Lecture Notes
Universal Algebra 2
Spring term 2020

Michael Kompatscher

Thursday 30th July, 2020 9:26

Contents

1 Equational logic **2**
1.1 The equational completeness theorem . . . 3
1.2 Term rewriting systems . . . 5
1.3 Knuth-Bendix algorithm . . . 10
1.4 Exercises . . . 11

2 Commutator theory **12**
2.1 Affine and Abelian algebras . . . 12
2.2 The fundamental theorem of Abelian algebras . . . 15
2.3 The term condition commutator . . . 16
2.4 Commutator and varieties . . . 21
2.5 Nilpotent algebras . . . 22
2.6 Exercises . . . 26

3 Finitely based algebras **28**
3.1 Park's conjecture and McKenzies DPC theorem . . . 31
3.2 Exercises . . . 35

4 CSPs and Maltsev conditions **36**
4.1 Constraint satisfaction problems (CSP) . . . 36
4.2 Clone and minion homomorphisms . . . 40
4.3 Taylor operations . . . 43
4.4 Exercises . . . 48

Chapter 1

Equational logic

In the first part of the lecture we discuss ways of deciding if a given identity holds in a variety or not. So, given a fixed finite set of identities $\mathcal{E}$ (in some type τ), we define the following computational problem:

Decide($\mathcal{E}$)
INPUT: An identity $s \approx t$ (in type τ)
QUESTION: Does $\mathcal{E} \models s \approx t$? (i.e. does $s \approx t$ hold in all models of $\mathcal{E}$?)

Our goal is to find an algorithm that solves the above problem.[1] Unfortunately there is no algorithm that solves Decide($\mathcal{E}$) in general - in fact, even for some sets $\mathcal{E}$ that define varieties of groups or semi-groups, the problem is undecidable (see e.g. [1] Chapter 12). However for other examples there are straightforward algorithms:

Example 1.1. Let $\mathcal{E}$ be the empty set. Then $\mathcal{E} \models s \approx t$, if and only if the two terms are equal (otherwise they induce different term operations in the totally free algebra). So return 'Yes' if and only if $s = t$ and 'No' otherwise.

Example 1.2. Let $\mathcal{E}$ be the set of identities in type $\tau = (+, 0, -, \cdot)$ that axiomatize the variety of commutative rings, so $\mathcal{E} = \{x + y \approx y + x, 0 + x \approx x, x \cdot (y + z) \approx x \cdot y + x \cdot z, \ldots\}$. Then Decide($\mathcal{E}$) is the problem of deciding whether an input identity $s(x_1, \ldots, x_n) \approx t(x_1, \ldots, x_n)$ holds in every commutative ring.

Clearly $\mathcal{E} \models s \approx t$ if and only if $\mathcal{E} \models s - t \approx 0$. By using the identities in $\mathcal{E}$ we can rewrite the left side into an equivalent term that is a sum of monomials (e.g. $a_1x_1 + a_{11}x_1^2 + a_2x_2 + a_{12}x_1x_2 + \ldots$). We output 'Yes' if all of the coefficients of all monomials are 0 and 'No' otherwise. This algorithm is correct, since in the 'No' case the two terms induce different operations in the ring $(\mathbb{Z}, +, 0, -, \cdot)$.

The algorithm described in Example 1.2 uses the ring identities in $\mathcal{E}$ as 'rewriting rules' to rewrite an arbitrary term into an equivalent 'normal form', namely the sum of monomials. We can then check if $\mathcal{E} \models s \approx t$, by simply checking if s and t have the same normal form. We are going to formalize this *term rewriting* approach and investigate for which identities/rewriting rules $\mathcal{E}$ it works.

[1] Decide($\mathcal{E}$) is often called the 'word problem' in literature, but we remark that 'word problem' has also other, slightly different meanings attached to it.

After that we discuss the *Knuth-Bendix algorithm*, a procedure to modify the sets of identities $\mathcal{E}$ into an equivalent set $\mathcal{F}$ for which the term rewriting approach works.

1.1 The equational completeness theorem

In the first lecture we are going to give a syntactic characterization of $\mathcal{E} \models s \approx t$, which will be the base of our algorithm.

Let us recall some definitions from UA1 (see also [2] Section 4.3-4.4).

Definition 1.3. Let τ be a similarity type. Then

- A *term* is an element of the *totally free algebra* in $\mathbf{F}_\tau(x_1, x_2, \ldots)$. For short we are going to use the notation $\mathbf{F} = \mathbf{F}_\tau(x_1, x_2, \ldots)$.
- An *identity* (of type τ) is an ordered pair $(s,t) \in \mathbf{F}^2$. We will also write $s \approx t$ for the identity (s,t). Note that according to this definition $s \approx t$ and $t \approx s$ are two different identities.
- An algebra $\mathbf{A}$ *satisfies* an identity $s \approx t$ (short $\mathbf{A} \models s \approx t$) iff for all $a_1, \ldots, a_n \in A$: $s^{\mathbf{A}}(a_1, \ldots, a_n) = t^{\mathbf{A}}(a_1, \ldots, a_n)$

The relation $\models$ induces a Galois connection between classes of algebras $\mathcal{K}$ and sets of identities $\mathcal{E}$ (of type τ), with the operators

$$\mathrm{Id}(\mathcal{K}) = \{s \approx t \colon \mathcal{K} \models s \approx t\}$$
$$\mathrm{Mod}(\mathcal{E}) = \{\mathbf{A} \colon \mathbf{A} \models \mathcal{E}\}$$

We saw in Universal Algebra 1 that the Galois closure $\mathrm{Mod}(\mathrm{Id}(\mathcal{K}))$ is $\mathsf{HSP}(\mathcal{K})$, the variety generated by $\mathcal{K}$.

The Galois-closed sets $\mathrm{Id}(\mathrm{Mod}(\mathcal{E})) \subseteq F^2$ on the identity side are called *equational theories*. For short, we write $\mathcal{E} \models s \approx t$ (and say $\mathcal{E}$ *semantically* implies $s \approx t$) if $(s \approx t) \in \mathrm{Id}(\mathrm{Mod}(\mathcal{E}))$, or in other words, if $s \approx t$ is satisfied in all models of $\mathcal{E}$.

Observation 1.4. An equational theory $\mathcal{D} \subseteq \mathbf{F}^2$ has three natural properties (which we discuss for the example $\tau = \{\cdot\}$)

1. $\mathcal{D}$ is an equivalence relation
2. $\mathcal{D}$ is even a *congruence* of $\mathbf{F}$ (since $\mathcal{D} \models s \approx t$ and $\mathcal{D} \models u \approx v$ implies $\mathcal{D} \models s \cdot u \approx t \cdot v$)
3. $\mathcal{D}$ is invariant under substitutions. If for instance $\mathcal{D} \models (x_1 \cdot x_2) \approx (x_1 \cdot x_1)$ then, after substituting variables via the map $x_1 \mapsto (x_1 \cdot x_2), x_2 \mapsto (x_3 \cdot x_1)$ we obtain $\mathcal{D} \models ((x_1 \cdot x_2) \cdot (x_3 \cdot x_1)) \approx ((x_1 \cdot x_2) \cdot (x_1 \cdot x_2))$.

Note that a substitution can be seen a map $\theta \colon \{x_1, x_2, \ldots\} \to \mathbf{F}$. By the properties of the totally free algebra $\mathbf{F}$, this map θ uniquely extends to a homomorphism $\theta \colon \mathbf{F} \to \mathbf{F}$. Thus (3) says that $\mathcal{D}$ is closed under endomorphisms $\theta \in \mathrm{End}(\mathbf{F})$. This motivates the following definition.

Definition 1.5. Let $\mathbf{A}$ be an algebra. We call a congruence $\mathcal{D} \in \mathrm{Con}(\mathbf{A})$ *fully invariant*, if for every endomorphism $\theta \in \mathrm{End}(\mathbf{A})$ we have $(s,t) \in \mathcal{D} \Rightarrow (\theta(s), \theta(t)) \in \mathcal{D}$.

We next show that the properties (1)-(3) from Observation 1.4 already completely characterize equational theories:

Proposition 1.6. *A binary relation $\mathcal{D} \subseteq \mathbf{F}^2$ is an equational theory if and only if it is a fully invariant congruence of* $\mathbf{F}$.

Proof. Every equational theory is a fully invariant congruence of $\mathbf{F}$. This can be shown as in Observation 1.4.

For the opposite direction, let $\mathcal{D}$ be a fully invariant congruence of $\mathbf{F}$. We then need to prove that $\mathcal{D} = \mathrm{Id}(\mathrm{Mod}(\mathcal{D}))$. Clearly $\mathcal{D} \subseteq \mathrm{Id}(\mathrm{Mod}(\mathcal{D}))$, so it is enough to prove the inclusion $\mathcal{D} \supseteq \mathrm{Id}(\mathrm{Mod}(\mathcal{D}))$.

Consider the quotient $\mathbf{A} = \mathbf{F}/\mathcal{D}$. This quotient is well-defined, since $\mathcal{D}$ is a congruence of $\mathbf{F}$. We claim that an identity is in $\mathcal{D}$, if and only if it satisfied in $\mathbf{A}$. Suppose $s(x_1, \dots, x_n) \approx t(x_1, \dots, x_n) \in \mathcal{D}$. Since $\mathcal{D}$ is fully invariant, this is equivalent to $s(u_1, \dots, u_n) \approx t(u_1, \dots, u_n) \in D$, for all terms $u_1, \dots, u_n \in \mathbf{F}$. By definition of $\mathbf{A}$, this is equivalent to $s^{\mathbf{A}}(u_1, \dots, u_n) = t^{\mathbf{A}}(u_1, \dots, u_n)$, for all $u_1, \dots, u_n \in A$. In other words, it is equivalent to $\mathbf{A} \models s(x_1, \dots, x_n) \approx t(x_1, \dots, x_n)$, which is what we claimed.

Since $\mathbf{A}$ is a model of $\mathcal{D}$, we obtain $\mathrm{Id}(\mathrm{Mod}(\mathcal{D})) \subseteq \mathrm{Id}(\{\mathbf{A}\}) = \mathcal{D}$. □

By Proposition 1.6, $\mathcal{E} \models s \approx t$, if $s \approx t$ lies in the fully invariant congruence generated by $\mathcal{E}$. If we only look at the identities that follow from directly applying one of the identities of $\mathcal{E}$ to a term (or a subterm of it) we get so called *immediate consequences* of $\mathcal{E}$:

Example 1.7. Let $l = (x_1 \cdot x_2) \approx (x_2 \cdot x_1) = r$ be the identity describing the commutativity law. Then, the identity $u = x_3 \cdot (x_1 \cdot (x_1 \cdot x_2)) \approx x_3 \cdot ((x_1 \cdot x_2) \cdot x_1) = v$ is an immediate consequence of $l \approx r$, since it follows from applying the commutativity law to the subterm $(x_1 \cdot (x_1 \cdot x_2))$ of u. Note that this subterm is equal to $\theta(l)$ for the substitution $\theta(x_1) = x_1, \theta(x_2) = x_1 \cdot x_2$.

We include a more formal definition (see also [1], Section 6.3):

Definition 1.8. Recall from UA1 that terms t can be represented by a term tree. If we label the vertices of a term tree t by sequences of natural numbers $a \in \omega^{<\omega}$ in a natural way (see Figure 1.1), every such label is called a *valid address of* t. The *subterm at address* a, or $t[a]$, is the term that corresponds to the subtree with root a.

Assume that $t[a] = s$. By $t[a \colon s \to r]$ we denote the term that we obtain by replacing the subterm s by r. Then every identity of the form $u \approx u[a \colon \theta(l) \to \theta(r)]$ for some $\theta \in \mathrm{End}(\mathbf{F})$ is called an *immediate consequence* of $l \approx r$.

We finish this section by proving that all identities in the equational theory generated by $\mathcal{E}$ lie in the equivalence relation generated by immediate consequences of $\mathcal{E}$. This result was shown by Birkhoff, and is also known as *equational completeness theorem* (in analogy to Gödel's completeness theorem in first order logic).

Definition 1.9. We say $\mathcal{E}$ *syntactically* implies $s \approx t$, and write $\mathcal{E} \vdash s \approx t$, if there is a sequence of terms $s = u_1, \dots, u_n = t$, such that for every $i = 1, \dots, n-1$, either (u_i, u_{i+1}), or (u_{i+1}, u_i) is a immediate consequence of an identity from $\mathcal{E}$.

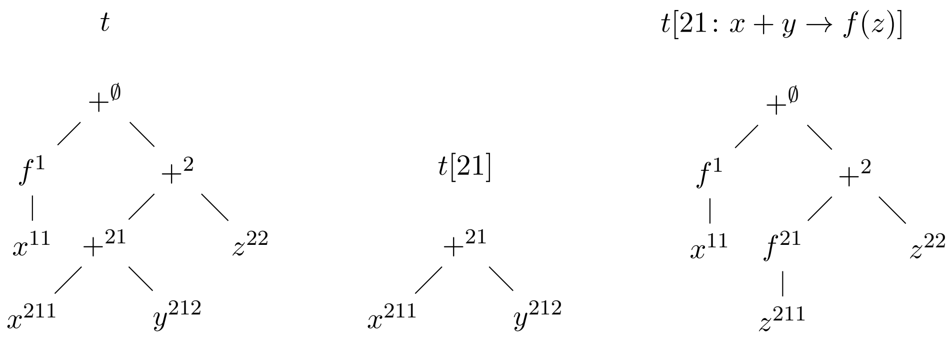

Figure 1.1: The term tree of $t = f(x) + ((x + y) + z)$; labelled its valid by addresses; its subterm $t[21] = (x + y)$, and the substitution $t[21\colon xy \to f(z)]$

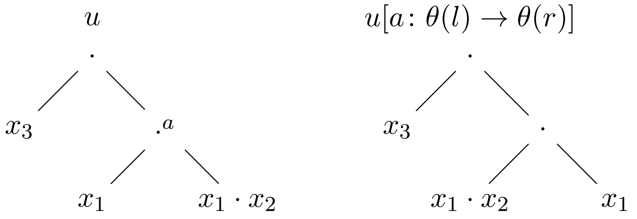

Figure 1.2: Example 1.7 revisited: $x_3 \cdot (x_1 \cdot (x_1 \cdot x_2)) \approx x_3 \cdot ((x_1 \cdot x_2) \cdot x_1)$ is an immediate consequence of $x_1 \cdot x_2 \approx x_2 \cdot x_1$ with $a = 2$ and $\theta(x_1) = x_1, \theta(x_2) = x_1 \cdot x_2$.

Theorem 1.10 (Birkhoff's equational completeness theorem)**.** *$\mathcal{E} \models s \approx t$ if and only if $\mathcal{E} \vdash s \approx t$.*

Proof. It follows straightforward from Observation 1.4 that $\mathcal{E} \vdash s \approx t$ implies $\mathcal{E} \models s \approx t$. For the opposite direction, we need to show that $\{s \approx t \colon \mathcal{E} \vdash s \approx t\}$ is a fully invariant congruence. This is left as a homework assignment (Exercise 1.7). By Proposition 1.6 $\{s \approx t \colon \mathcal{E} \vdash s \approx t\}$ is an equational theory, and since it contains $\mathcal{E}$, it also needs to contain all identities that semantically follow from $\mathcal{E}$. □

1.2 Term rewriting systems

Term rewriting systems are a tool that sometimes allows us to find an algorithm for Decide($\mathcal{E}$). In term rewriting algorithms we regard the identities in $\mathcal{E}$ as a set of rules on how to rewrite terms: For an input term s, if there is an immediate consequence $s \approx s'$ of $\mathcal{E}$, we replaces s by s', and repeat this process as long as possible. For certain rewriting rules $\mathcal{E}$ this algorithm terminates after finitely many steps and computes a unique term that is equivalent to s (its 'normal form'). If this is the case, we can solve Decide($\mathcal{E}$) by rewriting both s and t into their normal forms, and then check if they are equal (as in Example 1.2).

In the practicals we discussed some examples, in which this algorithm works and does not work (see Exercises 1.1, 1.2, 1.3). In this lecture we are going to formalize the observations we made.

We will use the following notation:

Definition 1.11. For a set of identities $\mathcal{E}$ (of type τ), we define a digraph $D(\mathcal{E})$ that has $\mathbf{F}$ (the terms of type τ) as vertex set and has an edge $s \to t$ if and only if $s \approx t$ is an immediate consequence of some identity $(l \approx r) \in \mathcal{E}$.

Definition 1.12. Let D be a digraph with edge relation $\to$.

- A vertex s is called *terminal* if there is no t with $s \to t$.
- We write $s \to^* t$, if (s,t) is in the transitive, reflexive closure of $\to$
- We write $s \leftrightarrow^* t$, if (s,t) is in the symmetric, transitive, reflexive closure of $\to$.

By definition $\mathcal{E} \vdash s \approx t$ if and only if $s \leftrightarrow^* t$ in $D(\mathcal{E})$.

Definition 1.13. We say that a digraph D

- is *finitely terminating* if it contains no infinite directed path $t_0 \to t_1 \to t_2 \to \cdots$ [2]
- is *normal*, if for all $s, t \in D$ such that $s \leftrightarrow^* t$ there is a $u \in D$ with $s \to^* u$ and $t \to^* u$.
- is *convergent*, if it is normal and finitely terminating

The following lemma shows that if $D(\mathcal{E})$ is convergent, then every vertex is is connected to a unique terminal vertex. And in fact it is connected to such vertex via a directed path.

Lemma 1.14. *Let D be a digraph that is convergent. Then*

1. *For every $t \in D$ there is a* unique *terminal vertex $NF(t)$, such that $t \to^* NF(t)$ (this is the* normal form *of t).*
2. *$s \leftrightarrow^* t$ if and only if $NF(s) = NF(t)$*

Proof. 1. Let $t \in D$. Since D is finitely terminating, there is a terminal vertex a, such that $t \to^* a$. For the uniqueness, let b be another terminal vertex, such that $t \to^* b$. So clearly $a \leftrightarrow^* b$. By normality, there is a u with $a \to^* u, b \to^* u$. Since a and b are terminal, this can only happen if $a = b = u$.

2. By normality $s \leftrightarrow^* t$ implies that there is an u with $s \to^* u, t \to^* u$. By the uniqueness in (1) we obtain $NF(u) = NF(s) = NF(t)$.

□

Note that by Lemma 1.14, the term rewriting algorithm described at the beginning of this section works, if $D(\mathcal{E})$ is convergent. In order to find out, for which sets of identities $\mathcal{E}$ this is true, we need to answer two questions: When is $D(\mathcal{E})$ is finitely terminating and when is it normal?

We start by discussing finite termination, by looking at some examples.

Example 1.15. Is $D(\mathcal{E})$ finitely terminating?

