} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$ *has at least two elements. Then* $\mathrm{HS}(\mathbf{B})$ *contains a subalgebra* $\mathbf{A}'$ *of* $\mathbf{A}$ *with at least two elements.*

Proof. Suppose that $\mathbf{C} \in S(\mathbf{B}^d)$ for some $d \in \mathbb{N}$ has a congruence K such that $\mathbf{A} := \mathbf{C}/K$ has at least two elements. We show the statement by induction on d. For $d = 1$ there is nothing to be shown because we can choose $\mathbf{A}' := \mathbf{A}$. If for any two equivalence classes E_1 and E_2 of K the intersection $\pi_1(E_1) \cap \pi_1(E_2)$ is empty, then let

$$\begin{aligned} C &:= \pi_1(C) \\ K' &:= \{(\pi_1(a), \pi_1(b)) \mid (a,b) \in K\}. \end{aligned}$$

Then C' is the universe of a subalgebra $\mathbf{C}'$ of $\mathbf{C}$, K' is a congruence of $\mathbf{C}'$, and $\mathbf{C}'/K'$ is isomorphic to $\mathbf{C}/K = \mathbf{A}$. Again, we have $\mathbf{A} \in \mathrm{HS}(\mathfrak{B})$.

Now suppose that K has two equivalence classes E_1 and E_2 such that $\pi_1(E_1) \cap \pi_1(E_2) \neq \emptyset$. Let $a \in \pi_1(E_1) \cap \pi_1(E_2)$ and define

$$\begin{aligned} C' &:= \pi_{\{2,\dots,n\}}(S \cap (\{a\} \times A^{d-1}) \\ K' &:= \{((b_2,\dots,b_d),(c_2,\dots,c_d)) \mid ((a,b_2,\dots,b_d),(a,c_2,\dots,c_d)) \in K\}. \end{aligned}$$

Since $\mathbf{B}$ and $\mathbf{C}$ are idempotent, the set C' is the universe of a subalgebra $\mathbf{C}'$ of $\mathbf{C}$, and K' is a congruence of $\mathbf{C}'$. The algebra $\mathbf{C}'/K'$ has at least two elements and is isomorphic to a subalgebra of $\mathbf{C}/K = \mathbf{A}$. Thus, the statement follows from the inductive assumption. □

Corollary 8.46. *Let* $\mathbf{B}$ *be an idempotent algebra. Then* $\mathrm{HSP}^{\mathrm{fin}}(\mathbf{B})$ *contains an algebra with at least two elements all of whose operations are projections if and only if* $\mathrm{HS}(\mathfrak{B})$ *does.*

Proof. If all operations $\mathbf{A}$ of an algebra are projections, then the same applies to all subalgebras of $\mathbf{A}$. Therefore the statement follows from Theorem 8.45. □

Since the size of the algebras in $\mathrm{HS}(\mathbf{B})$ is bounded by the size of $\mathbf{B}$, this leads to an algorithm that decides whether a given finite structure $\mathfrak{B}$ satisfies the equivalent conditions in Theorem 9.15. We summarise various equivalent conditions for finite idempotent algebras that were treated in this chapter.

Corollary 8.47. *Let* $\mathbf{B}$ *be a finite idempotent algebra. Then the following are equivalent.*

1. *There is no homomorphism from* $\mathrm{Clo}(\mathbf{B})$ *to* $\mathbf{Proj}$.
2. $\mathbf{B}$ *satisfies some non-trivial finite set of identities.*
3. $\mathrm{HSP}(\mathbf{B})$ *does not contain an at least 2-element algebra all of whose operations are projections.*
4. $\mathrm{HS}(\mathbf{B})$ *does not contain an at least 2-element algebra all of whose operations are projections.*

Proof. The equivalence of (1) and (2) follows from Corollary 8.41. The equivalence of (1) and (3) follows from Proposition 8.39. The equivalence of (3) and (4) follows from Theorem 8.20 combined with Theorem 8.45. □

9 Minions

(Abstract) minions generalise (abstract) clones, and function minions generalise operation clones. The name has been introduced in 2018 when it became clear that function minions over finite domains capture the complexity of so-called *promise CSPs*, which generalise CSPs [9].

This text does not cover promise CSPs; however, we still introduce minions, because minion homomorphisms play an important role when studying clones as well (see, e.g., [17, 29, 31]). In particular, minion homomorphisms can be used to characterise the operator HI from the second formulation of the tractability theorem, Theorem 5.28.

9.1 Minors and Minions

Let A, B be sets and $k, l \in \mathbb{N}_{\geq 1}$. Let $f \colon A^k \to B$ be a function and let $\alpha \colon [k] \to [m]$. Then f_α denotes the function $g \colon A^m \to B$ given by $g(x_1, \dots, x_m) := f(x_{\alpha(1)}, \dots, x_{\alpha(k)})$. A *minor* of f is a function of the form f_α, for some $\alpha \colon [k] \to [m]$. Note that

- if $\alpha \colon [k] \to [m]$ and $\beta \colon [m] \to [n]$, then
$$(f_\alpha)_\beta = f_{\beta \circ \alpha},$$
so the minor relation is transitive.
- if $\alpha \colon [k] \to [m]$ then $(\pi^k_i)_\alpha = \pi^m_{\alpha(i)}$.
- if f and g are idempotent operations on A, then for all $\alpha_1, \dots, \alpha_k \colon [m] \to [n]$ there exists $\beta \colon [km] \to [n]$ such that
$$f(g_{\alpha_1}, \dots, g_{\alpha_k}) = (f * g)_\beta \qquad \text{(recall Definition 8.32).}$$

Definition 9.1. A *function minion* is a subset $\mathscr{M}$ of $\bigcup_{k \geq 1} B^{A^k}$, where A and B be sets, which is closed under taking minors.

Note that every operation clone is a function minion where $A = B$ and where we additionally require the presence of the projections and closure under composition. A *minion* is the abstract version of a function minion, analogously as clones can be viewed as the abstract version of operation clones.

Definition 9.2. An *(abstract) minion* is a multi-sorted algebra $\mathbf{M}$ with sorts $M^{(1)}, M^{(2)}, \dots$ and for each $\alpha \colon [n] \to [m]$ the operation ${}_\alpha \colon M^{(n)} \to M^{(m)}$ such that for every $\sigma \colon [n] \to [m]$ and $\rho \colon [m] \to [k]$ and $f \in M^{(n)}$
$$(f_\sigma)_\rho = f_{\rho \circ \sigma}.$$

Clearly, every function minion gives rise to a minion in the obvious way. This statement has a converse, which is analogous to Cayley's theorem for groups (see Proposition 8.34 for the corresponding statement for clones); see Exercise 157.

Definition 9.3. Let $\mathbf{M}$ and $\mathbf{N}$ be minions.

- A *minion homomorphism* from $\mathbf{M}$ to $\mathbf{N}$ is a map $\xi \colon M \to N$ such that for every $n \in \mathbb{N}$ and $f \in M^{(n)}$ we have $\xi(f) \in D^{(n)}$ and for every $m, n \in \mathbb{N}$ and $\alpha \colon [n] \to [m]$ and $f \in M^{(n)}$ we have $\xi(f_\alpha) = \xi(f)_\alpha$.
- A *minion isomorphism* is a bijective minion homomorphism $\xi \colon \mathbf{M} \to \mathbf{N}$ such that ξ^{-1} is a minion homomorphism as well. In this case, $\mathbf{M}$ and $\mathbf{N}$ are called *isomorphic*.

Similarly as for clones in Lemma 8.40, we may apply the compactness theorem of first-order logic to characterise the existence of minion homomorphisms to function minions on finite sets.

Lemma 9.4. *Let* $\mathbf{M}$ *be a minion and let* $\mathbf{F}$ *be a minion with finitely many elements of each sort. If there is no minion homomorphism from* $\mathbf{M}$ *to* $\mathbf{F}$*, then there exists a primitive positive sentence over the signature of (abstract) minions that holds in* $\mathbf{M}$ *but not in* $\mathbf{F}$*.*

Exercises.

152. Let $\mathbf{C}$ and $\mathbf{D}$ be clones. Show that $\xi\colon C \to D$ is a minion homomorphism if and only if

- ξ preserves arities, i.e., $\xi(C^{(i)}) \subseteq D^{(i)}$ for all $i \in \mathbb{N}$, and
- ξ preserves composition with projections, that is, for all $n, k \geq 1$ and $f \in C^{(k)}$

$$\xi(f((\pi^n_{i_1})^{\mathbf{C}}, \ldots, (\pi^n_{i_k})^{\mathbf{C}})) = \xi(f)((\pi^n_{i_1})^{\mathbf{D}}), \ldots, (\pi^n_{i_k})^{\mathbf{D}}).$$

153. Let $\alpha\colon [k] \to [m]$. Write down a primitive positive formula $\phi(x, y)$ in the language of clones such that for every operation clone $\mathscr{C}$ and all $f, g \in \mathscr{C}$ we have $\mathscr{C} \models \phi(f, g)$ if and only if $g = f_\alpha$.

154. Let $\mathbf{M}$ be a minion. Show that for every injective $\alpha\colon [n] \to [m]$ the operation $f \mapsto f_\alpha$ from $M^{(n)}$ to $M^{(m)}$ is injective, and for every surjective $\alpha\colon [n] \to [m]$ the operation $f \mapsto f_\alpha$ from $M^{(n)}$ to $M^{(m)}$ is surjective.
155. (For readers familiar with basic category theory) Let **Set** denote the category of all set and mappings between them, and let **FinOrd** denote the category of finite ordinals and mappings between them.
    - Explain how minions can be viewed as functors from **FinOrd** to **Set**. Show that natural transformations between such functors are minion homomorphisms.
    - Explain how minions can be viewed as **Set**-endofunctors F which are *finitary*: that is, if X is a set, then $F(X)$ is the union of the sets $\{F(i)(u) \mid u \in U\}$, where U is a finite subset of X and $i\colon U \to X$ is the inclusion map. Show that natural transformations between **Set**-endofunctors are minion homomorphisms.

9.2 Reflections

In Section 8.4 we have seen that the $\mathrm{HSP}^{\mathrm{fin}}$ operator is the algebraic counterpart to full primitive positive interpretations. This section treats a relatively new universal-algebraic operator, for forming *reflections* (introduced in [17]), which can be used to characterise the structure-building operator HI. Recall from Section 5.9 that HI is the operator that is most relevant for constraint satisfaction.

Definition 9.5. Let $\mathbf{B}$ be a τ-algebra, let A be a set, and let $h\colon B \to A$ and $g\colon A \to B$ be two maps. Then the *reflection* of $\mathbf{B}$ with respect to (h, g) is the τ-algebra $\mathbf{A}$ with domain A where for all $a_1, \ldots, a_n \in A$ and $f \in \tau$ of arity n we define

$$f^{\mathbf{A}}(a_1, \ldots, a_n) := h(f^{\mathbf{B}}(g(a_1), \ldots, g(a_n)))\,.$$

The class of reflections of a class of τ-algebras $\mathcal{C}$ is denoted by $\mathrm{Refl}(\mathcal{C})$.

As for the other operators on algebras, we write $\mathrm{Refl}(\mathbf{B})$ instead of $\mathrm{Refl}(\{\mathbf{B}\})$. The following is an analog to the HSP-lemma (Lemma 8.17).

Lemma 9.6 (from [17]). *Let $\mathcal{C}$ be a class of τ-algebras.*

- *The smallest class of τ-algebras that contains $\mathcal{C}$ and is closed under* Refl, H, S, *and* P *equals* $\mathrm{Refl}\,\mathrm{P}(\mathcal{C})$.
- *The smallest class of τ-algebras that contains $\mathcal{C}$ and is closed under* Refl, H, S, *and* $\mathrm{P}^{\mathrm{fin}}$ *equals* $\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathcal{C})$.

Proof. For the first statement, it suffices to prove that $\mathrm{Refl}\,\mathrm{P}(\mathcal{C})$ is closed under Refl, H, S, P, and for the second that $\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathcal{C})$ is closed under Refl, H, S, $\mathrm{P}^{\mathrm{fin}}$. For the operator Refl this follows from the simple fact that $\mathrm{Refl}\,\mathrm{Refl}(\mathcal{K}) = \mathrm{Refl}(\mathcal{K})$ for any class $\mathcal{K}$.

To prove that $\mathrm{Refl}\,\mathrm{P}(\mathcal{C})$ and $\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathcal{C})$ are closed under H, we show that $\mathrm{H}(\mathcal{K}) \subseteq \mathrm{Refl}(\mathcal{K})$ for any class $\mathcal{K}$. Let $\mathbf{B} \in \mathcal{K}$ and $h \colon B \to A$ be a surjective homomorphism to an algebra $\mathbf{A}$. Pick any function g such that $h \circ g$ is the identity on A. Then h and g witness that $\mathbf{A}$ is a reflection of $\mathbf{B}$ since

$$\begin{aligned} h(f^{\mathbf{B}}(g(x_1), \dots, g(x_n)) &= f^{\mathbf{A}}(h \circ g(x_1), \dots, h \circ g(x_n)) && \text{(since h is a homomorphism)} \\ &= f^{\mathbf{A}}(x_1, \dots, x_n) && \text{(by the choice of g).} \end{aligned}$$

To prove that $\mathrm{Refl}\,\mathrm{P}(\mathcal{C})$ and $\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathcal{C})$ are closed under S, we show that $\mathrm{S}(\mathcal{K}) \subseteq \mathrm{Refl}(\mathcal{K})$ for any class $\mathcal{K}$. Let $\mathbf{B} \in \mathcal{K}$ and suppose that $\mathbf{A}$ is a subalgebra of $\mathbf{B}$. Let $g \colon A \to B$ be the identity on A, and $h \colon B \to A$ be any extension of g to B. Then h and g show that $\mathbf{A}$ is a reflection of $\mathbf{B}$ since

$$h(f^{\mathbf{B}}(g(x_1), \dots, g(x_n)) = f^{\mathbf{B}}(x_1, \dots, x_n) = f^{\mathbf{A}}(x_1, \dots, x_n)\,.$$

Let I be an arbitrary set, $(\mathbf{B}_i)_{i \in I}$ be algebras from $\mathrm{P}(\mathcal{C})$, and suppose that $\mathbf{A}_i$ is a reflection of $\mathbf{B}_i$ for every $i \in I$, witnessed by functions $h_i \colon B_i \to A_i$ and $g_i \colon A_i \to B_i$. Then the map $h \colon \prod_{i \in I} B_i \to \prod_{i \in I} A_i$ that sends $(b_i)_{i \in I}$ to $(h_i(b_i))_{i \in I}$ and the map $g \colon \prod_{i \in I} A_i \to \prod_{i \in I} B_i$ that sends $(a_i)_{i \in I}$ to $(g_i(a_i))_{i \in I}$ witness that $\prod_{i \in I} \mathbf{A}_i$ is a reflection of $\prod_{i \in I} \mathbf{B}_i$. This shows that $\mathrm{P}(\mathrm{Refl}\,\mathrm{P}(\mathcal{C})) \subseteq \mathrm{Refl}\,\mathrm{P}(\mathcal{C})$ and likewise that $\mathrm{P}^{\mathrm{fin}}(\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathcal{C})) \subseteq \mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathcal{C})$. $\square$

Theorem 9.7. *Let $\mathfrak{B}, \mathfrak{C}$ be finite relational structures and let $\mathbf{C}$ be a polymorphism algebra of $\mathfrak{C}$. Then*

1. $\mathfrak{B} \in \mathrm{H}(\mathfrak{C}')$ *for some structure $\mathfrak{C}'$ which is primitively positively definable in $\mathfrak{C}$ if and only if there is an algebra* $\mathbf{B} \in \mathrm{Refl}(\mathbf{C})$ *such that* $\mathrm{Clo}(\mathbf{B}) \subseteq \mathrm{Pol}(\mathfrak{B})$.
2. $\mathfrak{B} \in \mathrm{HI}(\mathfrak{C})$ *if and only if there is an algebra* $\mathbf{B} \in \mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{C})$ *such that* $\mathrm{Clo}(\mathbf{B}) \subseteq \mathrm{Pol}(\mathfrak{B})$.

Proof. To show (1), first suppose that $\mathfrak{B} \in \mathrm{H}(\mathfrak{C}')$ for some $\mathfrak{C}'$ which is pp definable in $\mathfrak{C}$; let $h \colon \mathfrak{C}' \to \mathfrak{B}$ and $g \colon \mathfrak{B} \to \mathfrak{C}'$ be homomorphisms witnessing homomorphic equivalence of $\mathfrak{B}$ and $\mathfrak{C}'$. Let $\mathbf{C}'$ be an expansion of $\mathbf{C}$ which is a polymorphism algebra of $\mathfrak{C}'$. Let $\mathbf{B}'$ be the reflection of $\mathbf{C}'$ with respect to (h, g). Every operation of $\mathrm{Clo}(\mathbf{B}')$ is obtained as a composition of homomorphisms, so preserves all the relations of $\mathfrak{B}$, so $\mathrm{Clo}(\mathbf{B}') \subseteq \mathrm{Pol}(\mathfrak{B})$. Let $\mathbf{B}$ be the reduct of $\mathbf{B}'$ where we only keep the operations for the signature of $\mathbf{C}$, and note that $\mathbf{B} \in \mathrm{Refl}(\mathbf{C})$ is such that $\mathrm{Clo}(\mathbf{B}) \subseteq \mathrm{Pol}(\mathfrak{B})$.

Conversely, suppose that the reflection $\mathbf{B}$ of $\mathbf{C}$ at $h\colon C \to B$ and $g\colon B \to C$ is such that $\mathrm{Clo}(\mathbf{B}) \subseteq \mathrm{Pol}(\mathfrak{B})$. Let $\mathfrak{C}'$ be the structure with domain C and the same signature as $\mathfrak{B}$ which contains for every k-ary relation symbol R of $\mathfrak{B}$ the relation

$$\begin{aligned} R^{\mathfrak{C}'} := \{ & (f(g(b_1^1), \dots, g(b_1^\ell)), \dots, f(g(b_k^1), \dots, g(b_k^\ell))) \\ & \mid f \in \mathrm{Pol}(\mathfrak{C}), (b_1^1, \dots, b_k^1), \dots (b_1^\ell, \dots, b_k^\ell) \in R^{\mathfrak{B}} \}. \end{aligned}$$

These relations are preserved by $\mathrm{Pol}(\mathfrak{C})$, so they are pp definable in $\mathfrak{C}$ by Theorem 6.6, and hence $\mathfrak{C}' \in \mathrm{Red}(\mathfrak{C})$. Clearly, g is a homomorphism from $\mathfrak{B}$ to $\mathfrak{C}'$. We claim that h is a homomorphism from $\mathfrak{C}'$ to $\mathfrak{B}$. Indeed, if $b_1, \dots, b_k \in B$ are such that $(f(g(b_1), \dots, g(b_k))) \in R^{\mathfrak{C}'}$, then $h(f(g(b_1), \dots, g(b_k))) \in R^{\mathfrak{B}}$ because the operation $(x_1, \dots, x_k) \mapsto h(f(g(x_1), \dots, g(x_k)))$ is an operation of $\mathbf{B}' \in \mathrm{Refl}(\mathbf{C})$ and hence a polymorphism of $\mathfrak{B}$ since $\mathrm{Clo}(\mathbf{B}') \subseteq \mathrm{Pol}(\mathfrak{B})$. Thus, $\mathfrak{B} \in \mathrm{H}(\mathfrak{C}')$.

Item (2) is a combination of item (1) with Theorem 8.18: First suppose that $\mathfrak{B} \in \mathrm{HI}(\mathfrak{C})$. Then there exists a structure $\mathfrak{D} \in I(\mathfrak{C})$ such that $\mathfrak{B} \in \mathrm{H}(\mathfrak{D})$. By Theorem 8.18 there is an algebra $\mathbf{D} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{C})$ such that $\mathrm{Clo}(\mathbf{D}) \subseteq \mathrm{Pol}(\mathfrak{D})$, and by item (1) there is an algebra $\mathbf{B} \in \mathrm{Refl}(\mathbf{D})$ such that $\mathrm{Clo}(\mathbf{B}) \subseteq \mathrm{Pol}(\mathfrak{B})$. This proves the statement since

$$\begin{aligned} \mathbf{B} \in \mathrm{Refl}(\mathbf{D}) & \subseteq \mathrm{Refl}\,\mathrm{HSP}^{\mathrm{fin}}(\mathbf{C}) \\ & = \mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{C}) && \text{(by Lemma 9.6).} \end{aligned}$$

Conversely, suppose that there exists $\mathbf{B} \in \mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{C})$ such that $\mathrm{Clo}(\mathbf{B}) \subseteq \mathrm{Pol}(\mathfrak{B})$. Then there exists $\mathbf{D} \in \mathrm{P}^{\mathrm{fin}}(\mathbf{C})$ such that $\mathbf{B} \in \mathrm{Refl}(\mathbf{D})$. Let $\mathfrak{D}$ be the structure with the same domain as $\mathbf{D}$ which contains all the relations preserved by all operations of $\mathrm{Clo}(\mathbf{D})$. Then $\mathfrak{D} \in \mathrm{I}(\mathfrak{C})$. Moreover, by item (1) there exists $\mathfrak{D}'$ which is pp definable in $\mathfrak{D}$ such that $\mathfrak{B} \in \mathrm{H}(\mathfrak{D}')$. Clearly, $\mathfrak{D}' \in \mathrm{I}(\mathfrak{C})$ and hence $\mathfrak{B} \in \mathrm{HI}(\mathfrak{C})$. $\square$

We now characterise in many different ways the hardness condition from the second formulation of the tractability theorem (Theorem 5.28); quite remarkably, we do not need to assume that the involved polymorphism algebra is idempotent.

Corollary 9.8. *Let $\mathfrak{B}$ be a structure with a finite domain and let $\mathbf{B}$ be a polymorphism algebra of $\mathfrak{B}$. Then the following are equivalent.*

1. $\mathrm{HI}(\mathfrak{B})$ *contains* K_3*;*
2. $\mathrm{HI}(\mathfrak{B})$ *contains all finite structures;*
3. $\mathrm{HI}(\mathfrak{B})$ *contains* $(\{0,1\}; \mathrm{NAE})$*;*
4. $\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{B})$ *contains an algebra of size at least 2 all of whose operations are projections.*
5. $\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{B})$ *contains for every finite set A an algebra on A all of whose operations are projections.*

If these condition apply then $\mathfrak{B}$ has a finite-signature reduct with an NP-hard CSP.

Proof. The implication from 1. to 2. follows from the fact that $\mathrm{I}(K_3)$ contains all finite structures (Theorem 5.17), and that $\mathrm{I\,H\,I}(\mathfrak{B}) \subseteq \mathrm{HI}(\mathfrak{B})$ by Theorem 5.26. The implication from 2. to 3. is trivial. The equivalence of 3. and 4. follows from the fact that all polymorphisms of $(\{0,1\}; \mathrm{NAE})$ are essentially unary (Exercise 121) and Theorem 9.7. We leave the proof of the equivalence of 5. to the reader. The final statement follows from Corollary 5.16. $\square$

Exercises.

156. Prove the equivalence of 5. with the other items of Corollary 9.8.

9.3 Birkhoff's Theorem for Height-One Identities

A *height-one identity* is an identity $s \approx t$ where the involved terms s and t have *height one*, i.e., each term involves exactly one function symbol. Three examples of properties that can be expressed by finite sets of height-one identities are listed below.

$$
\begin{aligned}
f(x,y) &\approx f(y,x) && (f \text{ is symmetric})\\
f(x,x,y) \approx f(x,y,x) &\approx f(y,x,x) \approx f(x,x,x) && (f \text{ is quasi majority})\\
f(x,x,y) &\approx f(y,x,x) \approx f(y,y,y) && (f \text{ is quasi Maltsev})
\end{aligned}
$$

A non-example is furnished by the Maltsev identities $f(x,x,y) \approx f(y,x,x) \approx y$ because the term y involves no function symbol. Identities where each term involves *at most* one function symbol are called *linear*; so the Maltsev identities are an example of a set of linear identities. An example of a non-linear identity is the associativity law

$$f(x, f(y,z)) \approx f(f(x,y),z).$$

If $\mathbf{A}$ is a τ-algebra, then we write $\mathrm{Minion}(\mathbf{A})$ for the smallest function minion that contains $\{f^{\mathbf{A}} \mid f \in \tau\}$. If $\mathbf{A}$ and $\mathbf{B}$ are τ-algebras then there exists a minion homomorphism $\xi\colon \mathrm{Minion}(\mathbf{B}) \to \mathrm{Minion}(\mathbf{A})$ that maps $f^{\mathbf{B}}$ to $f^{\mathbf{A}}$ if and only if for all $f,g \in \tau$ of arity k and l and all m-ary projections $p_1,\dots,p_k,q_1,\dots,q_l$ we have that $f^{\mathbf{A}}(p_1,\dots,p_k) = g^{\mathbf{A}}(q_1,\dots,q_l)$ whenever $f^{\mathbf{B}}(p_1,\dots,p_k) = g^{\mathbf{B}}(q_1,\dots,q_l)$. If this map exists it must be surjective and we call it the *natural minor-preserving map from* $\mathrm{Minion}(\mathbf{B})$ *to* $\mathrm{Minion}(\mathbf{A})$. The following theorem is a variant of Birkhoff's theorem (Theorem 8.20) for height-one identities.

Theorem 9.9 (cf. Proposition 5.3 of [17])**.** *Let* $\mathbf{A}$ *and* $\mathbf{B}$ *be* τ*-algebras such that* $\mathrm{Minion}(\mathbf{A})$ *and* $\mathrm{Minion}(\mathbf{B})$ *are operation clones. Then the following are equivalent.*

1. *The natural minor-preserving map from* $\mathrm{Minion}(\mathbf{B})$ *to* $\mathrm{Minion}(\mathbf{A})$ *exists.*
2. *All height-one identities that hold in* $\mathbf{B}$ *also hold in* $\mathbf{A}$*.*
3. $\mathbf{A} \in \mathrm{Refl}\,\mathrm{P}(\mathbf{B})$*.*

Moreover, if A and B are finite then we can add the following to the list:

4. $\mathbf{A} \in \mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{B})$*.*

Proof. The equivalence of 1. and 2. is straightforward from the definitions, as in the proof of Theorem 8.20.

The proof that 2. implies 3. is similar to the proof of Theorem 8.20. For every $a \in A$, let $\pi^A_a \in B^{B^A}$ be the function that maps every tuple in B^A to its a-th entry. Let $\mathbf{S}$ be the subalgebra of $\mathbf{B}^{B^A}$ generated by $\{\pi^A_a \mid a \in A\}$. Define $h\colon S \to A$ as

$$h(f^{\mathbf{B}}(\pi^A_{a_1},\dots,\pi^A_{a_n})) := f^{\mathbf{A}}(a_1,\dots,a_n).$$

Similarly as in the proof of Theorem 8.20 one can show that h is well defined using that all height-one identities that hold in $\mathbf{B}$ also hold in $\mathbf{A}$. Note that h is defined on all of S because $\operatorname{Minion}(\mathbf{B})$ is an operation clone.

Let $g\colon A \to S$ be the mapping which sends every $a \in A$ to $\pi_a^{\mathbf{A}}$. Then h and g show that $\mathbf{A} \in \operatorname{Refl}(\mathbf{S}) \subseteq \operatorname{Refl} \operatorname{S} \operatorname{P}(\mathbf{B}) = \operatorname{Refl} \operatorname{P}(\mathbf{B})$: for all $a_1, \dots, a_n \in A$

$$f^{\mathbf{A}}(a_1, \dots, a_n) = h(f^{\mathbf{B}}(g(a_1), \dots, g(a_n))).$$

If A and B are finite, then B^A is finite and hence $\mathbf{S} \in \operatorname{S} \operatorname{P}^{\mathrm{fin}}(\mathbf{B})$, so the proof implies that $\mathbf{A} \in \operatorname{Refl} \operatorname{P}^{\mathrm{fin}}(\mathbf{B})$.

3. implies 2. If $\mathbf{A} \in \operatorname{P}(\mathbf{B})$ then the statement follows from Theorem 8.20. Now suppose that $\mathbf{A}$ is a reflection of $\mathbf{B}$ via the maps $h\colon B \to A$ and $g\colon A \to B$. Let ϕ be the identity $\forall x_1, \dots, x_n\colon f_1(x_{i_1}, \dots, x_{i_k}) = f_2(x_{j_1}, \dots, x_{j_l})$ for $f_1, f_2 \in \tau$ and suppose that $\mathbf{B} \models \phi$. For all $a_1, \dots, a_n \in A$ we have

$$\begin{aligned} f_1^{\mathbf{A}}(a_{i_1}, \dots, a_{i_k}) &= h(f_1^{\mathbf{B}}(g(a_{i_1}), \dots, g(a_{i_k}))) \\ &= h(f_2^{\mathbf{B}}(g(a_{j_1}), \dots, g(a_{j_l}))) = f_2^{\mathbf{A}}(a_{j_1}, \dots, a_{j_l}). \end{aligned}$$

Since $a_1, \dots, a_n$ were chosen arbitrarily, we have that $\mathbf{A} \models \phi$. □

Exercises.

157. Prove a minion version of Cayley's theorem: show that every minion is isomorphic to a function minion.

158. Let Σ be a finite set of height-one identities.
Assume that there exists an algorithm with the following properties:

- it takes as input two finite τ-structures $\mathfrak{A}$ and $\mathfrak{B}$;
- if the algorithm returns 'no' then $\mathfrak{A} \not\to \mathfrak{B}$;
- it runs in polynomial time in the size of $\mathfrak{A}$ and $\mathfrak{B}$;
- if the polymorphisms of $\mathfrak{B}$ satisfy Σ, and the algorithm returns 'yes', then $\mathfrak{A} \to \mathfrak{B}$.

Show that:

(a) if Σ expresses the existence of a majority operation, then the path consistency procedure PC_H provides an example for such an algorithm A (viewing the graph H as part of the input of PC_H);

(b) if there is such an algorithm A for Σ, then there exists a polynomial-time algorithm that decides for a given finite τ-structure $\mathfrak{B}$ whether $\mathfrak{B}$ has polymorphisms that satisfy Σ.

9.4 Minion Homomorphisms and Primitive Positive Constructions

We have characterised primitive positive constructions in terms of polymorphism algebras and the reflection operator, and then we have characterised varieties that are additionally closed under reflection in terms of minion homomorphisms. In this section we present straightforward combinations of these links that are elegant and convenient for later use. The following is analogous to Corollary 8.43 for primitive positive constructions and clone homomorphisms.

(Abstract) Clone	(Abstract) Minion
Operation Clone	Function Minion
Identity	Height-one Identity
Clone Homomorphism	Minion Homomorphism
HSP	Refl P
$\mathrm{HSP}^{\mathrm{fin}}$	$\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}$
Primitive Positive Interpretation	Primitive Positive Construction
Corollary 8.35	Exercise 157

Figure 15: A dictionary between corresponding notions.

Corollary 9.10. *Let $\mathfrak{B}$ and $\mathfrak{A}$ be finite structures. Then $\mathfrak{A} \in \mathrm{HI}(\mathfrak{B})$ if and only if there exists a minion homomorphism from $\mathrm{Pol}(\mathfrak{B})$ to $\mathrm{Pol}(\mathfrak{A})$.*

Proof. Combine Theorem 9.7 with Theorem 9.9. □

Corollary 9.11. *Let $\mathfrak{B}$ be a finite structure. Then there exists a minion homomorphism from $\mathrm{Pol}(\mathfrak{B})$ to* **Proj** *if and only if $K_3 \in \mathrm{HI}(\mathfrak{B})$.*

Proof. Let $\mathbf{B}$ be an algebra such that $\mathrm{Minion}(\mathbf{B}) = \mathrm{Pol}(\mathfrak{B})$. Corollary 9.8 states that $K_3 \in \mathrm{HI}(\mathfrak{B})$ if and only if $\mathrm{Refl}\,\mathrm{P}^{\mathrm{fin}}(\mathbf{B})$ contains an algebra of size at least 2 all of whose operations are projections, and whose clone is therefore **Proj**. Theorem 9.9 implies that this is equivalent to the existence of a minor-preserving map to **Proj**. □

Proposition 9.12. *For every operation clone $\mathscr{C}$ on a finite set there exists an idempotent operation clone $\mathscr{D}$ on a finite set such that there exists a minion homomorphism from $\mathscr{C}$ to $\mathscr{D}$ and from $\mathscr{D}$ to $\mathscr{C}$.*

Proof. Let $\mathfrak{B}$ be a structure such that $\mathrm{Pol}(\mathfrak{B}) = \mathscr{C}$ (such a $\mathfrak{B}$ exists by Proposition 6.4). Let $\mathfrak{C}$ be the core of $\mathfrak{B}$ (which exists by the generalisation of Proposition 2.7 to relational structures). Let $\mathfrak{D}$ be the expansion of $\mathfrak{C}$ by all unary singleton relations. We have that $\mathfrak{D} \in \mathrm{HI}(\mathfrak{B})$ by Proposition 5.25. It follows from Corollary 9.10 that there exists a minion homomorphism from $\mathscr{C} = \mathrm{Pol}(\mathfrak{B})$ to the idempotent clone $\mathscr{D} := \mathrm{Pol}(\mathfrak{D})$. Conversely, we have that $\mathscr{D} \subseteq \mathrm{Pol}(\mathfrak{C})$, and $\mathrm{Pol}(\mathfrak{C})$ has a minion homomorphism to $\mathrm{Pol}(\mathfrak{B}) = \mathscr{C}$ since $\mathfrak{B} \in \mathrm{H}(\mathfrak{C})$ and again by Corollary 9.10. □

Exercises.