1. $\mathcal{E} = \{f(f(x)) \approx f(x)\}$ Yes: the number of f's in every immediate consequence decreases
2. $\mathcal{E} = \{x \approx x \cdot x\}$ No: $x \to xx \to (xx)(xx) \to \cdots$

[2] Note that this also forbids cycles or loops

3. $\mathcal{E} = \{x \cdot x \approx x\}$ Yes: the number of variables in every immediate consequence decreases

4. $\mathcal{E} = \{(x \cdot x) \cdot y \approx (y \cdot y)\}$ No: there is a loop $(x \cdot x) \cdot (x \cdot x) \to (x \cdot x) \cdot (x \cdot x)$

In the finitely terminating cases 1 and 3 in Example 1.15, we can see that for all pairs $s \to t$ in $D(\mathcal{E})$, the left side s is in some sense 'more complex' than t. This can be measured by so called reduction orders:

Definition 1.16. A strict (partial) order $<$ on the set of terms $\mathbf{F}$ is called a *reduction order* if

1. $s > t$ implies that for all $\theta \in \mathrm{End}(\mathbf{F})$: $\theta(s) > \theta(t)$
2. $s > t$ implies that, for all u with $u[a] = s$ we have $u > u[a\colon s \to t]$
3. $<$ is *well-founded*, so there is no infinite chain $t_1 > t_2 > t_3 > \cdots$

Observation 1.17. If there is a reduction order $<$ such that $l > r$ for all identities $(l \approx r) \in \mathcal{E}$, then $D(\mathcal{E})$ is finitely terminating. For the proof, note that (1) and (2) imply that $u > v$ for all immediate consequences of identities in $\mathcal{E}$. From (3) it follows that there is no infinite chain $t_1 \to t_2 \to t_3 \to \cdots$ in $D(\mathcal{E})$.

- Example 1.15 (1) is compatible with the reduction order defined by $s < t$ if and only if the number of f symbols in s is smaller than the number of f's in t.
- A reduction order compatible with Example 1.15 (3) is the order that sets $s < t$, if for all variables x_i, the number of x_i appearing is s is less or equal to the number for t and for some index i it is strictly less.
- Ordering terms by the total number of variables is NOT a reduction order, since it is not compatible with substitutions (as you can see in Example 1.15 (4), and the substitution $y \mapsto x \cdot x$).

Other reduction orders may take into account (weighted) number of function symbols and variables, but also their appearance/order within a term (see Exercise 1.3). A big family of rewriting orders to pick for a given $\mathcal{E}$ are the *Knuth-Bendix-orders*, which we are not going to define here (see [1], Section 13.7).

Remark 1.18. According to our definitions, the identity $x + y \approx y + x$ prohibits finite termination, and is also not compatible with any reduction order. So also $D(\mathcal{E})$ for Example 1.2 is not finitely terminating. This can be fixed by weakening the definitions and considering *pre-orders* instead of partial orders (see [1], Section 13.7). We are however not going to do so in our simplified setting.

We next discuss the normality of a digraph. By the following lemma, a finitely terminating digraph is normal, if and only if it satisfies the definition of normality in specific sub-cases.

Theorem 1.19. *Let D be a finitely terminating digraph. Then TFAE*

1. *D is normal.*
2. *D is* confluent: *For all $r, s, t \in D$ such that $r \to^* s$ and $r \to^* t$ there is an $u \in D$ with $s \to^* u$, $t \to^* u$.*
3. *D is* locally confluent: *For all $r, s, t \in D$ such that $r \to s$ and $r \to t$ there is an $u \in D$ with $s \to^* u$, $t \to^* u$.*

Proof. (1)→(2)→(3) holds by definition.

For (2)→(1), assume that there are vertices $a \leftrightarrow^* b$ such that there is no u with $a \to^* u$ and $b \to^* u$. By definition of $\leftrightarrow^*$ there are elements $t_0, \dots, t_{2n} \in D$, such that $a = t_0 \leftarrow^* t_1 \to^* t_2 \leftarrow^* \cdots \to^* t_{2n} = b$. Without loss of generality, let a, b be a pair such that the number n is minimal. By minimality there exists an u' with $a \to^* u'$ and $t_{2n-1} \to^* u'$. But by confluence there is also an u with $u' \to^* u$ and $b \to^* u$. Thus $a \to^* u' \to^* u$ and $b \to^* u$ - contradiction!

For (3)→(2), let D be locally confluent, and let P be the set of vertices of D, such that $v \in P$ if v has directed paths to more than one terminal vertex. If $P = \emptyset$, every vertex of D has (by finite termination) a directed path to exactly one terminal vertex. So in this case D is confluent.

Now assume that $P \neq \emptyset$. Then there is an element $v \in P$ such that for every $u \in P$, $v \to^* u$ implies $v = u$ (otherwise D would contain an infinite directed path, contradicting finite termination). Let s, t be two distinct terminal vertices with $v \to^* s$ and $v \to^* t$. By minimality of v there are distinct s_0, t_0 with $v \to s_0 \to^* s$, and $v \to t_0 \to^* t$. By local confluence, there is an u with $s_0 \to^* u$, $t_0 \to^* u$. Since D is finitely terminating we can pick u to be terminal. Since $s_0, t_0 \notin P$ we get $s = u$ and $u = t$, which is a contradiction to the assumption that s and t are distinct. □

Local confluence is an easier property to check than normality. In cases, where the rewriting affects 'disjoint' subterms, we always get local confluence (see Figure 1.3). The only obstacles to local confluence are immediate consequences $r \to s$ and $r \to t$ that affect 'overlapping' subterms. We give some examples.

Example 1.20. The example $\mathcal{E} = \{f(f(x)) \approx f(x)\}$ is locally confluent, while $\mathcal{E} = \{f(f(x)) \approx g(x)\}$ is not (see Exercise 1.2).

Example 1.21. Let $\mathcal{E} = \{x \cdot (y \cdot y) \approx x \cdot y\}$.

The term $s = (x(yy))(x(zz))$ rewrites to both $(xy)(x(zz))$ and $(x(yy))(xz)$. Both of these terms rewrite to $(xy)(xz)$. Note that corresponding immediate consequences affected the disjoint subterms $s[1] = (x(yy))$ respectively $s[2] = (x(zz))$, therefore we had no problems.

However $D(\mathcal{E})$ is not locally confluent: for this consider the the term $t = (xy)((uu)(uu))$. By applying the identity on the subterm $t[2]$ we obtain $(xy)((uu)(uu)) \to (xy)((uu)u)$, which is terminal. If we apply the identity on $t[\emptyset]$ we can only rewrite to $(xy)((uu)(uu)) \to (xy)(uu) \to (xy)u$.

The term $(xy)((uu)(uu))$ in Example 1.21 is a witness that $D(\mathcal{E})$ is not locally confluent. However it is not a 'minimal' example, since the same problem already appears when starting with the term $x((uu)(uu))$. We call the two terms $x((uu)u)$ and xu that result from a 'minimal' such example *critical pairs*. We give a more formal definition of critical pairs below.

Definition 1.22. Let t and s be two terms. A *unifier* of t and s is an endomorphism $\theta \in \mathrm{End}(\mathbf{F})$, such that $\theta t = \theta s$. A *most general unifier* of t and s is a unifier $\alpha \in \mathrm{End}(\mathbf{F})$, such that for all other unifiers θ, we have that $\theta = \beta \circ \alpha$, for some β.

If a unifier exists, there exists also a most general unifier, which is unique (up to renaming variables). We are not giving a proof, it can be found in [1], Section 13.1.

Example 1.23. Let $t(x, y) = (x \cdot x) \cdot (y \cdot y)$ and $s(u, v) = u \cdot (u \cdot v)$. Then any unifier θ of t and s must satisfy $\theta(u) = \theta(x) \cdot \theta(x)$ and $\theta(y) = \theta(u) = \theta(v)$. If we set $\theta(x) = x$, this gives us the most general unifier $\theta(t) = (xx)((xx)(xx)) = (xx)((xx)(xx))$.

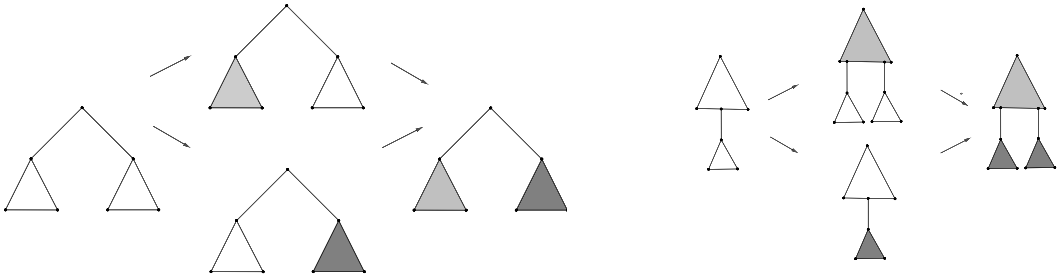

Figure 1.3: If the immediate consequences affect disjoint subterms, they are always confluent.

Example 1.24. The terms $t(x,y) = (x \cdot y) \cdot x$ and $s(u,v) = u \cdot (u \cdot v)$, do not have a unifier. If there was a unifier θ, it would need to satisfy $\theta(u) = \theta(x \cdot y) = \theta(x) \cdot \theta(y)$, but also $\theta(x) = \theta(u) \cdot \theta(v)$. This implies $\theta(u) = (\theta(u) \cdot \theta(v)) \cdot \theta(y)$, which is however not possible for any $\theta(u) \in \mathbf{F}$.

Definition 1.25. Let $l_1 \approx r_1$ and $l_2 \approx r_2$ be two identities in $\mathcal{E}$. Also let us assume that they contain no common variables (otherwise rename the variables of one identity). Let l' be a subterm of l_1, and let θ be the most general unifier of l' and l_2. Then the pair $(\theta r_1, \theta(l_1)[a\colon \theta l' \to \theta r_2])$ is called a *critical pair*. This corresponds to the output we get by applying the rewriting rules $l_1 \approx r_1$ and $l_2 \approx r_2$ to θl_1, respectively its subterm $\theta l' = \theta l_2$.

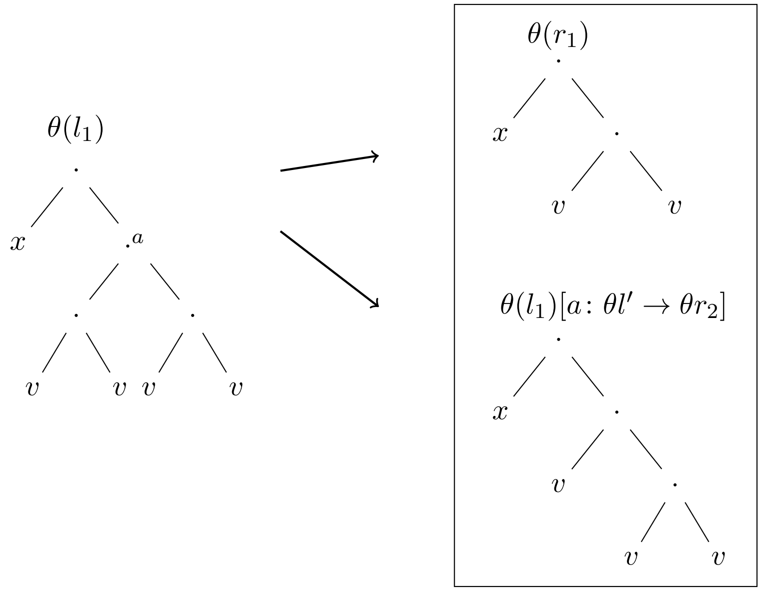

Figure 1.4: The critical pair in Example 1.26

Example 1.26. We discuss the definition of critical pair for Example 1.21. In both immediate consequences we apply the rule $x(yy) \approx xy$. So $l_1(x,y) = x(yy)$ and (after renaming variables) $l_2(u,v) = u(vv)$. The subterm l' of l_1 is equal to yy. The most general unifier of $l' = yy$ and

$l_2 = u(vv)$ is given by the map $\theta(x) = x$, $\theta(v) = v$ and $\theta(y) = \theta(u) = vv$. Thus the terms $\theta(r_1) = x(vv)$ and $\theta(l_1)[a\colon \theta l' \to \theta r_2] = x((vv)v)$ form a critical pair.

Without a proof we claim the following:

Theorem 1.27 (Theorem 3.1. in Chapter 13 of [1])**.** *Let $\mathcal{E}$ be a set of identities. Then $D(\mathcal{E})$ is locally confluent, if and only if it is locally confluent at all critical pairs.*

Note that, if $\mathcal{E}$ is finite, then there are only finitely many critical pairs. Thus by just checking local confluence for finitely many pairs, we can check if $D(\mathcal{E})$ is locally confluent.

1.3 Knuth-Bendix algorithm

The Knuth-Bendix algorithm is an algorithm that takes as input a set of identities $\mathcal{E}$ and tries to modify them in a way to obtain an equivalent set $\mathcal{F}$, such that $D(\mathcal{F})$ is convergent. The algorithm comes with a fixed reduction order $<$ as a parameter.

Knuth-Bendix algorithm
INPUT: A set of identities $\mathcal{E}$

- STEP 1: Check if for all identities $l \approx r \in \mathcal{E}$ either $l > r$ or $l < r$ holds.
    - If this is not the case, halt and output 'Failure'.
    - Otherwise rearrange the identities in a way such that $l > r$ for all $l \approx r \in \mathcal{E}$.
- STEP 2: Compute all critical pairs (s, t) of $\mathcal{E}$.
    - If $D(\mathcal{E})$ is locally confluent at all of them, halt and output $\mathcal{F} := \mathcal{E}$.
    - Otherwise pick a pair (s, t) that is not locally confluent. Compute two terminal vertices s_0, t_0 such that $s \to^* s_0$ and $t \to^* t_0$. Then add $s_0 \approx t_0$ to the set of identities $\mathcal{E}$ and return to Step 1.

Note that there are three possible outcomes:

1. The algorithm fails Step 1, since $\mathcal{E}$ is incompatible with the reduction order.
2. The algorithm enters an infinite loop, adding more and more identities to $\mathcal{E}$.
3. The algorithm stops after finitely many steps, and outputs a set of identities $\mathcal{F}$.

In case (3) we clearly obtain what we wanted, since then $D(\mathcal{F})$ is locally finite by Step (1), and locally confluent at all critical pairs by Step (2). Thus $D(\mathcal{F})$ is convergent. However both case (1) and case (2) can appear, even in situations where $\mathcal{E}$ is equivalent to some convergent system $\mathcal{F}$. In practice such problems can be sometimes overcome by wisely picking a reduction order, or specifying details of the algorithm (e.g. how to pick terminal vertices in Step 2). However we remind that Decide($\mathcal{E}$) is undecidable for some $\mathcal{E}$; for such $\mathcal{E}$ the Knuth-Bendix algorithm will always result in case (1) or (2).

1.4 Exercises

Exercise 1.1. Let $\mathcal{E} = \{f(x) \approx f(f(x))\}$. Draw $D(\mathcal{E})$, and explain why the term rewriting algorithm does not work. What is the problem and how can it be fixed?

Exercise 1.2. Let $\mathcal{E} = \{f(f(x)) \approx g(x)\}$.

1. Show that (with respect to $\mathrm{Id}(\mathrm{Mod}(\mathcal{E}))$) every term $t(x)$ is equivalent to a term of the form $g^n(x)$ or $g^n f(x)$
2. Show that the term rewriting algorithm does not work in computing the above normal from. What seems to be the problem?
3. Add an identity to $\mathcal{E}$ to fix this problem.

Exercise 1.3. Let $\mathcal{E} = \{f(x+y) \approx x + f(y), f(f(x)) \approx f(x)\}$.

1. Show that (with respect to $\mathrm{Id}(\mathrm{Mod}(\mathcal{E}))$) every term is equivalent to a term, which has f symbols only at its branches.
2. Show that every term is equivalent to exactly one such term.
3. Show that the 'term rewriting' algorithm solves $\mathrm{Decide}(\mathcal{E})$.
4. Find a rewriting order compatible with $\mathcal{E}$.

Exercise 1.4. Find an example of a digraph that is locally confluent, but not confluent.

Exercise 1.5. Find a convergent rewriting system equivalent to the identity $(x \cdot y) \cdot (y \cdot z) \approx y$ (Hint: Knuth-Bendix will produce 2 additional rules.)

Exercise 1.6. Consider the two rewriting rules

$$\mathcal{E} = \{f(x) + (y+z) \approx x + (f(f(y)) + z),\ f(x) + (y + (z+w)) \approx x + (y + (z+w))\}.$$

Show that $D(\mathcal{E})$ is finitely terminating.

Exercise 1.7. Finish the proof of the equational complete theorem (Theorem 1.10). For simplicities sake, you may assume that τ consists only of a binary operation $\cdot$.

Exercise 1.8. Let $\mathcal{E}_1 = \{(xy)z \approx x(yz)\}$, $\mathcal{E}_2 = \mathcal{E}_1 \cup \{(xy)z \approx x(yz), xx \approx x\}$ and $\mathcal{E}_3 = \mathcal{E}_2 \cup \{(xy)z \approx xz, x(yz) \approx xz\}$. (We remark that $\mathcal{E}_1$ defines the variety of all semigroups, $\mathcal{E}_2$ bands (idempotent semigroups), and $\mathcal{E}_3$ the so called rectangular bands and all $D(\mathcal{E}_i)$ are finitely terminating).

- Show that $D(\mathcal{E}_1)$ and $D(\mathcal{E}_3)$ are locally confluent, but $D(\mathcal{E}_2)$ is not.
- (*) Try to argue that our version of Knuth-Bendix algorithm will enter an infinite loop for $\mathcal{E}_2$.

Exercise 1.9. Try to come up with a convergent rewriting system for the variety of groups (Hint: Think of a normal form of group terms first; then write down the rewriting system. Don't get stuck in technical details).

Chapter 2

Commutator theory

Commutative groups, vector spaces and modules are examples of algebras that are well understood in both their structural properties (e.g. classification results), but also computational ones (e.g. solving equations by Gauss elimination). This stems mainly from the fact that all their term and polynomial operations are affine. More generally, many results in group theory are based on measuring how close a group or its subgroups are to being commutative (think of definitions like *commutator subgroups*, *centralizers*, *nilpotent* and *solvable* groups...).

Commutator theory is the part of universal algebra that tries to generalize such group theoretic definitions and results to arbitrary algebras. Its development started in the 1970ies with results of Gumm and Smith. A standard reference for commutator theory is the book of Freese and McKenzie that discusses commutator theory in congruence modular varieties [3]; for such varieties commutator theory works particularly well.