159. Show that every finite structure $\mathfrak{B}$ with totally symmetric polymorphisms of all arities can be pp-constructed in $(\{0,1\}; \{0,1\}^3 \setminus \{(1,1,0)\}, \{0\}, \{1\})$.

     Hints: Exercise 120, Lemma 6.26, Corollary 9.10, Theorem 3.15, Theorem 3.12.

9.5 Taylor Terms

The following goes back to Walter Taylor [92]. We slightly deviate from the historic definition in that we do not require idempotence – this allows us to give stronger formulations of several results in the following.

Definition 9.13 (Taylor operations)**.** A function $f\colon B^n \to B$, for $n \geq 2$, is called a *Taylor operation* if for every $i \in [n]$ there are $\alpha, \beta\colon [n] \to [2]$ such that $f_\alpha = f_\beta$ and $\alpha(i) \neq \beta(i)$.

Examples for Taylor operations are binary commutative operations, majority operations, and Maltsev operations. Since we do not insist on idempotence, also quasi majority operations (Exercise 62) are examples of Taylor operations.

A *Taylor term* of a τ-algebra $\mathbf{B}$ is a τ-term $t(x_1, \dots, x_n)$, for $n \geq 2$, such that $t^{\mathbf{B}}$ is a Taylor operation. Note that t is a Taylor term if and only if it satisfies a set of n height-one identities that can be written as

$$t \begin{pmatrix} x & ? & ? & \cdots & ? \\ ? & x & ? & & \vdots \\ \vdots & ? & \ddots & \ddots & \vdots \\ \vdots & & \ddots & x & ? \\ ? & \cdots & \cdots & ? & x \end{pmatrix} \approx t \begin{pmatrix} y & ? & ? & \cdots & ? \\ ? & y & ? & & \vdots \\ \vdots & ? & \ddots & \ddots & \vdots \\ \vdots & & \ddots & y & ? \\ ? & \cdots & \cdots & ? & y \end{pmatrix}$$

where t is applied row-wise and ? stands for either x or y.

Walter Taylor did not just introduce Taylor operations, but he also found a beautiful statement about their existence (Theorem 9.15). In the proof of this statement, we need the following important observation about the star product in idempotent algebras.

Lemma 9.14. *For $n \in \mathbb{N}$, let $(A; f_1, \dots, f_n)$ be an idempotent algebra. Then there exists $g \in \mathrm{Clo}(\mathbf{A})$ such that for every $f \in \{f_1, \dots, f_n\}$ there exists α such that $g_\alpha = f$.*

Proof. The statement is clear if $n \leq 1$. First consider the case that $n = 2$. Let m be the arity of f_1 and let l be the arity of f_2. Note that

$$\mathrm{Clo}(\mathbf{A}) \models \big(f_1 = \mathrm{comp}_m^{ml}(f_1 * f_2, \underbrace{\pi_1^m, \dots, \pi_1^m}_{l \text{ times}}, \underbrace{\pi_2^m, \dots, \pi_2^m}_{l \text{ times}}, \dots, \underbrace{\pi_m^m, \dots, \pi_m^m}_{l \text{ times}})\big) \tag{23}$$

$$\text{and } \mathrm{Clo}(\mathbf{A}) \models \big(f_2 = \mathrm{comp}_l^{ml}(f_1 * f_2, \underbrace{\pi_1^l, \dots, \pi_l^l}_{m \text{ times}}, \underbrace{\pi_1^l, \dots, \pi_l^l}_{m \text{ times}}, \dots, \underbrace{\pi_1^l, \dots, \pi_l^l}_{m \text{ times}})\big) \tag{24}$$

since $\mathbf{A}$ is idempotent. The general case can be shown easily by induction on n. $\square$

Theorem 9.15. *Let $\mathbf{B}$ be an idempotent algebra. Then the following are equivalent.*

(1) $\mathbf{B}$ has a Taylor term t.

(2) there is no minion homomorphism from $\mathrm{Clo}(\mathbf{B})$ to $\mathbf{Proj}$.

Proof. To show that (1) implies (2), suppose for contradiction that there is a minion homomorphism ξ from $\mathrm{Clo}(\mathbf{B})$ to $\mathbf{Proj}$. By definition of $\mathbf{Proj}$ we have $\xi(t^{\mathbf{B}}) = \pi_l^n$ for some $l \leq n$. By assumption, there are $\alpha, \beta\colon [n] \to [2]$ such that $(t^{\mathbf{B}})_\alpha = (t^{\mathbf{B}})_\beta$ and $\alpha(l) \neq \beta(l)$. Since

$\xi(t^{\mathbf{B}}) = \pi^n_l$ and ξ is a minion homomorphism, we therefore obtain that $\pi^2_1 = \pi^2_2$, which does not hold in **Proj**, a contradiction.

To show the converse implication, suppose that **B** does not have a Taylor term. We have to show that Clo(**B**) has a minion homomorphism to **Proj**. By Lemma 9.4, it suffices to show that every primitive positive sentence ϕ in the language of minions that holds in Clo(**B**) also holds in **Proj**. If $g_1, \dots, g_m$ are the existentially quantified variables in ϕ, then by Lemma 9.14 there exists $g \in \mathrm{Clo}(\mathbf{B})^{(n)}$, for some n, such that every g_i, $i \in [m]$, is a minor g_{α_i} of g. Hence, it suffices to define a minion homomorphism from Minion$(B; g)$ to **Proj**. By assumption, g is not a Taylor term, so there exists an argument i such that for all $\alpha, \beta \colon [n] \to [2]$ we have $\alpha(i) = \beta(i)$ or $g_\alpha \neq g_\beta$. For $\alpha \colon [n] \to [k]$, define $\xi(g_\alpha) := \pi^k_{\alpha(i)}$. This map is well-defined because if $g_\alpha = g_\beta$ for $\alpha, \beta \colon [n] \to [k]$, then $(g_\alpha)_\gamma = (g_\beta)_\gamma$ for all $\gamma \colon [k] \to [2]$, and hence $\gamma \circ \alpha(i) = \gamma \circ \beta(i)$ for all $\gamma \colon [k] \to [2]$, which implies that $\alpha(i) = \beta(i)$. Moreover, ξ is a minion homomorphism because $\xi(g_\alpha) = \pi^k_{\alpha(i)} = (\pi^n_i)_\alpha = \xi(g)_\alpha$. □

Remark 9.16. The original statement of Taylor is Theorem 9.15 with *clone homomorphism* instead of minion homomorphism; the version with minion homomorphisms in Theorem 9.15 is the statement we really care about in this course and leads to an easier proof.

The following lemma should be clear from the results that we have already seen.

Lemma 9.17. *Let $\mathfrak{B}$ and $\mathfrak{C}$ be homomorphically equivalent structures. Then $\mathfrak{B}$ has a Taylor polymorphism if and only if $\mathfrak{C}$ has a Taylor polymorphism.*

Proof. Let h be a homomorphism from $\mathfrak{B}$ to $\mathfrak{C}$, and g be a homomorphism from $\mathfrak{C}$ to $\mathfrak{B}$. Suppose that f is a Taylor polymorphism for $\mathfrak{B}$ of arity n. Then $(x_1, \dots, x_n) \mapsto h(f(g(x_1), \dots, g(x_n)))$ is a Taylor polymorphism of $\mathfrak{C}$. □

Corollary 9.18. *Let $\mathfrak{B}$ be a finite structure. Then the following are equivalent.*

1. $K_3 \notin \mathrm{HI}(\mathfrak{B})$.
2. *$\mathfrak{B}$ has a Taylor polymorphism.*
3. $\mathrm{Pol}(\mathfrak{B})$ *has no minion homomorphism to* **Proj**.

If these condition don't apply then $\mathfrak{B}$ has a finite-signature reduct with an NP-hard CSP.

Proof. The equivalence of 1. and 3. is Corollary 9.11. Now suppose that $K_3 \notin \mathrm{HI}(\mathfrak{B})$. If $\mathfrak{B}'$ is the core of $\mathfrak{B}$, and $\mathfrak{C}$ is the expansion of $\mathfrak{B}'$ by all unary singleton relations, then $\mathrm{C}(\mathrm{H}(\mathfrak{B})) \subseteq \mathrm{HI}(\mathfrak{B})$ implies that K_3 is not pp constructible in $\mathfrak{C}$. Hence, the idempotent clone $\mathrm{Pol}(\mathfrak{C})$ does not have a minion homomorphism to **Proj**. Theorem 9.15 shows that $\mathfrak{C}$ and thus also $\mathfrak{B}'$ must have a Taylor polymorphism. Lemma 9.17 implies that $\mathfrak{B}$ has a Taylor polymorphism.

Note that the existence of Taylor polymorphisms is preserved by minion homomorphisms, and since **Proj** does not have a Taylor operation we have that 2. implies 3.

The final statement follows from Corollary 5.16. □

Theorem 9.19 (Tractability Theorem, Version 3)**.** *Let $\mathfrak{B}$ be a relational structure with finite domain and finite signature. If $\mathfrak{B}$ has a Taylor polymorphism, then* CSP($\mathfrak{B}$) *is in P. Otherwise,* CSP($\mathfrak{B}$) *is NP-complete.*

Proof. An immediate consequence of Corollary 9.18 and Theorem 5.28. $\square$

A clone is said to be *Taylor* if it has a Taylor operation, and an algebra is called *Taylor* if it has a Taylor term operation.

Remark 9.20. We will from now on often use the formulation 'let $\mathbf{A}$ be a finite Taylor algebra' instead of 'let $\mathbf{A}$ be a finite algebra such that $\mathrm{Clo}(\mathbf{A})$ does not have a minion homomorphism to $\mathbf{Proj}$', even if we can avoid in the proofs the use of Taylor operations alltogether. The reason is that the assumption is shorter to state, and equivalent by the results of this section (see Exercise 159).

Exercises.

159. Show that the assumption in Theorem 9.15 that $\mathbf{B}$ is idempotent can be replaced by the assumption that its domain is finite.

9.6 Arc Consistency Revisited

In this section we revisit the arc consistency procedure from Section 3, generalised to arbitrary relational structures with finite domain and finite signature, in the light of minions and primitive positive constructions.

Definition 9.21. The minion $\mathbf{M}_{\mathrm{AC}}$ is defined as follows. For $n \geq 1$, the set $M_{\mathrm{AC}}^{(n)}$ consists of the set of all non-empty subsets of $\{1,\dots,n\}$. For $\alpha\colon [n] \to [m]$ and $f \in M_{\mathrm{AC}}^{(n)}$, we define f_α to be $\{\alpha(a) \mid a \in f\}$.

Note that $\mathbf{M}_{\mathrm{AC}}$ is isomorphic to $\mathrm{Pol}(\{0,1\};\{0\},\{1\},\{0,1\}^3 \setminus \{1,1,0\})$ (see Exercise 123). For a finite structure $\mathfrak{B}$, we write $\mathrm{AC}_{\mathfrak{B}}$ for the generalisation of the arc-consistency procedure AC_H from digraphs H to general relational structures $\mathfrak{B}$ (Exercise 84). We already know the following.

Theorem 9.22. *Let $\mathfrak{B}$ be a finite structure. Then the following are equivalent.*

1. $\mathrm{CSP}(\mathfrak{B})$ *is solved by* $\mathrm{AC}_{\mathfrak{B}}$.
2. $\mathrm{Pol}(\mathfrak{B})$ *has totally symmetric polymorphisms of all arities.*
3. $\mathbf{M}_{AC}$ *has a minion homomorphism to* $\mathrm{Pol}(\mathfrak{B})$.
4. $\mathfrak{B}$ *has a primitive positive construction in* $(\{0,1\};\{0\},\{1\},\{0,1\}^3 \setminus \{(1,1,0)\}$.

Proof. $1 \Leftrightarrow 2$ was already shown in Theorem 3.12.

$2 \Rightarrow 3$: let $s_k \in \mathrm{Pol}(\mathfrak{B})$ be a totally symmetric operation of arity k. Then the map which sends for every $n \in \mathbb{N}$ and $k \leq n$ the element $\{i_1,\dots,i_k\} \in M_{\mathrm{AC}}^{(n)}$ to the operation $(x_1,\dots,x_n) \mapsto s_k(x_{i_1},\dots,x_{i_k})$ is a minion homomorphism $\xi\colon \mathbf{M}_{\mathrm{AC}} \to \mathrm{Pol}(\mathfrak{B})$: if $\alpha\colon [n] \to [m]$ for some $n,m \in \mathbb{N}$, then

$$\begin{aligned}
\xi(\{i_1,\dots,i_k\}_\alpha) &= \xi(\{\alpha(i_1),\dots,\alpha(i_k)\}) \\
&= ((x_1,\dots,x_m) \mapsto s_k(x_{\alpha(i_1)},\dots,x_{\alpha(i_k)}) \\
&= ((x_1,\dots,x_n) \mapsto s_k(x_{i_1},\dots,x_{i_k})_\alpha \\
&= \xi(\{i_1,\dots,i_k\})_\alpha.
\end{aligned}$$

3 ⇒ 2: let ξ be the minion homomorphism from $\mathbf{M}_{\mathrm{AC}}$ to $\mathrm{Pol}(\mathfrak{B})$. Then the operation $s_n := \xi(\{1,\dots,n\}) \in \mathrm{Pol}(\mathfrak{B})$ is totally symmetric and of arity n. Indeed, suppose that $a_1,\dots,a_n,b_1,\dots,b_n \in B$ are such that $\{a_1,\dots,a_n\} = \{b_1,\dots,b_n\} = \{c_1,\dots,c_m\}$ for some $m \leq n$ such that $c_1,\dots,c_m$ are pairwise distinct. Then there are $\alpha_1,\alpha_2 \colon [n] \to [m]$ such that $\alpha_1(i) = \alpha_2(i) \in \{1,\dots,m\}$ for every $i \in [n]$. Note that $s_n(a_1,\dots,a_n) = (s_n)_{\alpha_1}(c_1,\dots,c_m) = (s_n)_{\alpha_2}(b_1,\dots,b_n)$.

3 ⇔ 4: Corollary 9.10. □

Exercises.

159. Verify that $\mathbf{M}_{\mathrm{AC}}$ (Definition 9.21) is indeed a minion.

10 Undirected Graphs

This section contains a proof of the dichotomy for finite undirected graphs of Hell and Nešetřil, Theorem 2.6. We prove something stronger, namely that the tractability theorem (Theorem 5.28) is true for finite undirected graphs $\mathfrak{B}$ [34]. More specifically, the following is true.

Theorem 10.1. *Let $\mathfrak{B}$ be a finite undirected graph. Then either*

- $\mathfrak{B}$ *is* bipartite *(i.e., homomorphic to K_2) or has a loop, or*
- $\mathrm{HI}(\mathfrak{B})$ *contains all finite structures.*

Note that in combination with Corollary 5.16, this theorem implies the tractability theorem (Theorem 5.28) for the special case of finite undirected graphs. This theorem also has a remarkable consequence in universal algebra, whose significance goes beyond the study of the complexity of CSPs, and which provides a strengthening of Taylor's theorem (Theorem 9.15), discovered by Siggers in 2010 (see Section 10.2).

10.1 The Hell-Nešetřil Theorem

The graph $K_4 - \{0,1\}$ (a clique with four vertices where one edge is missing) is called a *diamond*. A graph is called *diamond-free* if it does not contain a copy of a diamond as a (not necessarily induced) subgraph. For every $\ell \in \mathbb{N}$, the graph $(K_3)^\ell$ is an example of a diamond-free graph.

Lemma 10.2. *Let $\mathfrak{B}$ be a finite undirected loopless graph which is not bipartite. Then $\mathfrak{B}$ pp-constructs a diamond-free core containing a triangle.*

Proof. We may assume that

1. $\mathrm{HI}(\mathfrak{B})$ does not contain a non-bipartite loopless graph with fewer vertices than $\mathfrak{B}$, since otherwise we could replace $\mathfrak{B}$ by this graph. In particular, $\mathfrak{B}$ must then be a core.
2. $\mathfrak{B} = (V;E)$ contains a triangle: if the length of the shortest odd cycle in $\mathfrak{B}$ is k, then $(B;E^{k-2})$ is a graph and contains a triangle, so it can replace $\mathfrak{B}$.

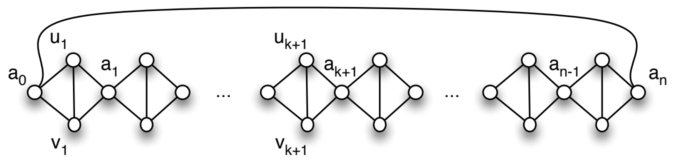

Figure 16: Diagram for the proof of Lemma 10.2.

Claim 1. Every vertex of $\mathfrak{B}$ is contained in a triangle: Otherwise, we can replace $\mathfrak{B}$ by the subgraph of $\mathfrak{B}$ induced by set defined by the primitive positive formula

$$\exists u, v \big(E(x,u) \wedge E(x,v) \wedge E(u,v)\big)$$

which still contains a triangle, contradicting our first assumption.

Claim 2. $\mathfrak{B}$ does not contain a copy of K_4. Otherwise, if a is an element from a copy of K_4, then the subgraph of $\mathfrak{B}$ induced by the set defined by the primitive positive formula $E(a,x)$ is a non-bipartite graph $\mathfrak{A}$, which has strictly less vertices than $\mathfrak{B}$ because $a \notin A$. Moreover, Theorem 5.26 implies that expansions of cores by constants can be pp-constructed, and hence that $\mathfrak{B}$ pp-constructs $\mathfrak{A}$, contrary to our initial assumption.

Claim 3. The graph $\mathfrak{B}$ must also be diamond-free. To see this, let R be the binary relation with the primitive positive definition

$$R(x,y) :\Leftrightarrow \exists u, v \big(E(x,u) \wedge E(x,v) \wedge E(u,v) \wedge E(u,y) \wedge E(v,y)\big)$$

and let T be the transitive closure of R. The relation T is clearly symmetric, and since every vertex of $\mathfrak{B}$ is contained in a triangle, it is also reflexive, and hence an equivalence relation of $\mathfrak{B}$. Since B is finite, for some n the formula $\exists u_1, \dots, u_n \big(R(x,u_1) \wedge R(u_1,u_2) \wedge \cdots \wedge R(u_n,y)\big)$ defines T, showing that T is primitively positively definable in $\mathfrak{B}$.

We claim that the graph $\mathfrak{B}/T$ (see Example 5.13) does not contain loops. It suffices to show that $T \cap E = \emptyset$. Otherwise, let $(a,b) \in T \cap E$. Choose (a,b) in such a way that the shortest sequence $a = a_0, a_1, \dots, a_n = b$ with $R(a_0,a_1)$, $R(a_1,a_2)$, $\dots$, $R(a_{n-1},a_n)$ in $\mathfrak{B}$ is shortest possible; see Figure 16. This chain cannot have the form $R(a_0,a_1)$ because $\mathfrak{B}$ does not contain K_4 subgraphs. Suppose first that $n = 2k$ is even. Let the vertices u_1, v_1, u_{k+1} and v_{k+1} be as depicted in Figure 16. Let S be the set defined by

$$\exists x_1, \dots, x_k \big(E(u_{k+1},x_1) \wedge E(v_{k+1},x_1) \wedge R(x_1,x_2) \wedge \cdots \wedge R(x_{k-1},x_k) \wedge E(x_k,x)\big).$$

Note that $a_0, u_1, v_1 \in S$ form a triangle. If $a_n \in S$ then we obtain a contradiction to the minimal choice of n. Hence, the subgraph induced by the primitively positively definable set S is non-bipartite and strictly smaller than $\mathfrak{B}$, in contradiction to the initial assumption.

If $n = 2k+1$ is odd, we can argue analogously with the set S defined by the formula

$$\exists x_1, \dots, x_k \big(R(a_{k+1},x_1) \wedge R(x_1,x_2) \wedge \cdots \wedge R(x_{k-1},x_k) \wedge E(x_k,x)\big)$$

and again obtain a contradiction. So we conclude that $\mathfrak{B}/T$ does not contain loops. It also follows that $\mathfrak{B}/T$ contains a triangle, because $\mathfrak{B}$ contains a triangle.

Thus, the initial assumption on $\mathfrak{B}$ then implies that T must be the equality relation on B, which in turn implies that $\mathfrak{B}$ does not contain any diamonds. $\square$

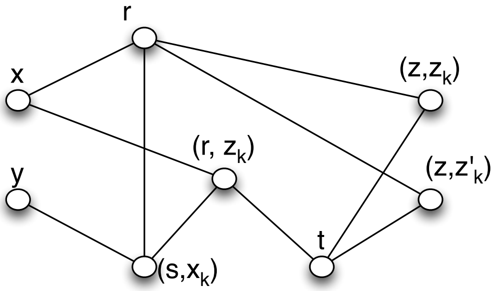

Figure 17: Diagram for the proof of Lemma 10.3.

Lemma 10.3 (from [34]). *Let $\mathfrak{B}$ be a diamond-free undirected graph and let $h\colon (K_3)^k \to \mathfrak{B}$ be a homomorphism. Then the image of h is isomorphic to $(K_3)^m$ for some $m \leq k$.*

Proof. Let $I \subseteq \{1, \ldots, k\}$ be maximal such that $\ker(h) \subseteq \ker(\pi_I)$. Note that π_I is defined even if $I = \emptyset$ (Definition 6.30). Such a set exists, because $\ker(\pi_\emptyset)$ is the total relation. We claim that $\ker(h) = \ker(\pi_I)$; this clearly implies the statement.

By the maximality of I, for every $j \in \{1, \ldots, k\} \setminus I$ there are $x, y \in (K_3)^k$ such that $h(x) = h(y)$ and $x_j \neq y_j$. We have to show that for all $z_1, \ldots, z_k, z'_j \in \{0, 1, 2\}$

$$h(z_1, \ldots, z_j, \ldots, z_k) = h(z_1, \ldots, z_{j-1}, z'_j, z_{j+1}, \ldots, z_k).$$

We may suppose that $z_j \neq x_j$ and $z'_j = x_j$. To simplify notation, we assume that $j = k$. As we have seen in Exercises 7 and 8, any two vertices in $(K_3)^k$ have a common neighbour.

- Let r be a common neighbour of x and $(z, z_k) := (z_1, \ldots, z_k)$. Note that r and (z, z'_k) are adjacent, too.
- For all $i \neq k$ we choose an element s_i of K_3 that is distinct from both r_i and y_i. Since x_k is distinct from r_k and y_k we have that $(s, x_k) := (s_1, \ldots, s_{k-1}, x_k)$ is a common neighbour of r and y.
- The tuple $(r, z_k) := (r_1, \ldots, r_{k-1}, z_k)$ is a common neighbour of both x and (s, x_k).
- Finally, for $i \neq k$ choose t_i to be distinct from z_i and r_i, and choose t_k to be distinct from z_k and from z'_k. Then $t := (t_1, \ldots, t_{k-1}, t_k)$ is a common neighbour of (z, z_k), of (z, z'_k), and of (r, z_k).

The situation is illustrated in Figure 17. Since $\mathfrak{B}$ is diamond-free, $h(x) = h(y)$ implies that $h(r) = h(r, z_k)$ and for the same reason $h(z, z_k) = h(z, z'_k)$ which completes the proof. $\square$

Lemma 10.4 (from [34]). *If a finite diamond-free graph $\mathfrak{B}$ contains a triangle, then for some $k \in \mathbb{N}$ there is a primitive positive interpretation of $(K_3)^k$ with constants in $\mathfrak{B}$.*

Proof. We construct a strictly increasing sequence of subgraphs $G_1 \subset G_2 \subset \cdots$ of $\mathfrak{B}$ such that G_i is isomorphic to $(K_3)^{k_i}$ for some $k_i \in \mathbb{N}$. Let G_1 be any triangle in $\mathfrak{B}$. Suppose now that G_i has already been constructed. If the domain of G_i is primitively positively definable in $\mathfrak{B}$ with constants, then we are done. Otherwise, there exists an idempotent polymorphism f of $\mathfrak{B}$ and $v_1, \dots, v_k \in G_i$ such that $f(v_1, \dots, v_k) \notin G_i$. The restriction of f to G_i is a homomorphism from $(K_3)^{k_i}$ to the diamond-free graph $\mathfrak{B}$. Lemma 10.3 shows that $G_{i+1} := f((G_i)^k)$ induces a copy of $(K_3)^{k_{i+1}}$ for some $k_{i+1} \leq k$. Since f is idempotent, we have that $G_i \subseteq G_{i+1}$, and by the choice of f the containment is strict. Since $\mathfrak{B}$ is finite, for some m the set G_m must have a primitive positive definition in $\mathfrak{B}$ with constants. $\square$

Proof of Theorem 10.1. Let $\mathfrak{B}$ be a finite undirected graph that is not bipartite. Then $\mathfrak{B}$ interprets primitively positively a graph that is homomorphically equivalent to a diamond-free core $\mathfrak{C}$ containing a triangle, by Lemma 10.2. We may now apply Lemma 10.4 to $\mathfrak{C}$ and obtain that for some $k \in \mathbb{N}$ there is a primitive positive interpretation of $(K_3)^k$ with constants in $\mathfrak{C}$. Since $\mathfrak{C}$ is a core, and since $(K_3)^k$ is homomorphically equivalent to K_3, it follows that there is a primitive positive interpretation of a structure that is homomorphically equivalent to K_3 in $\mathfrak{C}$. The structure K_3 interprets all finite structures primitive positively (Theorem 5.17), so Theorem 5.26 implies that $\mathrm{H}(\mathrm{I}(\mathfrak{B}))$ contains all finite structures. $\square$

10.2 Siggers Terms of Arity 6

An operation $s\colon B^6 \to B$ is called *Siggers operation* (of arity six[8]) if for all $x, y, z \in B$

$$s(x, y, x, z, y, z) = s(y, x, z, x, z, y).$$

As usual, if $\mathbf{A}$ is an algebra and t is a term such that $t^{\mathbf{A}}$ is a Siggers operation, we call t a *Siggers term.*

Theorem 10.5 (from [91]). *Let $\mathfrak{B}$ be a finite structure. Then either $\mathfrak{B}$ primitively positively interprets all finite structures up to homomorphic equivalence, or $\mathfrak{B}$ has a Siggers polymorphism.*

Proof. Pick $k \geq 1$ and $a, b, c \in B^k$ such that $\{(a_i, b_i, c_i) \mid i \leq k\} = B^3$. Let R be the binary relation on B^k such that $(u, v) \in R$ if and only if there exists a 6-ary $s \in \mathrm{Pol}(\mathfrak{B})$ such that $u = s(a, b, a, c, b, c)$ and $v = s(b, a, c, a, c, b)$. We make the following series of observations.

- The vertices $a, b, c \in B^k$ induce in $(B^k; R)$ a copy of K_3: each of the six edges of K_3 is witnessed by one of the six 6-ary projections from $\mathrm{Pol}(\mathfrak{B})$.
- The relation R is symmetric: Suppose that $(u, v) \in R$ and let $s \in \mathrm{Pol}(\mathfrak{B})$ be such that $u = s(a, b, a, c, b, c)$ and $v = s(b, a, c, a, c, b)$. Define $s' \in \mathrm{Pol}(\mathfrak{B})$ by $s'(x_1, \dots, x_6) := s(x_2, x_1, x_4, x_3, x_6, x_5)$; then

  $$v = s(b, a, c, a, c, b) = s'(a, b, a, c, b, c)$$
  $$u = s(a, b, a, c, b, c) = s'(b, a, c, a, c, b)$$

  and hence s' witnesses that $(v, u) \in R$.
- If the graph $(B^k; R)$ contains a loop $(w, w) \in R$, then there exists a 6-ary $s \in \mathrm{Pol}(\mathfrak{B})$ such that

  $$s(a, b, a, c, b, c) = w = s(b, a, c, a, c, b)\,.$$

[8]We stress the arity here since there is also a notion of Siggers operations for arity four, which is a similar but stronger result, see Section 14.3.

The operation s is Siggers: for all $x, y, z \in B$ there exists an $i \leq k$ such that $(x, y, z) = (a_i, b_i, c_i)$, and the above implies that

$$s(a_i, b_i, a_i, c_i, b_i, c_i) = s(b_i, a_i, c_i, a_i, c_i, b_i)$$

and we are done in this case.

So we may assume in the following that $(B^k; R)$ is a simple (i.e., undirected and loopless) graph that contains a copy of K_3. The relation R (as a $2k$-ary relation over B) is preserved by $\mathrm{Pol}(\mathfrak{B})$, and hence $(B^k; R)$ has a primitive positive interpretation in $\mathfrak{B}$. By Theorem 10.1 applied to the undirected graph $(B^k; R)$, there is a primitive positive interpretation in $(B^k; R)$ of all finite structures up to homomorphic equivalence, and hence also in $\mathfrak{B}$, and this concludes the proof. $\square$

Theorem 10.6 (Tractability Theorem, Version 4)**.** *Let $\mathfrak{B}$ be a relational structure with finite domain and finite signature. If $\mathfrak{B}$ has a Siggers polymorphism, then* $\mathrm{CSP}(\mathfrak{B})$ *is in P. Otherwise,* $\mathrm{CSP}(\mathfrak{B})$ *is NP-complete.*

Proof. An immediate consequence of Theorem 10.5 and Theorem 5.28. $\square$

Exercises.

160. A Taylor clone $\mathbf{C}$ is called *minimal Taylor* if every proper subclone of $\mathbf{C}$ is not Taylor. Show that every Taylor clone on a finite set contains a minimal Taylor clone.

11 Congruences

The study of congruences of algebras and varieties is one of the central topics in universal algebra. In Section 11.1 we present some basic facts about congruences that will be used later in the text. Section 11.2 about congruence permutability and Section 11.3 about congruence distributivity will not be used later in the text and can be skipped by the hasty reader.

11.1 The Congruence Lattice

Let $\mathbf{A}$ be a τ-Algebra. We write $\mathrm{Con}(\mathbf{A})$ for the set of all congruences of $\mathbf{A}$ (Definition 8.10). Clearly, $\mathrm{Con}(\mathbf{A})$ is closed under arbitrary intersections. On the other hand, the union of two congruences is in general not a congruence.

Lemma 11.1. *Let $\mathbf{A}$ be an algebra. Then $(\mathrm{Con}(\mathbf{A}), \subseteq)$ is a complete lattice (Example 8.5).*

Proof. Let $(E_i)_{i \in I}$ be a sequence of congruences of $\mathbf{A}$. Define (recall the definition of the relational product, Definition 5.7)

$$\bigvee_{i \in I} E_i = \bigcup \{E_{i_1} \circ \cdots \circ E_{i_k} \mid i_1, \dots, i_k \in I, k \in \mathbb{N}\}$$

Note that this is the smallest (with respect to inclusion) equivalence relation that contains all the E_i. Let $f \in \tau$ be n-ary and $(a_1, b_1), \dots, (a_n, b_n) \in E$. Then there are $i_1, \dots, i_k \in I$ such that for all $j \in \{1, \dots, n\}$

$$(a_j, b_j) \in E_{i_1} \circ \cdots \circ E_{i_k}.$$

Hence, $(f(a_1, \dots, a_n), f(b_1, \dots, b_n)) \in E_{i_1} \circ \cdots \circ E_{i_k}$ and $E \in \mathrm{Con}(\mathbf{A})$. $\square$

Every algebra has the following two congruences.

- Δ_A: the *diagonal relation* $\{(a,a) \mid a \in A\}$.
- ∇_A: the *universal relation* A^2.

Congruences that are different from ∇_A and Δ_A are called *proper*.

Definition 11.2. Algebras $\mathbf{A}$ without proper congruences are called *simple*.

Example 11.3. Groups that are simple in the sense of group theory are simple in the more general sense of Defintion 11.2. △

Example 11.4. Let G be a permutation group on the set A. Let $\mathbf{A}$ be the algebra with domain A and signature G, und define $g^{\mathbf{A}} := g$ for all $g \in G$. Then $\mathbf{A}$ is simple if and only if G is primitive as a permutation group. △

Definition 11.5. Let $\mathbf{A}$ be an algebra and let $\mathbf{B}$ be the expansion of $\mathbf{A}$ by all constant operations. A *polynomial over* $\mathbf{A}$ is a term in the signature of $\mathbf{B}$. A *polynomial operation of* $\mathbf{A}$ is a term operation of $\mathbf{B}$.

Definition 11.6. Two algebras $\mathbf{A}_1, \mathbf{A}_2$ with the same domain A are called *polynomially equivalent* if they have the same polynomial operations.

Note that polynomially equivalent algebras have the same congruences.

Lemma 11.7. *Let $\mathbf{B}$ be an algebra and $X \subseteq B^2$. Then the smallest congruence of $\mathbf{B}$ that contains X, denoted by $\mathrm{Cg}_{\mathbf{B}}(X)$, equals the symmetric transitive closure of*

$$T := \{(p(a), p(b)) \mid (a,b) \in X, p \text{ a unary polynomial operation of } \mathbf{B}\}. \tag{25}$$

Proof. Let C be the symmetric transitive closure of T. Clearly, if $(a,b) \in X$ and p is a unary polynomial operation of $\mathbf{B}$, then $(p(a), p(b)) \in \mathrm{Cg}_{\mathbf{B}}(X)$ since congruences are preserved by term operations and are reflexive. Since $\mathrm{Cg}_{\mathbf{B}}(X)$ is transitive and reflexive we obtain that $C \subseteq \mathrm{Cg}_{\mathbf{B}}(X)$. To prove that $\mathrm{Cg}_{\mathbf{B}}(X) \subseteq C$ it suffices to prove that S is a congruence of $\mathbf{B}$ that contains X. Since S is preserved by all constant operations, we have reflexivity of S, and we clearly have that $X \subseteq S$. We are left with the verification that S is a congruence. If $(u_1, v_1), \dots, (u_k, v_k) \in S$ and f is an operation of $\mathbf{B}$ of arity k, then for each $i \in \{1, \dots, k\}$ there exists a path of edges in T that starts in u_i and ends in v_i. By the reflexivity of T we may duplicate elements on these paths such these paths all have the same length, and for all ℓ, the ℓ-th edge is a forward edge on all paths, or it is a backward edge on all paths. Hence, we obtain a path of edges in T between $u_0 := f(u_1, \dots, u_k)$ and $v_0 := f(v_1, \dots, v_k)$, showing that $(u_0, v_0) \in S$ are we are done. □

Exercises.

161. Show that the algebra from Example 7.4 is simple.
162. Show that if $\mathbf{A}$ is an idempotent algebra and C a congruence of $\mathbf{A}$, then every congruence class of C is a subalgebra of $\mathbf{A}$.
163. Prove the remark after Definition 11.6.

164. Let $\mathbf{A}$ be an algebra. Show that an equivalence relation $R \subseteq A^2$ is a congruence of $\mathbf{A}$ if and only if it is preserved by all *unary* polynomials of $\mathbf{A}$. Which properties of an equivalence relation do you need in the proof?

165. Let $\mathbf{A}$ be an algebra. Then a relation $R \subseteq A^n$ is called

   - *reflexive* if $(a, \dots, a) \in R$ for every $a \in A$.
   - *transitive* if for all $(a_{ij})_{i,j \in \{1,\dots,n\}} \in A^{n \times n}$ whose rows and columns are from R, we have that $(a_{11}, \dots, a_{nn}) \in R$.

   Show that a reflexive and transitive R is preserved by $\mathrm{Clo}(\mathbf{A})$ if and only if it is preserved by all unary polynomial operations of $\mathbf{A}$.

166. Show that $\mathbf{A}$ is *polynomially complete*, i.e., the clone of polynomial operations of $\mathbf{A}$ equals the set of all operations on A, if and only if the *discriminator operation* $d \colon A^3 \to A$ is a polynomial operation of $\mathbf{A}$, which is defined as follows:

$$d(x, y, z) := \begin{cases} z & x = y \\ x & x \neq y. \end{cases}$$

167. Let $\mathbf{A}$ be an algebra on a finite set and $R \leq \mathbf{A}^2$ be subdirect. Then $\bigcup_{i \in \mathbb{N}} (R^{-1} \circ R)^i$ is a congruence of $\mathbf{A}$.

11.2 Congruence Permutability

Two congruences $C_1, C_2 \in \mathrm{Con}(\mathbf{A})$ *permute* if

$$C_1 \circ C_2 = C_2 \circ C_1.$$

An algebra $\mathbf{A}$ is called *congruence permutable* if all pairs of congruences of $\mathbf{A}$ permute. In the following, it will be convenient to write $a_0 C_1 a_1 C_2 a_2 \cdots a_n C_n a_n$ as a shortcut for $a_1 C_1 a_2$ and $a_2 C_2 a_3$ and $\cdots a_{n-1} C a_n$.

Lemma 11.8. *Let $\mathbf{A}$ be an algebra such that $\mathrm{Clo}(\mathbf{A})$ contains a Maltsev operation p. Then $\mathbf{A}$ is congruence permutable.*

Proof. Let $C, E \in \mathrm{Con}(\mathbf{A})$ and let $(a, b) \in C \circ E$. Then there exists $c \in A$ with $(a, c) \in C$ und $(c, b) \in E$. Note that

$$b = p^{\mathbf{A}}(c, c, b) \; C \; p^{\mathbf{A}}(a, c, b) \; E \; p^{\mathbf{A}}(a, b, b) = a$$

and thus $(b, a) \in C \circ E$ and $(a, b) \in E \circ C$. □

Note that most classical algebras, such as groups, rings, fields, etc., do have a Maltsev term operation, and hence are congruence permutable. A variety is congruence permutable if all algebras in the variety are congruence permutable.

Theorem 11.9 (Maltsev)**.** *Let $\mathcal{K}$ be a class of τ-algebras. Then $\mathrm{HSP}(\mathcal{K})$ is congruence permutable if and only if there exists a τ-term $t(x, y, z)$ such that every algebra in $\mathcal{K}$ satisfies $t(y, x, x) \approx t(x, x, y) \approx y$.*

Proof. "$\Leftarrow$". If every algebra in $\mathcal{K}$ has a Maltsev term operation, then so does $\mathrm{HSP}(\mathcal{K})$, and hence the statement follows from Lemma 11.8.

"$\Rightarrow$". Let $\mathbf{F} := F_{\mathcal{K}}(\{x,y,z\})$. For $\mathbf{F} := \mathbf{F}_{\mathcal{K}}(X)$ und $u,v \in \{x,y,z\}$ we write $C(u,v)$ for the smallest congruence of $\mathbf{F}$ that contains (u,v). Let $C_1 := C(x,y) \in \mathrm{Con}(\mathbf{F})$ and $C_2 := C(y,z) \in \mathrm{Con}(\mathbf{F})$. Since $(x,z) \in C_1 \circ C_2 = C_2 \circ C_1$ there exists $b \in F$ with $(x,b) \in C_2$ and $(b,z) \in C_1$. Since $\mathbf{F}$ is generated by $\{x,y,z\}$, there is a τ-term $p(x,y,z)$ with $b = p^{\mathbf{F}}(x,y,z)$. We will show that $\mathcal{K} \models \forall x,y.\, p(x,x,y) = y$. Let $\mathbf{A} \in \mathcal{K}$ and $u,v \in A$. Let $f \colon \mathbf{F} \to \mathbf{A}$ be a homomorphism with $f(x) = u$, $f(y) = u$, and $f(z) = v$. Then $f(p^{\mathfrak{F}}(x,y,z)) = p^{\mathbf{A}}(u,u,v)$. Since $(x,y) \in \mathrm{Ker}(f)$ we have $C_1 \subseteq \mathrm{Ker}(f)$. Thus, $(b,z) \in \mathrm{Ker}(f)$ and

$$v = f(z) = f(b) = f\big(p^{\mathbf{F}}(x,y,z)\big) = p^{\mathbf{A}}(u,u,v).$$

$\mathcal{K} \models p(y,x,x) \approx y$ can be shown similarly. $\square$

Recall that in Section 4.4 we proved that digraphs with a Maltsev polymorphism are rectangular, and in Theorem 4.19 we characterised the existence of Maltsev polymorphisms of digraphs using total rectangularity. The following corollary clarifies the connection between rectangularity and Maltsev terms.

Proposition 11.10. *Let $\mathbf{A}$ be an algebra. Then $\mathbf{A}$ has a Maltsev term if and only if every $R \leq \mathbf{B}^2$, for every algebra $\mathbf{B} \in \mathrm{HSP}(\mathbf{A})$, is rectangular.*

Proof. Clearly, if $\mathbf{A}$ has a Maltsev term, then every algebra $\mathbf{B}$ in $\mathrm{HSP}(\mathbf{A})$ has a Maltsev term, and hence every $R \leq \mathbf{B}^2$ is rectangular. Conversely, let $\mathbf{B} \in \mathrm{HSP}(\mathbf{A})$ be the free algebra for $\mathrm{HSP}(\mathbf{A})$ over $\{x,y\}$, and let R be the subuniverse of $\mathbf{B}^2$ generated by $\{(x,x),(x,y),(y,y)\}$. Since R is rectangular, we have that $(x,y) \in R$. Hence, there exists a term t such that $f((x,x),(x,y),(y,y)) = (y,x)$. Then t satisfies $t(x,x,y) \approx y$ and $t(x,y,y) \approx x$ in every algebra of $\mathrm{HSP}(\mathbf{A})$, (Lemma 8.24), and hence $\mathbf{A}$ has a Maltsev term. $\square$

11.3 Congruence Distributivity

A lattice $(P; \wedge, \vee, 0, 1)$ is called *distributive* if it satisfies

$$\begin{aligned} & (x \wedge y) \vee z \approx (x \vee z) \wedge (y \vee z) \\ \text{and} \quad & (x \vee y) \wedge z \approx (x \wedge z) \vee (y \wedge z). \end{aligned}$$

An example of a distributive lattice is the set of subsets of a set S, ordered by inclusion: $(\mathcal{P}(S); \subseteq)$. If the congruence lattice of $\mathbf{A}$ is distributive, we call $\mathbf{A}$ *congruence distributive.*

Lemma 11.11. *Every algebra with a majority term operation is congruence distributive.*

Proof. Let $C,D,E \in \mathrm{Con}(\mathfrak{A})$ and $(a,b) \in C \wedge (D \vee E)$. Then there are $c_1,\dots,c_n$ such that

$$a D c_1 E c_2 D c_3 \dots c_n E b.$$

Since $(a,b) \in C$, we have for all $c \in A$ that

$$m^{\mathfrak{A}}(a,c,b) \, C \, m^{\mathfrak{A}}(a,c,a) = a.$$

Thus, for all $c_1, c_2 \in A$ we obtain

$$m^{\mathfrak{A}}(a,c_1,b) \, C \, a \, C \, m^{\mathfrak{A}}(a,c_2,b). \tag{26}$$

Therefore,

$$\begin{aligned}
a = m^{\mathfrak{A}}(a,a,b)(C \wedge D)m^{\mathfrak{A}}&(a,c_1,b) && \text{(by (26))}\\
(C \wedge E)m^{\mathfrak{A}}&(a,c_2,b)\\
\dots&\\
(C \wedge D)m^{\mathfrak{A}}&(a,c_n,b)\\
(C \wedge E)m^{\mathfrak{A}}&(a,b,b) = b.
\end{aligned}$$

We conclude that $(a,b) \in (C \wedge D) \vee (C \wedge E)$. $\square$

As in the case of congruence permutability, there is even an equational characterisation of congruence distributivity of varieties.

Theorem 11.12 (Jónsson). *$\mathrm{HSP}(\mathcal{K})$ is congruence distributive if and only if there exists $n \in \mathbb{N}$ and τ-terms $p_0, \dots, p_n$ such that*

$$\begin{aligned}
\mathcal{K} \models p_i(x,y,x) &\approx x && \text{for } i \in \{1,\dots,n\}\\
p_0(x,y,z) &\approx x\\
p_i(x,x,y) &\approx p_{i+1}(x,x,y) && \text{for } i \text{ even}\\
p_i(x,y,y) &\approx p_{i+1}(x,y,y) && \text{for } i \text{ odd}\\
p_n(x,y,z) &\approx z
\end{aligned}$$

Proof. "$\Rightarrow$". Let $\mathbf{F} := F_{\mathcal{K}}(\{x,y,z\})$. We have

$$C(x,z) \wedge \big(C(x,y) \vee C(y,z)\big) = \big(C(x,z) \wedge C(x,y)\big) \vee \big(C(x,z) \wedge C(y,z)\big)$$

hence $(x,z) \in \big(C(x,z) \wedge C(x,y)\big) \vee \big(C(x,z) \wedge C(y,z)\big)$ in $\mathbf{F}$. Thus, there are $p_1, \dots, p_{n-1} \in \mathbf{F}$ such that each of $(x,p_1), (p_1,p_2), \dots, (p_{n-1},z)$ is in $C(x,z) \wedge C(x,y)$ or in $C(x,z) \wedge C(y,z)$. By padding the sequence $p_1, \dots, p_{n-1}$ with repeated entries, we may even suppose that

$$x\big(C(x,z) \wedge C(x,y)\big)p_1 \tag{27}$$
$$p_1\big(C(x,z) \wedge C(y,z)\big)p_2 \tag{28}$$
$$\vdots$$
$$p_{n-1}\big(C(x,z) \wedge C(y,z)\big)z \tag{29}$$

From (27) we obtain that

$$x = p_1(x,y,x) = p_1(x,x,y).$$

From (28) we obtain

$$\begin{aligned}
p_1(x,y,x) &= p_2(x,y,x)\\
\text{and} \quad p_1(x,y,y) &= p_2(x,y,y)
\end{aligned}$$

and from (29) that $p_{n-1}(x,y,x) = p_{n-1}(x,y,y) = z$. Similarly, all other identities from the statement can be derived from the sequence (28)-(29).

"$\Leftarrow$". Let $C_1, C_2, C_3 \in \mathrm{Con}(\mathbf{A})$ for $\mathbf{A} \in \mathrm{HSP}(\mathcal{K})$. It suffices to show that

$$C_1 \wedge (C_2 \vee C_3) \subseteq (C_1 \wedge C_2) \vee (C_1 \wedge C_3)$$

since the converse inclusion holds in every lattice. Let $(a,b) \in C_1 \wedge (C_2 \vee C_3)$. That is, there are $c_1, \dots, c_t$ with

$$aC_2c_1C_3c_2C_2 \cdots c_tC_3b.$$

For $i \in \{1, \dots, n\}$

$$p_i(a,a,b)C_2p_i(a,c_1,b)C_3p_i(a,c_2,b) \cdots C_3p_i(a,b,b)$$

and since $p_i(a,c,b)C_1p_i(a,c,a) = a$

$$p_i(a,a,b)(C_1 \wedge C_2)p_i(a,c_1,b)(C_1 \wedge C_3)p_i(a,c_2,b) \cdots (C_1 \wedge C_3)p_i(a,b,b).$$

Therefore,

$$p_i(a,a,b)\big((C_1 \wedge C_2) \vee (C_1 \wedge C_3)\big)p_i(a,b,b)$$

We conclude that $a\big((C_1 \wedge C_2) \vee (C_1 \wedge C_3)\big)b$. □

If the variety is generated by the polymorphism clone of a finite structure $\mathfrak{B}$ with finite relational signature, this condition has drastic consequences for CSP($\mathfrak{B}$), similarly as in the previous section for congruence permutable varities. Barto [7] proved that in this case $\mathfrak{B}$ must also have a near unanimity polymorphism and hence can be solved in polynomial time by the methods that will be presented in Section 15.

Exercises

168. Show that $\mathrm{SH}(\mathcal{K}) \subseteq \mathrm{HS}(\mathcal{K})$, $\mathrm{PS}(\mathcal{K}) \subseteq \mathrm{SP}(\mathcal{K})$, and $\mathrm{PH}(\mathcal{K}) \subseteq \mathrm{HP}(\mathcal{K})$.

169. Let $f \colon \mathbf{A} \to \prod_{i \in I} \mathbf{A}_i$ be a homomorphism. Show that

$$\mathrm{Ker}(f) = \bigcap_{i \in I} \mathrm{Ker}(\pi_i \circ f)$$

170. Show that the variety of all lattices is congruence distributive, but not congruence permutable.

171. Show that the variety of Boolean algebras is both congruence permutable and congruence distributive.

172. Show that a lattice satisfies the two identities

$$x \wedge (y \vee z) \approx (x \wedge y) \vee (x \wedge z)$$
$$x \vee (y \wedge z) \approx (x \vee y) \wedge (x \vee z)$$

if and only if it satisfies one of those identities.

12 Abelian Algebras

Modules (Example 8.3) have many strong properties and are very well understood. Affine algebras are 'essentially' modules and introduced in Section 12.1. The relevance of affine algebras in the context of constraint satisfaction is that core structures with more than one element and whose polymorphism algebra is idempotent and affine can pp-construct $(\mathbb{Z}_p; +, 1)$, for some prime p (Section 12.2). We have seen in Theorem 7.2 that the CSP of such structures cannot be solved by k-consistency, for any k.

Abelian algebras are defined more abstractly (Section 12.3). It turns out that under fairly general conditions, abelian algebras must be affine; this is the content of the fundamental theorem of abelian algebras which will be presented in Section 12 (Theorem 12.12) and generalised later in Section 13.4. Section 12.4 presents other useful characterisations of abelian algebras in terms of congruences.

12.1 Affine Algebras

An algebra $\mathbf{A}$ is called *affine* if $\mathbf{A}$ is polynomially equivalent (Definition 11.6) to a module (Example 8.3). Clearly, every module $\mathbf{M}$ has a Maltsev term operation, namely $(x, y, z) \mapsto x - y + z$, so every affine algebra has a Maltsev polynomial operation. Something stronger holds.

Lemma 12.1. *If* $\mathbf{A}$ *is affine, then* $(x, y, z) \mapsto x - y + z$ *is the* unique *Maltsev polynomial operation of* $\mathbf{A}$.

Proof. Suppose that $\mathbf{A}$ is polynomially equivalent to module $\mathbf{M}$ over the ring $\mathbf{R}$. Let $m(x, y, z) = \alpha x + \beta y + \gamma z + d$, for $\alpha, \beta, \gamma \in R$ and $d \in M$, be a Maltsev polynomial operation of $\mathbf{A}$. Since $m(0, 0, 0) = 0$ we must have $d = 0$. Moreover, for all $x \in M$ we have $x = m(x, 0, 0) = \alpha x$ and analogously we obtain $x = \gamma x$. Finally, $m(x, x, 0) = x + \beta x = 0$, and therefore $\beta x = -x$. We conclude that $m(x, y, z) = x - y + z$. □

Remark 12.2. The operation $(x, y, z) \mapsto x - y + z$ is an affine Maltsev operation as defined in Section 7.1. An algebra whose term operations are generated by an affine Maltsev operation is called an *affine Maltsev algebra* (also see Exercise 143). Note that affine Maltsev algebras are affine in the sense defined above.

Example 12.3. Every commutative group is affine: let $\mathbf{R}$ be the ring of integers $\mathbb{Z}$, and define scalar multiplication $n \cdot x$ as $\underbrace{x + \cdots + x}_{n \text{ times}}$. △

Definition 12.4. Let $\mathbf{A}$ be an algebra. An operation $m \colon A^k \to A$ is called *central in* $\mathbf{A}$ if m is a homomorphism from $\mathbf{A}^k \to \mathbf{A}$.

Remark 12.5. Note that m is central in $\mathbf{A}$ if and only if every operation of $\mathbf{A}$ preserves the graph of m (see Exercise 92). Hence, if $\mathbf{A} = \mathrm{Pol}(\mathfrak{A})$ for some finite structure $\mathfrak{A}$, then m is central in $\mathbf{A}$ if and only if its graph has a primitive positive definition in $\mathfrak{A}$.

Lemma 12.6. *Let* $\mathbf{A}$ *be affine. Then the operation* $(x, y, z) \mapsto x - y + z$ *is central.*

Proof. Let f be a basic operation of $\mathbf{A}$ of arity n. Since $\mathbf{A}$ is affine we can write f as $\sum_{i=1} \alpha_i x_i + c$. Then

$$\begin{aligned} f(\bar{x}) - f(\bar{y}) + f(\bar{z})) &= \sum_{i \in \{1,\dots,n\}} \alpha_i x_i + c - \left(\sum_{i=1} \alpha_i y_i + c \right) + \sum_{i=1} \alpha_i z_i + c \\ &= \sum_{i \in \{1,\dots,n\}} \alpha_i (x_i - y_i + z_i) + c \\ &= f(x_1 - y_1 + z_1, \dots, x_n - y_n + z_n). \end{aligned}$$

□

Exercises.

173. Let $\mathbf{A}$ be an affine conservative algebra. Show that $|A| = 2$.

12.2 Structures with an Idempotent Affine Polymorphism Clone

In this section we show that if a finite structure of size at least two has an idempotent affine polymorphism algebra, it can simulate systems of linear equations over a finite field.

Proposition 12.7. *Let $\mathfrak{B}$ be a finite structure with at least two elements and let $\mathbf{A}$ be an affine idempotent algebra such that* $\mathrm{Clo}(\mathbf{A}) = \mathrm{Pol}(\mathfrak{B})$*. Then there exists a prime number p such that $\mathfrak{B}$ pp-constructs $(\mathbb{Z}_p; +, 1)$.*

Proof. Since the operation $m\colon (x, y, z) \mapsto x - y + z$ is central in $\mathbf{A}$ (Lemma 12.6), and since $\mathbf{A}$ is idempotent, the addition operation $+\colon A^2 \to A$ which is given by $m(x, 0, y)$ is central as well, and hence primitively positively definable in $\mathfrak{B}$. Every element $a \in A$ of the abelian group $(A; +, -, 0)$ generates a cyclic group, and some $a \in A$ must have prime order p; choose $a \in A$ such that p is smallest possible. An element of $(A; +, -, 0)$ satisfies the formula

$$\underbrace{x + \cdots + x}_{p \text{ times}} = 0$$

if and only if it is 0 or has order p. The set of all these elements forms a subgroup $\mathbf{A}_p \leq (A; +, -, 0)$. By elementary group theory (see, e.g., Theorem 5 in Chapter 5 of [50]) there is an isomorphism i between $\mathbf{A}_p$ and $(\mathbb{Z}_p)^k$, for some $k \geq 1$. Let $b := i^{-1}(1, 0, \dots, 0)$. It suffices to show that $(A_p; +, b)$ is homomorphically equivalent to $(\mathbb{Z}_p, +, 1)$. The homomorphism from $(A_p; +, b)$ to $(\mathbb{Z}_p, +, 1)$ is given by $\pi_1 \circ i$, and the homomorphism from $(\mathbb{Z}_p, +, 1)$ to $(A_p, +, b)$ is given by $x \mapsto i(x, 0, \dots, 0)$. □

12.3 The Term Condition

We will see that affine algebras satisfy a general 'universal-algebraic' condition, abelianness.

Definition 12.8. An algebra $\mathbf{A}$ is *abelian* if it *satisfies the term condition*, i.e., for every term t of arity $k + 1$, all $a, b \in A$ and tuples $c, d \in A^k$

$$t(a, c) = t(a, d) \Rightarrow t(b, c) = t(b, d).$$

We also say that in Definition 12.8 the term condition is *applied to the first argument of* t; since we can permute arguments of t it is clear what is meant by applying the term condition to other arguments of t.

Example 12.9. Every algebra all of whose operations are unary is abelian. △

Example 12.10. A group $\mathbf{G} = (G; \circ, ^{-1}, e)$ (Example 8.1) is abelian if and only if multiplication is commutative, i.e., $\mathbf{G}$ satisfies $x \circ y \approx y \circ x$. Let us consider the term operation

$$[z_1, z_2] := z_1^{-1} \circ z_2^{-1} \circ z_1 \circ z_2$$

(the *commutator* from group theory) and let $x, y \in G$. Then $[e, y] = e = [e, e]$. The term condition implies that we can exchange e in the first argument of the term by x, and obtain $[x, y] = [x, e] = e$. Thus, $[x, y] = x^{-1}y^{-1}xy = e$ which implies that $xy = yx$. The converse direction follows from Lemma 12.11. △

Lemma 12.11. *Every affine algebra* $\mathbf{A}$ *is abelian.*

Proof. To verify the term condition of $\mathbf{A}$, let t be a term operation. By assumption, t can be written as $t(x, y_1, \dots, y_n) = \alpha_0 x + \sum_{i \in \{1,\dots,n\}} \alpha_i y_i + c$. Now, if $a, b \in A$ and $u, v \in A^n$ then

$$\begin{aligned} t(a, u) = t(a, v) &\Leftrightarrow \sum_{i \in \{1,\dots,n\}} \alpha_i u_i = \sum_{i \in \{1,\dots,n\}} \alpha_i v_i \\ &\Leftrightarrow t(b, u) = t(b, v). \end{aligned}$$

□

The following result was found by H. P. Gumm and, independently, J. D. H. Smith.

Theorem 12.12 (Fundamental theorem of abelian algebras; see [82])**.** *Let* $\mathbf{A}$ *be an algebra with a Maltsev term* m*. Then the following are equivalent.*

(1) $\mathbf{A}$ *is abelian.*

(2) $\mathbf{A}$ *is affine.*

(3) *There exists an abelian group* $(A; +, -, 0)$ *such that the operation* $(x, y, z) \mapsto x - y + z$ *is central in* $\mathbf{A}$ *and in* $\mathrm{Clo}(\mathbf{A})$*.*

(4) m *is central.*

Proof. We prove these implications in cyclic order.

For the implication from (1) to (2), we need to construct a module $\mathbf{M}$ that is polynomially equivalent to $\mathbf{A}$. Arbitrarily fix $0 \in A$. Define $x +^{\mathbf{M}} y := m(x, 0, y)$ and $-^{\mathbf{M}} x := m(0, x, 0)$. The ring $\mathbf{R}$ has the domain

$$R = \{r \in A^A \mid r \text{ unary polynomial operation such that } r(0) = 0\}$$

and the operations:

- $r_1 \cdot^{\mathbf{R}} r_2$ is defined as $x \mapsto r_1(r_2(x))$;
- $r_1 +^{\mathbf{R}} r_2$ is defined as $x \mapsto m(r_1(x), 0, r_2(x))$;
- $0^{\mathbf{R}}$ is the unary polynomial operation which is constant 0;
- $1^{\mathbf{R}}$ is the unary polynomial operation $x \mapsto x$.

For $r \in R$ we define $f_r^{\mathbf{M}}(a) := r(a)$; in the following, we just write ra instead. The algebra $\mathbf{M}$ thus defined is indeed a module:

- For every $x \in A = M$ we have $x + 0 = m(x, 0, 0) = x = m(0, 0, x) = 0 + x$, so 0 is the neutral element in $\mathbf{M}$.
- For associativity, consider the term $t(x_1, x_2, x_3, x_4)$ given by $((x_1 + x_2) + (x_3 + x_4))$. Note that $t(0, 0, b, c) = t(0, b, 0, c)$ for all $b, c \in A$. Applying the term condition to the first argument of t, we obtain $t(a, 0, b, c) = t(a, b, 0, c)$ for any $a \in A$. Hence, $a + (b + c) = (a + b) + c$.
- For any $a, b \in A$ we have $m(a, a, b) = m(b, a, a)$; the term condition applied to the middle argument yields $m(a, 0, b) = m(b, 0, a)$, showing that $a + b = b + a$.
- To see that $-a = m(0, a, 0)$ is the inverse of a, consider the polynomial $t(x, u_1, u_2) = u_1 + m(x, u_2, 0)$. For every $a \in A$ we then have $t(a, a, a) = t(a, 0, 0)$. Applying the term condition to the first argument, we get $t(0, a, a) = t(0, 0, 0)$; showing that $a + (-a) = 0$.
- To show (17), let $r \in R$ and consider the term $t(x, y) := r(x + y) - r(x) - r(y)$. Let $a, b \in A$. Note that $t(0, b) = 0 = t(0, 0)$, and applying the term condition to the first argument yields $t(a, b) = t(a, 0) = 0$, which proves that scalar multiplication by r distributes over addition.
- Let $r, s \in R$ and $a \in A$. Note that $(r + s)(a) = m(ra, 0, sa) = ra + sa$, showing (18). Moreover, $r(s(a)) = rs(b)$ by definition, showing (19).

To show that $\mathbf{A}$ and $\mathbf{M}$ are polynomially equivalent, first observe that every operation of $\mathbf{M}$ has been defined by a polynomial over $\mathbf{A}$. Conversely, let p be an operation of arity n of $\mathbf{A}$. We prove by induction on n that p is a polynomial operation of $\mathbf{M}$. If $n = 1$ then consider the unary polynomial operation $r(x) := p(x) - p(0)$. We have $r \in R$, and thus see that $p(x) = rx + p(0)$ is indeed a polynomial operation of $\mathbf{M}$. If $n > 1$, let $t(x_1, \dots, x_n)$ be the polynomial

$$p(x_1, x_2, \dots, x_n) - p(0, x_2, \dots, x_n) - p(x_1, 0, \dots, 0) + p(0, 0, \dots, 0).$$

We have $t(0, a_2, \dots, a_n) = 0 = t(0, 0, \dots, 0)$ for all $a_2, \dots, a_n \in A$, and by the term condition we get that $t(a_1, a_2, \dots, a_n) = t(a_1, 0, \dots, 0) = 0$. So

$$p(x_1, x_2, \dots, x_n) = p(0, x_2, \dots, x_n) + p(x_1, 0, \dots, 0) - p(0, 0, \dots, 0);$$

the three polynomials on the right have less variables and by the induction hypothesis can be written as polynomials over $\mathbf{M}$, which shows that p can be written as a polynomial over $\mathbf{M}$ as well.

(2) implies (3). We assume that there exists an abelian group $(A; +, -, 0)$ such that in particular the operation $m\colon (x, y, z) \mapsto x - y + z$ is a polynomial operation in $\mathbf{A}$. We have to show that m is not only a polynomial operation, but even a term operation of $\mathbf{A}$. Let t be a term such that $t^{\mathbf{A}}(x, y, z, a_1, \dots, a_n) = x - y + z$ for some constants $a_1, \dots, a_n \in A$. Since $\mathbf{A}$ is affine, $t^{\mathbf{A}}$ can be written in the form

$$(x, y, z, a_1, \dots, a_n) \mapsto x - y + z + \sum_{i=1}^{n} \lambda_i a_i + \lambda_0$$

for $\lambda_0, \lambda_1, \ldots, \lambda_n \in A$. Now consider the term

$$s(x, y, z) := t(x, t(y, y, y, y, \ldots, y), z, y, \ldots, y).$$

Note that

$$\begin{aligned} s^{\mathbf{A}}(x, y, z) &= x - (y + \sum \lambda_i y + \lambda_0) + z + \sum \lambda_i y + \lambda_0 \\ &= x - y + z. \end{aligned}$$

so indeed $m \in \mathrm{Clo}(\mathbf{A})$. Lemma 12.6 states that m is central.

(3) implies (4). We first prove that for every $f \in \mathrm{Clo}(\mathbf{A})$ there exist $a \in A$ and endomorphisms $e_1, \ldots, e_n$ of $(A; +, -, 0)$ such that for all $x_1, \ldots, x_n \in A$

$$f(x_1, \ldots, x_n) = \sum_{i=1}^{n} e_i(x_i) + a.$$

Define $a := f(0, \ldots, 0)$ and $e_i(x) := f(0, \ldots, 0, x, 0, \ldots, 0) - a$ for every $i \in \{1, \ldots, n\}$ and all $x \in A$. Then e_i is indeed an endomorphism of $(A; +, -, 0)$, because by the assumption that $(x, y, z) \mapsto x - y + z$ is central we have that

$$\begin{aligned} e_i(x - y) &= f(0, \ldots, 0, x - y, 0, \ldots, 0) - a \\ &= f(0, \ldots, 0, x, 0, \ldots, 0) - f(0, \ldots, 0, y, 0, \ldots, y) + f(0, \ldots, 0) - a \\ &= e_i(x) - e_i(y). \end{aligned}$$

Moreover,

$$f(x_1, \ldots, x_n) = f(x_1, 0, \ldots, 0) - f(0, \ldots, 0) + f(0, x_2, \ldots, x_n) = e_1(x_1) + f(0, x_2, \ldots, x_n)$$

and by induction it follows that $f(x_1, \ldots, x_n) = \sum e_i(x_i) + a$. In particular, $m(x, y, z) = e_1(x) + e_2(y) + e_3(z) + a$, for some endomorphisms e_1, e_2, e_3 of $(A; +, -, 0)$. We now proceed as in the proof of Lemma 12.1: we have $m(0, 0, 0) = 0$, and hence $a = 0$. Moreover, for all $x \in A$ we have $x = m(x, 0, 0) = e_1(x)$, and thus e_1 is the identity endomorphism. Analogously we have that e_3 must be the identity endomorphism. Finally, $m(x, x, 0) = x + e_2(x) = 0$, and therefore $e_2(x) = -x$. Hence, $m(x, y, z) = x - y + z$. So the assumptions imply that m is central.