In this chapter we are going to discuss some basics of commutator theory. We start with the definition of Abelianness for general algebras and prove Herrmann's fundamental theorem, which characterizes Abelian algebras in congruence modular varieties.

After that we define the commutator of two congruences and show that it generalizes the commutator subgroup (in groups) and the product of two ideals (in commutative rings). We further show how the commutator can be used to characterize congruence distributive varieties.

All presented results hold for algebras from congruence modular varieties; we are however only going to discuss algebras with a Maltsev operations, which is a stronger assumption that makes many proofs easier. Most of the results presented here can also be found in Chapter 7.2-7.4 of Bergman's book [2] (however in different notation).

2.1 Affine and Abelian algebras

We start by recalling some definitions from the Universal Algebra 1 lecture:

Definition 2.1. Let $\mathbf{A}$ be an algebra with universe A. A *polynomial operation* of $\mathbf{A}$ is an operation of the form $p(x_1, \dots, x_n) = t(x_1, \dots, x_n, a_1, \dots, a_k)$ where $t \in \mathsf{Clo}(\mathbf{A})$ is a term operation of $\mathbf{A}$ and $a_1, \dots, a_k \in A$. The set of all polynomial operations of an algebra forms a clone, which is sometimes denoted by $\mathrm{Pol}(\mathbf{A})$. [1]

[1] Attention: this should not be confused with the polymorphism clone $\mathrm{Pol}(\mathbb{A})$ of a relational structure $\mathbb{A}$ in Definition 4.7

Two algebras $\mathbf{A}$ and $\mathbf{A}'$ on the same universe are called *polynomially equivalent* if they have the same clone of polynomial operations.

So a polynomial operation is an operation that can be built from variables, constants from A, and basic operations of $\mathbf{A}$. If for instance $\mathbf{A} = (\{0, 1\}, +, \cdot)$ is the 2-element ring; then $p(x, y) = (1 + x) \cdot (y \cdot y + x) + 1$ is a polynomial operation of $\mathbf{A}$.

Next, recall from UA1 that we represented R-modules as algebras $\mathbf{A} = (A, +, 0, -, (r)_{r\in R})$, where $(A, +, 0, -)$ is the underlying group, and every ring element $r \in R$ corresponds to the unary scalar multiplication $r(x) = r \cdot x$. It is easy to see that every polynomial operation $p(x_1, \dots, x_n)$ of an R-module $\mathbf{A}$ is an affine combination of variables, so $p(x_1, \dots, x_n) = \sum_{i=1}^{n} \alpha_i x_i + c$ with $\alpha_i \in R$ and $c \in A$. This motivates the following definition:

Definition 2.2. We call an algebra $\mathbf{A}$ *affine* if it is polynomially equivalent to an R-module (for some ring R with identity).

Note that according to this definition also every commutative group is affine (by picking $R = \mathbb{Z}$ as the ring; and defining the scalar multiplication $n \cdot x := \underbrace{x + x + \cdots + x}_{n \text{ times}}$).

Our goal in this section is to characterize affine algebras in more ”‘universal algebraic”’ terms. We start by observing that affine algebras have two nice properties, namely that they they are *Abelian*, and that have a unique *Maltsev* polynomial, which is *central.*

Definition 2.3. A *Maltsev* (also *Mal'cev*) operation $m(x, y, z)$ is a ternary operation satisfying the identities $m(y, x, x) \approx m(x, x, y) \approx y$.

Maltsev operations were already defined in the UA1 lecture. Note that every group has a Maltsev term: $m(x, y, z) = xy^{-1}z$. Similarly also every ring and module has a Maltsev term $m(x, y, z) = x - y + z$.

Definition 2.4. Let $\mathbf{A}$ be an algebra. A polynomial $p \colon A^3 \to A$ of $\mathbf{A}$ is called *central* if it satisfies

$$p(f(x_1, \dots, x_n), f(y_1, \dots, x_n), f(z_1, \dots, z_n)) = f(p(x_1, y_1, z_1), \dots, p(x_n, y_n, z_n)),$$

for all basic operations f of $\mathbf{A}$. Note that this is equivalent to saying that p is a homomorphism $\mathbf{A}^3 \to \mathbf{A}$, or to the characterization in Figure 2.1.

$$\begin{array}{ccc|cc} x_1 & y_1 & z_1 & \to & p(x_1, y_1, z_1) \\ x_2 & y_2 & z_2 & \to & p(x_2, y_2, z_2) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ x_n & y_n & z_n & \to & p(x_n, y_n, z_n) \\ \hline \downarrow & \downarrow & \downarrow & & \downarrow \\ f(\bar{x}) & f(\bar{y}) & f(\bar{y}) & \to & \star \end{array}$$

Figure 2.1: If p is central, the result $\star$ is the same, no matter if we evaluate the rows (under p) or the columns (under f) first.

Observation 2.5. Let $\mathbf{A}$ be affine. Then $\mathbf{A}$ has a unique Maltsev polynomial operation $m(x, y, z)$. Moreover $m(x, y, z)$ is central.

Proof. By definition $\mathbf{A}$ is is polynomially equivalent to an R-module $\mathbf{A'} = (A, +, 0, -, (r)_{r \in R})$. Therefore it has a Maltsev polynomial $m(x, y, z) = x - y + z$. We claim that this is the only ternary polynomial satisfying the Maltsev identities.
So let $p(x, y, z)$ be some other polynomial of $\mathbf{A}$ satisfying the Maltsev identities. Since $\mathbf{A}$ is affine, we know that $p(x, y, z) = \alpha x + \beta y + \gamma z + c$ for some scalars $\alpha, \beta, \gamma \in R$ and $c \in A$. By the equality $p(0, 0, 0) = 0$, we obtain $c = 0$. For all x we further have $x = p(x, 0, 0) = \alpha x$; analogously we can prove $\gamma x = x$. At last note that for all $x \in A$: $p(x, x, 0) = x + \beta x = 0$; therefore $\beta x = -x$. We conclude that $p(x, y, z) = x - y + z = m(x, y, z)$.

For the centrality, let $f(x_1, \dots, x_n)$ be a basic operation of $\mathbf{A}$. Since $\mathbf{A}$ is affine we know that f is an affine combination of variables $f(x_1, \dots, x_n) = \sum_{i=1}^{n} \alpha_i x_i + c$. Therefore

$$\begin{aligned} m(f(\bar{x}), f(\bar{y}), f(\bar{z})) &= (\sum_{i=1}^{n} \alpha_i x_i + c) - (\sum_{i=1}^{n} \alpha_i y_i + c) + (\sum_{i=1}^{n} \alpha_i z_i + c) \\ &= \sum_{i=1}^{n} \alpha_i (x_i - y_i + z_i) + c \\ &= f(m(x_1, y_1, z_1), \dots, m(x_n, y_n, z_n)), \end{aligned}$$

so m is central. □

We next discuss Abelianness, which is defined by the following condition on the term operations of $\mathbf{A}$:

Definition 2.6. An algebra $\mathbf{A}$ is called *Abelian* if all term operations $t \in \mathsf{Clo}(\mathbf{A})$ (of arbitrary arity $n + 1$) satisfy the condition

$$\forall x, y \in A, \forall \bar{u}, \bar{v} \in A^n \colon t(x, \bar{u}) = t(x, \bar{v}) \Rightarrow t(y, \bar{u}) = t(y, \bar{v}). \tag{2.1}$$

Observation 2.7. Let $\mathbf{A}$ be affine. Then $\mathbf{A}$ is Abelian.

Proof. Let $t \in \mathsf{Clo}(\mathbf{A})$. Then t can be written as an affine combination $t(x, z_1, \dots, z_n) = \alpha_0 x + \sum_{i=1}^{n} \alpha_i z_i + c$. Now it is an easy computation that, for all $x, y \in A$ and $\bar{u}, \bar{v} \in A^n$:

$$t(x, u_1, \dots, u_n) = t(x, v_1, \dots, v_n) \Leftrightarrow \sum_{i=1}^{n} \alpha_i u_i = \sum_{i=1}^{n} \alpha_i v_i \Leftrightarrow t(y, u_1, \dots, u_n) = t(y, v_1, \dots, v_n),$$

thus $\mathbf{A}$ is Abelian. □

We give some more examples of Abelian algebras:

Example 2.8.

1. Every unary algebra $\mathbf{A} = (A, f(x))$ is Abelian.
2. every affine algebra is Abelian.
3. A semilattice $(L, \wedge)$ is Abelian if $|L| = 1$.
4. a group $(G, \cdot, e, ^{-1})$ is Abelian if and only if $x \cdot y \approx y \cdot x$.
5. a ring $(R, +, 0, -, \cdot)$ is Abelian if and only if $x \cdot y \approx 0$.

Proof (Solution to Exercise 2.2).

1. Every term is of the form $t(x_1, \dots, x_n) = f^k(x_i)$, and therefore only depends on one coordinate. Because of this the term condition (2.1) holds.
2. We showed this in Observation 2.7.
3. We apply the term condition (2.1) to the term $t(z_1, z_2) = z_1 \wedge z_2$. If $|L| > 1$ there are two elements $x, y \in L$ with $x < y$. Then we have $x = t(x, x) = t(x, y)$. By (2.1) we obtain $t(y, x) = t(y, y)$, which is equivalent to $x = y$ - a contradiction!
4. Let us consider the term operation $[z_1, z_2] := z_1^{-1} z_2^{-1} z_1 z_2$, and let $x, y \in G$. Then clearly $e = [e, y] = [e, e] = e$. By applying the term condition (2.1), we can exchange the first coordinate on both sides of this equation by x, obtaining $[x, y] = [x, e] = e$. Thus $[x, y] = e$ for all $x, y \in G$, or in other words, the multiplication is commutative. For the converse note, that we already saw before that commutative groups are affine, and therefore also Abelian.
5. For a ring $(R, +, 0, -, \cdot)$, let us look at the term $t(z_1, z_2) = z_1 \cdot z_2$. Let $x, y \in R$. Then clearly $t(0, y) = t(0, 0) = 0$. By applying (2.1) to this equation we get $t(x, y) = t(x, 0) = 0$. For the other direction, note that if $x \cdot y \approx 0$, then the ring is polynomially equivalent to the commutative group $(R, +, 0, -)$, which is Abelian.

□

2.2 The fundamental theorem of Abelian algebras

As you can see from Example 2.8 (1), Abelian algebras need not to be affine (for another example see Exercise 2.7). The fundamental theorem of Abelian algebras however states, that, under the additional assumption of $\mathbf{A}$ being in a congruence modular variety, Definition 2.6 is indeed equivalent to $\mathbf{A}$ being affine. We are only going to prove this theorem for the special case of algebras with a Maltsev polynomial m. In this case Abelianness is moreover equivalent to m being central.

Theorem 2.9 (Fundamental theorem (Herrmann '77))**.** *Let $\mathbf{A}$ be an algebra with a Maltsev polynomial $m(x, y, z)$. Then the following are equivalent:*

1. *$\mathbf{A}$ is Abelian*
2. *$\mathbf{A}$ is affine*
3. *m is central*

Proof. We already saw the implications (2) → (1),(3) in Observations 2.5 and 2.7.

For the implication (3)→(1) assume that m is central; let t be a term operation, and let $x \in A, \bar{u}, \bar{v} \in A^n$ such that $t(x, \bar{u}) = t(x, \bar{v})$. We then want to show that $t(y, \bar{u}) = t(y, \bar{v})$ for arbitrary y.

For this consider the expression $m(t(y, \bar{u}), t(x, \bar{u}), t(x, \bar{v}))$. Since $t(x, \bar{u}) = t(x, \bar{v})$ and m is Maltsev, it is equal to $t(y, \bar{u})$. On the other hand, using the centrality of m, we can rewrite the expression into $t(m(y, x, x), m(u_1, u_1, v_1), \dots, m(u_n, u_n, v_n))$, which is equal to $t(y, v_1, \dots, v_n)$ by the Maltsev identities. So we obtain $t(y, \bar{u}) = t(y, \bar{v})$, which is what we wanted to show.

For the implication (1)→(2), we need to construct a module that is polynomially equivalent to $\mathbf{A}$. Let us fix an arbitrary element $0 \in A$. We then define the group operations $+$ and $-$

by the polynomials $x+y=m(x,0,y)$ and $-x=m(0,x,0)$. The ring R is defined as $R=\{f$ unary polynomial : $f(0)=0\}$. We first verify the module axioms for $(A,+,0,-,(r)_{r\in R})$.

It is easy to see that 0 is the neutral element of $+$, since $x=m(0,0,x)=0+x=m(x,0,0)=x+0$ by the Maltsev identities. To prove associativity, consider the term $t(x_1,x_2,x_3,x_4)=(x_1+x_2)+(x_3+x_4)$. By the neutrality of 0, $t(0,0,b,c)=t(0,b,0,c)$ holds for all $b,c\in A$. Note that the first coordinate on both sides of the equation is set to 0. Thus we can apply the term condition (2.1) to this equation, and obtain that $t(a,0,b,c)=t(a,b,0,c)$ for arbitrary a. This is equivalent to $a+(b+c)=(a+b)+c$, hence $+$ is associative.

— *(Solution to Exercise 2.4)*

For commutativity, note that $m(a,a,b)=m(b,a,a)$ for every $a,b\in A$. By the term condition (2.1), this is still true if we substitute the middle a by 0, thus we obtain $m(a,0,b)=m(b,0,a)$, which is equivalent to $a+b=b+a$.

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 $t(a,a,a)=t(a,0,0)=a$. By the term condition (2.1) we obtain $t(0,a,a)=t(0,0,0)$; in other words $a+(-a)=0$.

Thus $(A,+,0,-)$ is a commutative group. Next let $r\in R$. Then, note that it is distributive with respect to $+$ iff $t(x,y)=r(x+y)-r(x)-r(y)$ is the constant 0 function. Let $a,b\in A$. Clearly $t(0,b)=t(0,0)=0$ for all $b\in A$. By the term condition we get $t(a,b)=t(a,0)=0$, which is what we wanted to prove.

—

It is still left to prove that every polynomial $p(x_1,\dots,x_n)$ of $\mathbf{A}$ is affine. We show it by an induction on the arity n. For $n=1$ the statement is clearly true (since by the definition of R, $p(x)=r(x)+p(0)$, for some $r\in R$).

For an induction step, let us look at the polynomial $t(x_1,\dots,x_n)=p(x_1,x_2,\dots,x_n)-p(0,x_2,\dots,x_n)-p(x_1,0,\dots,0)+p(0,\dots,0)$.Using the term condition (2.1) (as in the proof of the distributivity of $r\in R$) we see that t is the constant 0 function. Thus $p(x_1,x_2,\dots,x_n)=p(0,x_2,\dots,x_n)+p(x_1,0,\dots,0)-p(0,\dots,0)$; the latter three summands are affine because of the induction hypothesis, hence p is also affine. □

Using the observation that every affine algebra has a Maltsev term, we obtain the following corollary:

Corollary 2.10. *An algebra is affine if and only if it is Abelian and has a Maltsev polynomial.*

2.3 The term condition commutator

The actual strength of commutator theory is to give us tools for analyzing the congruences of an algebra and how they interact with each other. These tools are primarily the *centralizer relation C* and the *commutator* $[\alpha,\beta]$ of two congruences α,β.

In this section we are going to define them by a term condition similar to the condition (2.1) in the definition of Abelianness and show that in the group setting, the commutator of two congruences corresponds to the commutator subgroups of two normal subgroups. In the next section, however, we will see that the commutator is also very useful in discussing algebras that are far from being "group like".

Definition 2.11. Let $\mathbf{A}$ be an algebra, and let $\alpha, \beta, \delta \in \mathrm{Con}(\mathbf{A})$.

We then say that α *centralizes* β *modulo* δ, and write $C(\alpha, \beta; \delta)$, if, for all terms $t \in \mathsf{Clo}(\mathbf{A})$ (of arbitrary arity $n+1$), for all $x, y \in A$ and $\bar{u}, \bar{v} \in A^n$ with $x \mathrel{\alpha} y$ and $u_i \mathrel{\beta} v_i$ (for all indices $i = 1, \ldots, n$):

$$\begin{aligned} &t(x, \bar{u}) \mathrel{\delta} t(x, \bar{v}) \\ \Rightarrow\ &t(y, \bar{u}) \mathrel{\delta} t(y, \bar{v}) \end{aligned} \tag{2.2}$$

If $\delta = 0_A$, we say that α *centralizes* β.

Comparing Definition 2.11 with Definition 2.6 observe that $\mathbf{A}$ is Abelian, if and only if $C(1_A, 1_A; 0_A)$. If instead of single elements x and y, we also allow tuples $\bar{x}$, $\bar{y}$ such that $x_i \mathrel{\alpha} y_i$ holds for all entries, we still define the same relation C (see Exercise 2.15). For simplicity let us from now on write $\bar{u} \mathrel{\beta} \bar{v}$ if $\bar{u}$ and $\bar{v}$ are two tuples of the same length n, and for every index and $u_i \mathrel{\beta} v_i$ holds for every $i = 1, \ldots, n$.

We start by discussing the relation C for groups. Recall that the congruences of a group $(G, \cdot, e, ^{-1})$ correspond to its normal subgroups. We are going to show that in groups α centralizes β, if and only if the corresponding normal subgroups centralize each other in the classical group theoretical sense:

Proposition 2.12. *Let α, β be two congruences of a group $\mathbf{G} = (G, \cdot, e, ^{-1})$, and let $A = [e]_\alpha$ and $B = [e]_\beta$ be the corresponding normal subgroups. Then α centralizes β if and only if $a \cdot b = b \cdot a$, for all elements $a \in A$ and $b \in B$.*

Proof. To see that $C(\alpha, \beta; 0)$ implies $a \cdot b = b \cdot a$ for all $a \in A$ and $b \in B$, we can use the same proof idea as in Example 2.8 (4), applying the term condition (2.2) to the term $[x, y] = x^{-1}y^{-1}xy$.

Proving the converse is more work. Assume that $a \cdot b = b \cdot a$, for all elements $a \in A$ and $b \in B$. Our proof of $C(\alpha, \beta; 0)$ then uses the following two observations:

- **Observation 1:** for every term $s \in \mathsf{Clo}(\mathbf{G})$ and tuples $\bar{u}, \bar{v}$ such that $\bar{u} \mathrel{\beta} \bar{v}$ we have that $s(\bar{u})(s(\bar{v}))^{-1} \in B$. This follows straightforward from β being a congruence.
- **Observation 2:** For all $u, v \in G$ with $u \mathrel{\alpha} v$, and $b \in B$: $u^{-1}bu = v^{-1}bv$.
  This follows from $b = (u^{-1}v)(v^{-1}u)b = (u^{-1}v)b(v^{-1}u)$, where the last equality uses the assumption that elements of A and B commute.

To check the term condition for $C(\alpha, \beta; 0)$ let $t \in \mathsf{Clo}(\mathbf{G})$ be a group term, let $x, y \in G, \bar{u}, \bar{v} \in G^n$ such that $x \mathrel{\alpha} y$ and $\bar{u} \mathrel{\beta} \bar{v}$. We are going to prove that

$$t(x, \bar{u})(t(x, \bar{v}))^{-1} = t(y, \bar{u})(t(y, \bar{v}))^{-1}. \tag{2.3}$$

Note that the above equation directly implies the term condition (2.2); in that case both sides of the equation are equal to e.

Since t is a group term, it can be represented as product of variables or their inverses, so $t(x_1, x_2, \ldots, x_n) = z_1 \cdot z_2 \cdots z_m$, with $z_i \in \{x_1^d, \ldots, x_n^d \mid d \in \{-1, 1\}\}$. We prove (2.3) by induction on the length m.

For $m = 1$ the statement clearly holds: If $t(x, \bar{u}) = x$ we get that

$$t(x, \bar{u})(t(x, \bar{v}))^{-1} = xx^{-1} = e = yy^{-1} = t(y, \bar{u})(t(y, \bar{v}))^{-1};$$

the proof is analogue for x^{-1}. If $t(x,\bar{u}) = u_i$ for some i, we have $t(x,\bar{u})(t(x,\bar{v}))^{-1} = u_i v_i^{-1} = t(y,\bar{u})(t(y,\bar{v}))^{-1}$; the proof is analogous for u_i^{-1}.

For an induction step $m-1 \to m$ let $t = z_1 \cdot s$. Note that $s(x,\bar{u})(s(x,\bar{v}))^{-1}$ is in B by Observation 1, and equal to $s(y,\bar{u})(s(y,\bar{v}))^{-1}$ by the induction hypothesis (IH). We do a case distinction, depending on z_1.

If $t(x,\bar{u}) = u_i s(x,\bar{u})$, for some i, we get

$$t(x,\bar{u})(t(x,\bar{v}))^{-1} = u_i s(x,\bar{u})(s(x,\bar{v}))^{-1} v_i^{-1} \overset{IH}{=} u_i(s(y,\bar{u})(s(y,\bar{v}))^{-1})v_i^{-1} = t(y,\bar{u})(t(y,\bar{v}))^{-1}.$$

If $t(x,\bar{u}) = xs(x,\bar{u})$, we get that

$$\begin{aligned} t(x,\bar{u})(t(x,\bar{v}))^{-1} &= xs(x,\bar{u})(s(x,\bar{v}))^{-1}x^{-1} \\ &\overset{IH}{=} xs(y,\bar{u})(s(y,\bar{v}))^{-1}x^{-1} \\ &\overset{Ob2}{=} ys(y,\bar{u})(s(y,\bar{v}))^{-1}y^{-1} = t(y,\bar{u})(t(y,\bar{v}))^{-1}. \end{aligned}$$

If z_1 is an inverse of a variable the proof is analogous.

So we proved (2.3) for all $x \,\alpha\, y$ and $\bar{u} \,\beta\, \bar{v}$. This in turn implies that the term condition (2.2) holds; thus α centralizes β. □

There are a few observations that follow straightforward from the definition of $C(\alpha,\beta;\delta)$, and allow us to define the commutator.

Observation 2.13.

1. If $\alpha \geq \alpha'$, $\beta \geq \beta'$ and $C(\alpha,\beta;\delta)$, then $C(\alpha',\beta';\delta)$.
2. $C(\alpha,\beta;\delta)$ in $\mathbf{A}$ if and only if $C(\alpha/\delta,\beta/\delta;0_A)$ in $\mathbf{A}/\delta$
3. $C(\alpha,\beta;\alpha)$ and $C(\alpha,\beta;\beta)$
4. Let Δ be a set of congruences. If $C(\alpha,\beta;\delta)$ for all $\delta \in \Delta$ then $C(\alpha,\beta;\bigwedge\Delta)$

Proof. left as Exercise 2.12. □

Definition 2.14. The *commutator of* α *and* β, short $[\alpha,\beta]$, is the smallest $\delta \in \operatorname{Con}(\mathbf{A})$, such that $C(\alpha,\beta;\delta)$.

Note that the commutator is well defined: By Observation 2.13 (3), there is always a δ, such that α centralizes β modulo δ. By Observation 2.13 (4) the infimum γ of all such δ's has also the property that $C(\alpha,\beta;\gamma)$. In other words this infimum γ is the commutator $[\alpha,\beta]$. Note also that, by Observation 2.13 (3) and (4) we have $[\alpha,\beta] \leq \alpha \wedge \beta$.

In the literature, the commutator appears more commonly than the relation C; we are going to prove later that in the Maltsev case one can be defined via the other by $C(\alpha,\beta;\delta) \Leftrightarrow [\alpha,\beta] \leq \delta$.

We give some examples:

Example 2.15. Let α,β two congruences of a group $\mathbf{G} = (G,\cdot,e,^{-1})$, and let $A = [e]_\alpha$ and $B = [e]_\beta$ the corresponding normal subgroups. Then $[\alpha,\beta]$ corresponds to the *commutator subgroup* $[A,B]$, i.e. the normal subgroup generated by $\{a^{-1}b^{-1}ab : a \in A, b \in B\}$.

Proof. By Proposition 2.12 we have that $C(\alpha, \beta; \delta)$ if and only if $(a \cdot b) \, \delta \, (b \cdot a)$ for all $a \in A$, $b \in B$. If we set $D = [e]_\delta$, this is equivalent to $a^{-1}b^{-1}ab \in D$. Now the smallest normal subgroup D such that this holds is clearly the normal subgroup generated by all elements $a^{-1}b^{-1}ab$, which by definition is the commutator subgroup $[A, B]$. □

Example 2.16. Let $(R, +, 0, -, \cdot)$ be a commutative ring. Recall that its congruences α are uniquely determined by the ideals $I_\alpha = [0]_\alpha$. We claim that then α centralizes β if and only if $I_\alpha \cdot I_\beta = 0$. The congruence $[\alpha, \beta]$ corresponds to the ideal $I_\alpha \cdot I_\beta$.

Proof. Exercise 2.13. □

Example 2.17. Let $\mathbf{L} = (L, \wedge, \vee)$ be a lattice. Then $[\alpha, \beta] = \alpha \wedge \beta$ for all congruences $\alpha, \beta \in \mathrm{Con}(\mathbf{L})$.

Proof. We already know by Observation 2.13 that $[\alpha, \beta] \leq \alpha \wedge \beta$. So it is only left to show that $C(\alpha, \beta; \delta)$ implies $\delta \geq \alpha \wedge \beta$.

Let $x \, (\alpha \wedge \beta) \, y$. Without loss of generality we can assume that $x \leq y$ (since from UA1 we know that for any congruence $\gamma \in \mathrm{Con}(\mathbf{L})$ we have $(x, y) \in \gamma$ if and only if $(x \wedge y, x \vee y) \in \gamma$). Then, for the term $t(z_1, z_2) = z_1 \wedge z_2$ clearly $t(x, x) = t(x, y)$ holds. In particular this implies that $t(x, x) \, \delta \, t(x, y)$. By the term condition for $C(\alpha, \beta; \delta)$, we can exchange the first coordinate by y, and obtain $x = t(y, x) \, \delta \, t(y, y) = y$. Thus $\alpha \wedge \beta \leq \delta$, which concludes the proof. □

In general the commutator does not need to satisfy $[\alpha, \beta] = [\beta, \alpha]$, as one might think after Examples 2.15, 2.16, 2.17. However in the Maltsev case commutativity holds. In the following we give a proof and show a few other additional properties of the commutator in the Maltsev case.

We need the following helpful lemma first.

Lemma 2.18. *Assume* $\mathbf{A}$ *has a Maltsev polynomial* $m(x, y, z)$*, and let* $\alpha, \beta \in \mathrm{Con}(\mathbf{A})$ *such that* $[\alpha, \beta] = 0_A$*. Let* $t \in \mathsf{Clo}(\mathbf{A})$ *and* $\bar{x}, \bar{y}$ *such that* $\bar{x} \, \alpha \, \bar{y}$ *and* $\bar{u} \, \beta \, \bar{v}$*. Then* $t(\bar{y}, \bar{v}) = m(t(\bar{y}, \bar{u}), t(\bar{x}, \bar{u}), t(\bar{x}, \bar{v}))$.

Proof. Since m is a Maltsev operation, the following equation holds.

$$m(t(\bar{x}, \bar{v}), t(\bar{x}, \bar{u}), t(\bar{x}, \bar{u})) = m(t(\bar{x}, \bar{u}), t(\bar{x}, \bar{u}), t(\bar{x}, \bar{v}))$$

By applying the term condition for $C(\alpha, \beta; 0)$ we can exchange the first $\bar{x}$ by $\bar{y}$, and obtain:

$$m(t(\bar{y}, \bar{v}), t(\bar{x}, \bar{u}), t(\bar{x}, \bar{u})) = m(t(\bar{y}, \bar{u}), t(\bar{x}, \bar{u}), t(\bar{x}, \bar{v}))$$

But since m is a Maltsev operation, the left side is equal to $t(\bar{y}, \bar{v})$, which finishes the proof. □

Note that Lemma 2.18 can be seen as a generalization of the centrality of the Maltsev operation in the case where $[1_A, 1_A] = 0_A$.We can use Lemma 2.18 to prove some useful properties of the commutator:

Lemma 2.19. *Let* $\mathbf{A}$ *be an algebra with a Maltsev polynomial. Then*

1. $C(\alpha, \beta; \delta)$ *holds if and only if* $[\alpha, \beta] \leq \delta$
2. *The commutator is* commutative: $[\alpha, \beta] = [\beta, \alpha]$
3. *The commutator is* distributive*:* $[\bigvee \Gamma, \alpha] = \bigvee_{\gamma \in \Gamma} [\gamma, \alpha]$

Proof.

1. If follows from the definition of the commutator, that $C(\alpha,\beta;\delta)$ implies $[\alpha,\beta]\leq\delta$. For the opposite direction, assume that $[\alpha,\beta]\leq\delta$; we then want to verify the term condition for $C(\alpha,\beta;\delta)$.

   Without loss of generality we can assume that $[\alpha,\beta]=0$ (otherwise we work in the quotient $\mathbf{A}/[\alpha,\beta]$ by Observation 2.13(2)). Let $\bar{x}\,\alpha\,\bar{y}$ and $\bar{u}\,\beta\,\bar{v}$, and $t\in\mathsf{Clo}(\mathbf{A})$ such that $t(\bar{x},\bar{u})\,\delta\,t(\bar{x},\bar{v})$.

   By Lemma 2.18 we know that $t(\bar{y},\bar{v})=m(t(\bar{y},\bar{u}),t(\bar{x},\bar{u}),t(\bar{x},\bar{v}))$, which is δ-related to $m(t(\bar{y},\bar{u}),t(\bar{x},\bar{v}),t(\bar{x},\bar{v}))=t(\bar{y},\bar{u})$. Thus $t(\bar{y},\bar{u})\,\delta\,t(\bar{y},\bar{v})$, which is what we wanted to prove.

2. By (1) this is equivalent to showing $C(\alpha,\beta;\delta)\Leftrightarrow C(\beta,\alpha;\delta)$ in $\mathbf{A}$. By Observation 2.13 (2) we further can assume that $\delta=0$. So let us assume that $C(\alpha,\beta;0)$. Our goal is then to prove that for all terms t and for all $\bar{x},\bar{y}$ such that $\bar{x}\,\alpha\,\bar{y}$ and $\bar{u}\,\beta\,\bar{v}$, the equation $t(\bar{x},\bar{u})=t(\bar{y},\bar{u})$ implies $t(\bar{x},\bar{v})=t(\bar{y},\bar{v})$. By Lemma 2.18:

   $$t(\bar{y},\bar{v})=m(t(\bar{y},\bar{u}),t(\bar{x},\bar{u}),t(\bar{x},\bar{v}))=t(\bar{x},\bar{v}),$$

   where the second identity follows from the Maltsev identities.

3. Note that $[\bigvee\Gamma,\alpha]\geq[\gamma,\alpha]$ holds for every $\gamma\in\Gamma$ by Observation 2.13 (1). By taking the supremum over all such γ we get $[\bigvee\Gamma,\alpha]\geq\bigvee_{\gamma\in\Gamma}[\gamma,\alpha]$.

   For the other direction, let $\beta=\bigvee_{\gamma\in\Gamma}[\gamma,\alpha]$. To complete the proof we need to show $[\bigvee\Gamma,\alpha]\leq\beta$, which by (1) is equivalent to $C(\bigvee\Gamma,\alpha;\beta)$. So let $(x,y)\in\bigvee\Gamma$, and $\bar{u},\bar{v}$ such that $(u_i,v_i)\in\alpha$, and let $t\in\mathsf{Clo}(\mathbf{A})$. Since $(x,y)\in\bigvee\Gamma$ there is a chain of elements $a_0,\dots,a_k\in A$ and congruences $\gamma_1,\dots,\gamma_k\in\Gamma$ such that $x=a_0\gamma_1a_1\gamma_2a_2\cdots\gamma_ka_k=y$. By (1) $\beta\geq[\alpha,\gamma_i]$ is equivalent to $C(\alpha,\gamma_i;\beta)$. Using the corresponding term condition for every $i=1,\dots,k$, we get

   $$t(x,\bar{u})\,\beta\,t(x,\bar{v})\Rightarrow t(a_1,\bar{u})\,\beta\,t(a_1,\bar{v})\Rightarrow t(a_2,\bar{u})\,\beta\,t(a_2,\bar{v})\Rightarrow\ldots\Rightarrow t(y,\bar{u})\,\beta\,t(y,\bar{v}),$$

   which completes the proof of the term condition for $C(\bigvee\Gamma,\alpha;\beta)$.

□

Here is a short overview over all the properties of the commutator proved in this section:

Summary	
$\alpha'\leq\alpha,\ \beta'\leq\beta\Rightarrow[\alpha',\beta']\leq[\alpha,\beta]$	by Observation 2.13 (1)
$[\alpha,\beta]\leq\alpha\wedge\beta$	by Observation 2.13 (3), (4)
If A has Maltsev polynomial:	
$C(\alpha,\beta;\delta)\Leftrightarrow[\alpha,\beta]\leq\delta$	Lemma 2.19 (1)
$[\alpha/\gamma,\beta/\gamma]=[\alpha,\beta]/\gamma$ in $\mathbf{A}/\gamma$	Lemma 2.19 (1) + Obs. 2.13 (2)
$[\alpha,\beta]=[\beta,\alpha]$	Lemma 2.19 (2)
$[\bigvee\Gamma,\alpha]=\bigvee_{\gamma\in\Gamma}[\gamma,\alpha]$	Lemma 2.19 (3)

Remark 2.20. There are several other concepts from group theory that, building on this framework, can be directly generalized to arbitrary algebras:

- A congruence α is called *Abelian* if $[\alpha, \alpha] = 0_A$
- The *centralizer* of a congruence α is the biggest β, such that $C(\alpha, \beta; 0)$
- The *center* $\zeta_{\mathbf{A}}$ of an algebra $\mathbf{A}$ is the centralizer of 1_A.
- A *central series* is a series of congruences $0_A = \alpha_0 \leq \alpha_1 \leq \cdots \leq \alpha_n = 1_A$, such that $[1_A, \alpha_k] \leq \alpha_{k-1}$, for all $k = 1, \ldots, n$.
- An algebra is called *(n-)nilpotent* if it has a central series (of length n).
- An algebra is called *(n-)solvable* if there is a series $0_A = \alpha_0 \leq \alpha_1 \leq \cdots \alpha_n = 1_A$, with $[\alpha_k, \alpha_k] \leq \alpha_{k-1}$.

2.4 Commutator and varieties

Commutators allow us to characterize certain properties of varieties. We already saw an example in Corollary 2.10: A variety $\mathcal{V}$ with Maltsev term is affine, if and only if $[1_A, 1_A] = 0_A$ holds in every algebra $\mathbf{A} \in \mathcal{V}$.

A more involved result is the following theorem by Freese and McKenzie that discusses the subdirectly irreducible elements of an algebra:

Theorem 2.21 (Freese, McKenzie [3])**.** *Let* $\mathbf{A}$ *be an algebra, such that* $\mathcal{V} = \mathsf{HSP}(\mathbf{A})$ *is congruence modular. Then* $\mathcal{V}$ *contains only finitely many subdirectly irreducible elements, which are all finite* [2] *if and only if for all* $\alpha, \beta \in \mathrm{Con}(\mathbf{A})$*:* $\alpha \leq [\beta, \beta] \Rightarrow \alpha = [\alpha, \beta]$.

In this section we give another example of commutators neatly describing the properties of a variety. We already saw in Example 2.17 that in the variety of lattices $[\alpha, \beta] = \alpha \wedge \beta$ holds. By the same argument this is also true in the variety of Boolean algebras. Both are congruence-distributive variety; and in fact in the Maltsev case, congruence distributivity is equivalent to to the "commutator identity" $[\alpha, \beta] \approx \alpha \wedge \beta$:

Theorem 2.22 (Hagemann, Herrmann '79)**.** *Let* $\mathcal{V}$ *be a variety with Maltsev term. Then* $\mathcal{V}$ *is congruence-distributive if and only if* $[\alpha, \beta] = \alpha \wedge \beta$ *holds for all congruences* $\alpha, \beta \in \mathrm{Con}(\mathbf{A})$ *of all* $\mathbf{A} \in \mathcal{V}$.

Before reading the proof of the theorem, I suggest that you have a look at Exercise 2.18, which proves it for any group variety.

Proof. Suppose that $[\alpha, \beta] = \alpha \wedge \beta$ holds for all congruences of an algebra $\mathbf{A}$ with Maltsev polynomial. Then, by the distributivity of the commutator, this algebra has also a distributive congruence lattice: $(\alpha \vee \beta) \wedge \gamma = [\alpha \vee \beta, \gamma] = [\alpha, \gamma] \vee [\beta, \gamma] = (\alpha \wedge \gamma) \vee (\beta \wedge \gamma)$. Thus, if all congruences of all algebras in $\mathcal{V}$ satisfy $[\alpha, \beta] = \alpha \wedge \beta$ then $\mathcal{V}$ is congruence-distributive.

For the converse, suppose that there is an algebra $\mathbf{A} \in \mathcal{V}$ and congruences $\alpha, \beta \in \mathrm{Con}(\mathbf{A})$ such that $[\alpha, \beta] < \alpha \wedge \beta$. Our goal is to prove that $\mathcal{V}$ is not congruence distributive.