(4) implies (1). Suppose that m is central. We verify that $\mathbf{A}$ satisfies the term condition. Let t be a term operation of $\mathbf{A}$ and let $x, y \in A$ and $u, v \in A^n$ be such that $t(x, u) = t(x, v)$. We have to show that $t(y, u) = t(y, v)$. And indeed,

$$\begin{aligned} t(y, u) &= m(t(y, u), t(x, u), t(x, v)) && \text{(since } m \text{ is Maltsev)} \\ &= t(m(y, x, x), m(u_1, u_1, v_1), \ldots, m(u_n, u_n, v_n)) && \text{(centrality)} \\ &= t(y, v). \end{aligned}$$

□

Exercises.

174. Show that subalgebras of abelian algebras are abelian.

175. Show that a semilattice $(L; \wedge)$ (Example 8.4) is abelian if and only if $|L| = 1$.

176. Show that subalgebras of affine algebras are affine.

177. A ring **R** (Example 8.2) is abelian in the sense of Definition 12.8 if and only if for all $x, y \in R$ we have $x \cdot y = 0$.

178. Show that in the definition of the term condition (Definition 12.8), we could have equivalently phrased the condition for polynomial operations instead of term operations t. However, show that it is not sufficient to require the condition only for the operations of the algebra.

179. Show that $(\mathbb{Q}; (x,y) \mapsto \frac{x+y}{2})$ is idempotent abelian, but has no Maltsev polynomial.

180. Let $(A, +, -, 0)$ be a group. Show that $(x - y + z) \mapsto x - y + z$ is central in $\mathbf{A}$ if and only if $\{(x, y, u, v) \in A^4 \mid x + y = u + v\}$ is a subalgebra of $\mathbf{A}^4$.

181. Let $(A, +, -, 0)$ be a group. Show that $(x, y, z) \mapsto x - y + z$ is central in $\mathbf{A}$ if and only if for every $f \in \mathrm{Clo}(\mathbf{A})$ there exist $a \in A$ and endomorphisms $e_1, \dots, e_n$ of $(A; +, -, 0)$ such that for all $x_1, \dots, x_n \in A$

$$f(x_1, \dots, x_n) = \sum_{i=1}^{n} e_i(x_i) + a.$$

182. Let $(A, +, -, 0)$ be a group. Show that $(x, y, z) \mapsto x - y + z$ is a polymorphism of the relation $\{(a, b, c) \in A^3 \mid a + b = c\}$ if and only if the group is abelian.

12.4 The Congruence Condition

We close with a relational characterisation of abelianess (which for some authors is the official definition of abelianness). Recall that Δ_A denotes $\{(a, a) \mid a \in A\}$, and that $\mathrm{Cg}_{\mathbf{A}}(X)$ denotes the smallest congruence of $\mathbf{A}$ that contains X (see Lemma 11.7).

Theorem 12.13. *Let $\mathbf{A}$ be an algebra and $\Delta := \Delta_A$. Then the following are equivalent.*

1. *$\mathbf{A}$ is abelian.*
2. *Δ is a congruence class of $\mathrm{Cg}_{\mathbf{A}^2}(\Delta^2)$.*
3. *Δ is a congruence class of a congruence of $\mathbf{A}^2$.*

Proof. The implication (2) $\Rightarrow$ (3) is trivial. For the implication (3) $\Rightarrow$ (2), suppose that C is a congruence of $\mathbf{A}^2$ where Δ is a congruence class. Since $C' := \mathrm{Cg}_{\mathbf{A}^2}(\Delta^2)$ contains Δ^2 we have that Δ is contained in a congruence class of C'. But since $C' \subseteq C$, this congruence class must be Δ.

For the equivalence between (1) and (2), recall from Lemma 11.7 that $\mathrm{Cg}_{\mathbf{A}^2}(\Delta^2)$ equals that symmetric transitive closure of

$$\{(p(u), p(v)) \mid u, v \in \Delta, p \text{ a unary polynomial operation of } \mathbf{A}^2\}.$$

Note that every unary polynomial operation $p(x)$ of $\mathbf{A}^2$, can be written as

$$f(x, \begin{pmatrix} c_1 \\ d_1 \end{pmatrix}, \dots, \begin{pmatrix} c_n \\ d_n \end{pmatrix})$$

for some $c_1, d_1, \dots, c_n, d_n \in A, n \in \mathbb{N}$ and $f \in \mathrm{Clo}(\mathbf{A})$. Hence, Δ is a congruence class of $\mathrm{Cg}_{\mathbf{A}^2}(\Delta^2)$ if and only if for all $a, b \in A, c_1, d_1, \dots, c_n, d_n \in A, n \in \mathbb{N}$ and every term t

$$t^{\mathbf{A}^2}(\begin{pmatrix} a \\ a \end{pmatrix}, \begin{pmatrix} c_1 \\ d_1 \end{pmatrix}, \dots, \begin{pmatrix} c_n \\ d_n \end{pmatrix}) \in \Delta \text{ if and only if } t^{\mathbf{A}^2}(\begin{pmatrix} b \\ b \end{pmatrix}, \begin{pmatrix} c_1 \\ d_1 \end{pmatrix}, \dots, \begin{pmatrix} c_n \\ d_n \end{pmatrix}) \in \Delta;$$

this is exactly the term condition for $\mathbf{A}$ applied to the first argument of t. $\square$

Example 12.14. Let $n \geq 1$. Let $\mathbf{A}$ be the algebra $(\mathbb{Z}_n; m)$, where $m \colon \mathbb{Z}_n^3 \to \mathbb{Z}_n$ is given by $(x, y, z) \mapsto x - y + z$. Then $\mathbf{A}^2$ has the congruence θ defined as follows:

$$((x_1, x_2), (y_1, y_2)) \in \theta \ \Leftrightarrow \ (x_1 - x_2 = y_1 - y_2)$$

and clearly $\{(a, a) \mid a \in A\}$ is a congruence class of θ. △

Proposition 12.15. *Let* $\mathbf{A}$ *be an algebra with* $R \leq \mathbf{A}^3$ *such that for every* $a \in A$ *and* $i \in \{1, 2, 3\}$ *the binary relation defined by* $\exists x_i (R(x_1, x_2, x_3) \wedge x_i = a)$ *is the graph of an automorphism of* $\mathbf{A}$*. Then* $\mathbf{A}$ *is abelian.*

Proof. Note that the assumptions imply that R is the graph of a surjective binary operation $f \colon A^2 \to A$. Also note that f is central, i.e., $f \colon \mathbf{A}^2 \to \mathbf{A}$ is a homomorphism, because $R \leq \mathbf{A}^3$. Arbitrarily pick $a \in A$. Then $f^{-1}(a)$ is the graph of an automorphism α of $\mathbf{A}$. Then $(x, y) \mapsto f(x, \alpha(y))$ is central and its kernel C is a congruence of $\mathbf{A}^2$. The congruence class $(a, a)/C$ equals Δ_A and the statement follows from Theorem 12.13. □

13 Absorption

"The notion of absorption is, in a sense, complementary to abelianness"
(Barto and Kozik [13])

Absorption theory is an important topic in universal algebra, developed by Marcin Kozik and Libor Barto, which has powerful applications for the study of homomorphism problems. It can be seen as a tool to show the existence of certain solutions in instances of a CSP. This section covers material that stems from [8, 12, 15].

Definition 13.1 (Absorbing subalgebras)**.** Let $\mathbf{A}$ be an algebra and $f \in \mathrm{Clo}(\mathbf{A})$ of arity n. A subalgebra $\mathbf{B}$ of $\mathbf{A}$ is called an *absorbing subalgebra of* $\mathbf{A}$ *with respect to* f, in symbols $\mathbf{B} \lhd_f \mathbf{A}$, if for all $i \in \{1, \dots, n\}$

$$f(B \times B \times \underbrace{A}_{i} \times B \times \cdots \times B) \subseteq B,$$

i.e., if for all $a_1, \dots, a_n \in A$ we have $f(a_1, \dots, a_n) \in B$ whenever all but at most one out of $a_1, \dots, a_n$ are from B. If such an f exists we say that $\mathbf{B}$ *absorbs* $\mathbf{A}$, and write $\mathbf{B} \lhd \mathbf{A}$. A subalgebra $\mathbf{B}$ of $\mathbf{A}$ is called *n-absorbing* if $\mathbf{B} \lhd_f \mathbf{A}$ for some $f \in \mathrm{Clo}(\mathbf{A})$ of arity n.

Since subalgebras are uniquely determined by their domain, we also use the notation $B \lhd \mathbf{A}$ if B is the domain of an absorbing subalgebra $\mathbf{B}$ of $\mathbf{A}$. Note that if $\mathbf{A}$ is idempotent, then B is 1-absorbing if and only if $A = B$. We say that $\mathbf{A}$ is *absorption-free* if $\mathbf{A}$ has no proper absorbing subuniverse.

Example 13.2. The subuniverse $\{0\}$ of the algebra $(\{0, 1\}; \wedge)$ is absorbing with respect to the operation $\wedge$. More generally, if $\mathbf{A} = (A; \wedge)$ where A is finite and $\wedge$ is a semilattice operation, then $\{\bigwedge A\} \lhd \mathbf{A}$. △

Example 13.3. If $\mathbf{A} = (\{0, 1\}; \text{majority})$ then both $\{0\} \lhd \mathbf{A}$ and $\{1\} \lhd \mathbf{A}$. More generally, in any algebra $\mathbf{A}$ with a near unanimity term t (see Section 6.6), every one-element subalgebra is absorbing with respect to $t^{\mathbf{A}}$. △

Example 13.4. Let p be a prime, and let $\mathbf{A} = (\{0,\dots,p-1\}; m)$ be the algebra where $m(x,y,z) := x - y + z$ and where $+$ and $-$ are the usual addition and subtraction modulo p. Then the only absorbing subuniverses are $\emptyset$ and A, so $\mathbf{A}$ is absorption-free. Indeed, we already know that the only proper subuniverses of $\mathbf{A}$ are of the form $\{a\}$, for $a \in A$ (Exercise 144). However, if t it an $\{m\}$-term, then by induction on the term structure one can show that $t(a,\dots,a,x,a,\dots,a) = x$, which shows that $\{a\}$ is not absorbing with respect to $t^{\mathbf{A}}$, and hence $\{a\} \leq \mathbf{A}$ is not absorbing. △

Following the presentation in [13], we will prove in Proposition 13.7 below a converse to the statement from Example 13.3.

Lemma 13.5. *If* $\mathbf{B}, \mathbf{C} \lhd \mathbf{A}$ *then* $\mathbf{A}$ *has a term operation* f *such that* $\mathbf{B}, \mathbf{C} \lhd_f \mathbf{A}$.

Proof. If $\mathbf{B} \lhd_s \mathbf{A}$ and $\mathbf{C} \lhd_t \mathbf{B}$ for some $s,t \in \mathrm{Clo}(\mathbf{A})$, choose $f := s * t$ (Definition 8.32). □

Remark 13.6. If B is n-absorbing, then it is also $n+1$-absorbing.

Proposition 13.7. *Let* $\mathbf{A}$ *be a finite algebra. If every one-element subset is the domain of an absorbing subalgebra of* $\mathbf{A}$*, then* $\mathbf{A}$ *has a near unanimity term.*

Proof. Since $\mathbf{A}$ is finite we can use Lemma 13.5 to construct a single term operation h such that $\mathbf{B} \lhd_h \mathbf{A}$ for every one-element subalgebra $\mathbf{B}$. But then h must be a near unanimity operation. □

An operation $f\colon V^n \to V$ is called *(fully) symmetric* if $f(x_1,\dots,x_n) \approx f(x_{\rho(1)},\dots,x_{\rho(n)})$ for every permutation $\rho \in \mathrm{Sym}(\{1,\dots,n\})$. The following example from [72] shows a structure which has fully symmetric polymorphisms of all arities, but not totally symmetric polymorphisms of arity 3.

Example 13.8. Let $\mathbf{A}$ be the algebra on the domain $A = \{-1,0,1\}$ which has for every $k \geq 1$ the k-ary operation s_k defined as follows:

$$s_k(x_1,\dots,x_n) = \begin{cases} 0 & \text{if } \frac{x_1+\cdots+x_n}{n} \in (-\frac{1}{3}, \frac{1}{3}) \\ 1 & \text{if } \frac{x_1+\cdots+x_n}{n} \geq \frac{1}{3} \\ -1 & \text{otherwise.} \end{cases}$$

Note that s_k is fully symmetric and idempotent. Note that for every $k \geq 1$, the operation s_k preserves the following two relations

$$R_+ := \{(a_1,a_2,a_3) \in A \mid x_1 + x_2 + x_3 \geq 1\}$$
$$R_- := \{(a_1,a_2,a_3) \in A \mid x_1 + x_2 + x_3 \leq -1\}$$

Suppose for contradiction that there exists a totally symmetric operation $t\colon A^3 \to A$ which preserves R_+ and R_-. Then

$$t(1,1,-1) = t(1,-1,1) = t(-1,1,1) = t(1,-1,-1) = t(-1,1,-1) = t(-1,-1,1) =: c.$$

If we apply t to the three tuples $(1,1,-1),(1,-1,1),(-1,1,1) \in R_+$ we obtain $(c,c,c) \in R_+$, hence $c = 1$. The same argument applied to R_- instead of R_+ shows that $c = -1$, a contradiction.

Note that $\mathbf{A}$ has the proper subalgebras $\{-1\}, \{0\}, \{1\}, \{-1,0\}, \{0,1\}$. All of them are absorbing, witnessed by s_k for some large enough k. △

13.1 Absorption Transfer

We start with some warm-up exercises concerning absorption.

Lemma 13.9. *If* $\mathbf{C} \lhd \mathbf{B} \lhd \mathbf{A}$ *then* $\mathbf{C} \lhd \mathbf{A}$.

Proof. If $\mathbf{B} \lhd_t \mathbf{A}$ for some $t \in \mathrm{Clo}(\mathbf{A})$ and $\mathbf{C} \lhd_s \mathbf{B}$ for some $s \in \mathrm{Clo}(\mathbf{B})$, then $\mathbf{C} \lhd_{s*t} \mathbf{A}$. □

Corollary 13.10. *If* $\mathbf{B} \lhd \mathbf{A}$ *and* $\mathbf{C} \lhd \mathbf{A}$ *then* $(B \cap C) \lhd \mathbf{A}$.

Proof. Note that $(B \cap C) \lhd \mathbf{C}$ with respect to the same term as $B \lhd \mathbf{A}$. Now the statement follows from Lemma 13.9. □

Lemma 13.11. *Let* $\sim$ *be a congruence of* $\mathbf{A}$ *and suppose that* $B \lhd \mathbf{A}/{\sim}$. *Then* $\bigcup B \lhd \mathbf{A}$.

Proof. If $B \lhd_{t^{\mathbf{A}/\sim}} (\mathbf{A}/{\sim})$ for some term $t(x_1, \dots, x_n)$ and $b_1, \dots, b_n \in A$ are such that all but one are from $\cup B$, then $t^{\mathbf{A}}(b_1, \dots, b_n)/_{\sim} = t^{\mathbf{A}/\sim}(b_1/_{\sim}, \dots, b_n/_{\sim}) \in B$. Hence, $\bigcup B \lhd_{t^{\mathbf{A}}} \mathbf{A}$. □

Example 13.12. Consider the Maltsev operation m on the set $\{0, 1, 2\}$ from Example 7.4. We claim that $(\{0, 1, 2\}; m)$ is absorption-free. By symmetry, if $\{0, 1\}$ would be absorbing, then $\{1, 2\}$ would be as well, and by Lemma 13.11 the subalgebra $\{1\}$ would be absorbing. The argument that $\{1\}$ is not absorbing is as in Example 13.4. Again by symmetry, this shows that all proper subalgebras are not absorbing. △

Lemma 13.13. *If* $\mathbf{R} \leq \mathbf{A} \times \mathbf{B}$ *and* $S \lhd_f \mathbf{R}$, *then* $\pi_1(S) \lhd_f \pi_1(R)$.

Proof. Suppose that $a_1, \dots, a_n \in \pi_1(S)$, $c \in \pi_1(R)$, and $i \in \{1, \dots, n\}$. Then there are $b_1, \dots, b_n \in B$ and $d \in B$ such that $(a_1, b_1), \dots, (a_n, b_n) \in S$ and $(c, d) \in R$. Since $S \lhd_f \mathbf{R}$ we have

$$f((a_1, b_1), \dots, (a_{i-1}, b_{i-1}), (c, d), (a_{i+1}, b_{i+1}), \dots, (a_n, b_n)) \in S,$$

and hence $f(a_1, \dots, a_{i-1}, c, a_{i+1}, \dots, a_n) \in \pi_1(S)$, which proves that $\pi_1(S) \lhd_f \pi_1(R)$. □

Lemma 13.14. *Let* $\mathbf{A}$ *be an idempotent algebra and suppose that* $\mathbf{A}^2$ *has a proper* n*-absorbing subalgebra. Then* $\mathbf{A}$ *has a proper* n*-absorbing subalgebra as well.*

Proof. Suppose that B is a proper non-empty subset of A^2 such that $B \lhd_f \mathbf{A}^2$ for some $f \in \mathrm{Clo}(\mathbf{A})^{(n)}$. Note that $\pi_1(B) \lhd_f \mathbf{A}$ by Lemma 13.13, and that $\pi_1(B)$ is non-empty. Hence, if $\pi_1(B) \neq A$ then we are done, so suppose that $\pi_1(B) = A$. Since $B \neq A^2$, there exists an $a \in A$ such that $B' := \pi_2\big(B \cap (\{a\} \times A)\big) \neq A$. Since $\pi_1(B) = A$ we have that $B' \neq \emptyset$, so it suffices to show that $B' \lhd_f \mathbf{A}$. Let $b_1, \dots, b_n \in B'$, $i \leq n$, and $c \in A$. We have to show that $d := f(b_1, \dots, b_{i-1}, c, b_{i+1}, \dots, b_n) \in B'$. By the definition of B' we have $(a, b_1), \dots, (a, b_n) \in B$. Then

$$\begin{aligned} &f((a, b_1), \dots, (a, b_{i-1}), (a, c), (a, b_{i+1}), \dots, (a, b_n)) \\ &= (f(a, \dots, a), f(b_1, \dots, b_{i-1}, c, b_{i+1}, \dots, b_n)) = (a, d) \in B \end{aligned}$$

since $B \lhd_f \mathbf{A}$, hence $d \in B'$. □

Exercises.

183. Let $\mathbf{A}$ and $\mathbf{B}$ be two algebras of the same signature, and let R be the domain of a subalgebra of $\mathbf{A} \times \mathbf{B}$. For $X \subseteq A$ we define

$$X + R := \{b \in B \mid \exists a \in X \colon R(a, b)\}.$$

Prove that

- if $X \leq \mathbf{A}$ then $(X + R) \leq \mathbf{B}$;
- if $R \leq \mathbf{A} \times \mathbf{B}$ is subdirect and $X \lhd_f \mathbf{A}$, then $(X + R) \lhd_f \mathbf{B}$.

184. Let $\mathfrak{A}$ be a finite relational τ-structure and ϕ a primitive positive τ-formula. Let $\mathfrak{A}'$ be a τ-structure on the same domain such that for each $R \in \tau$ we have $R^{\mathfrak{A}'} \lhd R^{\mathfrak{A}}$. If ϕ defines S in $\mathfrak{A}$ and defines S' in $\mathfrak{A}'$, then $S' \lhd S$.

13.2 Essential Relations

This section presents a relational characterisation of absorption that will be needed in Section 13.6 and in Section 14.2. The following material is mostly from Barto and Kazda [10].

Definition 13.15. Let $B \leq \mathbf{A}$ and $n \geq 1$. Then $R \leq \mathbf{A}^n$ is *B-essential* if for every $i \in \{1, \dots, n\}$

$$R \cap (B \times \cdots \times B \times \underbrace{A}_{i} \times B \times \cdots \times B) \neq \emptyset$$

$$\text{and} \qquad R \cap B^n = \emptyset.$$

Note that if $B \leq \mathbf{A}$ is a proper subuniverse and $\mathbf{A}$ is idempotent, then $\{a\}$ is B-essential for every $a \in A \setminus B$.

Lemma 13.16. *Let $B \leq \mathbf{A}$. If there is no B-essential relation of arity m, then for every $n \geq m$ there is no B-essential relation of arity n.*

Proof. If $R \leq \mathbf{A}^n$ is B-essential, then $\pi_{1,\dots,n-1}(R \cap (A^{n-1} \times B))$ is B-essential. $\square$

Lemma 13.17. *Let $\mathbf{A}$ be an algebra with a term operation t of arity m such that $B \lhd_t \mathbf{A}$. Then there are no B-essential relations $R \leq \mathbf{A}^m$.*

Proof. Suppose for contradiction that $R \leq \mathbf{A}^m$ is B-essential. Then there are $a^1, \dots, a^m \in A^m$ such that $\{a^i_1, \dots, a^i_m\} \setminus \{a^i_i\} \subseteq B$ for every $i \in \{1, \dots, m\}$. Therefore, $t(a^1, \dots, a^m) \in R \cap B^m$, because $B \lhd_t \mathbf{A}$, contrary to our assumptions. $\square$

Proposition 13.18. *Let $B \leq \mathbf{A}$ and $R \leq \mathbf{A}^n$ for $n \geq m - 1$. Suppose that $\mathbf{A}$ has no B-essential relation of arity m and for every $I \in \binom{\{1,\dots,n\}}{m-1}$ we have $\pi_I(R) \cap B^{m-1} \neq \emptyset$. Then*

$$R \cap B^n \neq \emptyset.$$

Proof. The proof is by induction on $n \geq m-1$. The base case $n = m-1$ is immediate by the assumption applied for $I = \{1, \dots, n\}$. For the inductive step, suppose that $n \geq m$. For every $i \in \{1, \dots, n\}$ define

$$R_i := \pi_{[n]\setminus\{i\}}(R) \leq \mathbf{A}^{n-1}$$

and note that $\pi_I(R_i) \cap B^{m-1} \neq \emptyset$ for every $I \subseteq \binom{\{1,\dots,n\}\setminus\{i\}}{m-1}$. Hence, by the inductive assumption we have that

$$R_i \cap B^{n-1} \neq \emptyset.$$

Since R is not essential by Lemma 13.16 we therefore must have $R \cap B^n \neq \emptyset$. □

Corollary 13.19. *Let $t \in \mathrm{Clo}(\mathbf{A})^{(m)}$ be such that $B \lhd_t \mathbf{A}$ and let $R \leq \mathbf{A}^n$ for $n \geq m-1$. Suppose that for every $I \in \binom{\{1,\dots,n\}}{m-1}$ we have $\pi_I(R) \cap B^{m-1} \neq \emptyset$. Then*

$$R \cap B^n \neq \emptyset.$$

Proof. Combine Lemma 13.18 with Lemma 13.17. □

Lemma 13.17 has a converse (Proposition 16 in [10]); also see [94, 95].

Theorem 13.20. *Let $m \geq 1$. A subalgebra $\mathbf{B} \leq \mathbf{A}$ m-absorbs $\mathbf{A}$ if and only if there are no B-essential relations $R \leq \mathbf{A}^m$.*

Proof. The forward implication is Lemma 13.17. For the converse, suppose that $\mathbf{A}$ has no B-essential relations of arity m. Let $\mathbf{F} \leq \mathbf{A}^{A^m}$ be the free algebra generated by $x_1, \dots, x_m$ in $\mathrm{HSP}(\mathbf{A})$ (see Section 8.5). For $i \in \{1, \dots, m\}$, let $X_i := B^{i-1} \times (A \setminus B) \times B^{m-i}$, and let $X := X_1 \cup \dots \cup X_m$, which will be used as index set of the relation R defined as follows:

$$R := \pi_X(F) \leq \mathbf{A}^{X_1} \times \dots \times \mathbf{A}^{X_m}.$$

Let $I \subseteq \binom{X}{m-1}$. We claim that $\pi_I(R) \cap B^{m-1} \neq \emptyset$: indeed, by the pigeon-hole principle there exists $i \in \{1, \dots, m\}$ such that $I \cap X_i = \emptyset$. Since $\mathbf{F}$ contains π_i^m we have

$$R \cap (B^{X_1} \times \dots \times B^{X_{i-1}} \times A^{X_i} \times B^{X_{i+1}} \times \dots \times B^{X_m}) \neq \emptyset$$

which shows the claim.

Therefore, Proposition 13.18 implies that $R \cap B^n \neq \emptyset$. By definition, any element of $R \cap B^X$ can be extended to an element $t \in \mathfrak{F}$, and any such t is a term operation of $\mathbf{A}$ of arity m which absorbs B. □

13.3 The Absorption Theorem

The presentation of this section is based on the lecture notes of Libor Barto. The goal of this section is to show that finite idempotent Taylor algebras (see Remark 9.20) must have some form of absorption; this idea will be formalised in the absorption theorem, Theorem 13.33 below, which is from [12].

Definition 13.21. Let $\mathbf{A}$ be an algebra. A subset $B \subseteq A$ is called *projective in* $\mathbf{A}$ (in some papers called *cube term blocker* [81]) if for every $f \in \mathrm{Clo}(\mathbf{A})$ of arity n there exists $i \in \{1, \dots, n\}$ such that

$$f(A, A, \dots, A, \underbrace{B}_{\text{position } i}, A, \dots, A) \subseteq B;$$

as usual, the term on the left stands for $\{f(x_1, \dots, x_n) \mid x_1, \dots, x_n \in A, x_i \in B\}$.

Note that subsets of A that are projective in $\mathbf{A}$ are subuniverses of $\mathbf{A}$. Recall the definition of minion homomorphisms from Section 9.1. Our starting point is the following theorem.

Theorem 13.22. *Let $\mathbf{A}$ be an algebra such that there is no minion homomorphism from* $\mathrm{Clo}(\mathbf{A})$ *to* $\mathbf{Proj}$ *and let $B \subseteq A$. Then B 2-absorbs $\mathbf{A}$, or is not projective.*

Proof. Suppose that B is projective, so for every $f \in \mathrm{Clo}(\mathbf{A})$ of arity n there exists $i_f \in [n]$ such that

$$f(A, \ldots, A, \underbrace{B}_{i_f}, A, \ldots, A) \subseteq B. \tag{30}$$

If i_f is *unique* for every $f \in \mathrm{Clo}(\mathbf{A})$, then $\mathrm{Clo}(\mathbf{A}) \to \mathbf{Proj}$ given by $f \mapsto \pi^n_{i_f}$ is a minion homomorphism: indeed, let $\alpha\colon [n] \to [k]$ and $f \in \mathrm{Clo}(\mathbf{A})$ and suppose that there exists a unique $i \in [n]$ such that (30) holds. Then f_α is an operation of arity k such that

$$f_\alpha(A, \ldots, A, \underbrace{B}_{\alpha(i)}, A, \ldots, A) \subseteq f(A, \ldots, A, \underbrace{B}_{i}, A, \ldots, A) \subseteq B$$

and by assumption $\alpha(i)$ is the only index $j \in [k]$ such that $f_\alpha(A, \ldots, A, \underbrace{B}_{j}, A, \ldots, A) \subseteq B$.

Hence,

$$\xi(f_\alpha) = \pi^k_{\alpha(i)} = (\pi^n_i)_\alpha = \xi(f)_\alpha$$

and ξ is a minion homomorphism.

So there exists $f \in \mathrm{Clo}(\mathbf{A})$ and $i \neq j$ such that

$$f(A, \ldots, A, \underbrace{B}_{i}, A, \ldots, A) \subseteq B \qquad \text{and} \qquad f(A, \ldots, A, \underbrace{B}_{j}, A, \ldots, A) \subseteq B$$

Define $r(x, y) := f(x, \ldots, x, \underbrace{y}_{j}, x, \ldots, x)$ and observe that $B \lhd_r \mathbf{A}$. □

If $\mathbf{A}$ does not have proper projective subuniverses, then in exchange it must have a term operation satisfying the following strong condition.

Definition 13.23. An operation $t\colon A^n \to A$ is called *transitive* if for every $a \in A$ and $i \in \{1, \ldots, n\}$ we have

$$t(A, \ldots, A, \underbrace{\{a\}}_{i}, A, \ldots, A) = A.$$

Clearly, if $|A| > 1$, then a transitive operation must have arity at least two.

Theorem 13.24. *Let $\mathbf{A}$ be a finite idempotent algebra without proper projective subuniverses. Then* $\mathrm{Clo}(\mathbf{A})$ *contains a transitive operation.*

Proof. By assumption, for every proper subset B of A there exists $t_B \in \mathrm{Clo}(\mathbf{A})$ of arity n such that for every $i \in \{1, \ldots, n\}$

$$t_B(A, \ldots, A, \underbrace{B}_{i}, A, \ldots, A) \not\subseteq B.$$

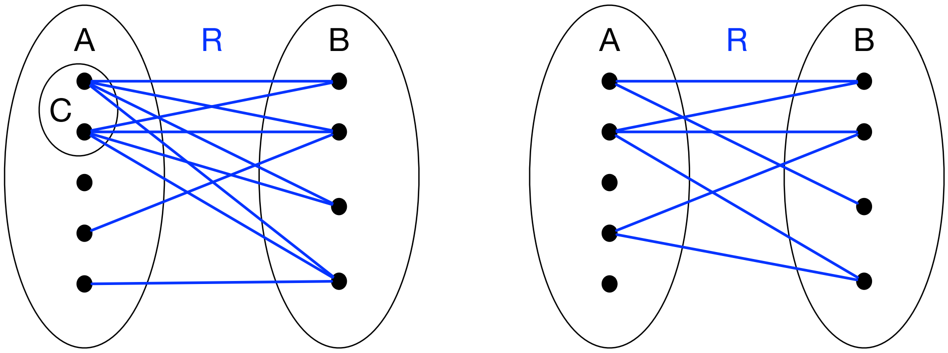

Figure 18: Illustrations of $\mathbf{R} \leq \mathbf{A} \times \mathbf{B}$ with non-empty left center C (left side), and of $\mathbf{R} \leq \mathbf{A} \times \mathbf{B}$ which is linked (right side); none of the two examples is subdirect.

Using the star product and the idempotence of $\mathbf{A}$, Lemma 9.14 implies that we may suppose that there exists a single term t that works for all proper $B \subset A$. Then $u := \underbrace{t * \cdots * t}_{|A| \text{ times}}$ is transitive, because for every $a \in A$ and $j \in \{1, \dots, |A|\}$

$$|\underbrace{t * \cdots * t}_{j \text{ times}}(A, \dots, A, \underbrace{\{a\}}_{i}, A, \dots, A)| \geq j$$

and hence $u(A, \dots, A, \underbrace{\{a\}}_{i}, A, \dots, A) = A$. □

Corollary 13.25. *Let* $\mathbf{A}$ *be a finite idempotent Taylor algebra. Then* $\mathbf{A}$ *has a proper 2-absorbing subuniverse, or* $\mathrm{Clo}(\mathbf{A})$ *contains a transitive operation.*

Proof. Since $\mathbf{A}$ is Taylor, there is no minion homomorphism from $\mathrm{Clo}(\mathbf{A})$ to $\mathbf{Proj}$ (Theorem 9.15; Remark 9.20). If $\mathbf{A}$ has a proper projective subuniverse B, then B 2-absorbs $\mathbf{A}$ by Theorem 13.22 and we are done. Otherwise, all proper subuniverses of $\mathbf{A}$ are not projective, and $\mathrm{Clo}(\mathbf{A})$ contains a transitive operation by Theorem 13.24. □

We will now explore consequences of having a transitive term operation for the existence of proper absorbing subuniverses. Let $\mathbf{A}$ and $\mathbf{B}$ be algebras, and let $R \subseteq A \times B$ be a relation.

Definition 13.26 (left centre[9])**.** The *left centre of* R is the set

$$\{a \in A \mid (a, b) \in R \text{ for every } b \in B\}.$$

See Figure 18, left side.

Proposition 13.27. *Let* $\mathbf{A}$ *and* $\mathbf{B}$ *be idempotent algebras with the same signature and let* $R \leq \mathbf{A} \times \mathbf{B}$ *with left centre* C *be such that for every* $a \in A$ *there exists* $b \in B$ *such that* $(a, b) \in R$*. If there exists a term such that* $t^{\mathbf{B}}$ *is transitive, then* $C \lhd_{t^{\mathbf{A}}} \mathbf{A}$.

[9]There is no connection with the notion of centrality from Definition 12.4.