Let us define $\gamma' = \alpha \wedge \beta$. Then, by the monotonicity of the commutator we have $[\gamma', \gamma'] \leq [\alpha, \beta] < \gamma'$. The quotient $\mathbf{B} = \mathbf{A}/[\alpha, \beta]$ also lies in the variety $\mathcal{V}$. Let $\gamma = \gamma'/[\alpha, \beta]$. Then γ is a non-trivial congruence of $\mathbf{B}$ with $[\gamma, \gamma] = 0_B$.

[2] Such varieties are called *residually finite*

We can regard γ as a subalgebra of $\mathbf{B}^2$. Hence also γ is an element of the variety $\mathcal{V}$. To make it clear that we consider γ as an algebra, we write $\mathbf{B}(\gamma)$ instead of γ, and write its element by column vectors, so $\mathbf{B}(\gamma) = \{\binom{a}{b} : a, b \in B, a\,\gamma\,b\}$.

We then claim that the congruence lattice of $\mathbf{B}(\gamma)$ contains the following sublattice:

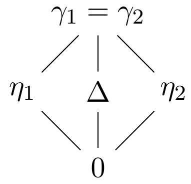

If we can prove this claim, $\operatorname{Con}(\mathbf{B}(\gamma))$ is not distributive, since this sublattice is not distributive.

We define the congruences $\binom{a}{b}\eta_1\binom{c}{d}$ if and only if $a = c$, and $\binom{a}{b}\eta_2\binom{c}{d}$ if and only if $b = d$. Similarly we define $\binom{a}{b}\gamma_1\binom{c}{d}$, if $a\gamma c$. Note that this is equivalent to all four elements a, b, c, d being in the same γ-class. It is easy to see that $\eta_1 \wedge \eta_2 = 0$ and $\eta_1 \vee \eta_2 = \gamma_1$.

The last congruence Δ is defined by

$$\Delta = \left\{ \left(\begin{pmatrix} t(\bar{x}, \bar{u}) \\ t(\bar{y}, \bar{u}) \end{pmatrix}, \begin{pmatrix} t(\bar{x}, \bar{v}) \\ t(\bar{y}, \bar{v}) \end{pmatrix} \right) : \bar{x}\,\gamma\,\bar{y},\ \bar{u}\,\gamma\,\bar{v},\ t \in \mathsf{Clo}(\mathbf{B}) \right\}.$$

By a lemma of Maltsev (Exercise 2.19), Δ is indeed a congruence of $B(\gamma)$.

If $\binom{a}{b}(\eta_1 \wedge \Delta)\binom{c}{d}$, there is a term t and tuples $\bar{x}, \bar{y}, \bar{u}, \bar{v}$ such $\bar{x}\,\gamma\,\bar{y}, \bar{u}\,\gamma\,\bar{v}$ and $a = t(\bar{x}, \bar{u}) = t(\bar{x}, \bar{v}) = c$, $b = t(\bar{y}, \bar{u})$ and $d = t(\bar{y}, \bar{v})$. By the term condition for $[\gamma, \gamma] = 0$ it follows that $b = t(\bar{y}, \bar{u}) = t(\bar{y}, \bar{v}) = d$, and therefore $\binom{a}{b} = \binom{c}{d}$. So $\eta_1 \wedge \Delta = 0$. Symmetrically we can prove that $\eta_2 \wedge \Delta = 0$.

To see that $\eta_1 \vee \Delta = \gamma_1$, simply observe that $\binom{a}{a}\Delta\binom{c}{c}$ holds for all $a\,\gamma\,c$. Therefore, for all pairs $\binom{a}{b}, \binom{c}{d} \in \gamma$ we have:

$$\begin{pmatrix} a \\ b \end{pmatrix} \eta_1 \begin{pmatrix} a \\ a \end{pmatrix} \Delta \begin{pmatrix} c \\ c \end{pmatrix} \eta_1 \begin{pmatrix} c \\ d \end{pmatrix},$$

showing that $\eta_1 \vee \Delta = \gamma_1$. Symmetrically $\eta_2 \vee \Delta = \gamma_1$ holds. Thus $\mathbf{B}(\gamma)$ is not congruence distributive. Since it is an element of $\mathcal{V}$, $\mathcal{V}$ is not congruence distributive, which finishes our proof. □

Remark 2.23. Our version of Theorem 2.22 does not cover varieties of lattices (Example 2.17), since lattice varieties usually don't have a Maltsev term. But it is known that Theorem 2.22 holds for all congruence modular varieties $\mathcal{V}$ (and thus also for lattices), see [3].

Even more generally, Andrew Moorehead (who was until recently a member of our department) showed in [4], that the commutator identity $[\alpha, \beta]_H \approx \alpha \wedge \beta$ is equivalent to a variety $\mathcal{V}$ being congruence meet-semidistributive, where his 'hypercommutator' $[\alpha, \beta]_H$ coincides with our normal 'term condition commutator' $[\alpha, \beta]$ on congruence modular varieties.

2.5 Nilpotent algebras

In this last section we shortly discuss a few open problems in commutator theory. This is not part of the core (or exam) material of the lecture, but meant to show you that commutator theory (despite being developed in the 80ies) is still an active field of research today. Recent development indicates that it can be a useful tool in tackling certain problems from theoretical computer science. One such example is the concept of 'similarity' / 'bridges', which appears in the Zhuk's proof of the CSP dichotomy conjecture [6].

In this section we are discussing another example, namely the *(polynomial) equivalence problem* over a fixed algebra $\mathbf{A}$, which is defined as follows:

PolEQV($\mathbf{A}$)
Input: Two polynomials $s(x_1, \dots, x_n), t(x_1, \dots, x_n)$ over $\mathbf{A}$
Question: Does $\mathbf{A} \models s(x_1, \dots, x_n) \approx t(x_1, \dots, x_n)$?

Note that, if the algebra $\mathbf{A}$ is finite, there is always an algorithm solving PolEQV($\mathbf{A}$): Evaluate s and t at every tuple $\bar{a} \in A^n$, and check if the result is the same. However this algorithm is not very efficient, since its running time is $\mathcal{O}(|A|^n)$, exponential in the input size!

It would be nice to have a characterization of all finite algebras $\mathbf{A}$, for which PolEQV($\mathbf{A}$) is in P, so for which there is a deterministic polynomial time algorithm solving PolEQV($\mathbf{A}$). But at the current state we are far from such a classification.

However for algebras with Maltsev term much more is known. There seems to be a very close connection between the complexity of PolEQV($\mathbf{A}$) and commutator properties. It is for instance easy to see, that affine algebras have an easy equivalence problem:

Example 2.24. Let $\mathbf{A}$ be an affine algebra. Then $\mathbf{A} \models s \approx t$ is equivalent to $\mathbf{A} \models m(s, t, 0) \approx 0$. So we can reduce the equivalence problem PolEQV($\mathbf{A}$) to the problem of deciding whether an input polynomial is equivalent to the constant 0 function.

By Exercise 2.5, a polynomial operation of $\mathbf{A}$ of arity $n \geq 2$ is constant 0, if and only if it is 0-absorbing. This implies that $p(x_1, x_2, \dots, x_n) \approx 0$, if and only if all of the unary polynomials $p(x_1, 0, \dots, 0), p(0, x_2, 0, \dots, 0), \dots, p(0, \dots, 0, x_n)$ describe the constant 0 function. Thus $p(x_1, x_2, \dots, x_n) \not\approx 0$ if and only if there is a tuple $\bar{a}$, that has at most one non-0 entry, such that $p(\bar{a}) \neq 0$. There are only $n \cdot |A|$ many such tuples; evaluating p at them takes only linear time. So PolEQV($\mathbf{A}$) is in P.

On the other hand, by results of Idziak and Krzaczkowski [5], finite Maltsev algebras $\mathbf{A}$ with trivial center $\zeta_{\mathbf{A}} = 0_A$ are polynomially equivalent to an $\mathbf{A}'$, for which PolEQV($\mathbf{A}'$) is coNP-complete (and therefore not in P, unless $\mathsf{P} = \mathsf{NP}$).

This leaves essentially only nilpotent algebras as possible candidates, for which PolEQV is in P. Recall that an algebra is n-nilpotent if there is a series of congruences, such that $0_A = \alpha_0 \leq \alpha_1 \leq \cdots \leq \alpha_n = 1_A$, such that $[1_A, \alpha_k] \leq \alpha_{k-1}$, for all $k = 1, \dots, n$. Note that this is equivalent to $\underbrace{[\cdots[[1_A, 1_A], 1_A], \dots, 1_A]}_{n+1} = 0_A$.

We are going to show that similar algorithms to Example 2.24 work to for *some* nilpotent algebras, for instance nilpotent rings and groups.

Example 2.25. We saw in Example 2.16, that in a commutative ring, the commutator of two congruences is given by the product $I \cdot J$ of the corresponding ideals. Therefore a commutative ring R is k-step nilpotent, if and only if $R^{k+1} = 0$, or in other words, if and only if the identity $x_1 \cdot x_2 \cdots x_{k+1} \approx 0$ holds.

Because of this, a polynomial in a nilpotent ring R is always equivalent to a sum of monomials of total degree less than k. This in turn implies that every $k+1$-ary 0-absorbing polynomial of R is constant. Thus, in order to check if a polynomial $p(x_1, \dots, x_n)$ is equivalent to the constant 0 function, we only need to evaluate it at all tuples with at most k-many non-0 entries. This can be done in time $\mathcal{O}(\binom{n}{k} \cdot |A|^k) = \mathcal{O}(n^k)$. Thus PolEQV$(R)$ is in P.

It is further known (by Burris, Lawrence '93) that for non-nilpotent rings R, PolEQV(R) is coNP-complete (and therefore not in P, unless $\mathsf{P} = \mathsf{NP}$). Therefore nilpotent rings are exactly those rings, in which checking identities is computationally easy.

Example 2.26. By Proposition 2.12 a group $(G, \cdot, e, ^{-1})$ is n-nilpotent, if and only if the identity $[x_1, [x_2, \dots [x_n, x_{n+1}] \cdots] \approx e$ holds, with $[x, y] = x^{-1}y^{-1}xy$. A similar, but more technical argument than for rings shows that in n-nilpotent groups all 0-absorbing polynomials of arity bigger than n are constant. Thus also for nilpotent groups the above algorithm shows that PolEQV(G) is in P.

So both k-nilpotent groups and k-nilpotent rings have the property that 0-absorbing polynomials of arity $> k$ are constant. Maltsev algebras with this special property are called *k-supernilpotent*.[3] From the above examples one might jump to the conclusion that all k-nilpotent Maltsev algebras have this property, but this is false! Every supernilpotent Maltsev algebra is also nilpotent; but the opposite is not true:

Example 2.27. Let $\mathbf{A} = (\mathbb{Z}_9, +, 0, -, (f_n(x_1, \dots, x_n))_{n \in \mathbb{N}})$, where $f_n(x_1, \dots, x_n) = 3 \cdot x_1 \cdot x_2 \cdots x_n$ for every n. This algebra is 2-step nilpotent but not k-supernilpotent, for any k.

Example 2.28. Let $\mathbf{A} = (\mathbb{Z}_2 \times \mathbb{Z}_3, +, (0,0), -, f(x))$, where f is the unary operation defined by

$$f((l, u)) = \begin{cases} (1,0) \text{ if } u = 0 \\ (0,0) \text{ else.} \end{cases}$$

This algebra is 2-step nilpotent but not k-supernilpotent, for any k.

You may try to verify Example 2.27 as an exercise.

So what makes a nilpotent algebra? By the following result, one can regard the operations of a k-nilpotent algebras as a $k-1$-nilpotent algebra 'acting' on an affine algebra.

Proposition 2.29 ([3]). *Let $\mathbf{A}$ be a k-nilpotent algebra with a Maltsev term. Then there exists a $(k-1)$ nilpotent algebra $\mathbf{U}$ and an affine algebra $\mathbf{L}$ of the same signature, such that $A = L \times U$, and every basic operation of $\mathbf{A}$ is of the form*

$$f^{\mathbf{A}}((l_1, u_1), \dots, (l_n, u_n)) = (f^{\mathbf{L}}(l_1, \dots, l_n) + \hat{f}(u_1, \dots, u_n), f^{\mathbf{U}}(u_1, \dots, u_n)),$$

where $\hat{f} \colon U^n \to L$.

Proof idea. You may try to prove the result for algebras $\mathbf{A}$ containing Abelian group operations $+, 0, -$. Let ζ be the center of $\mathbf{A}$. Then set $\mathbf{U} = \mathbf{A}/\zeta$, and construct $\mathbf{L}$ as an extension of $([0]_\zeta, +, 0, -)$. □

[3] In general supernilpotent algebras are defined in a different way, but for Maltsev algebras those two definitions coincide.

We saw that in Abelian algebras $x \cdot y = m(x, 0, y)$ is a commutative group operation. More generally one can prove that in nilpotent algebras $x \cdot y = m(x, 0, y)$ is a loop multiplication. This follows from the following lemma:

Lemma 2.30. *Let* $\mathbf{A}$ *be a nilpotent algebra with a Maltsev term* m*. Then* $x \mapsto m(x, a, b)$ *is a bijection for all constants* $a, b \in A$*.*

Proof. The lemma follows from Exercise 2.14! □

Proposition 2.29 and Lemma 2.30 give us a really good idea how the polynomials of nilpotent Maltsev algebras look like. But nevertheless many questions about them are still unanswered.

Question 2.31. For which nilpotent $\mathbf{A}$ is the equivalence problem solvable in polynomial time?

By Lemma 2.30, we can reduce this always to the problem, of deciding whether a polynomial is equivalent to the constant 0 function. As we saw, there is a polynomial time algorithm for this problem in supernilpotent algebras. Also for 2-nilpotent algebras the problem is known to be in P. However, assuming the exponential time hypothesis (an open conjecture in computer science) Idziak et al. constructed very recently (2020) some examples that are not in P.

Another computational question concerns the subpower membership problem:

Question 2.32. The *subpower membership problem* of $\mathbf{A}$ is the computational problem where the input consists of tuples $\bar{b}, \bar{a}_1, \ldots, \bar{a}_n \in A^k$, and the question is if $\bar{b}$ is generated by $\bar{a}_1, \ldots, \bar{a}_n$ in $\mathbf{A}^k$. Is the subpower membership problem in P for finite nilpotent Maltsev algebras $\mathbf{A}$?

This question is also only verified for the supernilpotent case (by results of Peter Mayr); for general Maltsev algebras we only know that the problem is in NP.

At last we mention an open question that is connected to the next chapter of this lecture:

Question 2.33. Let $\mathbf{A}$ be a finite nilpotent algebra with Maltsev term. Is $\mathbf{A}$ *finitely based*? So is there a finite set of identities $\mathcal{E}$, such that $\mathsf{HSP}(\mathbf{A}) = \mathrm{Mod}(\mathcal{E})$?

This is also true for supernilpotent algebras (by results in [3]), but unknown in general.

2.6 Exercises

Exercise 2.1. Recall from UA1 the proof that a variety $\mathcal{V}$ is congruence permutable if and only if it has a Maltsev term.

Exercise 2.2. (Without using Theorem 2.9) show that

1. a semilattice $(L, \wedge)$ is Abelian if and only if $|L| = 1$.
2. a group $(G, \cdot, e, ^{-1})$ is Abelian if and only if $x^{-1}y^{-1}xy \approx e$.
3. a ring $(R, +, 0, -, \cdot)$ is Abelian if and only if $x \cdot y \approx 0$.

(*) What does Abelianness mean in your favorite variety? (E.g.: Show that $x \cdot y \approx y \cdot x$ is not sufficient for loops, quasigroups, monoids, semigroups...)

Exercise 2.3.

- Show that Definition 2.6 is equivalent to satisfying the term condition (2.1) for all polynomial operations of $\mathbf{A}$.
- Show that it is however not sufficient for Abelianness to satisfy (2.1) only for basic operations of $\mathbf{A}$. (Hint: groups)

Exercise 2.4. Complete the missing proof steps in Theorem 2.9: In the proof of (1)→(2) why is $(A, +, 0, -)$ a commutative group? And why are the scalar multiplications distributive with respect to $+$?

Exercise 2.5. Let A be a set and $0 \in A$. An operation $f\colon A^n \to A$ is called 0-*absorbing*, if $f(0, x_2, \ldots, x_n) \approx f(x_1, 0, x_3, \ldots, x_n) \approx \cdots \approx f(x_1, x_2, x_3, \ldots, x_{n-1}, 0) \approx 0$. Show that, if $\mathbf{A}$ is Abelian, all of its 0-absorbing polynomials operations (of arity $n \geq 2$) must be constant.

Compare this with Exercise 2.2, and the proof steps of Theorem 2.9.

Exercise 2.6. Alternatively, try to prove (3)→(2) in Theorem 2.9 with $x + y = m(x, 0, y)$ and $-x = m(0, x, 0)$. Use that m is central with respect to itself. (You can find the solution as Theorem 7.34 in [2])

Exercise 2.7. Show that $(\mathbb{Q}, +)$ is Abelian, but has no Maltsev polynomial. (Hint: Show that $+$ preserves the order $\leq$ on $\mathbb{Q}$ and that Malcev operations do not preserve linear orders). Observe that every polynomial operation $f(x_1, \ldots, x_k)$ of $(\mathbb{Q}, +)$ is equal to some affine combination $\sum_{i=1}^{k} n_i x_i + c$ with $n_i \in \mathbb{Z}$. Why is this not a counterexample to Corollary 2.10?

Exercise 2.8.

- Show that, if $\mathcal{V}$ is a variety with Maltsev term, then also the subclass $\mathcal{V}_{ab}$ containing all Abelian algebras in $\mathcal{V}$ forms a variety.
- Let $\mathcal{E}$ be a set of identities such that $\mathcal{V} = \mathrm{Mod}(\mathcal{E})$. Try to find a 'nice' extension of $\mathcal{E}$ that defines $\mathcal{V}_{ab}$ (hint: m is central)
- (*) What could prevent $\mathcal{V}_{ab}$ from being a variety for general varieties $\mathcal{V}$?

Exercise 2.9. Let A be a 4-element set, and $0 \in A$. There are two ways of defining a commutative group operation on A; Let $+_1$ be such that $(A, +_1) \cong \mathbb{Z}_4$, and $+_2$ be such that $(A, +_2) \cong \mathbb{Z}_2 \times \mathbb{Z}_2$. Show that $(A, +_1, +_2)$ is not Abelian (Hint: this can be done in one line, using the results of Section 2.2).

Exercise 2.10. Recall the definition of a loop $(L, \cdot, /, \backslash, 0)$ (see also Bergman [2], page 5). Show that a loop is Abelian if and only if $\cdot$ is an Abelian group operation.