Proof. If C is empty, then the statement is trivial, so suppose that C is non-empty. Since $\mathbf{B}$ is idempotent, $C \leq \mathbf{A}$. Let $i \in [n]$, $z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_n \in C$, $a \in A$, and

$$a' := t^{\mathbf{A}}(z_1, \dots, z_{i-1}, a, z_{i+1}, \dots, z_n).$$

To show that $C \lhd_{t^{\mathbf{A}}} \mathbf{A}$ we need to show that $a' \in C$, i.e., $(a', b) \in R$ for every $b \in B$. Arbitrarily choose $b \in B$. By assumption, there exists $c \in B$ such that $(a, c) \in R$. By the transitivity of $t^{\mathbf{B}}$ there are $d_1, \dots, d_{i-1}, d_{i+1}, \dots, d_n \in B$ such that $t^{\mathbf{B}}(d_1, \dots, d_{i-1}, c, d_{i+1}, \dots, d_n) = b$. Note that $(z_1, d_1), \dots, (z_{i-1}, d_{i-1}), (z_{i+1}, d_{i+1}), \dots, (z_n, d_n) \in R$ since $z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_n$ are from the left center of R. Since $R \leq \mathbf{A} \times \mathbf{B}$, we have that

$$(a', b) = \big(t^{\mathbf{A}}(z_1, \dots, z_{i-1}, a, z_{i+1}, \dots, z_n), t^{\mathbf{B}}(d_1, \dots, d_{i-1}, c, d_{i+1}, \dots, d_n)\big) \in R$$

and the proof is complete. □

Corollary 13.28. *Let $\mathbf{A}$ and $\mathbf{B}$ be finite idempotent algebras with the same signature such that $\mathbf{B}$ is Taylor. Let $R \leq \mathbf{A} \times \mathbf{B}$ with left centre C be such that for every $a \in A$ there exists $b \in B$ such that $(a, b) \in R$. Then $\mathbf{B}$ has a proper 2-absorbing subuniverse or $C \lhd \mathbf{A}$.*

Proof. Suppose that $\mathbf{B}$ does not have a proper 2-absorbing subuniverse. Then Corollary 13.25 implies that $\mathbf{B}$ has a transitive term operation t. Hence, Proposition 13.27 implies that $C \lhd_t \mathbf{A}$. □

The relation R can be viewed as the edge relation of a bipartite graph G_R with color classes A and B (this perspective was already presented in Section 8.3).

Definition 13.29 (Linked relations). $R \subseteq A \times B$ is *linked* if G_R is connected after removing isolated vertices.

See Figure 18 on the right. Note that if R is a subdirect subalgebra of $\mathbf{A} \times \mathbf{B}$ (Definition 8.16) then G_R has no isolated vertices. Also note that if R has a non-empty left centre, then it is linked (but not necessarily subdirect). Recall the definition of R^{-1} from Exercise 94.

Proposition 13.30. *Let $\mathbf{A}$ and $\mathbf{B}$ be finite idempotent algebras with the same signature and let $R \leq \mathbf{A} \times \mathbf{B}$ be with empty left centre and such that $R^{-1} \circ R = B^2$. Then there exists a subdirect $R' \leq \mathbf{B}^2$ whose left centre is a proper subuniverse of $\mathbf{B}$.*

Before we go into the proof we consider an example.

Example 13.31. Suppose that $\mathbf{B}$ has domain $B = \{1, 2, 3\}$ and $\mathbf{B}^2$ has the subuniverse $R := \{(u, v) \in B^2 \mid u \neq v\}$. Clearly, the left centre of R is empty. Then

$$R' := \{(x, y) \mid \exists a(R(a, x) \wedge R(a, y) \wedge R(a, 1))\} = \{(1, 1), (1, 2), (1, 3), (2, 1), (3, 1)\}$$

is a subuniverse of $\mathbf{B}^2$ (we use that $\{1\}$ is a subuniverse), is subdirect, and has the left centre $\{1\}$. △

Proof of Proposition 13.30. For $D = \{d_1, \dots, d_k\} \subseteq B$ define

$$S_D := \big\{(x, y) \in B^2 \mid \exists a\big(R(a, x) \wedge R(a, y) \wedge R(a, d_1) \wedge \dots \wedge R(a, d_k)\big)\big\}.$$

Then

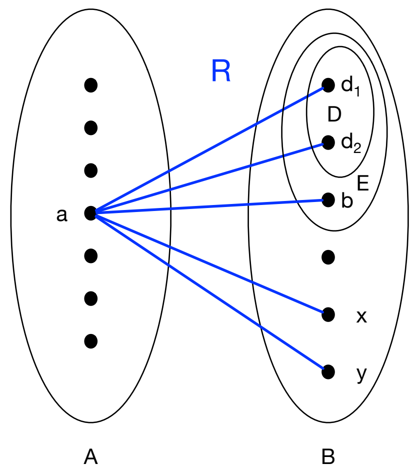

Figure 19: An illustration for the definition of S_E in the proof Proposition 13.30.

- $S_D \leq \mathbf{B}^2$ by the idempotence of $\mathbf{B}$,
- $S_\emptyset = R^{-1} \circ R = B^2$ by assumption, and
- $S_B = \emptyset$ because the left centre of R is empty.

Let D be maximal such that $S_D = B^2$, and let $E \subseteq B$ and $b \in B$ be a set such that $E \setminus D = \{b\}$. See Figure 19. Let C be the left centre of S_E.

- C contains b and hence is non-empty: indeed, for any $y \in B$ there exists $a \in A$ witnessing that $(b, y) \in S_D = B^2$, i.e., $R(a, b)$, $R(a, y)$, and $R(a, d)$ for every $d \in D$. Hence, $(b, y) \in S_E$ and $b \in C$.
- $S_E \leq \mathbf{B}^2$ is subdirect: since C is non-empty, for every $y \in B$ there is $c \in C$ such that $(y, c) \in S_E$ and $(c, y) \in S_E$.
- C is a proper subset of B. Otherwise, the centrality of C would imply that $S_E = B^2$, contrary to the choice of D and E.

Therefore, $R' := S_E$ meets the requirements. $\square$

Proposition 13.32. *Let $\mathbf{A}, \mathbf{B}$ be finite idempotent algebras with the same signature and let $R \leq \mathbf{A} \times \mathbf{B}$ be subdirect and linked such that $R \neq A \times B$. Then at least one of the following cases applies.*

- *R has a non-empty left centre.*
- *there exists a subdirect $R' \leq \mathbf{B}^2$ whose left centre is a proper subuniverse of $\mathbf{B}$.*

Proof. Suppose that the left centre C of R is empty. If $R^{-1} \circ R = B^2$ then the statement follows from Proposition 13.30. Otherwise, $R^{-1} \circ R \leq \mathbf{B}^2$ is subdirect, proper, and linked (Exercise 186), so we may replace $\mathbf{A}$ by $\mathbf{B}$ and R by $R^{-1} \circ R$. Since R is linked and subdirect, we have that $(R^{-1} \circ R)^n = B^2$ for some $n \in \mathbb{N}$. Hence, if we repeat the argument, we eventually find a proper, subdirect, and linked subuniverse R of $\mathbf{B}^2$ such that $R^{-1} \circ R = B^2$. If the left centre of R is non-empty it is a proper subalgebra of $\mathbf{B}$ and we are done. Otherwise, the statement again follows from Proposition 13.30. □

Theorem 13.33 (Absorption theorem [12]). *Let $\mathbf{A}, \mathbf{B}$ be finite idempotent algebras such that $\mathbf{B}$ is Taylor. Then for every linked and subdirect $R \leq \mathbf{A} \times \mathbf{B}$ one of the following is true:*

1. *$R = A \times B$;*
2. *$\mathbf{A}$ has a proper absorbing subuniverse.*
3. *$\mathbf{B}$ has a proper absorbing subuniverse.*

Proof. Suppose that $R \neq A \times B$ because otherwise item 1 of the theorem holds and we are done. Let C be the left centre of R. Note that $C \neq A$ because $R \neq A \times B$.

Suppose also that $\mathbf{B}$ is absorption-free, because otherwise item 3 of the theorem holds. Corollary 13.28 then implies that $C \lhd \mathbf{A}$. If C is non-empty, then we have found a proper absorbing subuniverse of $\mathbf{A}$ and item 2 of the theorem holds. Otherwise, Proposition 13.32 implies that there exists a subdirect $R' \leq \mathbf{B}^2$ whose left centre is a proper subuniverse of $\mathbf{B}$. In this case, $\mathbf{B}$ has a proper absorbing subuniverse by Corollary 13.28, in contradiction to the assumption above. □

Exercises.

185. Let $\mathbf{A}$ and $\mathbf{B}$ be idempotent algebras and $R \leq \mathbf{A} \times \mathbf{B}$. Show that the left center of R is a subalgebra of $\mathbf{A}$.

186. Show that if $R \leq \mathbf{A} \times \mathbf{B}$ is linked, then $R^{-1} \circ R \leq \mathbf{B}^2$ is linked, too.

187. Suppose that $\mathbf{A}$ is a simple algebra. Then every subdirect $R \leq \mathbf{A}^2$ is linked or the graph of an automorphism of $\mathbf{A}$.
Hint. Consider $\bigcup_{i \in \mathbb{N}} (R \circ R^{-1})^i$.

188. Let $\mathbf{A}$ be a finite algebra and $B \subseteq A$. Then B is a projective in $\mathbf{A}$ if and only if $\mathrm{Clo}(\mathbf{A})$ preserves for every n the relation $B[n] := A^n \setminus (A \setminus B)^n$.

189. Let $\mathbf{A}, \mathbf{B}$ be a algebras and let $\mathbf{R} \leq \mathbf{A} \times \mathbf{B}$ be subdirect. Let θ_A be the kernel of $\pi_1 \colon R \to A$ and let θ_B be the kernel of $\pi_2 \colon R \to B$. Show that R is linked if and only if $\theta_A \vee \theta_B = 1_R$.

13.4 Abelianness Revisited

The fundamental theorem of abelian algebras (Theorem 12.12) implies that every abelian algebra with a Maltsev term is affine. In this section we considerably strengthen this theorem for finite idempotent algebras by replacing the assumption of having a Maltsev term by having a Taylor term (Corollary 13.39). This result follows from tame congruence theory (Hobby and McKenzie [63]; see the discussion in [15]); the new proof based on absorption that we present here is from [15]; the presentation follows lecture notes of Libor Barto.

Definition 13.34. An algebra $\mathbf{A}$ is called *hereditarily absorption-free (HAF)* if no subalgebra of $\mathbf{A}$ has a proper absorbing subalgebra, i.e., whenever $\mathbf{C}$ is a non-empty absorbing subalgebra of an subalgebra $\mathbf{B}$ of $\mathbf{A}$, then $C = B$.

We will first prove that 'HAF and Taylor implies Maltsev' (Theorem 13.37), and then that 'abelian implies HAF' (Theorem 13.38). First we prove that HAF is closed under taking direct products.

Lemma 13.35. *Let* $\mathbf{A}, \mathbf{B}$ *be idempotent hereditarily absorption-free algebras with the same signature. Then* $\mathbf{A} \times \mathbf{B}$ *is hereditarily absorption-free.*

Proof. Suppose that $S \lhd \mathbf{R} \leq \mathbf{A} \times \mathbf{B}$ is non-empty. We have to show that $S = R$. Let $(a, b) \in R$ and consider the subalgebras $\mathbf{D} \leq \mathbf{C} \leq \mathbf{B}$ given by

$$\begin{aligned} C &:= \{b' \mid (a, b') \in R\} \leq \mathbf{B} && \text{(since } \mathbf{B} \text{ is idempotent)} \\ \text{and } D &:= \{b' \mid (a, b') \in S\} \leq \mathbf{C} && \text{(since } \mathbf{B} \text{ is idempotent).} \end{aligned}$$

Claim 1. $D \neq \emptyset$. Note that $\pi_1(S) \leq \pi_1(R) \leq \mathbf{A}$ (for the notation, see the comments after Definition 6.30). We even have $\pi_1(S) \lhd \pi_1(R)$ (Lemma 13.13). Since $\mathbf{A}$ is HAF, we get $\pi_1(S) = \pi_1(R)$. Since $a \in \pi_1(R)$, there must be $b' \in B$ such that $(a, b') \in S$. Therefore, $b' \in D \neq \emptyset$.

Claim 2. $D \lhd \mathbf{C}$. By assumption, there exists a term operation $f \in \mathrm{Clo}(\mathbf{R})^{(n)}$ such that $S \lhd_f \mathbf{R}$. Let $b_1, \ldots, b_n \in C$ be such that all but one of them are from D. Then $f((a, b_1), \ldots, (a, b_n)) = (a, f(b_1, \ldots, b_n)) \in S$ since $\mathbf{A}$ is idempotent and $S \lhd_f \mathbf{R}$. It follows that $f(b_1, \ldots, b_n) \in D$.

By the assumption that $\mathbf{B}$ is HAF, we must have $D = C$ and hence $(a, b) \in S$. Since $(a, b) \in R$ was chosen arbitrarily, this implies that $R = S$. □

Corollary 13.36. *The class of idempotent HAF algebras of fixed signature* τ *forms a pseudo-variety.*

Proof. By definition of HAF, the class is closed under subalgebras. Closure under finite products has been established in Lemma 13.35. Closure under homomorphic images is by Lemma 13.11. □

Theorem 13.37 (Theorem 1.4 in [15])**.** *Let* $\mathbf{A}$ *be a finite idempotent Taylor algebra. If* $\mathbf{A}$ *is hereditarily absorption-free, then* $\mathbf{A}$ *has a Maltsev term.*

Proof. Let $\mathbf{F} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})$ be the free algebra over two generators x, y (see Section 8.5). It follows from Corollary 13.36 that $\mathbf{F}$ is HAF. Let $\mathbf{R}$ be the subalgebra of $\mathbf{F}^2$ generated by (x, y), (x, x), and (y, x).

Claim 1. $R \leq \mathbf{F}^2$ is subdirect. Every element of F can be written as $t(x, y)$ for some term t, and since $(x, y) \in R$ and $(y, x) \in R$ we have that $(t^{\mathbf{F}}(x, y), t^{\mathbf{F}}(y, x)) \in R$. A similar statement holds for the second argument of R. This shows that $\mathbf{R}$ is a subdirect subalgebra of $\mathbf{F}^2$.

Claim 2. R is linked. Every element of R can be written as

$$s^{\mathbf{R}}((x, y), (x, x), (y, x)) = (s^{\mathbf{F}}(x, x, y), s^{\mathbf{F}}(y, x, x))$$

for some term s. Since $(x, x), (y, x) \in R$ we have that $(s^{\mathbf{F}}(x, x, y), s^{\mathbf{F}}(x, x, x)) \in R$ and since $(x, y), (x, x) \in R$ we have $(s^{\mathbf{F}}(x, x, x), s^{\mathbf{F}}(y, x, x)) \in R$. Note that $s^{\mathbf{F}}(x, x, x) = x$ by the idempotence of $\mathbf{A}$ and $\mathbf{F}$, and thus between any two elements of F there is a path of length at most three in the bipartite graph G_R of R, which proves the claim.

Since $\mathbf{F} \in \mathrm{SP}(\mathbf{A})$ (Proposition 8.23) and $\mathrm{Clo}(\mathbf{A})$ has no minion homomorphism to **Proj**, neither has $\mathbf{F}$ (Proposition 8.39). Since $\mathbf{F}$ has no proper non-empty absorbing subalgebra, Theorem 13.33 implies that $R = F \times F$. Let m be a term such that $m^{\mathbf{F}}((x,y),(x,x),(y,x)) = (y,y)$. Then $m^{\mathbf{A}}$ is a Maltsev operation. □

The following is a special case of Lemma 4.1 in [15].

Theorem 13.38. *Let $\mathbf{A}$ be a finite idempotent algebra. If $\mathbf{A}$ is abelian then $\mathbf{A}$ is hereditarily absorption-free.*

Proof. Since every subalgebra of an abelian algebra is abelian (Exercise 174), it suffices to show that if $\mathbf{B} \lhd \mathbf{A}$, then $B = A$. We will show that if $\mathbf{B} \lhd_t \mathbf{A}$ for some n-ary term operation $t \in \mathrm{Clo}(\mathbf{A})$, for $n \geq 2$, then $\mathbf{B} \lhd_s \mathbf{A}$ for some $n-1$-ary s. This is enough, because if $\mathbf{B} \lhd_s \mathbf{A}$ and s is unary, then s must be the identity by the idempotence of $\mathbf{A}$, hence $B = A$. Define the term $t_m(\bar{x}, y)$, where $\bar{x} = (x_1, \dots, x_{n-1})$, as follows

$$t_m(\bar{x}, y) := \underbrace{t(\bar{x}, t(\bar{x}, \dots, t(\bar{x}}_{m \text{ times}}, y))).$$

Note that $\mathbf{B} \lhd_{t_m} \mathbf{A}$ for every $m \geq 1$.

Claim 1. For $m = |A|!$ we have

$$t_m(\bar{x}, t_m(\bar{x}, y)) = t_m(\bar{x}, y). \tag{31}$$

To see this, define $r_{\bar{x}} \colon A \to A$ by $r_{\bar{x}}(y) := t(\bar{x}, y)$. Then note that

$$t_m(\bar{x}, y) = \underbrace{r_{\bar{x}} \circ \cdots \circ r_{\bar{x}}}_{m \text{ times}}(y)$$

and observe that (see Exercise 115)

$$\underbrace{r_{\bar{x}} \circ \cdots \circ r_{\bar{x}}}_{2m \text{ times}}(y) = \underbrace{r_{\bar{x}} \circ \cdots \circ r_{\bar{x}}}_{m \text{ times}}(y).$$

Claim 2. $\mathbf{B} \lhd_s \mathbf{A}$ for $s \colon A^{n-1} \to A$ defined by

$$s(x_1, \dots, x_{n-1}) := t_m(x_1, \dots, x_{n-1}, x_{n-1}).$$

Let $a \in A$ and $b_1, \dots, b_{n-2} \in B$. Clearly, $s(b_1, \dots, b_i, a, b_{i+1}, \dots, b_{n-2}) \in B$ for all $i < n-1$, since $B \lhd_t \mathbf{A}$. We have to verify that $s(b_1, \dots, b_{n-2}, a) \in B$. From (31) we obtain that

$$t_m(b_1, \dots, b_{n-2}, b_{n-1}, t_m(b_1, \dots, b_{n-2}, b_{n-1}, a)) = t_m(b_1, \dots, b_{n-2}, b_{n-1}, a)$$

and since $\mathbf{A}$ is abelian, we may apply the term condition to the term t_m at the $(n-1)$-st argument and obtain

$$t_m(b_1, \dots, b_{n-2}, a, t_m(b_1, \dots, b_{n-2}, b_{n-1}, a)) = t_m(b_1, \dots, b_{n-2}, a, a).$$

The right hand side of this equation equals $s(b_1, \dots, b_{n-2}, a)$, and the left hand side is contained in B since $B \lhd_{t_m} \mathbf{A}$. □

Corollary 13.39. *Let $\mathbf{A}$ be a finite idempotent abelian Taylor algebra. Then $\mathbf{A}$ is affine.*

Proof. If $\mathbf{A}$ is abelian, then by Theorem 13.38 it is hereditarily absorption-free. Theorem 13.37 implies that $\mathbf{A}$ has a Maltsev term m. Then Theorem 12.12 implies that $\mathbf{A}$ is affine. □

Exercises.

190. Let $\mathbf{A} := (A, +, -, 0)$ be a group for $A = \{a_1, \dots, a_n\}$. Show that $\mathrm{CSP}(A; +, a_1, \dots, a_n)$ is in P if $\mathbf{A}$ is abelian, and is NP-hard otherwise.[10]

13.5 Paper, Scissors, Stone

This section describes a fundamental example of a three-element algebra which shows some interesting behaviour and which provides important intuition for the abstract results in the following sections. On the one hand, it is absorption-free, but on the other hand it is not hereditarily absorption-free.

Definition 13.40 (Paper-Scissors-Stone algebra)**.** Let $\mathbf{A}$ be the algebra with the domain $A := \{0, 1, 2\}$ and let $\cdot \colon A^2 \to A$ be the binary operation given by the multiplication table on the right.

$\cdot$	0	1	2
0	0	1	0
1	1	1	2
2	0	2	2

Note that $\mathbf{A}$ has the automorphism

$$\rho \colon x \mapsto x + 1 \mod 3.$$

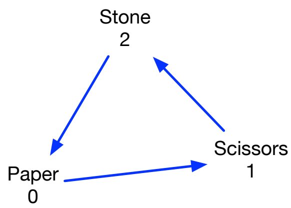

Let $C_3 := \{(a, \rho(a)) \mid a \in A\}$ be the binary relation on A which denotes the graph of ρ. All three 2-element subsets of A are subuniverses of the algebra $\mathbf{A} = (A; \cdot)$, and in each of the corresponding subalgebras the operation $\cdot$ denotes a semilattice operation; however, $\cdot$ itself is not a semilattice operation. We will see below (see Remark 13.44) that none of the proper subalgebras of $\mathbf{A}$ is absorbing. However, $\{1\}$ is a proper absorbing subuniverse of the subalgebra of $\mathbf{A}$ with domain $\{0, 1\}$, so $\mathbf{A}$ is not HAF, and in particular not Abelian 13.38. Note that for any $a, b \in A$

$$(a \cdot \rho^{-1}(b)) \cdot b = b. \tag{32}$$

The algebra $\mathbf{A}$ is simple. Indeed, if C is a congruence which contains $(0, 1)$, then it must also contain $(0, 1) \cdot (2, 2) = (0, 2)$, and therefore also $(1, 0)$ and $(2, 0)$ by symmetry. By similar reasoning we conclude that $C = A^2$, which shows that $\mathbf{A}$ has no proper congruences.

We first present an interesting relational description of $\mathrm{Clo}(\mathbf{A})$. Note that $\mathrm{Inv}(\mathbf{A})$ also contains the relation

$$R_3^{=} := \big\{(x, y, z) \in A^3 \mid x \in \{0, 1\} \wedge (x = 0 \Rightarrow y = z)\big\}.$$

The relation $R_3^{=} \leq \mathbf{A}^3$ is not subdirect, because the first argument cannot take value 2.

Lemma 13.41. *Let $R \leq \mathbf{A}^n$, for $n \geq 1$, be subdirect. Then R can be defined by a conjunction of atomic formulas over $(A; C_3)$.*

Proof. Our proof is by induction on n. For $n = 1$ we have $R = A$ and hence R can be defined by $x = x$. If $n = 2$, then R is the graph of an automorphism of $\mathbf{A}$ or linked, because $\mathbf{A}$ is simple (Exercise 187). In the first case, either $C_3(x, y)$, $C_3(y, x)$, or $x = y$ defines R, and we are done, so let us assume that R is linked.

[10] This result is due to Goldmann and Russell [58]; it can be derived relatively easily from the results in this text. Hint: Combine Exercise 182, Theorem 7.1, Proposition 12.15, Corollary 9.18, and Corollary 13.39.

Claim. $R = A^2$. Indeed, if $(u,v) \in A^2$, then the linkedness of R implies that $(u,v) \in R$ or there exists a path $p_1, \dots, p_{2k} \in A$ for $k \geq 1$ such that $(u,p_1),(p_2,p_1),(p_2,p_3),\dots,(p_{2k},v) \in R$. In the first case there is nothing to be shown. Otherwise, choose k as small as possible. If $k = 1$, we may assume that $p_2 \neq u$ and $p_1 \neq v$, because otherwise we are in the first case. Let $a, b \in A$ be such that $\{u, p_2, a\} = A = \{v, p_1, b\}$. Since R is subdirect, there exist $a', b' \in A$ such that $(a,a'),(b',b) \in R$.

Note that $(u,p_1) \cdot (p_2,v) \in \{(u,v),(u,p_1),(p_2,p_1),(p_2,v)\}$. If $(u,p_1) \cdot (p_2,v) = (u,v)$ then $(u,v) \in R$, contrary to the minimal choice of k. If $(u,p_1) \cdot (p_2,v) = (p_2,v)$, we consider the following subcases.

1. $a' = v$. Then $(a,a') \cdot (u,p_1) = (u,a') = (u,v) \in R$, contradiction.

2. $a' = p_1$. Then $(a,a') \cdot (p_2,v) = (a,v) \in R$, and we are in the first subcase.

3. $a' = b$. Then $(a,a') \cdot (p_2,p_1) = (a,p_1) \in R$ and we are in the second subcase.

The case that $(u,p_1) \cdot (p_2,v) = (u,p_1)$ is similar. Finally, suppose that $(u,p_1) \cdot (p_2,v) = (p_2,p_1)$. Again, we break into subcases.

1. $a' = v$ and $b' = u$. Then $(a,a') \cdot (b',b) = (u,v) \in R$, a contradiction.

2. $a' = p_1$ and $b' = p_2$. Then $(a,a') \cdot (b',b) = (a,b)$ and $(a,b) \cdot (p_2,v) = (a,v) \in R$. Moreover, $(u,p_1) \cdot (a,b) = (u,b) \in R$. Hence, $(a,v) \cdot (u,b) = (u,v) \in R$, and we are done.

3. $a' = v$ and $b' = p_2$. Then $(p_2,p_1) \cdot (a,a') = (a,p_1) \in R$ and we are in subcase 2.

4. $a' = p_1$ and $b' = u$. Then $(b',b) \cdot (p_2,p_1) = (p_2,b) \in R$ and we are again in subcase 2.

5. $a' = b$. Then $(a,a') \cdot (p_2,v) = (a,v) \in R$ and $(a,a') \cdot (u,p_1) = (u,a') \in R$, and we are in subcase number one.

Finally, if $k \geq 2$, then we may assume that $\{u,p_2,p_4\} = A = \{v,p_1,p_3\}$. Then either $(u,p_1) \cdot (p_4,p_3)$ or $(u,p_1) \cdot (p_4,v)$ is from $\{(u,v),(u,p_3),(p_4,p_1)\}$, and in each case we obtain a contradiction to the minimal choice of k. This concludes the proof of the claim.

Now consider the case $n \geq 3$. If $R(x_1,\dots,x_n)$ implies $C_3(x_i,x_j)$ for some $\{i,j\} \in \binom{[n]}{2}$, then R has the definition $C_3(x_i,x_j) \wedge \psi$ in $\mathfrak{A}$, where ψ is the definition of $\pi_{[n]\setminus\{j\}}(R)$ in $\mathfrak{A}$, which exists by inductive assumption. Similarly we can treat the case that $x_i = x_j$ is implied instead of $C_3(x_i,x_j)$. Otherwise, we will show that $R = A^n$. Let $t \in A^n$. For any $a \in A$, the $(n-1)$-ary relation $R_a := \{\bar{x} \mid (\bar{x},a) \in R\}$ is is preserved by $\cdot$. Moreover, we will prove that it is subdirect. Indeed, let $b \in A$. Note that the binary relation $\pi_{2,3}(R)$ equals A^2 by the case $n = 2$, and hence in particular contains (b,a). Therefore, there exists $c' \in A^{n-2}$ such that $(c',b,a) \in R$, and thus $(c',b) \in R_a$. Similar arguments apply to the other arguments of R_a, showing that $R_a \leq \mathbf{A}^{n-1}$ is subdirect.

First consider the case $n = 3$. Since $R_{t_3} \leq \mathbf{A}^3$ is subdirect, by the case $n = 2$ the formula R_{t_3} equals A^2, $=_A$, C_3, or $\{(y,x) \mid (x,y) \in C_3\}$. In any case, R_{t_3} contains (t_1,t_2') and (t_1',t_2) for some $t_1',t_2' \in A$. If $t_1' = t_1$ or $t_2' = t_2$, then $t \in R$ and we are done. If $t_1' = \rho^{-1}(t_1)$ and $t_2' = \rho^{-1}(t_2)$, then $(t_1',t_2,t_3) \cdot (t_1,t_2',t_3) = t \in R$ and we are again done. Hence, up to reordering the arguments of R we may assume that $t_2' = \rho(t_2)$, and since R_{t_3} is preserved by ρ^{-1} we get that $(\rho^{-1}(t_1),t_2) \in R_{t_3}$. Therefore, $t' := (\rho^{-1}(t_1),t_2,t_3) \in R$.

By similar reasoning with the relation $\{(x,y) \mid (t_1,x,y) \in R\}$ instead of R_{t_3} we obtain that $t'' := (t_1, \rho^{-1}(t_2), t_3) \in R$ or $t'' := (t_1, t_2, \rho^{-1}(t_3)) \in R$. Applying $\cdot$ to $t', t'' \in R$, we again obtain $t \in R$.

Finally, consider the case $n > 3$. Then for $i, j \in \binom{[n-1]}{2}$ we have that $\pi_{i,j,n}(R) = A^3$ by the case $n = 3$. Hence, for any $a \in A$ we have $\pi_{i,j}(R_a) = A^2$, and it follows from the case $n-1$ that $R_a = A^{n-1}$. This means that $R = A^n$. $\square$

Definition 13.42. A *pss-Horn clause* is a formula of the form

$$\bigwedge_{i \in [k]} x_i \in \{a_i, \rho(a_i)\} \wedge ((\bigwedge_{i \in [k]} x_i = a_i) \Rightarrow \psi)$$

where $a_1, \dots, a_k \in A$ are constants, $x_1, \dots, x_k$ are variables, and where ψ is

- of the form $y \in \{c, d\}$, for $c, d \in A$ and a variable y,
- of the form $C_3(y, z)$ for variables y and z, or
- of the form $y = z$ for variables y and z.

Note that $k = 0$ is permitted and some variables may be equal.

Proposition 13.43. *For every $R \subseteq A^n$, the following are equivalent.*

1. *R is preserved by $\cdot$;*
2. *R can be defined by a conjunction of pss-Horn clauses;*
3. *R has a primitive positive definition in the structure $(\{0,1,2\}; C_3, R_3^{=})$.*

Proof. For the implication from 1. to 2., suppose that R is preserved by $\cdot$. Let $\phi(x_1, \dots, x_n)$ be the conjunction over all pss-Horn clauses that are implied by $R(x_1, \dots, x_n)$. We prove that ϕ defines R. Suppose that t satisfies ϕ. Let $\{i_1, \dots, i_k\} \subseteq [n]$ be maximal such that $R(x_1, \dots, x_n)$ implies

$$\begin{aligned} & x_{i_1} \in \{t_{i_1}, \rho(t_{i_1})\} \\ & \wedge x_{i_1} = t_{i_1} \Rightarrow x_{i_2} \in \{t_{i_2}, \rho(t_{i_2})\} \\ & \wedge (x_{i_1} = t_{i_1} \wedge x_{i_2} = t_{i_2}) \Rightarrow x_{i_3} \in \{t_{i_3}, \rho(t_{i_3})\} \\ \cdots & \wedge (x_{i_1} = t_{i_1} \wedge \cdots \wedge x_{i_{k-1}} = t_{i_{k-1}}) \Rightarrow x_{i_k} \in \{t_{i_k}, \rho(t_{i_k})\}. \end{aligned}$$

For the sake of notation, we assume that $i_1 = n, \dots, i_k = n - k + 1$, which is without loss of generality, because otherwise we may reorder the arguments of R accordingly. Define

$$R' := \{(x_1, \dots, x_{n-k}) \mid (x_1, \dots, x_{n-k}, t_{n-k+1}, \dots, t_n) \in R\} \leq \mathbf{A}^{n-k}.$$

Note that if $\pi_i(R') = \{a, \rho(a)\}$, for $a \in A$ and $i \in \{m+1, \dots, n-k\}$, then $R(x_1, \dots, x_n)$ implies that $(x_{n-k+1} = t_{n-k+1} \wedge \cdots \wedge x_n = t_n) \Rightarrow x_i \in \{a, \rho(a)\}$. If $a = t_i$ we obtain a contradiction to the maximality of k. Hence, we must have $t_i = \rho(a)$.

Also note that $R' \neq \emptyset$ because otherwise the following pss-Horn clause is implied by $R(x_1, \dots, x_n)$.

$$\bigwedge_{i \in \{n-k+1, \dots, n\}} x_i \in \{t_i, \rho(t_i)\} \wedge \left(\bigwedge_{i \in \{n-k+1, \dots, n\}} x_i = t_i \right) \Rightarrow C(x_n, x_n)$$

But this formula is false for t, contrary to the assumptions that t satisfies ϕ. By further reordering the arguments of R we may additionally assume that there exists $m \in \{0,\dots,n-k\}$ such that $\pi_i(R') = A$ for $i \in [m]$ and $|\pi_i(R')| \leq 2$ for $i \in \{m+1,\dots,n-k\}$.

If $m \geq 1$, then $\pi_{[m]}(R')$ is subdirect in $\mathbf{A}^m$. Hence, Lemma 13.41 implies that $\pi_{[m]}(R')$ can be defined by a conjunction of atomic formulas ψ over $(A; C_3)$. Then for every conjunct χ of ψ we have that $R(x_1,\dots,x_n)$ implies the pss-Horn clause

$$\bigwedge_{j\in\{n-k+1,\dots,n\}} x_j \in \{t_j, \rho(t_j)\} \wedge \left(\bigwedge_{j\in\{n-k+1,\dots,n\}} x_j = t_j \right) \Rightarrow \chi.$$

It follows that $(t_1,\dots,t_m) \in \pi_{[m]}(R')$ satisfies ψ, and hence $(\rho^{-1}(t_1),\dots,\rho^{-1}(t_m)) \in \pi_{[m]}(R')$ since C_3 is preserved by ρ^{-1}. So $(t_1,\dots,t_m)$ and $(\rho^{-1}(t_1),\dots,\rho^{-1}(t_m))$ can be extended to tuples $p,q \in R'$, respectively. If $m = 0$, then we pick $p,q \in R'$ arbitrarily.