Exercise 2.11. We saw in UA 1 that for every group G, the commutator subgroup $[G, G]$ is the smallest normal subgroup such that $G/[G, G]$ is Abelian. Can you generalize this statement to arbitrary algebras?

Exercise 2.12. Prove Observation 2.13.

Exercise 2.13. Prove Example 2.16.
Hint: Use that polynomials in rings can be written as sums of monomials.

Exercise 2.14. Let $\mathbf{A}$ be an algebra with Maltsev polynomial $m(x, y, z)$ and let $\alpha, \beta \in \mathrm{Con}(\mathbf{A})$, such that α centralizes β. Let $a, b \in A$ such that $a\,\alpha\,b$. Then prove that $[a]_\beta$ and $[b]_\beta$ have the same size.
Hint: show that the map $x \mapsto m(x, a, b)$ is a bijection from $[a]_\beta$ to $[b]_\beta$.

Exercise 2.15. Show that $C(\alpha, \beta; \delta)$ in Definition 2.11 is equivalent to

$$t(\bar{x}, \bar{u})\,\delta\,t(\bar{x}, \bar{v}) \Rightarrow t(\bar{y}, \bar{u})\,\delta\,t(\bar{y}, \bar{v})$$

for all terms $t \in \mathsf{Clo}(\mathbf{A})$ and all *tuples* $\bar{x}\alpha\bar{y}$ and $\bar{u}\beta\bar{v}$.

Exercise 2.16. Look up the definition of the lower central series and upper central series of a group. How would you generalize the two definitions to arbitrary algebras?

Exercise 2.17. Let $\mathcal{V}$ be a variety with Maltsev term. Then prove that all of its 2-nilpotent members, i.e. all of the algebras satisfying $[[1, 1], 1] = 0$ form a subvariety.

Exercise 2.18.

1. Let $\mathbb{Z}_p$ be the cyclic group of prime order p. How does the congruence lattice of $\mathbb{Z}_p \times \mathbb{Z}_p$ look like? Show in particular that it is not distributive.
2. Derive that no non-trivial group variety is congruence distributive.
3. After reading the proof of Theorem 2.22: Which congruences of $\mathbb{Z}_p \times \mathbb{Z}_p$ correspond to the congruences η_1, η_2, Δ?

Exercise 2.19. Recall from UA1 the proof of the *Maltsev chain lemma*:

- Let $\mathbf{A}$ be an algebra and let $X \subseteq A^2$. Then $Cg_{\mathbf{A}}(X)$, the congruence generated by the pairs in X, is equal to the equivalence relation generated by $\{(p(x), p(y)) : p$ unary polynomial of $\mathbf{A}\}$.
- If $\mathbf{A}$ has additionally a Maltsev polynomial and let $X \subseteq A^2$, then $Cg_{\mathbf{A}}(X)$ is equal to $\{(p(x_1, \ldots, x_n), p(y_1, \ldots, y_n)) : p$ polynomial of $\mathbf{A}$ and $(x_i, y_i) \in X$ for all $i\}$. (This can be found as Theorem 4.65 in [2].)
- Derive that Δ in the proof of Theorem 2.22 is a congruence.

Exercise 2.20. Let $\mathcal{W}$ be a variety with Maltsev term. Show that

1. If two subvarieties $\mathcal{V}_1$ and $\mathcal{V}_2$ of $\mathcal{W}$ are congruence distributive, then so is $\mathcal{V}_1 \vee \mathcal{V}_2$
2. If $\mathcal{W}$ is locally finite, then there is a largest congruence distributive subvariety.

Chapter 3

Finitely based algebras

An *equational basis* (or short, just *basis*) of a variety $\mathcal{V}$ is a set of identities $\mathcal{E}$, such that $\mathcal{V} = \mathrm{Mod}(\mathcal{E})$. We say that $\mathcal{E}$ is an *equational basis of an algebra* $\mathbf{A}$ if $\mathcal{E}$ is a basis of $\mathsf{HSP}(\mathbf{A})$. A variety (or algebra) is called *finitely based* if it has a finite basis.

Many familiar varieties, such as the variety of groups, rings, lattices, or Boolean algebras are finitely based, since they are defined by finitely many identities. An example of a non-finitely based variety is the variety of vector spaces $(V, +, 0, -, (q)_{q \in \mathbb{Q}})$ over $\mathbb{Q}$. It is also defined by identities, but requires infinitely many (to describe the distributivity and composition of the infinitely many scalar multiplications $q \in \mathbb{Q}$). For this reason we generally exclude varieties of infinite type when discussing finite basis results.

For varieties that are given as the HSP closure of an algebra $\mathbf{A}$ (or a class of algebras) it is however in general not clear, if they are finitely based or not. If $\mathbf{A}$ is a finite algebra (so $\mathbf{A}$ has finite universe and finite type), it can be completely specified by a finite amount of information. So one might expect that also its equational theory is generated by finitely many identities - but this is not true in general!

The first counterexample was a 7-element algebra with a binary operation found by Lyndon in 1954. The smallest non-finitely based algebra is a 3-element algebra with a binary operation constructed by Murskiĭ in 1965. We are going to discuss a non-finitely based 4-element algebra with a commutative (but not associative) binary operation later in this section.

On the side of positive results, it is known since the 70ies that every finite group and every finite ring is finitely based. It is still open until today, whether this generalizes to finite algebras with a Maltsev term (see also Question 2.33). Baker showed in 1977 that if $\mathbf{A}$ generates a congruence distributive variety, $\mathbf{A}$ is finitely based. We are going to discuss this in Section 3.1 and prove a result of McKenzie that was a stepping stone to it.

In general there is little hope for a total classification of finitely based algebras: McKenzie showed in 1996 that there is no algorithm that decides whether a finite algebra is finitely based or not.

But why should we care about finitely based varieties in the first place? Having a finite basis is a desirable property for a variety $\mathcal{V}$, since satisfaction of the identities in the basis provides a 'nice' test condition to check if an arbitrary algebra belongs to $\mathcal{V}$.

Example 3.1. Let S_3 the symmetric group on a 3-element set. Then $G \in \mathsf{HSP}(S_3)$ if and only if:

1. G is a quotient group of a subgroup of a power of S_3.

2. G has a normal subgroup N such that N is an elementary abelian 3-group and G/N is an elementary abelian 2-group.
3. G is a group satisfying the identities $x^6 \approx e$, $[x^2, y^2] \approx e$ and $[x, y]^3 \approx e$.

The first characterization, while true, offers little information. The second characterization is more informative, and perhaps the most useful to a group theorist. However it is not an efficient condition to test, as it does not specify how to find the normal subgroup N. The third characterization shows that S_3 is finitely based; and testing if a finite G satisfies the identities can be done in polynomial time (in the size of G and its operation tables).

In fact being finitely based is equivalent to having a 'testing condition' that is definable in first-order logic:

Observation 3.2. A variety $\mathcal{V}$ is finitely based if and only if it has an axiomatization by a first-order sentence.

Proof. If $\mathcal{V}$ has a finite equational basis $\mathcal{E}$ then the logical conjunction of the identities in $\mathcal{E}$ clearly is a first-order sentence, which defines $\mathcal{V}$.

So for the opposite direction, assume that ϕ is a first-order sentence such that $\mathcal{V}$ consist of all models of ϕ. For contradiction assume that $\mathcal{V}$ is not finitely based. So for every finite subset $\mathcal{E} \subseteq \mathrm{Id}(\mathcal{V})$, $\mathcal{E} \cup \{\neg\phi\}$ has a model. By the compactness theorem of first-order logic, also $\mathrm{Id}(\mathcal{V}) \cup \{\neg\phi\}$ has a model, which is a contradiction. □

Before discussing the non-finitely based example, we prove a theorem of Birkhoff, which allows us to 'approximate' finite bases for finite $\mathbf{A}$ in a certain sense. Let us denote by $\mathrm{Id}_n(\mathbf{A})$ the subset of all identities of $\mathrm{Id}(\mathbf{A})$ that only contain variables from $\{x_1, \dots, x_n\}$. Birkhoff's theorem then states as follows:

Theorem 3.3 (Birkhoff '35). *Let $\mathbf{A}$ be a finite algebra, and let $n \in \mathbb{N}$. Then there exists a finite subset $\mathcal{E} \subseteq \mathrm{Id}_n(\mathbf{A})$ such that $\mathcal{E} \models \mathrm{Id}_n(\mathbf{A})$.*

So in other words, the variety generated by $\mathrm{Id}_n(\mathbf{A})$ is always finitely based for finite $\mathbf{A}$. Note that this implies that $\mathbf{A}$ is finitely based, if and only if there is some natural number n such that $\mathrm{Id}_n(\mathbf{A})$ is an equational basis of $\mathbf{A}$.

Proof of Theorem 3.3. Let $\mathbf{T}(x_1, \dots, x_n)$ be the totally free algebra of the same type as $\mathbf{A}$, generated by variables $x_1, \dots, x_n$. Further let $\mathbf{F} = \mathbf{F_A}(x_1, \dots, x_n)$ be the free algebra in $\mathsf{HSP}(\mathbf{A})$, generated by n-many elements.

By definition, $\mathbf{F}$ is isomorphic to $\mathbf{T}(x_1, \dots, x_n)/\mathrm{Id}_n(\mathbf{A})$. On the other hand we know from Universal Algebra 1 (see also Exercise 4.34.3 in [2]) that $\mathbf{F}$ is finite (as it can be considered as a subalgebra of $\mathbf{A}^{A^n}$). Therefore $\mathrm{Id}_n(\mathbf{A})$ has only finitely many equivalence classes.

We fix a set $Q = \{q_1, \dots, q_k\}$ of terms of $\mathbf{T}(x_1, \dots, x_n)$ that represent the equivalence classes of $\mathrm{Id}_n(\mathbf{A})$. Then we define $\mathcal{E}$ as the set of all the identities in $\mathrm{Id}_n(\mathbf{A})$ that are of one of the following forms:

$$\begin{aligned} x_i &\approx x_j \\ x_i &\approx q_l \\ f(q_{l_1}, \dots, q_{l_m}) &\approx q_l \end{aligned}$$

where f is in the type of $\mathbf{A}$ and $1 \leq i, j \leq n$ and $1 \leq l, l_1, \dots, l_m \leq k$. Since Q is finite, and the type of $\mathbf{A}$ is finite, also $\mathcal{E}$ is finite.

We claim that $\mathcal{E} \models p \approx q$ if and only if $(p \approx q) \in \mathrm{Id}_n(\mathbf{A})$. The left to right implication holds by definition of $\mathcal{E}$. For the opposite direction, assume that $(p \approx q) \in \mathrm{Id}_n(\mathbf{A})$. We then prove that $\mathcal{E} \models p \approx q$. Note that it is enough to prove this for $q \in Q$, as every term is equivalent to a unique $q \in Q$ modulo $\mathrm{Id}_n(\mathbf{A})$.

We prove the statement by induction on the size of p. If $p = x_i$, then $(p \approx q) \in \mathcal{E}$ by definition. For an induction step, let $p = f(p_{i_1}, \dots, p_{i_m})$. Then, for each p_{i_k} there is a $q_{i_k} \in Q$, such that $p_{i_k} \approx q_{i_k} \in \mathrm{Id}_n(\mathbf{A})$. By the induction hypothesis $\mathcal{E} \models p_{i_k} \approx q_{i_k}$. By its definition, $\mathcal{E}$ contains the identity $f(q_{i_1}, \dots, q_{i_m}) \approx q$. All together, this implies that $\mathcal{E} \models f(p_1, \dots, p_m) \approx q$, which is what we wanted to prove. □

Example 3.4 (Park '80)**.** The algebra $\mathbf{A} = (A, \cdot)$ with $A = \{0, 1, 2, u\}$ and the following operation table is not finitely based.

$\cdot$	0	1	2	u
0	0	1	u	u
1	1	1	2	u
2	u	2	2	u
u	u	u	u	u

Proof. Note that $\cdot$ is commutative, but not associative. When we write products without brackets, we compute the multiplications left to right, so $x \cdot y \cdot z$ is a shorthand for $(x \cdot y) \cdot z$. We can describe the multiplication table of this algebra nicely via a directed graph with vertices $A \setminus \{u\}$ as depicted in Figure 3.1, by setting $x \cdot y = y \cdot x = y$ if (x, y) is an edge of the graph, and $x \cdot y = u$ else.

For an arbitrary integer $n \geq 2$ we define the two n-ary terms

$$s(x_1, \dots, x_n) \approx x_1 \cdot x_2 \cdots x_{n-1} \cdot x_n \cdot x_1 \cdots x_n$$
$$t(x_1, \dots, x_n) \approx x_2 \cdots x_{n-1} \cdot x_n \cdot x_1 \cdots x_n \cdot x_1$$

We are going to show that $\mathbf{A} \models s \approx t$, but $\mathrm{Id}_{n-1}(\mathbf{A}) \not\models s \approx t$. Therefore there is no n such that the $(n-1)$-ary identities form a basis of $\mathbf{A}$, and so $\mathbf{A}$ cannot be finitely based.

It is straightforward to see that $\mathbf{A} \models s \approx t$, by a case distinction: If $a_i = u$ for some i, then $s(a_1, \dots, a_n) = t(a_1, \dots, a_n) = u$. The subalgebras $(\{0, 1\}, \cdot)$, and $(\{1, 2\}, \cdot)$ are semilattices, so restricted to those values the identity also clearly holds. The only remaining case is when both 0 and 2 appear among the input to s and t. Then, both of $s(a_1, \dots, a_n)$ and $t(a_1, \dots, a_n)$ are of the form $c_1 c_2 \cdots 2 \cdots 0 \cdots c_{2n}$. As we evaluate this expression, multiplying from left to right, the result must be in $\{2, u\}$ once we reach 2, and remain there. When we reach 0 the result must be u. Thus $s(a_1, \dots, a_n) = t(a_1, \dots, a_n) = u$.

We proceed by constructing an algebra $\mathbf{B}$ that satisfies $\mathrm{Id}_{n-1}(\mathbf{A})$, but not $s \approx t$. We define $\mathbf{B}$ to be the algebra given by the graph in Figure 3.1, so its universe is $\{0, 1, \dots, n-1\} \cup \{u\}$, and the multiplication is $x \cdot y = y \cdot x = y$ if $x = y$ or $(y = x + 1 \mod n)$ and $x \cdot y = u$ otherwise.

It is easy to see that

$$s^{\mathbf{B}}(0, 1, 2, \dots, n-1) = n - 1 \neq 0 = t^{\mathbf{B}}(0, 1, 2, \dots, n-1),$$

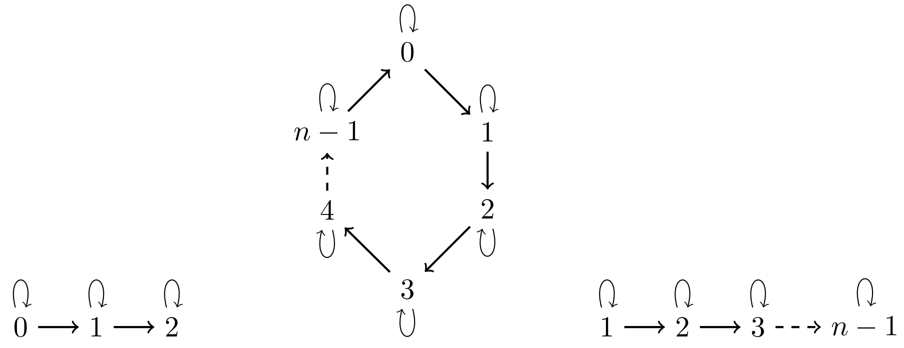

A as graph

B as graph

C as graph

Figure 3.1: The algebras in the proof of Example 3.4.

showing that $\mathbf{B} \not\models s \approx t$.

So it is only left to show that $\mathbf{B}$ is a model of $\mathrm{Id}_{n-1}(\mathbf{A})$. Assume otherwise. Thus there is an identity $(p \approx q) \in \mathrm{Id}_{n-1}(\mathbf{A})$, and $b_1, \dots, b_{n-1} \in B$, such that $p^{\mathbf{B}}(b_1, \dots, b_{n-1}) \neq q^{\mathbf{B}}(b_1, \dots, b_{n-1})$. By a counting argument, there is an element of $\{0, 1, \dots, n-1\}$ that is not among $b_1, \dots, b_{n-1}$. By the symmetry of $\mathbf{B}$ we can assume without loss of generality that $0 \notin \{b_1, \dots, b_{n-1}\}$.

Let $C = B \setminus \{0\}$. Then C is a subuniverse of $\mathbf{B}$. The corresponding subalgebra $\mathbf{C}$ can be described by the graph in Figure 3.1. It follows from the above that $\mathbf{C} \not\models p \approx q$. But $\mathbf{C}$ is in $\mathsf{HSP}(\mathbf{A})$, which can be seen by defining the elements $\bar{c}_i \in A^{n-3}$ by

$$\begin{aligned}
\bar{c}_1 &= (0, 0, \dots, 0, 0, 0) \\
\bar{c}_2 &= (0, 0, \dots, 0, 0, 1) \\
\bar{c}_3 &= (0, 0, \dots, 0, 1, 2) \\
\bar{c}_4 &= (0, 0, \dots, 1, 2, 2) \\
&\vdots \\
\bar{c}_{n-1} &= (2, 2, \dots, 2, 2, 2)
\end{aligned}$$

Then $D = \{\bar{c}_1, \dots, \bar{c}_{n-1}\} \cup \{\bar{y} \in A^{n-3} : u \in \{y_1, \dots, y_{n-3}\}\}$ forms a subuniverse of $\mathbf{A}^{n-3}$. The map that sends each $\bar{c}_i$ to i and all other elements to u is a homomorphisms from $\mathbf{D}$ to $\mathbf{C}$, thus $\mathbf{C} \in \mathsf{HSP}(\mathbf{A})$. This is a contradiction to $\mathbf{C} \not\models p \approx q$.

We conclude that $\mathbf{B}$ is a model of $\mathrm{Id}_{n-1}(\mathbf{A})$, but $\mathbf{B} \not\models s \approx t$, and therefore $\mathbf{B} \notin \mathsf{HSP}(\mathbf{A})$. Hence $\mathbf{A}$ is not finitely based. □

3.1 Park's conjecture and McKenzies DPC theorem

As already mentioned, it is not likely that there will ever be a complete characterization of all finitely based algebras. However we can still try to look for *sufficient* criteria that work in many situations. A promising candidate for such a condition is residual finiteness:

Definition 3.5. A variety $\mathcal{V}$ is *residually finite*, if it contains (up to isomorphism) only finitely many subdirectly irreducible algebras, which are all finite.

Recall from Universal Algebra 1 that an algebra is subdirectly irreducible, if it has a non-trivial congruence that is contained in all other non-trivial congruences. We denote the subdirectly irreducible algebras in $\mathcal{V}$ by $\mathcal{V}_{SI}$. We saw in Universal Algebra 1 that every algebra $\mathbf{A} \in \mathcal{V}$ can be written as a subdirect product of elements from $\mathcal{V}_{SI}$.

So if $\mathcal{V}$ is residually finite, the subdirectly irreducible elements $\mathcal{V}_{SI}$ give us a description of $\mathcal{V}$ by a finite amount of information. Park conjectured in his thesis that this can be used to show the existence of a finite equational basis.

Conjecture 3.6 (Park's conjecture '75)**.** *Let* $\mathbf{A}$ *be a finite algebra, such that* $\mathsf{HSP}(\mathbf{A})$ *is residually finite. Then* $\mathbf{A}$ *is finitely based.*

Until now no counterexample to Park's conjecture was found [1] and it is verified for algebras $\mathbf{A}$ that additionally satisfy one of the following conditions:

1. $\mathsf{HSP}(\mathbf{A})$ is congruence distributive (Baker '77),
2. $\mathsf{HSP}(\mathbf{A})$ is congruence modular (McKenzie '87),
3. $\mathsf{HSP}(\mathbf{A})$ has a difference term (Kearnes, Szendrei, Willard '15).

Without specifying what a difference term is, we remark that congruence modular varieties have always a difference term. So (3) is the strongest known result up to date. Note also that varieties with a Maltsev term are congruence modular, so by (2) Park's conjecture holds for all algebras $\mathbf{A}$ with a Maltsev term.

We are going to prove a result of McKenzie, which states that Park's conjecture is true for algebras with *definable principle congruences (DPC)*. This result is not as strong as the three results above, but nicely illustrates the core idea of the proofs.

Definition 3.7. A variety $\mathcal{V}$ has *definable principle congruences (DPC)* if there is a first-order formula $\phi(u, v, x, y)$ such that for every $\mathbf{A} \in \mathcal{V}$ and all $a, b, c, d \in A$:

$$(c, d) \in \mathrm{Cg}^{\mathbf{A}}(a, b) \Leftrightarrow \mathbf{A} \models \phi(c, d, a, b).$$

We then also say that ϕ defines principle congruences on $\mathcal{V}$.

Having DPC is a relatively strong assumption on a variety. Here are a few examples and non-examples:

Example 3.8.

1. The variety of commutative rings with a unit has DPC. For this recall that congruences in rings $\mathbf{A}$ correspond to ideals; so $(c, d) \in \mathrm{Cg}^{\mathbf{A}}(a, b)$ if and only if $c - d$ lies in the principle ideal generated by $a - b$. Thus $\phi(c, d, a, b) := \exists x(c - d = x \cdot (a - b))$ is equivalent to $(c, d) \in \mathrm{Cg}^{\mathbf{A}}(a, b)$.

[1] You can win a price of 50 Euro from Ross Willard if you find one ;)

2. For any positive integer n the variety $\mathcal{A}_n$ of Abelian groups of exponent n has DPC. In this case $(c,d) \in \mathrm{Cg}^{\mathbf{A}}(a,b)$ if and only if $c-d$ lies in the cyclic subgroup generated by $a-b$. Thus the formula

   $$\phi_n(c,d,a,b) := (c=d) \vee (c-d = a-b) \vee (c-d = 2\cdot(a-b)) \vee \cdots \vee (c-d = n\cdot(a-b))$$

   is equivalent to $(c,d) \in \mathrm{Cg}^{\mathbf{A}}(a,b)$.

3. On the other hand, the variety of *all* Abelian groups $\mathcal{A}$ does not have DPC.

   Note that in any Abelian group $\mathbf{A} \in \mathcal{A}$ we have $(c,d) \in \mathrm{Cg}^{\mathbf{A}}(a,b)$ if and only if there is an n such that $\psi_n(c,d,a,b) := \phi_n(c,d,a,b) \vee \phi_n(d,c,a,b)$ holds. But any individual such formula is not strong enough to define principle congruences (for instance in $\mathbb{Z}$, $\psi_n(2(n+1),0,2,0)$ does not hold, although $2(n+1)$ is in the subgroup generated by 2).

   Now for contradiction, assume that there is some formula $\psi(c,d,a,b)$ that defines principle congruences on $\mathcal{A}$. Let $\mathbf{c},\mathbf{d},\mathbf{a},\mathbf{b}$ be some new constant symbols. By the above paragraph, every finite subset of $\psi(\mathbf{c},\mathbf{d},\mathbf{a},\mathbf{b}) \cup \{\neg\psi_n(\mathbf{c},\mathbf{d},\mathbf{a},\mathbf{b}) : n \in \mathbb{N}\}$ is satisfiable in an Abelian group extended by constant symbols $\mathbf{c},\mathbf{d},\mathbf{a},\mathbf{b}$. By the compactness theorem of first-order logic, also $\psi(\mathbf{c},\mathbf{d},\mathbf{a},\mathbf{b}) \cup \{\neg\psi_n(\mathbf{c},\mathbf{d},\mathbf{a},\mathbf{b}) : n \in \mathbb{N}\}$ has a model. So there is an Abelian group $\mathbf{A}$ with elements $c,d,a,b \in A$ that satisfy $\psi(c,d,a,b)$, but $\neg\psi_n(c,d,a,b)$ for *every* $n \in \mathbb{N}$. But this implies that $\psi(c,d,a,b)$ and $(c,d) \notin \mathrm{Cg}^{\mathbf{A}}(a,b)$ - contradiction!

4. The variety of distributive lattices has DPC (Exercise 3.5).

5. The variety of semilattices has DPC (Exercise 3.6).

6. If a variety has DPC, also all of its subvarieties have DPC (by the same formula).

Interestingly the formulas in Example 3.8 (1) and (2) do not contain universal quantifiers $\forall$ or negations $\neg$. Therefore they are preserved under homomorphisms. In fact it can be proven that, whenever there is is a formula ϕ that defines principle congruences, then there is also an equivalent formula ψ with the following properties:

1. If $h\colon \mathbf{A} \to \mathbf{B}$ is a homomorphism, then $\mathbf{A} \models \psi(c,d,a,b)$ implies $\mathbf{B} \models \psi(h(c),h(d),h(a),h(b))$.
2. For every algebra $\mathbf{A}$: $\mathbf{A} \models (\psi(u,v,x,x) \to u=v)$.

Let us call a 4-ary formula ϕ *nice* [2] if it has the two above properties.

Lemma 3.9. *Let $\mathcal{V}$ be a variety and suppose that $\phi(u,v,x,y)$ defines principle congruences on $\mathcal{V}$. Then, there is a nice formula $\psi(u,v,x,y)$ such that $\mathcal{V} \models \phi(u,v,x,y) \leftrightarrow \psi(u,v,x,y)$.*

Proof idea. This lemma can be proven by a compactness argument, similar to the one in Example 3.8 (3). We only give a short proof sketch here, and refer to Lemma 5.31. and 5.32. in [2] for the interested readers.

The Maltsev chain lemma (see also Exercise 2.19) states that, for every algebra $\mathbf{A} \in V$ and $c,d,a,b \in A$ such that $(c,d) \in \mathrm{Cg}(a,b)$, there is a chain of elements c $= z_0, z_1, \ldots, z_n = d$, and unary polynomials $f_0,\ldots,f_{n-1}$, such that $\{z_i, z_{i+1}\} = \{f_i(a), f_i(b)\}$. Every individual such chain can be described by a nice formula $\psi_{c,d,a,b}(c,d,a,b)$. Now a compactness argument similar to the one in Example 3.8 (3), shows that $\phi(u,v,x,y)$ needs to be equivalent to a finite disjunction of formulas $\psi_{c,d,a,b}$ modulo $\mathrm{Id}(\mathcal{V})$. Such a disjunction is also a nice formula. □

[2] In Bergman's book [2] this is called a 'congruence formula', which in my opinion is a misleading notation, since not every nice formula defines congruences.

On an arbitrary algebra, a nice formula ψ does not need to define principle congruences. But we can describe the algebras, in which it does, by another first-order sentence:

Lemma 3.10. *Suppose that $\psi(u,v,x,y)$ is a nice formula (in a finite language). Then there is a first-order formula $\alpha_\psi(x,y)$, such that $\mathbf{A} \models \alpha_\psi(a,b)$ if and only if $\theta_{a,b} = \{(c,d) \in A^2 : \psi(c,d,a,b)\}$ is equal to $\operatorname{Cg}(a,b)$.*

Proof. Let's consider the following formulas:

1. $\forall u\colon \psi(u,u,x,y)$
2. $\forall u,v\colon \psi(u,v,x,y) \to \psi(v,u,x,y)$
3. $\forall u,v,w\colon \psi(u,v,x,y) \wedge \psi(v,w,x,y) \to \psi(u,w,x,y)$
4. for every n-ary function symbol f: $\forall \bar{u},\bar{v} \bigwedge_{i=1}^n \psi(u_i,v_i,x,y) \to \psi(f(\bar{u}),f(\bar{v}),x,y)$
5. $\psi(x,y,x,y)$,

and let $\alpha_\psi(x,y)$ be the conjunction of all of them. If $\mathbf{A} \models \alpha_\psi(a,b)$ then $\theta_{a,b}$ is an equivalence relation by (1)-(3). By (4) it preserves all operation of $\mathbf{A}$ and is therefore also a congruence of $\mathbf{A}$. By (5) $\theta_{a,b}$ contains (a,b), and thus $\operatorname{Cg}(a,b) \subseteq \theta_{a,b}$.

To see the opposite inclusion, assume that $(c,d) \in \theta_{a,b}$. Let $\mathbf{B} = \mathbf{A}/\operatorname{Cg}(a,b)$ and $h\colon \mathbf{A} \to \mathbf{B}$ be the quotient map. Since ψ is nice, it is preserved under homomorphisms, so $\mathbf{B} \models \psi(h(c),h(d),h(a),h(b))$. But $h(a) = h(b)$ implies that $h(c) = h(d)$ (because ψ is nice) and therefore $(c,d) \in \operatorname{Cg}(a,b)$. □

Theorem 3.11. *Let $\mathcal{V}$ be a variety of finite type, and assume that $\mathcal{V}$ has DPC, and $\mathcal{V}_{SI}$ is (up to isomorphism) a finite set of finite algebras. Then $\mathcal{V}$ is finitely based.*

Proof. Let $\{\mathbf{C}_1,\dots\mathbf{C}_n\}$ be an enumeration of the elements of $\mathcal{V}_{SI}$. For every $\mathbf{C}_i$ there is a first-order sentence γ_i such that $\mathbf{A} \models \gamma_i$ if and only if $\mathbf{A}$ is isomorphic to $\mathbf{C}_i$ (see also Exercise 3.7). So $\gamma = \gamma_1 \vee \cdots \vee \gamma_n$ holds in $\mathbf{A}$ if and only if $\mathbf{A}$ is isomorphic to an element of $\mathcal{V}_{SI}$.

By Lemma 3.9 there is a nice formula $\psi(u,v,x,y)$ defining principle congruences in $\mathcal{V}$. Then let α be the sentence $\forall x,y\, \alpha_\psi(x,y)$, where α_ψ is as in Lemma 3.10. An algebra satisfies α, if and only if ψ defines principle congruences on it. Therefore every element of $\mathcal{V}$ satisfies α.

Next let β be the sentence $\exists u \neq v \forall x \neq y(\psi(u,v,x,y))$. If an algebra $\mathbf{A}$ satisfies α and β, this implies that there exists $u \neq v \in A$ such that $(u,v) \in \operatorname{Cg}^{\mathbf{A}}(x,y)$, for all pairs of distinct x,y. In other words, $\operatorname{Cg}^{\mathbf{A}}(u,v)$ is contained in all other non-trivial congruences of $\mathbf{A}$, so $\mathbf{A}$ is subdirectly irreducible.

By the above observations, all algebras in $\mathcal{V}$ satisfy the formula $\alpha \wedge (\beta \leftrightarrow \gamma)$. So $\operatorname{Id}(\mathcal{V}) \models \alpha \wedge (\beta \leftrightarrow \gamma)$. By the compactness theorem of first order logic, there is a finite subset $\mathcal{E} \subset \operatorname{Id}(\mathcal{V})$, such that $\mathcal{E} \models \alpha \wedge (\beta \leftrightarrow \gamma)$.

We claim that $\mathcal{E}$ is a finite basis of $\mathcal{V}$. Clearly every element of $\mathcal{V}$ is in $\operatorname{Mod}(\mathcal{E})$. For the opposite direction it is sufficient to show that the subdirectly irreducible elements of $\operatorname{Mod}(\mathcal{E})$ are in $\mathcal{V}_{SI}$. But this follows from the fact that $\mathcal{E} \models \alpha \wedge (\beta \leftrightarrow \gamma)$, so every subdirectly irreducible algebra in $\operatorname{Mod}(\mathcal{E})$ needs to satisfy γ. □

Note that the proof of Theorem 3.11 is not constructive. So although Theorem 3.11 states that some varieties are finitely bounded it gives us no method to find an explicit equational basis.

We finish by giving two examples, in which the theorem holds:

Example 3.12. Let $\mathbf{R}$ be a finite commutative ring with a unit satisfying $x^n \approx x$. We saw in Example 3.8 (1) that $\mathsf{HSP}(\mathbf{R})$ has definable principle congruences. On the other hand, we know from Bergman's book (see Theorem 3.28 in [2]) that the only subdirectly irreducible rings satisfying $x^n \approx x$ are the finite fields whose order divides $n-1$. There are only finitely many such fields, so $\mathsf{HSP}(\mathbf{R})$ is finitely based by Theorem 3.11.

Example 3.13. Jónsson's Lemma (which was discussed in Universal Algebra 1) implies that for finite $\mathbf{A}$, such that $\mathsf{HSP}(\mathbf{A})$ is congruence distributive, $\mathsf{HSP}(\mathbf{A})$ is residually finite. So, whenever $\mathsf{HSP}(\mathbf{A})$ is congruence distributive, and has DPC, $\mathbf{A}$ is finitely based by Theorem 3.11.

This is for example true if $\mathbf{A}$ is a distributive lattice (by Exercise 3.5).

3.2 Exercises

Exercise 3.1. (*) Try to prove the statements in Example 3.1 (showing (3)→(1) or (2)→(1) is quite challenging... So don't worry, if you don't manage to prove it).

Exercise 3.2. Let $\mathbf{A}$ and $\mathbf{A}'$ be two algebras in finite type with $\mathsf{Clo}(\mathbf{A}) = \mathsf{Clo}(\mathbf{A}')$. Show that then $\mathbf{A}$ is finitely based if and only if $\mathbf{A}'$ is finitely based. Is this still true if we don't assume finite type?

Exercise 3.3. Try to prove the following (using the compactness theorem): If $\mathcal{V}$ is finitely based, every basis of $\mathcal{V}$ contains a finite subset that is already a basis.

Exercise 3.4. For any directed graph $G = (V, E)$, we can define a binary algebra $\mathbf{A} = (V \cup \{u\}, \cdot)$ as in Example 3.4. Does $\mathsf{HSP}(\mathbf{A})$ then contain the algebras corresponding to

- (induced) subgraphs of G?
- powers of G?
- homomorphic images of G?

(This and other versions of 'graph algebras' are often used to construct interesting examples in UA - see `https://en.wikipedia.org/wiki/Graph_algebra`)

Exercise 3.5. Let $\mathbf{L}$ be a distributive lattice. Prove that $(c, d) \in \mathrm{Cg}^{\mathbf{L}}(a, b)$ if and only if

$$c \wedge (a \wedge b) = d \wedge (a \wedge b) \text{ and } c \vee (a \vee b) = d \vee (a \vee b).$$

Thus the variety of distributive lattice has DPC.

Exercise 3.6. Prove that the variety of semilattices has DPC.

Exercise 3.7. In the proof of Theorem 3.11 we use that, for every finite algebra $\mathbf{C}$ there is a first-order formula γ such that $\mathbf{B} \models \gamma$ if and only if $\mathbf{B}$ is isomorphic to $\mathbf{C}$. Take a moment to think why this is true (Hint: up to isomorphism $\mathbf{C}$ is determined by having exactly $|C|$-many distinct elements, and its operation tables).

Exercise 3.8. Let $\mathbf{A}$ be a finite nilpotent algebra with Maltsev term. Show that $\mathsf{HSP}(\mathbf{A})$ is residually finite, if and only if $\mathbf{A}$ is Abelian (use Theorem 2.21).

Chapter 4

CSPs and Maltsev conditions

4.1 Constraint satisfaction problems (CSP)

Constraint satisfaction problems (short *CSPs*) are a broad class of computational problems that have many theoretical and real-life applications. Because of this CSPs are an important research topic in computer science (in particular in artificial intelligence and operations research).

But CSPs also had a major impact on universal algebra in the last two decades, and lead to many new results in the study of Maltsev conditions. In this section we define CSPs over a fixed relational structure, illustrate them by some basic examples, and discuss how the $\mathrm{Pol} - \mathrm{Inv}$ Galois-connection can be used in studying the complexity of CSPs.

Definition 4.1. A *relational structure* $\mathbb{A} = (A; R_1, \ldots, R_t)$ consists of a set A and a family of relations $R_i \subseteq A^{k_i}$ on it. We are only going to study finite relational structures, so $\mathbb{A}$ will always have a finite domain A, and finitely many relations $R_1, \ldots, R_t$.

A *primitive positive sentence* (or *pp-sentence*) over $\mathbb{A}$ is a first order sentence that can be constructed using the relations of $\mathbb{A}$, equality $=$, conjunctions $\wedge$ and existential quantification $\exists$ (for example $\exists x, y, z\, R_1(x) \wedge R_2(y, x, x) \wedge R_2(z, z, y)$).

Now the constraint satisfaction problem $\mathrm{CSP}(\mathbb{A})$ of a relational structure $\mathbb{A}$ is defined as:

$\mathrm{CSP}(\mathbb{A})$
INPUT: A primitive positive sentence ϕ over $\mathbb{A}$.
QUESTION: Is ϕ true in $\mathbb{A}$?

A lot of important computational problems fit this framework:

Example 4.2. 3-COLOR is the computational problem, where the input is a finite graph (V, E), and the question is if it can be colored by 3 colors (so if there is a map $c\colon V \to \{1, 2, 3\}$ such that for every edge $(x, y) \in E$ we have $c(x) \neq c(y)$). This problem is equal to $\mathrm{CSP}((\{1, 2, 3\}; \neq))$.

To illustrate this consider for example the following graph:

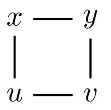

This graph is 3-colorable if and only if the pp-sentence $\exists x, y, u, v\, (x \neq y) \wedge (x \neq u) \wedge (y \neq v) \wedge (u \neq v)$ is satisfied in $\{1, 2, 3\}$; here the values of the variables x, y, u, v correspond to the colors of the vertices; the constraints $(x \neq y), (x \neq u), \ldots$ describe the condition that two connected vertices have different colors.