For $i \in [n-k]$, let $s^i \in R'$ be such that $s^i_i = t_i$. Define $s := s^{m+1} \cdot (\cdots (s^{n-k-1} \cdot s^{n-k}) \cdots) \in R'$, and note that $s_i = s^i_i = t_i$ for all $i \in \{m+1,\dots,n-k\}$. Observe that $(s_i \cdot \rho^{-1}(t_i)) \cdot t_i = t_i$ for $i \in [m]$ using (32). Also observe that $s_i \cdot q_i = s_i = t_i$ and that $s_i \cdot p_i = s_i = t_i$ for $i \in \{m+1,\dots,n-k\}$, because $p_i, q_i \in \{\rho^{-1}(t_i), t_i\} = \pi_i(R')$. Therefore,

$$\begin{aligned}(s\cdot q)\cdot p &= \big((s_1\cdot\rho^{-1}(t_1))\cdot t_1,\dots,(s_m\cdot\rho^{-1}(t_m))\cdot t_m, s_{m+1},\dots,s_{n-k}\big)\\ &= (t_1,\dots,t_m,t_{m+1},\dots,t_{n-k}) \in R'.\end{aligned}$$

This in turn shows that $t \in R$ and concludes the proof of the implication from 1 to 2.

For the implication from 2. to 3. it suffices to show that every pss-Horn clause has a primitive positive definition in $(\{0,1,2\}; C_3, R^{=}_3)$. Note that

- $\{0,1\}$ has the primitive positive definition $\psi(x)$ given by
$$\exists y,z. R^{=}_3(x,y,z);$$
- $\{1,2\}$ has the primitive positive definition
$$\exists y \big(C_3(y,x) \wedge y \in \{0,1\}\big);$$
- similarly, $\{2,0\}$ and hence also $\{0\}$, $\{1\}$, and $\{2\}$ are primitively positively definable.

Next, for every $k \geq 1$, the relation

$$R^{=}_{k+2}\{(x_1,\dots,x_k,y,z) \in \{0,1\}^k \times A^2 \mid x_1 = \cdots = x_k = 0 \Rightarrow y = z\}$$

has the following primitive positive definition

$$\exists u_1,\dots,u_{k+1}\big(R^{=}_3(x_1,y,u_1) \wedge R^{=}_3(x_2,u_1,u_2) \wedge \cdots \wedge R^{=}_3(x_{k-1},u_{k-1},u_k) \wedge R^{=}_3(x_k,u_k,z)\big).$$

This allows us to define for every $a_1,\dots,a_k \in A$ the relation

$$R^{=}_{a_1,\dots,a_k} := \left\{ (x_1,\dots,x_k,y,z) \in A^{k+2} \mid \bigwedge_{i=1}^{k} x_i \in \{a_i, \rho(a_i)\} \wedge \left(\left(\bigwedge_{i=1}^{k} x_i = a_i \right) \Rightarrow y = z \right) \right\}$$

by the formula

$$\exists u_1,\dots u_k \left(R^{=}_{k+2}(u_1,\dots,u_k,y,z) \wedge \bigwedge_{i=1}^{k} \phi_i(x_i,u_i) \right),$$

where

$$\phi_i := \begin{cases} x_i = u_i & \text{if } a_i = 0; \\ C_3(u_i,x_i) & \text{if } a_i = 1; \\ C_3(x_i,u_i) & \text{if } a_i = 2. \end{cases}$$

Finally, let $c,d \in A$. The pss-Horn clause

$$\bigwedge_{i=1}^{k} x_i \in \{a_i,\rho(a_i)\} \wedge \left(\left(\bigwedge_{i=1}^{k} x_i = a_i \right) \Rightarrow y \in \{c,d\} \right)$$

can be defined by

$$\exists u \; (R^{=}_{a_1,\dots,a_k}(x_1,\dots,x_k,y,u) \wedge u \in \{c,d\})$$

and the pss-Horn clause

$$\bigwedge_{i=1}^{k} x_i \in \{a_i,\rho(a_i)\} \wedge \left(\left(\bigwedge_{i=1}^{k} x_i = a_i \right) \Rightarrow C_3(y,z) \right)$$

by

$$\exists u \; \left(R^{=}_{a_1,\dots,a_k}(x_1,\dots,x_k,y,u) \wedge C_3(u,z)\right).$$

Finally, for the implication from 3. to 1. we verify that C_3 and $R_3^{=}$ are preserved by $\cdot$. For C_3, this is immediate from the fact that ρ is an automorphism of $\mathbf{A}$. If $(x_1,y_1,z_1),(x_2,y_2,z_2) \in R_3^{=}$, we have to show that $(x_0,y_0,z_0) := (x_1 \cdot x_2, y_1 \cdot y_2, z_1 \cdot z_2) \in R_3^{=}$. We have $x_1,x_2 \in \{0,1\}$ and hence $x_0 = x_1 \cdot x_2 \in \{0,1\}$. If $x_0 = 1$, then $(x_0,y_0,z_0) \in R_3^{=}$ and we are done. Otherwise, we must have that $x_1 = x_2 = 0$, and hence $y_1 = z_1$ and $y_2 = z_2$. But then $y_0 = y_1 \cdot y_2 = z_1 \cdot z_2 = z_0$ and again $(x_0,y_0,z_0) \in R_3^{=}$. Hence, if R has a primitive positive definition in $(\{0,1,2\}; C_3, R_3^{=})$, it is preserved by $\cdot$, proving that 3. implies 1. $\square$

Remark 13.44. The algebra $\mathbf{A}$ is absorption free. First note that $B = \{0,1\}$ is not absorbing. Indeed, for every $n \geq 1$ the relation

$$R := \{(x_1,\dots,x_{n-1},y) \in A^n \mid x_1,\dots,x_{n-1} \in \{1,2\} \wedge x_1 = \cdots = x_{n-1} = 1 \Rightarrow y = 2\}$$

can be defined by a pss-Horn clause and hence is a subalgebra of $\mathbf{A}^n$, and is B-essential: for every $i \in \{1,\dots,n-1\}$ we have

$$\begin{aligned} R \cap (B^i \times A \times B^{n-i-1}) &= \{(1,\dots,1,2,1,\dots,1,0),(1,\dots,1,2,1,\dots,1,1)\} \\ R \cap (B^{n-1} \times A) &= \{(1,\dots,1,2)\}, \text{ but} \\ R \cap B^n &= \emptyset. \end{aligned}$$

Hence, Lemma 13.17 implies that B is not absorbing. The same argument shows that $C = \{1\}$ is not absorbing. Every other proper subuniverse is symmetric to B or C via ρ.

Algorithms to solve $\mathrm{CSP}(\{0,1,2\}; C_3, R_3^{=})$ will be discussed in Section 15.

Exercises.

190. Show that CSP($\{0,1,2\}; C_3, R_3^{=}$) can be solved by the 3-consistency procedure (see Exercise 87).

191. Let $\mathbf{A} = (\{0,1,2,3,4\}; \circ)$ be the idempotent algebra where $\circ$ is given by the rock, paper, scissors, lizzard, spock game:

   - spock smashes scissors and vaporises rock,
   - scissors cuts paper and decapitates lizard,
   - paper disproves spock and covers rock,
   - rock crushes scissors and rock, and
   - lizard eats paper and poisons spock.

   Determine the proper subalgebras and proper congruences of $\mathbf{A}$. Which subalgebras are absorbing? Is $\mathbf{A}$ Taylor, Abelian, absorption-free? Is $\mathbf{A}$ subdirectly complete?

192. Is there a structure $\mathfrak{A}$ with a finite relational signature such that a relation $R \subseteq A^n$ is preserved by the operation $\circ$ from the previous exercise if and only if R has a primitive positive definition in $\mathfrak{A}$?

13.6 Ternary Absorption

It will be convenient later to work with ternary absorbing (i.e., 3-absorbing) subalgebras instead of absorbing subalgebras with respect to terms of unbounded arity; this will in particular help in some applications of the absorption theorem in Section 14. The results in this section are from [95] and the presentation is based on [94].

Definition 13.45. We say that an absorbing subalgebra $\mathbf{C}$ of $\mathbf{A}$ is *centrally absorbing*, written $\mathbf{C} \lhd_Z \mathbf{A}$, if

$$(a,a) \notin \langle (\{a\} \times C) \cup (C \times \{a\}) \rangle_{\mathbf{A}^2}$$

for every $a \in A \setminus C$.

Example 13.46. The absorbing subuniverse $\{0\}$ of $\mathbf{A} := (\{0,1\}; \text{majority})$ (Example 13.2) is centrally absorbing, because

$$(1,1) \notin \langle (0,1),(1,0) \rangle_{\mathbf{A}^2} = \{(0,1),(1,0)\}.$$

The absorbing subuniverse $\{0\}$ of $\mathbf{B} := (\{0,1\}; \wedge)$ (Example 13.3) is centrally absorbing, because $(1,1) \notin \langle (0,1),(1,0) \rangle_{\mathbf{B}^2} = \{(0,0),(0,1),(1,0)\}$. △

The next example shows an algebra with an absorbing subuniverse which is *not* centrally absorbing.

Example 13.47. Let $\mathbf{A} := (\{0,1\}; \wedge, \vee)^2$ be the square of the 2-element lattice (Example 8.5). Note that $\{(0,0)\}$ is absorbing with respect to $\wedge$; however, $\langle \{(0,1),(1,0)\} \rangle = \{0,1\}^2$, so $\{(0,0)\}$ is not centrally absorbing. △

One source of centrally absorbing subalgebras is the following proposition.

Proposition 13.48. *Let* $\mathbf{A}$ *and* $\mathbf{B}$ *be finite idempotent and such that* $\mathbf{B}$ *is Taylor and has no proper 2-absorbing subuniverses. Let* $R \leq \mathbf{A} \times \mathbf{B}$ *with left center* C *be such that for every* $a \in A$ *there exists* $b \in B$ *such that* $(a, b) \in R$*. Then* $C \lhd_Z \mathbf{A}$*.*

Proof. By Corollary 13.28 we have that $C \lhd \mathbf{A}$. If C is not centrally absorbing, then for some $n, m \in \mathbb{N}$ there exists $a \in A \setminus C$ and a term operation t of $\mathbf{A}$ of arity $n + m$ and $c_1, \dots, c_m, d_1, \dots, d_n \in C$ such that

$$t(a, \dots, a, c_1, \dots, c_m) = a = t(d_1, \dots, d_n, a, \dots, a).$$

Note that $\{a\} + R$ (using the terminology from Exercise 183) is a proper subalgebra of $\mathbf{B}$: we have $\{a\} + R \neq B$ because $a \notin C$, and $\{a\} + R \neq \emptyset$ by assumption. Moreover, $\{a\} + R$ is 2-absorbing with respect to f given by

$$f(x, y) := t(\underbrace{x, \dots, x}_{n}, \underbrace{y, \dots, y}_{m}):$$

if $b \in \{a\} + R$ and $u \in B$, note that $(c_1, u), \dots, (c_m, u) \in R$ and $(a, b) \in R$, and hence

$$\begin{aligned} f(b, u) &= t(b, \dots, b, u, \dots, u) \\ &\in \{t(a, \dots, a, c_1, \dots, c_m)\} + R = \{a\} + R. \end{aligned}$$

Similarly, $f(u, b) = t(u, \dots, u, b, \dots, b) \in \{t(d_1, \dots, d_n, a, \dots, a)\} + R = \{a\} + R$. This contradicts the assumption that $\mathbf{B}$ has no proper 2-absorbing subuniverses. $\square$

Interestingly, centrally absorbing subalgebras are 3-absorbing (Proposition 13.50). To prove this result, we need the following lemma about essential relations (Definition 13.15).

Lemma 13.49 (Essential doubling)**.** *Let* $\mathbf{A}$ *be finite idempotent and let* $C \lhd_Z \mathbf{A}$*. Suppose that* $R \leq \mathbf{A}^n$*, for* $n \geq 3$*, is* C*-essential. Then there exists* $R' \leq \mathbf{A}^{2n-2}$ *which is* C*-essential.*

Proof. From all relations R that satisfy the assumptions given in the lemma for fixed n, choose R such that $B \leq \mathbf{A}$ given by

$$B := \pi_n(R \cap (C^{n-1} \times A))$$

has minimal size. Note that B is non-empty and disjoint from C by the assumption that R is C-essential. Also note that for every $b \in B$ we have that $B' := \langle C \cup \{b\} \rangle_{\mathbf{A}}$ contains B. To see this, suppose for contradiction that $d \in B \setminus B'$. Then we could replace R by the relation $\tilde{R} := R \cap (A^{n-1} \times B')$. Note that $\tilde{R}$ is C-essential, and that $\pi_n(\tilde{R} \cap (C^{n-1} \times A)) \subseteq B'$ does not contain d, in contradiction to the choice of R and B.

Pick $b \in B$ and define

$$S := \langle (\{b\} \times C) \cup (C \times \{b\}) \rangle_{\mathbf{A}^2}.$$

Finally, let $R' \leq \mathbf{A}^{2n-2}$ be given as the set of all tuples $(x_1, \dots, x_{n-1}, y_1, \dots, y_{n-1})$ that satisfy

$$\exists x_n, y_n \big(R(x_1, \dots, x_n) \wedge S(x_n, y_n) \wedge R(y_1, \dots, y_n) \big). \tag{33}$$

We verify that R' is C-essential. First, for $i \in [2n-2]$, we need to show that

$$R' \cap (C^{i-1} \times A \times C^{2n-2-i}) \neq \emptyset. \tag{34}$$

If $i \in [n-1]$, we may choose $(x_1, \ldots, x_n) \in R \cap (C^{i-1} \times A \times C^{n-i})$ since R is C-essential, and we may choose $(y_1, \ldots, y_n) \in R \cap (C^{n-1} \times \{b\})$ since $b \in B$. Note that $(x_n, y_n) = (x_n, b) \in S$, so $(x_1, \ldots, x_n, y_1, \ldots, y_n)$ satisfies the quantifier-free part of (33), which proves (34). If $i \in \{n, \ldots, 2n-2\}$ then the proof is analogous.

Second, we need to show that

$$R' \cap C^{2n-2} = \emptyset,$$

which is equivalent to

$$S \cap B^2 = \emptyset.$$

Suppose for contradiction that there is $(b_1, b_2) \in S \cap B^2$. As we have observed earlier, $b \in \langle C \cup \{b_1\} \rangle_{\mathbf{A}}$. Hence, for some $k \geq 1$ there exists $f \in \mathrm{Clo}(\mathbf{A})^{(k)}$ and $d_2, \ldots, d_k \in C$ such that $b = f(b_1, d_2, \ldots, d_k)$. Let $b_2' := f(b_2, b, \ldots, b) \in B$. Since $(b_1, b_2) \in S$ and $(d_2, b), \ldots, (d_k, b) \in S$ we have that $(b, b_2') \in S$. We also have $b \in \langle C \cup \{b_2'\} \rangle_{\mathbf{A}}$ by the same argument as above, and hence there exists $f' \in \mathrm{Clo}(\mathbf{A})^{(\ell)}$ and $e_2, \ldots, e_\ell \in C$ such that $b = f'(b_2', e_2, \ldots, e_\ell)$. Since $(b, b_2'), (b, e_2), \ldots, (b, e_\ell) \in S$ we have that $(f'(b, b, \ldots, b), f'(b_2', e_2, \ldots, e_\ell)) = (b, b) \in S$. This contradicts the assumption that $C \lhd_Z \mathbf{A}$, which completes the verification that R' is C-essential. □

Proposition 13.50. *Let* $\mathbf{A}$ *be finite idempotent such that* $\mathbf{C} \lhd_Z \mathbf{A}$*. Then* $\mathbf{C}$ *3-absorbs* $\mathbf{A}$*.*

Proof. Suppose for contradiction that $\mathbf{C}$ does not 3-absorb $\mathbf{A}$. Then there exists a C-essential relation $R \leq \mathbf{A}^3$ by Theorem 13.20. Applying Lemma 13.49 sufficiently many times we may obtain C-essential relations of arbitrarily large arity. Then Theorem 13.20 implies that $\mathbf{C}$ is not absorbing, contrary to our assumptions. □

Combining these results with the proof from Section 13.3, we obtain a strengthened form of the Absorption Theorem (Theorem 13.33).

Theorem 13.51. *Let* $\mathbf{A}$ *and* $\mathbf{B}$ *be finite idempotent algebras with the same signature such that* $\mathbf{B}$ *is Taylor. Then for every linked and subdirect* $R \leq \mathbf{A} \times \mathbf{B}$ *one of the following is true:*

1. $R = A \times B$;
2. $\mathbf{A}$ *has a proper* 3*-absorbing subuniverse.*
3. $\mathbf{B}$ *has a proper* 3*-absorbing subuniverse.*

Proof. Suppose that $R \neq A \times B$, because otherwise item 1 of the theorem holds and we are done. Let C be the left centre of R. Note that $C \neq A$ because $R \neq A \times B$.

If $\mathbf{B}$ has a proper 2-absorbing subuniverse, then item 3 of the theorem holds, so suppose that it does not. Corollary 13.28 implies that $C \lhd \mathbf{A}$. If C is non-empty, then by Proposition 13.48 we have found a proper absorbing subuniverse of $\mathbf{A}$ which is centrally absorbing. In this case, Proposition 13.50 implies that C 3-absorbs $\mathbf{A}$ and item 2 of the theorem holds.

Otherwise, if $C = \emptyset$, then Proposition 13.32 implies that there exists a subdirect $R' \leq \mathbf{B}^2$ whose left centre C' is a proper absorbing subuniverse of $\mathbf{B}$. Then Proposition 13.48 implies that C' is centrally absorbing, Hence, C' 3-absorbs $\mathbf{B}$ by Proposition 13.50 and item 3 of the theorem holds. □

13.7 Zhuk's Cases

The property of the paper-scissors-stone algebra established in Lemma 13.41 is of general importance when studying finite idempotent Taylor algebras (i.e., $\mathrm{Clo}(\mathbf{A})$ does not have a minion homomorphism to **Proj**, see Theorem 9.15).

Definition 13.52. Let $\mathbf{A}$ be an algebra and let $\mathfrak{A}$ be the relational structure with the same domain as $\mathbf{A}$ whose relations are the graphs of the automorphisms of $\mathbf{A}$. Then $\mathbf{A}$ is called *subdirectly complete* if every subdirect $R \leq \mathbf{A}^n$, for every $n \in \mathbb{N}$, can be defined by a conjunction of atomic formulas over $\mathfrak{A}$.

We state a consequence of the Absorption Theorem for finite *simple* algebras. This result is known as 'Zhuk's four cases' but we have combined two cases with absorption into one, so we only show three cases here. The presentation is based on the lecture notes of Brady Zarathustra [94], who cites [95].

Theorem 13.53. *Let $\mathbf{A}$ be a simple finite idempotent Taylor algebra. Then at least one of the following cases applies.*

1. $\mathbf{A}$ *has a proper 3-absorbing subuniverse.*
2. $\mathbf{A}$ *is affine.*
3. $\mathbf{A}$ *is subdirectly complete.*

Proof. Suppose that $\mathbf{A}$ has no proper 3-absorbing subuniverse and is not affine.

We prove by induction on $n \geq 1$ that every subdirect $R \leq \mathbf{A}^n$ has a definition by a conjunction of atomic formulas in the structure $\mathfrak{A}$ whose relations are the graphs of the automorphisms of $\mathbf{A}$. If $n = 1$, then $R = A$ since R is subdirect, and there is nothing to be shown. If $n = 2$, then R is linked or the graph of an automorphism of $\mathbf{A}$, by the simplicity of $\mathbf{A}$ (see Exercise 187). In the latter case we are done, so suppose that R is linked. If $R = A^2$, then we are also done. Otherwise, the Absorption Theorem in its strengthend version (Theorem 13.51) implies that $\mathbf{A}$ has a proper 3-absorbing subuniverse, which is a contradiction to our assumptions.

Now suppose that $n \geq 3$. We first consider the case that for some $\{i,j\} \in \binom{[n]}{2}$ the relation $R' := \pi_{i,j}(R)$ is the graph of an automorphism of $\mathbf{A}$. For the sake of notation, suppose that $j = n$. Then the relation $\pi_{[n-1]}(R)$ is subdirect and by the inductive assumption has a definition $\phi(x_1,\dots,x_{n-1})$ by a conjunction of atomic formulas over $\mathfrak{A}$. Then $\phi(x_1,\dots,x_{n-1}) \wedge R'(x_i,x_j)$ is a definition of R over $\mathfrak{A}$ and we are done. Therefore, we may assume that for every $\{i,j\} \in \binom{[n]}{2}$ the relation $R' := \pi_{i,j}(R)$ is *not* the graph of an automorphism, and hence $R' = A^2$ by the case $n = 2$. We have to show that $R = A^n$.

Note that for every $a \in A$ the $(n-1)$-ary relation

$$R_a := \{\bar{x} \mid (a,\bar{x}) \in R\}$$

is a subuniverse of $\mathbf{A}^{n-1}$ because $\mathbf{A}$ is idempotent. Moreover, R_a is subdirect. Indeed, let $b \in A$. Note that the binary relation $\pi_{1,2}(R)$ equals A^2 by the case $n = 2$, and hence in particular it contains (a,b). Hence, there exists $c' \in A^{n-2}$ such that $(a,b,c') \in R$, and thus $(b,c') \in R_a$. Similar arguments apply to the other arguments of R_a, showing that $R_a \leq \mathbf{A}^{n-1}$ is subdirect.

We first consider the case $n = 3$. Since $R_a \leq \mathbf{A}^2$ is subdirect, by the case $n = 2$ it is either the graph of an automorphism of $\mathbf{A}$ or it equals A^2. First suppose that there exists an $a \in A$ such that $R_a = A^2$. Then a is an element of the left center C of R if R is considered as a subalgebra of $\mathbf{A} \times \mathbf{A}^2$. If $C = A$ then $R = A^3$ and we are done. Otherwise, C is a proper subuniverse of $\mathbf{A}$. Clearly, $\mathbf{A}^2$ is Taylor since $\mathbf{A}$ is Taylor. Lemma 13.14 implies that $\mathbf{A}^2$ does not have proper 2-absorbing subuniverses because $\mathbf{A}$ has no 3-absorbing, and hence no 2-absorbing, subuniverses. Since R is subdirect, Proposition 13.48 implies that $C \lhd_Z \mathbf{A}$. Therefore, C is a (proper) 3-absorbing subuniverse of $\mathbf{A}$ by Proposition 13.50, contrary to our assumptions.

Otherwise, for every $a \in A$ the relation R_a is the graph of an automorphism of $\mathbf{A}$. A similar argument applies to R after permuting the arguments. So we may assume that for every $a \in A$ and every $i \in \{1, 2, 3\}$ the relation defined by $\exists x_i (R(x_1, x_2, x_3) \wedge x_i = a)$ is the graph of an automorphism of $\mathbf{A}$. Then $\mathbf{A}$ is abelian by Proposition 12.15, contrary to our assumptions.

Finally, suppose that $n > 3$. Then for $i, j \in \binom{\{2,\dots,n\}}{2}$ we have that $\pi_{\{1,i,j\}}(R) = A^3$ by the case $n = 3$. Hence, for any $a \in A$ we have $\pi_{\{i,j\}}(R_a) = A^2$, and it follows from the case $n - 1$ that $R_a = A^{n-1}$. This means that $R = A^n$. $\square$

Exercises.

193. Use Theorem 13.53 to give another proof of Lemma 13.41.

194. Show that a subdirect relation $R \subseteq \{0, 1, 2\}^n$ is preserved by the Maltsev operation m in Example 7.4 if and only if R can be defined by a conjunction of graphs of permutations $\{(0, 1), (1, 2), (2, 3)\}$ and $\{(0, 0), (1, 2), (2, 1)\}$.

     Hint. Use a similar proof architecture as in the proof of Theorem 13.53. In the situation where we can apply Proposition 12.15, we obtain that $\mathbf{A}$ is abelian and by Theorem 12.12, $\mathbf{A}$ is affine with a central Maltsev operation m'. However, $m' \neq m$, a contradiction (Exercise 132).

195. Show that a subdirect relation $R \subseteq \{0, 1, 2\}^n$ is preserved by the Maltsev operation m in Example 7.4 if and only if R can be defined primitively positively from the graphs of permutations $\{(0, 1), (1, 2), (2, 3)\}$ and $\{(0, 0), (1, 2), (2, 1)\}$ and from $\{(x, y, z) \in \{0, 1\}^3 \mid x + y + z = 0 \bmod 2\}$.

     Hints. First prove the following substeps:

     - Reduce the general case to the case that $|\pi_i(R)| \geq 2$ for every $i \in \{1, \dots, n\}$.
     - Reduce the general case to the case that $\pi_i(R) = \{0, 1\}$ whenever $|\pi_i(R)| = 2$ for some $i \in \{1, \dots, n\}$.
     - Reduce the general case to the case where $R = R' \times \{0, 1, 2\}^m$ where $R' \subseteq \{0, 1\}^n$.
     - Show by induction on m that a relation R as in the previous item can be defined by a conjunction of linear equations over $\{0, 1\}$ (the most interesting step).
     - Express linear equations over two-element subsets with primitive positive formulas over the given relations.

196. Show that a relation $R \subseteq \{0, 1, 2\}^n$ is preserved by the Maltsev operation m in Example 7.5 if and only if it can be defined from the unary relations, H, and L.

14 Cyclic Terms

An operation $c\colon A^n \to A$, for $n \geq 2$, is *cyclic* if it satisfies for all $a_1, \dots, a_n \in A$ that $c(a_1, \dots, a_n) = c(a_2, \dots, a_n, a_1)$. Cyclic operations are in particular Taylor operations. Conversely, a result of Barto and Kozik (Theorem 14.4 below) implies that every Taylor operation on a finite set generates a cyclic operation.

We start with some easy but useful observations about cyclic terms. The *cyclic composition* $s \circlearrowleft t$ of s and t is the operation (or term) of arity q defined by

$$(x_1, \dots, x_q) \mapsto s\big(t(x_1, \dots, x_q), t(x_2, \dots, x_q, x_1), \dots\big).$$

The following is easy to see.

Lemma 14.1. *Let $s\colon A^k \to A$ and $t\colon A^l \to A$ be operations.*

- *If s is arbitrary and t is cyclic then $s \circlearrowleft t$ is cyclic.*
- *If s is cyclic, t is arbitrary, and l divides k then $s \circlearrowleft t$ is cyclic.*

Exercises.

197. Show that if $\mathbf{A} = (\{0,1\}; \min)$ and $f \in \mathrm{Clo}(\mathbf{A})^{(k)}$ is cyclic, then

$$f(x_1, \dots, x_k) = \min(x_1, \dots, x_k).$$

198. If s and t are cyclic operations or arity k and l, respectively, then star product $s * t$ (Definition 8.32) is cyclic after reordering the arguments, i.e., there is a permutation α of $[kl]$ such that $(s * t)^\alpha$ is cyclic.

199. Suppose that $\mathbf{A} = (\{0,1\}; \text{majority})$ and $f \in \mathrm{Clo}(\mathbf{A})^{(k)}$ is cyclic. Show that

   - $k \geq 3$;
   - if $r > k/2$ and $c \in A^k$ is such that $c_i = a$ for $i \leq r$ and $c_i = b$ otherwise, then $f(c) = a$;
   - if r, s, t are such that $r + s > t$, $s + t > r$, and $t + r > s$, then the function

$$(x, y, z) \mapsto f(\underbrace{x, \dots, x}_{r}, \underbrace{y, \dots, y}_{s}, \underbrace{z, \dots, z}_{t}) \tag{35}$$

   is the ternary majority operation on $\{0, 1\}$.

200. Suppose that p is a prime and $\mathbf{A} = (\{0, \dots, p-1\}; m)$ where $m\colon A^3 \to A$ is given by $m(x,y,z) = x - y + z \mod p$ and that $f \in \mathrm{Clo}(\mathbf{A})^{(k)}$ is cyclic. Show that if r, s, t are such that $r = t = k \mod p$ and $s = -k \mod p$, then the ternary function defined in (35) equals $x - y + z \mod p$.

201. Does the previous exercise remain true if we drop the assumption that p is prime?

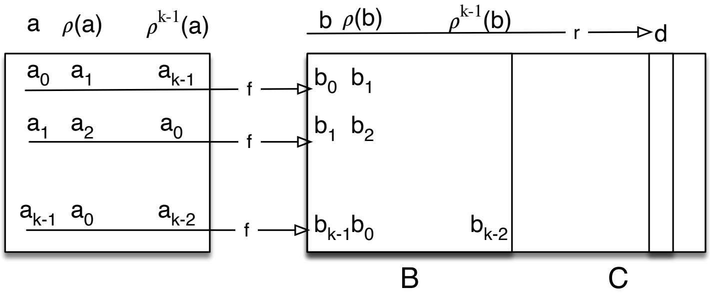

Figure 20: Diagram for the proof of Lemma 14.3.

14.1 Cyclic Relations

When $a = (a_0, a_1, \dots, a_{k-1})$ is a k-tuple, we write $\rho(a)$ for the k-tuple $(a_1, \dots, a_{k-1}, a_0)$.

Definition 14.2. An n-ary relation R on a set A is called *cyclic* if for all $a \in A^k$

$$a \in R \Rightarrow \rho(a) \in R\,.$$

Lemma 14.3 (from [12])**.** *A finite idempotent algebra* $\mathbf{A}$ *has a* k*-ary cyclic term if and only if every nonempty cyclic subalgebra of* $\mathbf{A}^k$ *contains a constant tuple.*

Proof. Let τ be the signature of $\mathbf{A}$. For the easy direction, suppose that $\mathbf{A}$ has a cyclic τ-term $t(x_1, \dots, x_k)$. Let $a = (a_0, a_1, \dots, a_{k-1})$ be an arbitrary tuple in a cyclic subalgebra $\mathbf{R}$ of $\mathbf{A}^k$. As R is cyclic, $\rho(a), \dots, \rho^{k-1}(a) \in R$, and since $\mathbf{R}$ is a subalgebra

$$b := t^{\mathbf{A}}(a, \rho(a), \dots, \rho^{k-1}(a)) \in R.$$

Since t is cyclic, the k-tuple b is constant.

To prove the converse direction, we assume that every nonempty cyclic subalgebra of $\mathbf{A}^k$ contains a constant tuple. For a τ-term $f(x_0, x_1, \dots, x_{k-1})$, let $S(f)$ be the set of all $a \in A^k$ such that $f^{\mathbf{A}}(a) = f^{\mathbf{A}}(\rho(a)) = \cdots = f^{\mathbf{A}}(\rho^{k-1}(a))$. Choose f such that $|S(f)|$ is maximal (here we use the assumption that A is finite). If $|S(f)| = |A^k|$, then $f^{\mathbf{A}}$ is cyclic and we are done. Otherwise, arbitrarily pick $a = (a_0, a_1, \dots, a_{k-1}) \in A^k \setminus S(f)$. For $i \in \{0, \dots, k-1\}$, define $b_i := f(\rho^i(a))$, and let $B := \{b, \rho(b), \dots, \rho^{k-1}(b)\}$.

We claim that the smallest subalgebra $\mathbf{C}$ of $\mathbf{A}^k$ that contains B is cyclic. So let $c \in C$ be arbitrary. Since $\mathbf{C}$ is generated by B, there exists a τ-term $s(x_0, x_1, \dots, x_{k-1})$ such that $c = s^{\mathbf{A}}(b, \rho(b), \dots, \rho^{k-1}(b))$. Then $\rho(c) = s^{\mathbf{A}}(\rho(b), \rho^2(b), \dots, \rho^{k-1}(b), b) \in C$, proving the claim.

Since C is cyclic, by our assumption it contains a constant tuple d. Then there exists a τ-term $r(x_0, \dots, x_{k-1})$ such that $d = r^{\mathbf{C}}(b, \rho(b), \dots, \rho^{k-1}(b))$. Note that

$$r^{\mathbf{A}}(b) = r^{\mathbf{A}}(\rho(b)) = \cdots = r^{\mathbf{A}}(\rho^{k-1}(b))$$

since d is constant. It follows that $b \in S(r)$.

Now consider the τ-term $t(x_0, x_1, \dots, x_{k-1})$ defined by

$$t(x) := r \circlearrowleft f = r(f(x), f(\rho(x)), \dots, f(\rho^{k-1}(x))).$$

where $x := (x_0, x_1, \dots, x_{k-1})$. We claim that $S(f) \subseteq S(t)$. Let $e \in S(f)$. To show that $e \in S(t)$, note that for all $i \in \{0, \dots, k-1\}$

$$\begin{aligned} t^{\mathbf{A}}(\rho^i(e)) &= r^{\mathbf{A}}\big(f^{\mathbf{A}}(\rho^i(e)), f^{\mathbf{A}}(\rho^{i+1}(e)), \dots, f^{\mathbf{A}}(\rho^{i-1}(e))\big) \\ &= r^{\mathbf{A}}\big(f^{\mathbf{A}}(e), f^{\mathbf{A}}(\rho^1(e)), \dots, f^{\mathbf{A}}(\rho^{k-1}(e))\big) && \text{(since } e \in S(f)) \\ &= t^{\mathbf{A}}(e). \end{aligned}$$

Moreover, $a \in S(t)$, because

$$\begin{aligned} t^{\mathbf{A}}(\rho^i(a)) &= r^{\mathbf{A}}\big(f^{\mathbf{A}}(\rho^i(a)), f^{\mathbf{A}}(\rho^{i+1}(a)), \dots, f^{\mathbf{A}}(\rho^{i-1}(a))\big) \\ &= r^{\mathbf{A}}(b_i, b_{i+1}, \dots, b_{i-1}) \\ &= r^{\mathbf{A}}(\rho^i(b)) \end{aligned}$$

is constant for all i by the choice of r. We obtain a contradiction to the maximality of $|S(f)|$. $\square$

Exercises.

202. Show that the digraph C_2^{++} from Exercise 76 has a ternary cyclic polymorphism.

14.2 The Cyclic Terms Theorem

In this section we prove the following theorem of Barto and Kozik [12].

Theorem 14.4 (of [12])**.** *Let* $\mathbf{A}$ *be a finite algebra. Then the following are equivalent.*

1. $\mathbf{A}$ *has a Taylor term;*
2. $\mathbf{A}$ *has a cyclic term;*
3. *for all prime numbers* $p > |A|$*, the algebra* $\mathbf{A}$ *has a* p*-ary cyclic term.*

Proof. The implication from 3 to 2 and from 2 to 1 are trivial. For the implication from 1 to 3, let $p > |A|$ be prime. Our proof is by induction on $|A|$. We may assume that $\mathbf{A}$ is idempotent (see Lemma 9.12). For $|A| = 1$ the statement is trivial. For the induction step, we use Lemma 14.3. Let $R \leq \mathbf{A}^p$ be non-empty and cyclic. We have to show that R contains a constant tuple. We may assume that R is subdirect: indeed, if $\pi_i(R)$ is a proper subuniverse of $\mathbf{A}$, for some $i \in [p]$, then $R \leq \pi_i(R)^p$ contains a constant tuple by the inductive assumption.

If there is $\{i, j\} \in \binom{[p]}{2}$ such that $\pi_{i,j}(R)$ (Definition 6.30) is the graph of an automorphism α of $\mathbf{A}$, then $\pi_{j,2j-i}(R)$ is the graph of α as well, because R is cyclic, and the same applies to $\pi_{2j-i,3j-2i}(R)$, etc. Moreover, $\alpha^p = \mathrm{id}_A$ since R is cyclic and of arity p. Since $p > |A|$ is a prime, we must have $\alpha = \mathrm{id}_A$. This shows that R contains for every $a \in A$ the constant tuple $(a, \dots, a)$: by subdirectness, there exists a tuple $t = (t_1, \dots, t_p) \in R$ such that $t_i = a$; by what we have seen above, $t_j = a$, $t_{2j-i} = a$, etc, and since p is prime we obtain that $t = (a, \dots, a)$.

So we may suppose that for every $\{i, j\} \in \binom{[p]}{2}$, the relation $\pi_{i,j}(R)$ is not the graph of an automorphism of $\mathbf{A}$.

Suppose that $\mathbf{A}$ has a proper congruence C. Let h be the homomorphism from $\mathbf{A}$ to $\mathbf{A}/C$. Since C is proper, $|A/C|$ is strictly smaller than $|A|$. Let h^* be the homomorphism from $\mathbf{A}^p$ to $(\mathbf{A}/C)^p$ obtained by applying h componentwise. Then $h^*(R) \leq (\mathbf{A}/C)^p$

(Lemma 8.15) is cyclic, so the inductive assumption implies that $h^*(R)$ contains a constant tuple $(a/C, \dots, a/C)$. Note that a/C is a nonempty proper subalgebra of $\mathbf{A}$ since $\mathbf{A}$ is idempotent and C is proper. Then $R \cap (a/C)^p \leq \mathbf{A}^p$ is non-empty and cyclic, and hence contains a constant tuple by the inductive assumption. Thus, if $\mathbf{A}$ is not simple we are done.

Now suppose that $\mathbf{A}$ is simple. Therefore, one of the Zhuk cases from Theorem 13.53 applies. First we consider the case that $\mathbf{A}$ is subdirectly complete. By assumption, for every $\{i,j\} \in \binom{[p]}{2}$, the relation $\pi_{i,j}(R)$ is not the graph of an automorphism of $\mathbf{A}$, and by subdirect completeness equals A^2. In particular, R contains a constant tuple.

If $\mathbf{A}$ is affine with underlying abelian group $(A; +, -, 0)$, then for all $k \geq 1$ and $a_1, \dots, a_k \in \mathbb{Z}$ such that $a_1 + \cdots + a_k \equiv 1 \mod |A|$ the operation $(x_1, \dots, x_k) \mapsto a_1 x_1 + \cdots + a_k x_k$ is a term operation of $\mathbf{A}$ (combine Theorem 12.12 and Exercise 143). Since $p > |A|$ is prime, there exists i such that $ip = 1 \mod |A|$. In particular, $(x_1, \dots, x_p) \mapsto i x_1 + \cdots + i x_p$ is a term operation of $\mathbf{A}$, and clearly cyclic.

Finally, suppose that $\mathbf{A}$ has a proper 3-absorbing subalgebra U.[11] Define a directed graph $\mathfrak{D}$ whose vertices are pairs (i, B) where $i \in [p]$ and $B \lhd \mathbf{A}$ is proper 3-absorbing, and whose edges are the pairs $((i,B),(j,B'))$ for distinct $i,j \in [p]$ with $B + \pi_{i,j}(R) \subseteq B'$ (the notation has been introduced in Exercise 183). Clearly, the edge relation of $\mathfrak{D}$ is transitive. Also note that if $B \lhd \mathbf{A}$ is 3-absorbing, then $B + \pi_{i,j}(R) \lhd \mathbf{A}$ is 3-absorbing as well (Exercise 183).

Claim. If $\mathfrak{D}$ contains the edge $((i,B),(j,B'))$, then $((j,B'),(i,B))$ is not an edge. So suppose that $i,j \in [p]$ are distinct such that $B + \pi_{i,j}(R) \subseteq B'$. If $B' + \pi_{j,i}(R) \subseteq B$, then this is in contradiction to the fact that $\pi_{i,j}(R)$ is linked: indeed, by assumption, $\pi_{i,j}(R)$ is not the graph of an automorphism of $\mathbf{A}$. Since R is subdirect, so is $\pi_{i,j}(R)$. Hence, the simplicity of $\mathbf{A}$ implies that $\pi_{i,j}(R)$ is linked (Exercise 187).

The claim together with the transitivity of the edge relation implies that $\mathfrak{D}$ is acyclic. Since $\mathfrak{D}$ is finite, non-empty, and acyclic, it must contain a sink (i,B). Note that B is a proper 3-absorbing subuniverse of $\mathbf{A}$ such that $B + \pi_{i,j}(R) = A$ for all $j \in [p]$. Since R is cyclic, this implies that $\pi_I(R) \cap B^2 \neq \emptyset$ for all $I \in \binom{[p]}{2}$. Hence, $R' := R \cap B^p$ is non-empty by Corollary 13.19, because B is 3-absorbing. Since $|B| < |A|$, we obtain a constant tuple in $R' \subseteq R$ by the inductive assumption. $\square$

Theorem 14.5 (Tractability Theorem, Version 5)**.** *Let $\mathfrak{B}$ be a relational structure with finite domain and finite signature. If $\mathfrak{B}$ has a cyclic polymorphism, then* $\mathrm{CSP}(\mathfrak{B})$ *is in P. Otherwise,* $\mathrm{CSP}(\mathfrak{B})$ *is NP-complete.*

Proof. An immediate consequence of Theorem 14.4 and Theorem 5.28. $\square$

Exercises.

203. Show that if $\mathbf{A}$ and $\mathbf{B}$ are finite algebras, each with a cyclic term, then $\mathbf{A} \times \mathbf{B}$ has a cyclic term as well. How about the same statement for Taylor terms?

204. Use the results presented in the text to show that a finite idempotent algebra $\mathbf{A}$ has a cyclic term if and only if it has a *weak near unanimity term* of arity $n \geq 2$, i.e., a an idempotent term t such that $\mathbf{A}$ satisfies

$$f(x, \dots, x, y) \approx f(x, \dots, x, y, x) \approx \cdots \approx f(y, x, \dots, x).$$

[11]The author thanks Michael Pinsker for a suggestion how to simplify the argument in this case.

205. Show that a finite structure has a cyclic term if and only if it has a *quasi weak near unanimity term* of arity $n \geq 2$, which is defined exactly as *weak near unanimity term* except that we drop the idempotence assumption.

206. Give an immediate proof (without using results from the text) that K_3 does not have quasi weak near unanimity polymorphisms.

207. Find a cyclic term in the algebra $(\mathbb{Z}_p; m)$ where m is given by $(x, y, z) \mapsto x - y + z$.

14.3 Siggers Terms of Arity 4

Interestingly, whether a finite algebra has a Taylor term (equivalently: a weak near unanimity term, or a cyclic term) can be tested by searching for a single 4-ary term s that satisfies

$$s(x, x, y, z) \approx s(y, z, z, x),$$

a so-called *4-ary Siggers term*. Note that this definition comes in numerous variants, because we may permute the arguments of s and rename the variables of the identity and obtain equivalent conditions. One such variant is

$$t(a, r, e, a) \approx t(r, a, r, e)\,.$$

Siggers originally found a 6-ary term (see Section 10.2), which has been improved later to the 4-ary term given above. The observation that this condition can be obtained by equating variables of a cyclic term of sufficiently high arity is from [71]; the proof below is based on a variant from [94] of their proof.

Theorem 14.6. *A finite algebra has a cyclic term if and only if it has a 4-ary Siggers term.*

Proof. Suppose that $\mathbf{A}$ has a cyclic term. Let $c(x_1, \ldots, x_p)$ be a cyclic term of $\mathbf{A}$ for some $p \geq 2$. Then there are numbers $a, b \in \mathbb{N}$ be such that $2a + 3b = m$, and we define $s(x, y, z, w)$ to be the term

$$s(x, y, z, w) := c(\underbrace{x, \ldots, x}_{b}, \underbrace{y, \ldots, y}_{a}, \underbrace{z, \ldots, z}_{b}, \underbrace{w, \ldots, w}_{a+b}).$$

Then

$$\begin{aligned} s(x, x, y, z) &= c(\underbrace{x, \ldots, x}_{b}, \underbrace{x, \ldots, x}_{a}, \underbrace{y, \ldots, y}_{b}, \underbrace{z, \ldots, z}_{a+b}) \\ &\approx c(\underbrace{y, \ldots, y}_{b}, \underbrace{z, \ldots, z}_{a}, \underbrace{z, \ldots, z}_{b}, \underbrace{x, \ldots, x}_{a+b}) = s(y, z, z, x). \end{aligned}$$

Conversely, a Siggers term is a Taylor term, and therefore the other direction follows from Theorem 14.4. $\square$

The previous result is optimal in the sense that there is no equivalent characterisation using a single ternary Taylor term [69, 71]. However, there is also a system of equations involving only ternary terms that characterises the existence of a Taylor term [67]. Computationally, checking whether a given finite structure has polymorphisms satisfying these identities is easier than checking for a 4-ary Siggers polymorphism (for computer experiments, see [25]).

Proposition 14.7 (from [67]). *Let* $\mathbf{A}$ *be a finite algebra. Then* $\mathbf{A}$ *has a Taylor term if and only if* $\mathbf{A}$ *has terms* p, q *satisfying the following identities ('p-q-terms').*

$$q(y,x,x) \approx q(x,x,y) \tag{36}$$
$$q(x,x,y) \approx p(x,y,y) \tag{37}$$
$$p(x,y,x) \approx q(x,y,x) \tag{38}$$

Proof. First suppose that $\mathbf{A}$ has a Taylor term and therefore a 4-ary Siggers term. Define

$$p(x,y,z) := s(x,x,y,z) \text{ and } q(x,y,z) := s(y,z,x,x)$$

and observe that they p and q satisfy the equations from the statement.

$$\begin{aligned} q(y,x,x) &= s(y,y,x,x) \approx s(x,x,x,y) = q(x,x,y) \\ q(x,x,y) &= s(x,x,x,y) \approx s(y,y,x,x) = p(x,y,y) \\ p(x,y,x) &= s(y,x,x,x) \approx s(x,x,y,x) = q(x,y,x) \end{aligned}$$

Conversely, let $\mathbf{A}$ be a algebra that satisfies (36), (37) and (38). Then there is no $\xi \colon \mathrm{Clo}(\mathbf{A}) \to$ **Proj**, because otherwise

$$\begin{aligned} \xi(q) &= \pi^3_2 && \text{(because of 36)} \\ \xi(p) &= \pi^3_1 && \text{(because of 37)} \\ \xi(p) &= \pi^3_2 && \text{(because of 38)} \end{aligned}$$

which is a contradiction unless $|A| = 1$. □

Exercises.

208. Show that every algebra with a Maltsev term has a 4-ary Siggers term (directly, without using other results).

14.4 Summary of Equivalent Dichotomy Formulations

In the following we list all the equivalent conditions on a finite structure $\mathfrak{B}$ with finite relational signature that imply that CSP($\mathfrak{B}$) is in P. If $\mathfrak{B}$ does not satisfy these conditions, then CSP($\mathfrak{B}$) is NP-complete.

Corollary 14.8. *Let* $\mathfrak{B}$ *be a finite structure with core* $\mathfrak{C}$*, let* $\mathfrak{D}$ *be the expansion of* $\mathfrak{C}$ *by all singleton unary relations, and let* $\mathbf{D}$ *be a polymorphism algebra of* $\mathfrak{D}$*. Then the following are equivalent.*

1. K_3 *does not have a primitive positive interpretation in* $\mathfrak{D}$*;*
2. HSP($\mathbf{D}$) *does not contain an at least 2-element algebra all of whose operations are projections;*
3. HS($\mathbf{D}$) *does not contain an at least 2-element algebra all of whose operations are projections (Corollary 8.47);*
4. *there is no clone homomorphism from* Pol($\mathfrak{D}$) *to* **Proj** *(Corollary 8.47);*

5. **D** *satisfies some non-trivial finite set of identities (Corollary 8.47);*

6. $\mathfrak{B}$ *does not pp-construct* K_3*, i..e.,* $K_3 \notin \mathrm{HI}(\mathfrak{B})$ *(Corollary 8.47);*

7. *there is no minion homomorphism from* $\mathrm{Pol}(\mathfrak{B})$ *to* **Proj** *(Corollary 9.18);*

8. $\mathrm{Pol}(\mathfrak{B})$ *contains a Taylor operation (Theorem 9.15);*

9. $\mathrm{Pol}(\mathfrak{B})$ *has a 6-ary Siggers operation (Theorem 10.5);*

10. $\mathrm{Pol}(\mathfrak{B})$ *has a cyclic operation (Theorem 14.4);*

11. $\mathrm{Pol}(\mathfrak{B})$ *has for all prime numbers* $p > |B|$ *a cyclic operation of arity* p *(Theorem 14.4);*

12. $\mathrm{Pol}(\mathfrak{B})$ *has a 4-ary Siggers operation (Theorem 14.6);*

13. $\mathrm{Pol}(\mathfrak{B})$ *has p-q-operations (Proposition 14.7).*

14.5 Undirected Graphs Revisited

As another application of the cyclic term theorem, we obtain another proof (from [16]) of the classification of the complexity of H-colouring for finite undirected graphs H (Theorem 2.6).

Proof. If the core G of H equals K_2 or has just one vertex, then $\mathrm{CSP}(H)$ can be solved in polynomial time, e.g. by the Path Consistency Procedure, Section 4. Otherwise, G is not bipartite and there exists a cycle $a_0, a_1, \ldots, a_{2k}, a_0$ of odd length in H. If H has no Taylor polymorphism, then by Theorem 9.18 $\mathrm{CSP}(H)$ is NP-hard.

Otherwise, if H has a Taylor polymorphism, then Theorem 14.4 asserts that there exists a p-ary cyclic polymorphism c of H where p is a prime number greater than $\max\{2k, |A|\}$. Since the edges in H are undirected, we can also find a cycle $a_0, a_1, \ldots, a_{p-1}, a_0$ in H. Then $c(a_0, a_1, \ldots, a_{p-1}) = c(a_1, \ldots, a_{p-1}, a_0)$, which implies that H contains a loop, a contradiction to the assumption that the core of H has more than one element. □

This proof naturally generalises to smooth digraphs that are strongly connected. In fact, the assumption that H is strongly connected can be dropped.

Theorem 14.9 (Barto, Kozik, Nieven [14])**.** *Let H be a smooth digraph. If H has a Taylor polymorphism, then H is homomorphically equivalent to a cycle.*

In the proof we need the concept of *algebraic length* of a graph. It is the minimum number $k \geq 1$ such that the graph contains a cycle of net length k.

Proof. We only present a proof for the special case where H is strongly connected. Let p be a prime larger than $|V(H)|$. If any two paths in G that start and end in the same vertex have the same net length modulo n, then $H \to \vec{C}_n$ (Exercise 14) and we are done. Otherwise, H has algebraic length one, and since H is strongly connected we find a directed cycle of length p. Theorem 14.4 asserts that there exists a p-ary cyclic polymorphism c of H. As in the proof above, we have $c(a_0, a_1, \ldots, a_{p-1}) = c(a_1, \ldots, a_{p-1}, a_0)$, which implies that H contains a loop and hence is homomorphically equivalent to a loop. □

Corollary 14.10 (Loop Lemma)**.** *Let* $\mathbf{A}$ *be a finite Taylor algebra and let* $R \leq \mathbf{A}^2$ *be subdirect. If the digraph* (A, R) *has a connected component of algebraic length one, then* R *has a loop.*

Figure 21: The smooth digraph leading to the Siggers terms of arity four.

Proof. The digraph (A, R) has no sources and sinks because R is subdirect and has the Taylor term operation of $\mathbf{A}$ as a polymorphism. Hence, Theorem 14.9 implies that (A, R) is homomorphically equivalent to a disjoint union of cycles. The only cycle that is homomorphically equivalent to a digraph of algebraic length one is the loop, so (A, R) is homomorphically equivalent to a structure that contains a loop, so it must contain a loop. $\square$

If a graph H is homomorphically equivalent to a disjoint union of cycles, then $\mathrm{CSP}(H)$ is in P (e.g., we can use the algorithm PC_H to solve it; see Section 4). On the other hand, a digraph without a Taylor polymorphism has an NP-hard CSP. Therefore, Theorem 14.9 shows that the Feder-Vardi conjecture is true for digraphs without sources and sinks: their CSPs are in P or NP-complete.

As another consequence we present a second proof that Taylor algebras have a 4-ary Siggers term, which is the original proof from [69].

Second proof of Theorem 14.6. Let $\mathbf{F}$ be the free algebra with three generators x, y, z in the variety generated by $\mathbf{A}$ (see Section 8.5). Let $\mathbf{R} \leq \mathbf{F}^2$ be generated by

$$\left\{ \begin{pmatrix} x \\ y \end{pmatrix}, \begin{pmatrix} y \\ z \end{pmatrix}, \begin{pmatrix} z \\ x \end{pmatrix}, \begin{pmatrix} x \\ z \end{pmatrix} \right\}.$$

Then $\mathbf{R}$ is subdirect and (F, R) is a smooth digraph of algebraic length 1 (see Figure 21 for the restriction of this digraph to $\{x, y, z\}$). Hence, the Loop Lemma (Corollary 14.10) implies that R contains a pair (f, f), so there exists a term $t(x_1, x_2, x_3, x_4)$ such that

$$t^{\mathbf{F}^2}(\begin{pmatrix} x \\ y \end{pmatrix}, \begin{pmatrix} y \\ z \end{pmatrix}, \begin{pmatrix} z \\ x \end{pmatrix}, \begin{pmatrix} x \\ z \end{pmatrix}) = \begin{pmatrix} f \\ f \end{pmatrix}.$$

Thus, $t^{\mathbf{F}}$ satisfies $t(x, y, z, x) = f = t(y, z, x, z)$. $\square$

Exercises.

209. Let G and H be finite smooth digraphs. Show that if $\mathrm{CSP}(G \times H)$ can be solved in polynomial time, then $\mathrm{CSP}(G)$ or $\mathrm{CSP}(H)$ can be solved in polynomial time as well (so we cannot use them to solve Exercise 27).

210. Let $G := (\{1, 2, 3, 4\}; E)$ be the digraph given by

$$E := \{(1, 2), (1, 3), (2, 3), (3, 2), (2, 4), (3, 4)\}.$$

Show that every finite structure has a primitive positive interpretation in G.

15 Bounded Width

This section is under construction. Equipped with the universal-algebraic approach, we come back to one of the questions that occupied us at the beginning of the course: which H-colouring problems can be solved by the path-consistency procedure (PC_H, introduced in Section 4)? We have seen in Section 4.2 that if H has a majority or a semilattice polymorphism, then PC_H solves the H-colouring problem. But these were just sufficient, not necessary conditions.

A necessary and sufficient polymorphism condition for solvability by PC_H has been found by Barto and Kozik [11]. Their result is much stronger: it characterises not just the strength of PC_H, but more generally of k-consistency (introduced in Section 4), and not just for H-colouring, but for CSPs of finite structures in general. Before we state their result in Theorem 15.8 below, it is convenient to use a more flexible terminology to discuss the idea of k-consistency for general relational structures more precisely (Section 15.1). It turns out that if $\mathrm{CSP}(\mathfrak{B})$ for a finite structure $\mathfrak{B}$ can be solved by k-consistency, for some k, then it can already be solved by a particularly natural algorithm, namely by the singleton arc consistency (SAC) procedure. SAC is weaker than 3-consistency in the sense that if the SAC procedure detects that an instance is unsatisfiable, then so does the 3-consistency procedure.[12]

15.1 k-Consistency

When generalising 3-consistency for the H-colouring problem to k-consistency for CSPs of arbitrary finite structures $\mathfrak{B}$, there are two essential parameters:

- the first is the arity l of the relations maintained for all l-tuples of variables in the instance. For PC_H, for instance, we have $l = 2$.
- the second is the number of variables considered at a time within the main loop of the algorithm. For PC_H, for instance, we have $k = 3$.

Hence, for each pair $(l, k) \in \mathbb{N}^2$, we obtain a different form of consistency, called (l, k)-consistency.

Note that it is easy to come up with finite structures $\mathfrak{B}$ whose CSP cannot be solved by (l, k)-consistency when $\mathfrak{B}$ might contain relations of arity larger than k (there is no possibility of the (l, k)-consistency algorithm to take constraints into account that are imposed on more than k variables). We say that $\mathrm{CSP}(\mathfrak{B})$ has *width* (l, k) if it can be solved by (l, k)-consistency, and that is has *bounded width*) if it has width (l, k) for some $l, k \in \mathbb{N}$. We mention that a CSP has bounded width if and only if unsatisfiability of an instance of $\mathrm{CSP}(\mathfrak{B})$ can be detected by a *Datalog program* (see [55]).

The following lemma suggests that the universal-algebraic approach can be used to study the question for which structures $\mathfrak{B}$ the problem $\mathrm{CSP}(\mathfrak{B})$ has bounded width.

Lemma 15.1. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures with finite relational signature such that $\mathfrak{A} \in \mathrm{HI}(\mathfrak{B})$. If $\mathrm{CSP}(\mathfrak{B})$ has bounded width, then so does $\mathrm{CSP}(\mathfrak{A})$.*

Proof. Let τ be the signature of $\mathfrak{A}$ and σ the signature of $\mathfrak{B}$. Suppose that $\mathrm{CSP}(\mathfrak{B})$ has width (l, k). Let d be the dimension of the primitive positive interpretation I of $\mathfrak{A}$ in $\mathfrak{B}$, let $\delta_I(x_1, \dots, x_d)$ be the domain formula, and let $h \colon D \to A$ be the coordinate map where

[12] The converse of this statement in general does not hold, so SAC is *strictly* weaker than 3-consistency.

$D := \{(b_1, \ldots, b_d) \in B^d \mid \mathfrak{B} \models \delta_I(b_1, \ldots, b_d)\}$. Let ϕ be an unsatisfiable instance of CSP($\mathfrak{A}$) with variable set $U = \{x_1, \ldots, x_n\}$. From ϕ we construct an unsatisfiable instance ψ of CSP($\mathfrak{B}$). This instance will be used as a "guide" when we inductively show that (l, k)-consistency derives false on ϕ.

For fresh and pairwise distinct variables $V := \{y^i_j \mid 1 \leq i \leq d, 1 \leq j \leq n\}$ let ψ_1 be

$$\bigwedge_{1 \leq i \leq n} \delta_I(y^1_i, \ldots, y^d_i) \; .$$

Let ψ_2 be the conjunction of the formulas $\theta_I(y^1_{i_1}, \ldots, y^d_{i_1}, \ldots, y^1_{i_k}, \ldots, y^d_{i_k})$ over all conjuncts $\theta = R(x_{i_1}, \ldots, x_{i_k})$ of ϕ. By moving existential quantifiers to the front, the sentence

$$\exists y^1_1, \ldots, y^d_n \; (\psi_1 \wedge \psi_2)$$

can be re-written to a primitive positive σ-formula ψ.

We claim that ψ is unsatisfiable in $\mathfrak{B}$. Suppose for contradiction that $f \colon V \to B$ satisfies all conjuncts of ψ in $\mathfrak{B}$. By construction of ψ, if ϕ has a conjunct $\theta = R(x_{i_1}, \ldots, x_{i_k})$, then

$$\mathfrak{B} \models \theta_I((f(y^1_{i_1}), \ldots, f(y^d_{i_1})), \ldots, (f(y^1_{i_k}), \ldots, f(y^d_{i_k}))) \; .$$

By the definition of interpretations, this implies that

$$\mathfrak{A} \models R(h(f(y^1_{i_1}), \ldots, f(y^d_{i_1})), \ldots, h(f(y^1_{i_k}), \ldots, f(y^d_{i_k}))) \; .$$

Hence, the mapping $g \colon U \to A$ that sends x_i to $h(f(y^1_i), \ldots, f(y^d_i))$ satisfies all conjuncts of ϕ in $\mathfrak{A}$, in contradiction to the assumption that ϕ is unsatisfiable.

Since CSP($\mathfrak{B}$) has width (l, k) we consequently have that the (l, k)-consistency procedure applied to ψ derives *false*. This derivation can be used to show that the (l, k)-consistency procedure applied to ϕ derives *false*, too. We leave the details to the reader. $\square$

15.2 Singleton AC

For $n, l \in \mathbb{N}$, a matrix $M \in \{0,1\}^{n \times l}$ is called *skeleton* if for every $j \in [n]$, either the j-th row only contains 0 entries, or some column of M is the j-th standard unit vector.

Definition 15.2. The minion $\mathbf{M}_{\mathrm{SAC}}$ is defined as follows. For $n \geq 1$, the set $M^{(n)}_{\mathrm{SAC}}$ consists of all columns of matrices $M \in \{0,1\}^{n \times l}$, for some $l \in \mathbb{N}$, such that the following properties are satisfied:

- there is at least one 1 in every column of M, and
- M is skeleton.

For $\alpha \colon [n] \to [m]$ and $f = (a_1, \ldots, a_n) \in M^{(n)}_{\mathrm{AC}}$, we define f_α to be $(a_{\alpha(1)}, \ldots, a_{\alpha(n)})$.

Note that if M is skeleton, then for every $\alpha \colon [n] \to [m]$, the matrix with the columns $\{f_\alpha \mid f \text{ column of } M\}$ is skeleton as well. Moreover, the property to contain at least one 1 is preserved as well.

Definition 15.3. A tuple $t \in A^{nl}$ is called *conservative* if for every entry a there exists a block consisting of a entirely.

Definition 15.4. An operation $f\colon B^{nl} \to B$ is called *weak l-block totally symmetric polymorphism* if

- it is symmetric on consecutive blocks of size l, and
- for every conservative tuple $(\bar{a}_1, \dots, \bar{a}_n)$ we have that

$$f(\bar{a}_1, \dots, \bar{a}_n) = g(\{a_1^i \mid i \in [l]\}, \dots, \{a_n^i \mid i \in [l]\})$$

for some function $g\colon \mathcal{P}(A)^n \to A$.

Theorem 15.5. *Let $\mathfrak{B}$ be a relational structure with finite signature and finite domain. Then the following are equivalent.*

- *SAC solves* $\mathrm{CSP}(\mathfrak{B})$.
- *there exists a minion homomorphism from* $\mathbf{M}_{SAC}$ *to* $\mathrm{Pol}(\mathfrak{B})$.
- *For every $n, l \in \mathbb{N}$ there exists a weak totally symmetric polymorphism of $\mathfrak{B}$.*

Exercises.

211. Verify that $\mathbf{M}_{\mathrm{SAC}}$ (Definition 15.4) is indeed a minion.

15.3 Weak Near Unanimity Operations

A *weak near unanimity* operation is an operation that satisfies

$$\forall x, y.\, w(x, \dots, x, y) = w(x, \dots, y, x) = \dots = w(y, x, \dots, x)\,.$$

We write $\mathrm{WNU}(k)$ for the k-ary weak near unanimity operations. Again, we warn the reader that many authors additionally assume that weak near unanimity operations are idempotent; we do not make this assumption since it gives us more flexibility of the terminology.

Example 15.6. The algebra $\mathbf{A}_n := (\{0, \dots, n-1\}; m)$ where $m(x, y, z) := x - y + z$ (see Exercise 143) has an $\mathrm{WNU}(k)$ term if and only if $\gcd(k, n) = 1$:

- if $\gcd(k, n) = 1$ then there is an $a \in \{0, \dots, n-1\}$ such that $ak \equiv 1 \mod n$. Hence, $\sum_i ax_i \in \mathrm{WNU}(k)$ and we have $\sum_i a = ka = 1$ so this operation is in $\mathrm{Clo}(\mathbf{A}_n)$.
- Conversely, let $g \in \mathrm{WNU}(k)$. In particular, we have

$$\begin{aligned} g(0, \dots, 0, 1) &= a_k \\ \equiv g(1, 0, \dots, 0) &= a_1 \end{aligned}$$

and it follows that $a := a_0 \equiv \dots \equiv a_{k-1} \mod n$. But $1 = \sum_i a_i = ka \mod n$, which implies that n and m are pairwise prime.

For example, $\mathrm{Clo}(\mathbf{A}_6)$ has a $\mathrm{WNU}(5)$ term, but not $\mathrm{WNU}(k)$ term for $k \leq 4$. △

Theorem 15.7. *Let $\mathfrak{A}$ be a finite idempotent algebra. Then the following are equivalent.*

- *$\mathfrak{A}$ has for every $k \geq 3$ a weak near unanimity operation of arity k.*
- $\mathrm{HS}(\mathfrak{A})$ *does not contain an affine algebra.*

Exercises.

212. Let $(A; +, -, 0)$ be a finite Abelian group. Then every idempotent weak near-unanimity operation that preserves the relation defined by $y_1 + y_2 = y_3 + y_4$ is of the form $(x_1, \dots, x_n) \mapsto t \cdot (x_1 + \cdots + x_n)$ for some $t \in \mathbb{N}_{\geq 1}$.

15.4 The Bounded Width Theorem

Theorem 15.8. *Let $\mathfrak{B}$ be a finite structure. Then the following are equivalent.*

1. $\mathrm{CSP}(\mathfrak{B})$ *has width (l, k) for some $l, k \in \mathbb{N}$.*

2. *for every prime number p, the structure $(\mathbb{Z}_p; +, 1)$ does not have a primitive positive construction in $\mathfrak{B}$.*

3. *$\mathfrak{B}$ does not pp-construct a structure $\mathfrak{C}$ with at least two elements such that there exists an idempotent affine algebra $\mathbf{A}$ with $\mathrm{Clo}(\mathbf{A}) = \mathrm{Pol}(\mathfrak{C})$.*

4. $\mathrm{CSP}(\mathfrak{B})$ *can be solved by singleton linear arc-consistency (SLAC) which will be introduced below.*

5. $\mathrm{CSP}(\mathfrak{B})$ *has width $(2, k)$ where k is the maximal arity of $\mathfrak{B}$.*

6. *$\mathfrak{B}$ has* 3-4 weak near unanimity *polymorphisms, i.e., operations $f \in \mathrm{WNU}(3)$ and $g \in \mathrm{WNU}(4)$ satisfying*
$$\forall x, y.\ f(y, x, x) = g(y, x, x, x)\,.$$

7. *$\mathfrak{B}$ has a binary polymorphism f_2 and polymorphisms $f_n \in \mathrm{WNU}(n)$ for every $n \geq 3$ and*
$$\forall x, y.\ f_n(x, y, \dots, y) = f_2(x, y)\,.$$

8. *$\mathfrak{B}$ has ternary polymorphisms p, q such that $p \in \mathrm{WNU}(3)$ and*
$$\forall x, y\ \big(p(x, x, y) = q(x, y, x) \wedge q(x, x, y) = q(x, y, y)\big)\,.$$

Item 3. mentions a procedure that we only introduce informally here, called *Singleton Linear Arc Consistence* (SLAC). It comes close to strategies that humans perform when solving Sudoku puzzles. First, *Linear Arc Consistency* (LAC) is the restriction of the arc consistency procedure for arbitrary relational signatures where, informally, each inference uses at most one fact that has been derived previously. *SLAC* is the extension of LAC which performs the following with an instance I of $\mathrm{CSP}(\mathfrak{B})$:

1. Run LAC on I; if LAC derives *false*, return **No**.

2. Create a copy I' of I.

3. Pick some variable x of I' and some value v from B; set $x := v$.

4. If LAC derives false on I', remove v from the list for x in I.

5. Otherwise, do nothing (I is unchanged).

We repeat these steps until for no pair (x, v) a value can be removed from I, in which case we return **Yes**.

We do not give a complete proof of the important Theorem 15.8, but only show some of the easy implications, and we explain how to deduce the remaining implications from statements that can be found explicitly in the literature.

Proof. 1. $\Rightarrow$ 2.: By Lemma 15.1, it suffices to show that $\mathrm{CSP}(\mathbb{Z}_p; R_+, \{1\})$ does not have width (l, k). Two proof sketches of this fact can be found in [55]. A stronger non-expressibility can be found in [5] (the given CSP is not even expressible in least fixed point logic with counting quantifiers).

The implication from 2. $\Rightarrow$ 1. was open for a while, has been conjectured by Larose and Zadori [76] (and in equivalent form, by Feder and Vardi [55]; also see [75]), and was proven by Barto and Kozik [11]; the proof requires several important concepts, e.g., *Prague strategies* and *absorption theory* (Section 13).

For the implication from 2. to 3., suppose that $\mathfrak{B}$ pp-constructs a structure $\mathfrak{C}$ with at least two elements such that there exists an idempotent affine algebra $\mathbf{A}$ with $\mathrm{Clo}(\mathbf{A}) = \mathrm{Pol}(\mathfrak{B})$. Since pp-constructions compose, it suffices to show that $\mathfrak{C}$ pp-constructs $(\mathbb{Z}_p; +, 1)$ for some prime p; this has been shown in Proposition 12.7.

The (unexpected) implication from 2. $\Rightarrow$ 4. is from [70]. There it is shown that if $\mathbf{A}$ is finite, idempotent, and has no affine factors with more than one element, and $\mathrm{Clo}(\mathbf{A}) = \mathrm{Pol}(\mathfrak{C})$, then $\mathrm{CSP}(\mathfrak{C})$ can be solved by SLAC. To use their result, let $\mathfrak{C}$ be the expansion of the core of $\mathfrak{B}$ by all singleton unary sets; by Proposition 5.25, $\mathfrak{C}$ has a pp-construction in $\mathfrak{B}$. If $\mathbf{A}$ has an affine factor with more than one element, then $\mathfrak{C}$ pp-constructs a structure with at least two elements and an idempotent affine polymorphism algebra. Composing pp-constructions, we obtain a contradiction to 2. Otherwise, $\mathrm{CSP}(\mathfrak{C})$ can be solved by SLAC, and it immediately follows that $\mathrm{CSP}(\mathfrak{B})$ can be solved by SLAC as well.

The implication from 4. to 5. is easy: a derivation of *false* by SLAC can be simulated by a derivation of *false* by $(2, k)$-consistency where k is the maximal arity of the relations in $\mathfrak{B}$.

The implication from 5. to 6. has an elegant short proof, see [71]. The equivalences between the final three items when we additionally require idempotency for the terms has been shown in [67] (e.g., for 7., see Proposition 4.1 in [67]). But since a structure has polymorphism satisfying an equation without nesting (and the equations under consideration are of this type) if and only if its core does, and since a core has such a polymorphism if and only if it has a polymorphism that is additionally idempotent, the idempotent case implies the statement as given in the theorem.

The implication from 7. to 2. is easy: first note that if $\mathfrak{B}$ has a $\mathrm{WNU}(k)$-polymorphism, then so do have all structures in $\mathrm{H}(\mathrm{I}(\mathfrak{B}))$ (recall that we do not require that the operations in $\mathrm{WNU}(k)$ are idempotent). But the structure $(\mathbb{Z}_p; +, 1)$ does not have $\mathrm{WNU}(k)$ polymorphisms for both $k = 3$ and $k = 4$ (see Example 15.6). ☐

In the following we point out some immediate consequences of Theorem 15.8.

Corollary 15.9. *Let H be a finite digraph. Then strong path consistency solves* $\mathrm{CSP}(H)$ *if and only if H has weak near unanimity polymorphisms f and g satisfying*

$$\forall x, y.\ g(y, x, x) = f(y, x, x, x)$$

Another remarkable consequence is that for the H-colouring problem, $(2, 3)$-consistency is as powerful as $(2, k)$-consistency for all $k \geq 3$ (we already stated this in Theorem 4.2). One technical step of the proof of Theorem 15.8 is to reduce the argument to an argument about the strength of $(2,3)$-consistency via Corollary 5.24.

16 Open Problems

The Feder and Vardi dichotomy conjecture [55] has been the outstanding open problem in the field; it was solved in 2017 by Bulatov [36] and, independently, by Zhuk [96]. The boarder between polynomial and NP-hard cases has numerous equivalent logical and algebraic characterisations, for example characterisations based on primitive positive constructability and characterisations based on identities that are satisfied by the polymorphism clone (see Corollary 14.8).

There are many interesting problems in the field that are still left open. We start with open research problems where all the relevant concepts have already been introduced in the course.

1. Is there a *uniform* algorithm for finite-domain CSPs, i.e., is there a polynomial-time algorithm that takes as input a pair $(\mathfrak{A}, \mathfrak{B})$ of finite structures with the same finite relational signature, where $\mathfrak{B}$ is promised to have a Taylor polymorphism (for equivalent promises, see Corollary 14.8), that decides whether there is a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$. This question is also open if we replace 'Taylor polymorphism' by 'Maltsev polymorphism'.

2. Is there a polynomial-time algorithm to determine whether a given core structure $\mathfrak{B}$ has a Siggers polymorphism? Is this true for the special case where $\mathfrak{B}$ is a digraph or an orientation of a tree? This problem is known to be NP-complete if $\mathfrak{B}$ is not required to be a core structure [44]. This problem is known to be in P if $\mathfrak{B}$ is a core and there is a uniform algorithm for finite-domain CSPs (as in the previous question; see [44]).

3. Is the class of all finite structures, ordered by pp-constructability and factored by the respective equivalence relation, a lattice [28–30]? Is it countably infinite or uncountably infinite? Are there infinite ascending chains?

4. What is the computational complexity of determining whether a given finite core structure H has tree duality? Is this problem in P? Is it in P if H is a digraph or even an orientation of a tree?

5. (Bulín [40]) Is is true that the CSP of an orientation of a tree is in P if and only if it can be solved by Datalog?

6. Is it true that *most* orientations of finite trees are hard, i.e., is it true that the probability that an orientation of a tree drawn uniformly at random from the set of all such trees with vertex set $\{1, \ldots, n\}$ is NP-hard tends to 1 as n tends to infinity [25]? The answer is yes if we ask the question for random labelled digraphs instead of random labelled trees [79].

7. Determine the smallest trees whose CSP is P-hard (assuming that NL $\neq$ P). It is known that they must have at least 16 vertices, since all smaller trees have a majority polymorphism and thus are in NL [25].

8. Is the CSP of most digraphs with a Taylor polymorphism P-hard?

9. Is the CSP of most digraphs with a Taylor polymorphism not in Datalog?

We continue with some open problems that require knowledge of concepts that have not been covered in this course; however, references are provided where these concepts are defined formally.

1. Prove that a finite-domain CSP is in P if and only if it can be expressed in *Choiceless Polynomial Time* [20].

2. (Dalmau [47]) Is it true that if CSP(H) is in NL, then CSP(H) is in linear Datalog? Is this at least true for digraphs H? The same question would already be interesting for orientations of trees.

3. (Egri-Larose-Tesson [51]) Is it true that if CSP(H) is in L, then CSP(H) is in symmetric Datalog? Is this at least true for digraphs H? It would already be interesting for orientations of trees.

4. (Larose-Tesson [74]) Is it true that if the polymorphism algebra of H generates a congruence join-semidistributive variety, then CSP(H) is in linear Datalog? Is this at least true for digraphs H? It would already be interesting for orientations of trees.

5. Is it true that if CSP(H) is not P-hard under logspace reductions, then it is in NC? It is known that NC is closed under logspace reductions, and it is believed that P is different from NP. Moreover, the CSP for the structure $(\{0,1\}; \{0,1\}^3 \setminus \{(1,1,0)\}, \{0\}, \{1\})$ is P-hard (see Exercise 120). Is it true that if CSP(H) does not pp-construct this structure then CSP(H) is in NC?

Finally some curious questions for concrete finite digraphs where we do not know the answer.

1. Is the CSP of the orientation of a tree displayed on the right in NL [25]?

2. What are the smallest digraphs with a Taylor polymorphism that cannot be solved by Datalog?

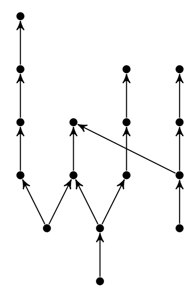

For CSPs over infinite domains, there are numerous open problems, and I invite the reader to have a look at [21].

References

[1] S. Aaronson. P=$^?$NP. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:4, 2017.

[2] G. Ahlbrandt and M. Ziegler. Quasi-finitely axiomatizable totally categorical theories. *Annals of Pure and Applied Logic*, 30(1):63–82, 1986.

[3] E. Aichinger, P. Mayr, and R. McKenzie. On the number of finite algebraic structures. *Journal of the European Mathematical Society*, 16(8):1673–1686, 2014.

[4] B. Aspvall, M. F. Plass, and R. E. Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Information Processing Letters*, 8(3):121–123, 1979.

[5] A. Atserias, A. A. Bulatov, and A. Dawar. Affine systems of equations and counting infinitary logic. *Theoretical Computer Science*, 410(18):1666–1683, 2009.

[6] L. Barto. The dichotomy for conservative constraint satisfaction problems revisited. In *Proceedings of the Symposium on Logic in Computer Science (LICS)*, Toronto, Canada, 2011.

[7] L. Barto. Finitely related algebras in congruence distributive varieties have near unanimity terms. *Canadian Journal of Mathematics*, 65(1):3–21, 2013.

[8] L. Barto, Z. Brady, A. Bulatov, M. Kozik, and D. Zhuk. Minimal Taylor algebras as a common framework for the three algebraic approaches to the CSP. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–13. IEEE, 2021.

[9] L. Barto, J. Bulín, A. A. Krokhin, and J. Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):28:1–28:66, 2021.

[10] L. Barto and A. Kazda. Deciding absorption. *Int. J. Algebra Comput.*, 26(5):1033–1060, 2016.

[11] L. Barto and M. Kozik. Constraint satisfaction problems of bounded width. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 595–603, 2009.

[12] L. Barto and M. Kozik. Absorbing subalgebras, cyclic terms and the constraint satisfaction problem. *Logical Methods in Computer Science*, 8/1(07):1–26, 2012.

[13] L. Barto and M. Kozik. Absorption in universal algebra and CSP. In *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 45–77, 2017.

[14] L. Barto, M. Kozik, and T. Niven. The CSP dichotomy holds for digraphs with no sources and no sinks (a positive answer to a conjecture of Bang-Jensen and Hell). *SIAM Journal on Computing*, 38(5), 2009.

[15] L. Barto, M. Kozik, and D. Stanovský. Mal'tsev conditions, lack of absorption, and solvability. *Algebra Universalis*, 74:185–206, 2015.

[16] L. Barto, A. A. Krokhin, and R. Willard. Polymorphisms, and how to use them. In *The Constraint Satisfaction Problem: Complexity and Approximability*, pages 1–44. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.

[17] L. Barto, J. Opršal, and M. Pinsker. The wonderland of reflections. *Israel Journal of Mathematics*, 223(1):363–398, 2018.

[18] N. Biggs. Constructions for cubic graphs with large girth. *Electronic Journal of Combinatorics*, 5, 1998.

[19] G. Birkhoff. On the structure of abstract algebras. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):433–454, 1935.

[20] A. Blass, Y. Gurevich, and S. Shelah. Choiceless polynomial time. *Annals of Pure and Applied Logic*, 100(1-3):141–187, 1999.

[21] M. Bodirsky. *Complexity of Infinite-Domain Constraint Satisfaction*. Lecture Notes in Logic (52). Cambridge University Press, Cambridge, United Kingdom; New York, NY, 2021.

[22] M. Bodirsky. Model theory, 2022. Course Notes, TU Dresden, `https://wwwpub.zih.tu-dresden.de/~bodirsky/Model-theory.pdf`.

[23] M. Bodirsky. Diskrete Strukturen, 2023. Skript zur Vorlesung, TU Dresden, `https://wwwpub.zih.tu-dresden.de/~bodirsky/Diskrete-Strukturen.pdf`.

[24] M. Bodirsky. Introduction to mathematical logic, 2023. Course notes, TU Dresden, `https://wwwpub.zih.tu-dresden.de/~bodirsky/Logic.pdf`.

[25] M. Bodirsky, J. Bulín, F. Starke, and M. Wernthaler. The smallest hard trees. *Constraints*, abs/2205.07528, 2022.

[26] M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.

[27] M. Bodirsky and M. Pinsker. Topological Birkhoff. *Transactions of the American Mathematical Society*, 367(4):2527–2549, 2015.

[28] M. Bodirsky and F. Starke. Maximal digraphs with respect to primitive positive constructability. *Combinatorica*, 42:997–1010, 2022.

[29] M. Bodirsky, F. Starke, and A. Vucaj. Smooth digraphs modulo primitive positive constructability and cyclic loop conditions. *International Journal on Algebra and Computation*, 31(5):939–967, 2021. Preprint available at ArXiv:1906.05699.

[30] M. Bodirsky and A. Vucaj. Two-element structures modulo primitive positive constructability. *Algebra Universalis*, 81(20), 2020. Preprint available at ArXiv:1905.12333.

[31] M. Bodirsky, A. Vucaj, and D. Zhuk. The lattice of clones of self-dual operations collapsed. *Accepted for publication in the International Journal on Algebra and Computation*, 2023. Preprint available at https://arxiv.org/abs/2109.01371.

[32] V. G. Bodnarčuk, L. A. Kalužnin, V. N. Kotov, and B. A. Romov. Galois theory for Post algebras, part I and II. *Cybernetics*, 5:243–539, 1969.

[33] A. A. Bulatov. Tractable conservative constraint satisfaction problems. In *Proceedings of the Symposium on Logic in Computer Science (LICS)*, pages 321–330, Ottawa, Canada, 2003.

[34] A. A. Bulatov. H-coloring dichotomy revisited. *Theoretical Computer Science*, 349(1):31–39, 2005.

[35] A. A. Bulatov. Conservative constraint satisfaction re-revisited. *Journal Computer and System Sciences*, 82(2):347–356, 2016. ArXiv:1408.3690.

[36] A. A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17*, pages 319–330, 2017.

[37] A. A. Bulatov and V. Dalmau. A simple algorithm for Mal'tsev constraints. *SIAM Journal on Computing*, 36(1):16–27, 2006.

[38] A. A. Bulatov and P. Jeavons. Algebraic structures in combinatorial problems. Technical report MATH-AL-4-2001, Technische Universität Dresden, 2001.

[39] A. A. Bulatov, A. A. Krokhin, and P. G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.

[40] J. Bulín. On the complexity of H-coloring for special oriented trees. *Eur. J. Comb.*, 69:54–75, 2018.

[41] J. Bulín, D. Delic, M. Jackson, and T. Niven. A finer reduction of constraint problems to digraphs. *Log. Methods Comput. Sci.*, 11(4), 2015.

[42] S. N. Burris and H. P. Sankappanavar. *A Course in Universal Algebra*. Springer Verlag, Berlin, 1981.

[43] C. Carvalho, L. Egri, M. Jackson, and T. Niven. On Maltsev digraphs. *Electr. J. Comb.*, 22(1):P1.47, 2015.

[44] H. Chen and B. Larose. Asking the metaquestions in constraint tractability. *TOCT*, 9(3):11:1–11:27, 2017.

[45] D. A. Cohen, M. C. Cooper, P. G. Jeavons, and S. Živný. Binarisation via dualisation for valued constraints. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 3731–3737, 2015.

[46] B. Csákány. Minimal clones. *Algebra Universalis*, 54(1):73–89, 2005.

[47] V. Dalmau. Linear Datalog and bounded path duality of relational structures. *Logical Methods in Computer Science*, 1(1), 2005.

[48] V. Dalmau and J. 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.

[49] R. Dechter. *Constraint Processing*. Morgan Kaufmann, 2003.

[50] Dummit and Foote. *Abstract Algebra*. Wiley, 2004. Third edition.

[51] L. Egri, B. Larose, and P. Tesson. Symmetric Datalog and constraint satisfaction problems in logspace. In *Proceedings of the Symposium on Logic in Computer Science (LICS)*, pages 193–202, 2007.

[52] M. M. El-Zahar and N. Sauer. The chromatic number of the product of two 4-chromatic graphs is 4. *Combinatorica*, 5(2):121–126, 1985.

[53] T. Feder. Classification of homomorphisms to oriented cycles and of k-partite satisfiability. *SIAM Journal on Discrete Mathematics*, 14(4):471–480, 2001.

[54] T. Feder and M. Y. Vardi. Monotone monadic SNP and constraint satisfaction. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 612 – 622, 1993.

[55] T. Feder and M. 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, 1999.

[56] M. Garey and D. Johnson. *A guide to NP-completeness.* CSLI Press, Stanford, 1978.

[57] D. Geiger. Closed systems of functions and predicates. *Pacific Journal of Mathematics*, 27:95–100, 1968.

[58] M. Goldmann and A. Russell. The complexity of solving equations over finite groups. *Information and Computation*, 178(1):253–262, 2002.

[59] G. H. Hardy and E. M. Wright. *An introduction to the theory of numbers.* Oxford University Press, 2008. Sixth edition.

[60] P. Hell and J. Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48:92–110, 1990.

[61] P. Hell and J. Nešetřil. The core of a graph. *Discrete Mathematics*, 109:117–126, 1992.

[62] P. Hell and J. Nešetřil. *Graphs and Homomorphisms.* Oxford University Press, Oxford, 2004.

[63] D. Hobby and R. McKenzie. *The structure of finite algebras*, volume 76 of *Contemporary Mathematics.* American Mathematical Society, 1988.

[64] W. Hodges. *Model theory.* Cambridge University Press, Cambridge, 1993.

[65] W. Hodges. *A shorter model theory.* Cambridge University Press, Cambridge, 1997.

[66] P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, 1997.

[67] J. Jovanović, P. Marković, R. McKenzie, and M. Moore. Optimal strong Mal'cev conditions for congruence meet-semidistributivity in locally finite varieties. *Algebra Universalis*, 76:305–325, 2016.

[68] A. Kazda. Maltsev digraphs have a majority polymorphism. *European Journal of Combinatorics*, 32:390–397, 2011.

[69] K. A. Kearnes, P. Marković, and R. McKenzie. Optimal strong Mal'cev conditions for omitting type 1 in locally finite varieties. *Algebra Universalis*, 72(1):91–100, 2015.

[70] M. Kozik. Weak consistency notions for all the CSPs of bounded width. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 633–641, 2016.

[71] M. Kozik, A. Krokhin, M. Valeriote, and R. Willard. Characterizations of several Maltsev conditions. *Algebra universalis*, 73(3):205–224, 2015.

[72] G. Kun and M. Szegedy. A new line of attack on the dichotomy conjecture. *Eur. J. Comb.*, 52:338–367, 2016.

[73] B. Larose, C. Loten, and C. Tardif. A characterisation of first-order constraint satisfaction problems. *Logical Methods in Computer Science*, 3(4:6), 2007.

[74] B. Larose and P. Tesson. Universal algebra and hardness results for constraint satisfaction problems. *Theoretical Computer Science*, 410(18):1629–1647, 2009.

[75] B. Larose, M. Valeriote, and L. Zádori. Omitting types, bounded width and the ability to count. *International Journal of Algebra and Computation*, 19(5), 2009.

[76] B. Larose and L. Zádori. Bounded width problems and algebras. *Algebra Universalis*, 56(3-4):439–466, 2007.

[77] D. Lau. *Function algebras on finite sets : a basic course on many-valued logic and clone theory / Dietlinde Lau.* Springer monographs in mathematics. Springer, Berlin Heidelberg New York, cop. 2006.

[78] L. Libkin. *Elements of Finite Model Theory.* Springer, 2004.

[79] T. Łuczak and J. Nešetřil. When is a random graph projective? *Eur. Journal Comb.*, 27(7), 2006.

[80] A. K. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 8:99–118, 1977.

[81] P. Marković, M. Maróti, and R. McKenzie. Finitely related clones and algebras with cube-terms. *Order*, 29:345–359, 2012.

[82] R. N. McKeznie, G. F. McNulty, and W. F. Taylor. *Algebras, Lattices, Varieties (Volume 1).* American Mathematical Society, 1987.

[83] U. Montanari. Networks of constraints: Fundamental properties and applications to picture processing. *Information Sciences*, 7:95–132, 1974.

[84] C. H. Papadimitriou. *Computational Complexity.* Addison-Wesley, 1994.

[85] J. Płonka. r-prime idempotent reducts of groups. *Arch. Math. (Basel)*, 24:129–132, 1973.

[86] E. L. Post. *The two-valued iterative systems of mathematical logic*, volume 5. Princeton University Press, Princeton, 1941.

[87] I. G. Rosenberg. Minimal clones I: the five types. *Lectures in Universal Algebra (Proc. Conf. Szeged, 1983), Colloq. Math. Soc. J. Bolyai*, 43:405–427, 1986.

[88] T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 216–226, 1978.

[89] U. Schöning. *Logic for Computer Scientists.* Springer, 1989.

[90] Y. Shitov. Counterexamples to Hedetniemi's conjecture, 2019. arXiv:1905.02167.

[91] M. H. Siggers. A strong Mal'cev condition for varieties omitting the unary type. *Algebra Universalis*, 64(1):15–20, 2010.

[92] W. Taylor. Varieties obeying homotopy laws. *Canadian Journal of Mathematics*, 29:498–527, 1977.

[93] Y. I. Yanov and A. A. Muchnik. On the existence of k-valued closed classes without a finite basis. *Dokl. Akad. Nauk SSSR*, 127:44–46, 1959. in Russian.

[94] B. Zarathustra. Notes on CSPs and polymorphisms. *CoRR*, abs/2210.07383, 2022.

[95] D. Zhuk. Strong subalgebras and the constraint satisfaction problem. *CoRR*, abs/2005.00593, 2020.

[96] D. N. Zhuk. A proof of CSP dichotomy conjecture. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17*, pages 331–342, 2017. https://arxiv.org/abs/1704.01914.

A O-notation

The letters o and O stand for the *order* of growth of the function. The *big-O notation* is used to express upper bounds, and the *little-o notation* to express lower bounds. We mention that there exists related notation to describe other kinds of bounds on asymptotic growth, e.g., Θ, Ω, ω, of which we only need Θ in this text, so we skip the definitions of the others.

Let $g\colon \mathbb{R} \to \mathbb{R}$ (we use $\mathbb{R}$ for convenience; similar definitions exist for other domains such as $\mathbb{N}$ and $\mathbb{Q}$, etc). Then $O(g)$ is the set of all functions $f\colon \mathbb{R} \to \mathbb{R}$ such that there exists $c, x_0 \in \mathbb{R}$ such that $|f(x)| \leq cg(x)$ for all $x \geq x_0$. Note that

$$f \in O(g) \Leftrightarrow \limsup_{x\to\infty} \left|\frac{f(x)}{g(x)}\right| < \infty.$$

In typical usage, the formal definition of $O(g)$ is not used directly; rather, we first use the following simplification rules:

- if $g(x)$ is a sum of several terms, if there is one with largest growth rate, then we drop all other terms;
- if $g(x) = c \cdot f(x)$ and c is a constant that does not depend on x, then c can be omitted.

When we write $O(g)$, we typically choose g to be as simple as possible. O-notation can also be used within arithmetic terms. For example, $h + O(g)$ denotes the set of functions of the form $h + f$ for $f \in O(g)$. In other words, $k \in h + O(g)$ is equivalent to $k - h \in O(G)$.

We write $o(g)$ for the set of all functions $f\colon \mathbb{R} \to \mathbb{R}$ such that for every $\epsilon \in \mathbb{R}_{>0}$ there exists $x_0 \in \mathbb{R}$ such that $|f(x)| \leq \epsilon g(x)$ for all $x \geq x_0$. Informally, $f \in o(g)$ means that g grows much faster than f. For example, $x \mapsto 2x$ is in $o(x \mapsto x^2)$, and $x \mapsto 1/x$ is in $o(1)$. Note that $o(g) \subseteq O(g)$, and that

$$f \in o(g) \Leftrightarrow \lim_{x\to\infty} \frac{f(x)}{g(x)} = 0.$$

Similarly as in the case of the O-notation we may use the o-notation in arithmetic expressions. Note that if $f \in o(g)$ and c is a constant, then $cf \in o(g)$. Frequent notation is to write $f \ll g$ (or $g \gg f$) if $f \in o(g)$.

We write $\Theta(g)$ for the set of all functions f such that there are constants c, C and $x_0 \in \mathbb{R}$ such that $cg(x) \leq f(x) \leq Cg(x)$ for every $x \geq x_0$. In other words, $f \in \Theta(g)$ if $f \in O(g)$ and $g \in O(f)$.

Finally, we write $f(x) \sim g(x)$ if

$$\lim_{x\to\infty} \frac{f(x)}{g(x)} = 1$$

and we say that f and g are *asymptotically equivalent* (for $x \to \infty$).

B Basics of Complexity Theory

For a set A, we write A^* for the set of all words over the alphabet A. A *word over A* can be seen as a function from $\{1, \dots, n\} \to A$, for some $n \in \mathbb{N}$. We write ϵ for the empty word (i.e., for the function with the empty domain).

The most classical setting of complexity theory is the study of the computational complexity of functions f from $\{0,1\}^* \to \{0,1\}$. Alternatively, we may view f as a set of words, namely that set of words w such that $f(w) = 1$; such sets are also called *formal languages*. There are several mathematically rigorous machine models to formalise the set of such functions that are computable or efficiently computable. The first insight is that most of these machine models lead to the same, or to closely related classes of functions. Complexity theory maps out the landscape of the resulting classes of functions. Typically the first machine model that is introduced in introductory courses are *Turing machines*. They strike a good balance between the following two (almost contradictory!) requirements that a theoretician has for these machine models:

- the model should be relatively simple, so that it is easy to show that it can be simulated by many other machine models.
- the model should be relatively powerful, so that it is easy to show that it can simulate many other machine models.

Turing machines are simple, but still the definition does not easily fit into a few lines. On the other hand, today academics are most likely to already have a very good idea of what a computer program can do (in polynomially many steps); and this coincides with what a Turing machine M can do (in polynomially many computational steps). In a nutshell, a Turing machine

- has an unboundedly large memory containing values from $\{-1, 0, 1\}$ (the symbol -1 will be called the *blank* symbol);
- has finitely many *states* Q;

- has a *read/write* head;
- has a finite transition function $\delta\colon Q \times \{-1,0,1\} \to \Sigma \times Q \times \{l,r\}$;
- has a *accept* state $y \in Q$.
- has a *start* state $s \in Q$.

Initially, the memory just contains the word $w \in \{0,1\}^*$, i.e., in the first cell there is w_1, in the second cell there is w_2, etc, and in all further memory cells there is -1, and the machine *is in state* s. Depending on its state $u \in Q$ and the tape content c under the read-write head, let $(v,d,m) := \delta(u,c)$; then

1. the machine changes to state v;
2. the tape content under the read-write head is changed from c to d,
3. the read-write tape moves one cell to the left if $m = l$, and one to the right if $m = r$.

If the machine reaches state y it accepts. Every Turing machine describes a formal language, namely the function $f\colon \{0,1\}^* \to \{0,1\}$ such that $f(w) = 1$ if and only if when running the machine on input w it eventually accepts. We also say that M *computes* f, and we then sometimes write $M(f)$ instead of $f(w)$. More generally, Turing machines can be used to describe functions f from $\{0,1\}^*$ to $\{0,1\}^*$ where $f(w)$, for a given word w, is the string that is written on the output tape when the Turing machine accepts (here we require that the machine terminates on every input after finitely many steps, and again we say that M *computes* f).

So we will pretend in the following that the reader already knows what Turing machines M are. It turns out that despite the simplicity of Turing machines, they can simulate most of the other machine models, and they can simulate any machine that humans ever constructed (even when neglecting the restriction that we one have some fixed finite maximal memory size in this universe).

In complexity theory we are interested in the number of computation steps that M needs to perform to compute $f(w)$, which corresponds to computation time. For example, we say that a Turing machine runs *in polynomial time* if the number of computation steps is in $O(|w|^k)$ for some $k \in \mathbb{N}$. The class of such functions is denoted by P.

Coding. In the main text we have met computational complexity for example for computational problems for finite graphs, whereas in the above we have only treated formal languages. But this is just a matter of coding. We first observe that we can simulate any alphabet by our alphabet $\{0,1\}$, by just grouping bits together to represent a richer alphabet. In particular, we will typically use the letter # to separate different numbers in the input. One way to represent a graph as a word is to first write the number n of vertices, followed by the symbol #, followed by a sequence of n^2 bits for the adjacency matrix.

The second most important complexity class is NP.

Definition B.1. NP (for *nondeterministic polynomial time*) stands for the class of all functions $f\colon \{0,1\}^* \to \{0,1\}$ such that there exists a polynomial-time Turing machine M and a $d \in \mathbb{N}$ such that for every $w \in \{0,1\}^*$ there exists a $a \in \{0,1\}^*$ with $|a| \in O(n^d)$ such that $f(w) = M(w\#a)$.

It is a famous open problem whether P = NP, and it is widely conjectured that P $\neq$ NP. To explain the significance of this conjecture, we need a couple of more concepts. Let $f_1, f_2 \colon \{0,1\}^* \to \{0,1\}$. A *reduction* from f_1 to f_2 is a function $g \colon \{0,1\}^* \to \{0,1\}^*$ such that $f_1(w) = f_2(g(w))$. A reduction g is *polynomial-time* if g can be computed a Turing machine that runs in polynomial time.

Definition B.2. A function $f \colon \{0,1\}^* \to \{0,1\}$ is *NP-hard* if every function g in NP has a polynomial-time reduction to f. A function is called *NP-complete* if it is in NP and NP-hard.

The class *coNP* is dual to NP: it is the class of all functions f such that $1 - f$ is in NP. There is an analogous definition for any complexity class K: a function is in co-K if $1 - f$ is in K. Clearly, every function in P is both in NP and in co-NP.

A class of finite graphs $\mathcal{C}$ is in NP if there exists a formal language in NP such that each word in the language codes a graph in $\mathcal{C}$ (say in the way we described above), and every graph in $\mathcal{C}$ is coded by some word in the language. Unlike the class P, it is possible to define the class of all graph classes in NP transparently and fully formally in a few lines (without any reference to Turing machines).

Theorem B.3 (Fagin)**.** *A class of finite graphs $\mathcal{C}$ is in NP if and only if there exists an existential second-order sentence Φ such that for every finite graph G we have*

$$G \in \mathcal{C} \text{ if and only if } G \models \Phi.$$

We do not define *existential second-order logic* here. The interested reader is referred to a textbook on finite model theory to learn more about such connections between logic and complexity theory, e.g. [78].

We now return to the question why most researchers believe that P $\neq$ NP. In order to show that P=NP is suffices to provide for *any* of the known NP-complete problems a polynomial-time algorithm. There are many NP-complete problems that are of central importance in optimisation, scheduling, cryptography, bioinformatics, artificial intelligence and many more areas. If P=NP, then this would mean a simultaneous breakthrough in all of these areas. It is fair to say that every day, thousands of researchers are directly or indirectly working on proving that P=NP (since they work on things that are related to the better understanding of some NP-complete problem). The fact that nobody has succeeded (not even came close to) is one of the reasons why we believe that P cannot be equal to NP. A world where P = NP would probably be drastically different from the world we live in. On the other hand, we also have no clue on how to possibly prove that P $\neq$ NP. Quite a bit is known about approaches to proving P $\neq$ NP that must fail (see [1]; great read, free download at `https://www.scottaaronson.com/papers/pnp.pdf`).