Example 4.3. In 3-SAT the input is a Boolean formula ϕ, given as conjunction of 3-clauses (for example $(x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_2 \vee \neg x_4 \vee x_1) \wedge (x_1 \vee x_3 \vee x_3))$. The question then is if ϕ is satisfiable in $\{0, 1\}$.

Note that every 3-clause can be modelled by a ternary relation on $\{0, 1\}$ (for example $x_1 \vee \neg x_2 \vee x_3$ is equivalent to $(x_1, x_2, x_3) \in \{0, 1\}^3 \setminus \{(0, 1, 0)\})$. Therefore 3-SAT is equivalent to $\mathrm{CSP}((\{0, 1\}; R_{000}, R_{001}, R_{011}, R_{111}))$, where $R_{ijk} = \{0, 1\}^3 \setminus \{(i, j, k)\}$.

Example 4.4. Let p be a fixed prime. 3-LIN-p is defined as the problem whose input is a system of linear equations over $\mathbb{Z}_p$, such that each equation contains at most 3 variables. The question is: Does this system have a solution or not?

3-LIN-p can be modelled by $\mathrm{CSP}(\mathbb{Z}_p; (R_{\bar{a}})_{\bar{a} \in \mathbb{Z}_p^4})$, where $R_{\bar{a}}$ is a ternary relation defined by $R_{\bar{a}}(x, y, z) \Leftrightarrow a_1 x + a_2 y + a_3 z = a_4$, for every $\bar{a} = (a_1, a_2, a_3, a_4)$.

We remark that there are many related problems (finding explicit solutions to some input, counting the number of solutions, finding 'approximations' to solutions, infinite domain CSP, promise CSPs, etc.) that the CSP community is studying. But we are only going to focus on 'classical' CSPs here.

Then the question is: **What is the computational complexity of** $\mathrm{CSP}(\mathbb{A})$ **for a given** $\mathbb{A}$**?** We are primarily interested in two complexity classes:

- P: problems that can be solved in polynomial time (with respect to the input size). This includes 3-LIN-p in Example 4.4, since any system of linear equations can be solved by Gaussian elimination in polynomial time. Problems in P are usually considered to be computationally 'easy'.
- NP: problems, for which solutions can be verified in polynomial time. 3-COLOR in Example 4.2 is in NP, since for every map $c\colon V \to \{1, 2, 3\}$ we can check in polynomial time, if it indeed is a 3-coloring of the graph.

Actually for *every* finite relational structure $\mathbb{A}$, $\mathrm{CSP}(\mathbb{A}) \in \mathsf{NP}$, since it only takes polynomial time to check, if an assignment of values ot the variables is a solution. Now $\mathsf{P} \subseteq \mathsf{NP}$, and it is widely believed that this is a proper inclusion. So a major goal in studying CSPs is (or was) to classify which of them are in P (and therefore 'easy'), and which are not (assuming that $\mathsf{P} \neq \mathsf{NP}$).

There is a natural way of comparing the complexity of two problems. We say that a problem S *reduces* to a problem Q in polynomial time (and write $S \leq Q$), if there is a polynomial time algorithm that transforms an input I of S in polynomial time to an input I_Q of Q, such that I is a YES-instance of S if and only if I_Q is a YES-instance of Q.

Finally, a problem Q is called NP-complete, if $Q \in \mathsf{NP}$, and $S \leq Q$ holds for every other problem $S \in \mathsf{NP}$. So NP-complete problems are, in some sense, the 'hardest' problem in NP. [1]

[1] The definitions of P, NP and NP-completeness can be made much more precise using Turing machines, as most of you probably know from some introduction to theoretical computer science / complexity theory. But for our purposes this is not necessary.

3-SAT and 3-COLOR are well-known examples of NP-complete problems. Starting from the beginning of CSP research it was observed that all studied examples of CSPs are either in P (like 3-LIN-p) or NP-complete (like 3-SAT):

- In 1978 Schaefer gave a full classification of structure $\mathbb{A}$ with $|A| = 2$ and showed that CSP($\mathbb{A}$) is either in P or NP-complete.
- In 1990 Hell and Nešetřil showed that for all symmetric graphs $\mathbb{A}$, CSP($\mathbb{A}$) is either in P or NP-complete.

These observations lead to the conjecture of Feder and Vardi (in 1993) that CSP($\mathbb{A}$) is always either in P or NP-complete.

But what does universal algebra have to do with all this? In 1998 Jeavons observed that the clone of polymorphisms of $\mathbb{A}$, Pol($\mathbb{A}$) actually determines the complexity of CSP($\mathbb{A}$). This lead to conjectures about *algebraic* conditions on Pol($\mathbb{A}$) that separate the problems in P or NP-complete (around 2000), followed by many partial results on CSPs and many new results in universal algebra (also by Libor and members of our group). Ultimately two proofs of the conjecture were given by Bulatov [7] and Zhuk [6] in 2017.

In the following we discuss the very basic of this 'universal algebraic approach' to CSPs.

Observation 4.5. Let $\mathbb{A} = (A; R_1, \dots, R_n)$ and $\mathbb{B} = (A; S_1, \dots, S_k)$ be two structures on the same set A, and assume that every relation of $\mathbb{B}$ is *pp-definable* from $\mathbb{A}$. So every S_i can be defined using relation of $\mathbb{A}$, conjunctions and existential quantifiers. Then CSP($\mathbb{B}$) reduces to CSP($\mathbb{A}$).

Proof. To see this it is enough to see that every instance of CSP($\mathbb{B}$) can be turned into an equivalent instance of CSP($\mathbb{A}$) by substituting every constraint S_i by its definition. □

Example 4.6. In the structure $(\{0,1\}; R_{000}, R_{001}, R_{011}, R_{111})$ in Example 4.3 we can pp-define the unary relation $C_0 = \{0\}$ by the formula $C_0(x) \leftrightarrow R_{111}(x,x,x)$; also $XOR = \{(0,1),(1,0)\}$ has a pp-definition by $XOR(x,y) \leftrightarrow R_{000}(x,y,y) \wedge R_{111}(x,y,y)$, or alternatively $XOR(x,y) \leftrightarrow \exists a,b\colon R_{000}(x,y,a) \wedge R_{111}(x,y,b) \wedge C_0(a) \wedge C_1(b)$. In fact, it can be shown that every relation on the two element set has a pp-definition (maybe you already know this by some logic or CS lecture).

Observation 4.5 allows to compare CSPs, and was already used by Schaefer to prove his result on 2-element CSPs. But this is not very deep in itself. As it turns out, the reductions from Observation 4.5 have an algebraic counterpart. Let recall the definition of polymorphisms from Universal Algebra 1:

Definition 4.7. A function $f\colon A^n \to A$ *preserves* a relation $R \subseteq A^k$, if for all tuples $\bar{r}_1, \dots, \bar{r}_n \in R$, also $f(\bar{r}_1, \dots, \bar{r}_n) \in R$. Here $f(\bar{r}_1, \dots, \bar{r}_n)$ is computed row-wise, so:

$$\begin{pmatrix} r_{11} \\ r_{12} \\ \vdots \\ r_{1k} \end{pmatrix}, \begin{pmatrix} r_{21} \\ r_{22} \\ \vdots \\ r_{2k} \end{pmatrix}, \dots, \begin{pmatrix} r_{n1} \\ r_{n2} \\ \vdots \\ r_{nk} \end{pmatrix} \in R \Rightarrow \begin{pmatrix} f(r_{11}, r_{21}, \dots, r_{n1}) \\ f(r_{12}, r_{22}, \dots, r_{n2}) \\ \vdots \\ f(r_{1k}, r_{2k}, \dots, r_{nk}) \end{pmatrix} \in R$$

An operation $f\colon A^n \to A$ is called a *polymorphism* of $\mathbb{A}$, if it preserves all the relations of $\mathbb{A}$. The set of all polymorphisms forms a clone, which is denoted by Pol($\mathbb{A}$).

The relations that are preserved under some set of operations $\mathcal{F}$ are also called invariant relations $\mathrm{Inv}(\mathcal{F})$. As discussed in Universal Algebra 1, Pol and Inv forms a Galois connection between operations and relations on a set A. The Galois-closed subsets of operations are the (polymorphism) clones. We are going to prove that the Galois-closure on the relational side corresponds exactly to the pp-definable relations:

Proposition 4.8. *Let $\mathbb{A} = (A; R_1, \dots, R_t)$ be a finite structure, and $S \subseteq A^k$. Then S is preserved by* $\mathrm{Pol}(\mathbb{A})$ *if and only if S is pp-definable in $\mathbb{A}$.*

Proof. It is easy to see that, pp-definable relations of $\mathbb{A}$ are preserved by $\mathrm{Pol}(\mathbb{A})$.

For the opposite direction assume that S is invariant under $\mathrm{Pol}(\mathbb{A})$. Let $S = \{\bar{a}_1, \dots \bar{a}_l\} \subseteq A^k$ so that $\bar{a}_1, \dots \bar{a}_l$ are the column vectors of the following matrix:

$$\begin{pmatrix} a_{11} & a_{21} & \cdots & a_{l1} \\ a_{12} & a_{22} & \cdots & a_{l2} \\ \vdots & \vdots & & \vdots \\ a_{1k} & a_{2k} & \cdots & a_{lk} \end{pmatrix}$$

Without loss of generality we can assume that this matrix has no duplicate rows (if for instance the first and second row are equal, then $S(x_1, x_2, x_3, \dots, x_k)$ is equivalent to $S'(x_1, x_3, \dots, x_k) \wedge x_1 = x_2$, for an invariant relation S').

Next, we extend this matrix in a way, such that the rows enumerate all tuples in A^l:

$$\left(\begin{array}{cccc} a_{11} & a_{21} & \cdots & a_{l1} \\ a_{12} & a_{22} & \cdots & a_{l2} \\ \vdots & \vdots & & \vdots \\ a_{1k} & a_{2k} & \cdots & a_{lk} \\ \hline \vdots & \vdots & & \vdots \\ a_{1m} & a_{2m} & \cdots & a_{lm} \end{array}\right)$$

with $m = |A|^l$. Let then $\bar{a}_1^+, \dots, \bar{a}_l^+$ the resulting column vectors and $Q = \{f(\bar{a}_1^+, \dots, \bar{a}_l^+) \colon f \in \mathrm{Pol}(\mathbb{A})\}$. We claim that Q has a pp-definition in $\mathbb{A}$. If we can prove this claim, then also S is pp-definable in in $\mathbb{A}$, since S is the projection of Q to the first k coordinates; in other words $S(x_1, \dots, x_k)$ is equivalent to $\exists x_{k+1}, \dots, x_m Q(x_1, \dots, x_k, x_{k+1}, \dots, x_m)$.

To prove the claim, note first that whenever we have a tuple $\bar{b} \in A^m$, there is a unique map $f \colon A^l \to A$ with $f(\bar{a}_1^+, \dots, \bar{a}_l^+) = \bar{b}$. This follows from the fact that the rows of the above matrix enumerate A^l. Let $\phi(x_1, \dots, x_m)$ be the conjunction of all atomic formulas that hold for all tuples $\bar{a}_i^+$. (For instance, if $R_1(a_{i1}, a_{i3}, a_{i1})$ holds for every $i = 1, \dots, l$, then $R_1(x_1, x_3, x_1)$ is part of the conjunction ϕ.)

We are going to show that ϕ is a pp-definition of Q. It is easy to see that every tuple in Q satisfies ϕ. For the opposite direction, assume $\bar{b}$ satisfies ϕ. Since the rows of $\bar{a}_1^+, \dots, \bar{a}_l^+$ enumerate A^l, there is a unique map $f \colon A^l \to A$, such that $f(\bar{a}_1^+, \dots, \bar{a}_l^+) = \bar{b}$. By definition of ϕ, f preserves all relations of $\mathbb{A}$. So f is a polymorphism of $\mathbb{A}$ and therefore $\bar{b} = f(\bar{a}_1^+, \dots, \bar{a}_l^+) \in Q$. □

Proposition 4.8 together with Observation 4.5 now directly implies:

Corollary 4.9. *Let $\mathbb{A}$ and $\mathbb{B}$ be finite structures such that* $\mathrm{Pol}(\mathbb{A}) \subseteq \mathrm{Pol}(\mathbb{B})$. *Then $\mathbb{B}$ is pp-definable in $\mathbb{A}$, and therefore* $\mathrm{CSP}(\mathbb{B})$ *reduces to* $\mathrm{CSP}(\mathbb{A})$.

So structures with small polymorphism clones have computationally hard CSPs. As an immediate consequence we get the following corollary:

Corollary 4.10. *Let $\mathbb{A}$ be a relational structure such that $|A| \geq 2$, and* $\mathrm{Pol}(\mathbb{A})$ *consists only of projections. Then* $\mathrm{CSP}(\mathbb{A})$ *is* NP*-complete.*

Proof. For every other relational structure $\mathbb{B}$ with the same domain A, we have that $\mathrm{Pol}(\mathbb{A}) \subseteq \mathrm{Pol}(\mathbb{B})$. By Proposition 4.8, $\mathbb{B}$ is pp-definable in $\mathbb{A}$, and therefore $\mathrm{CSP}(\mathbb{B})$ reduces to $\mathrm{CSP}(\mathbb{A})$. Since there is a $\mathbb{B}$ with NP-complete CSP, also $\mathrm{CSP}(\mathbb{A})$ needs to be NP-complete. □

More generally we can ask: How small does the polymorphism clone of $\mathbb{A}$ need to be, such that $\mathrm{CSP}(\mathbb{A})$ is NP-complete? Is the only possible reason that $\mathrm{Pol}(\mathbb{A})$ is the projection clone? The answer to this question is NO: In the example of 3-COLOR, $\mathrm{Pol}(\{1,2,3\},\neq)$ does not only consist of projections, since all the permutations of $\{1,2,3\}$ are unary polymorphisms of $(\{1,2,3\},\neq)$.

Next time we are going to discuss, which other algebraic conditions imply NP-completeness, and lead to the algebraic dichotomy conjecture.

4.2 Clone and minion homomorphisms

We saw in the last section, that $\mathrm{CSP}(\mathbb{B})$ reduces to $\mathrm{CSP}(\mathbb{A})$ if $\mathrm{Pol}(\mathbb{A}) \subseteq \mathrm{Pol}(\mathbb{B})$. This allowed us to compare the complexity of CSPs on the same set. There are some more general reductions between CSPs on (possibly) different sets, that correspond to special maps between $\mathrm{Pol}(\mathbb{A})$ and $\mathrm{Pol}(\mathbb{B})$, namely clone and and minion homomorphisms. We introduce clone and and minion homomorphisms in this section.

Recall from UA1 that a *clone* $\mathscr{A}$ is a set of operations on a set A, such that

- $\mathscr{A}$ contains all projections $\pi_i^n \colon A^n \to A$ defined by $\pi_i^n(x_1,\dots,x_n) = x_i$, for all $1 \leq i \leq n$
- $\mathscr{A}$ is closed under composition. So if $g \in \mathscr{A}$ is n-ary, $f_1,\dots,f_n \in \mathscr{A}$, then also $g \circ (f_1,\dots,f_n) \in \mathscr{A}$, where $g \circ (f_1,\dots,f_n)$ is defined by

$$g \circ (f_1,\dots,f_n)(x_1,\dots,x_m) = g(f_1(x_1,\dots,x_m),\dots,f_n(x_1,\dots,x_m)).$$

For every algebra $\mathbf{A}$, the set of its term operation $\mathsf{Clo}(\mathbf{A})$ forms a clone. The polymorphisms $\mathrm{Pol}(\mathbb{A})$ of a relational structure $\mathbb{A}$ also naturally form a clone.

Clone homomorphisms are defined as follows:

Definition 4.11. Let $\mathscr{A},\mathscr{B}$ be clones. A map $\xi \colon \mathscr{A} \to \mathscr{B}$ is called a *clone homomorphism* if

1. $\xi(f)$ has the same arity as f, for every $f \in \mathscr{A}$,
2. $\xi(\pi_i^n) = \pi_i^n$ for all projections [2]
3. ξ preserves compositions:

$$\xi(g \circ (f_1,\dots,f_n)) = \xi(g) \circ (\xi(f_1),\dots,\xi(f_n)).$$

[2] Here we are lazy with notation, and do not distinguish between projections on different sets.

Let us write $\mathscr{A} \leq \mathscr{B}$ if there is a clone homomorphism $\xi \colon \mathscr{A} \to \mathscr{B}$, and $\mathscr{A} \sim \mathscr{B}$ if $\mathscr{A} \leq \mathscr{B} \leq \mathscr{A}$.

Note that a clone homomorphism is a map from the operations in $\mathscr{A}$ to the operations in $\mathscr{B}$; this definition does not mention the universes of $\mathscr{A}$ or $\mathscr{B}$ at all. So clone homomorphisms should not be confused with homomorphisms between algebras.

Observation 4.12. By definition, clone homomorphisms preserve arbitrary compositions of functions in $\mathscr{A}$. Let for instance $f, g \in \mathscr{A}$ and $h(x,y) = f(g(x,y,x),x)$, which is equivalent to $h = f \circ (g \circ (\pi_1^2, \pi_2^2, \pi_1^2), \pi_1^2)$. If $\xi \colon \mathscr{A} \to \mathscr{B}$ is a clone homomorphism, then $\xi(h) = \xi(f) \circ (\xi(g) \circ (\pi_1^2, \pi_2^2, \pi_1^2), \pi_1^2)$; or in other words $\xi(h)(x,y) = \xi(f)(\xi(g)(x,y,x),x)$ in $\mathscr{B}$.

Observation 4.13. Clone homomorphisms preserve identities. For example, let $\xi \colon \mathscr{A} \to \mathscr{B}$ be a clone homomorphism and assume that $\mathscr{A}$ has a Maltsev operation. So there is an $m \in \mathscr{A}$, such that $x \approx m(x,y,y) \approx m(y,y,x)$ for all $x, y \in A$. Since ξ is a clone homomorphism, also $\xi(m)(x,y,y) \approx \xi(m)(y,y,x) \approx x$ for all $x, y \in B$. So if $\mathscr{A}$ has a Maltsev operation, then also $\mathscr{B}$ has one.

By Observation 4.13, clone homomorphisms are exactly those maps between clones that preserve identities. Therefore $\mathscr{A} \leq \mathscr{B}$ holds if and only if all the identities that hold in $\mathscr{A}$ hold also in $\mathscr{B}$. Note also, that we use 'identities' here in a different way than for algebras, since clones do not have a fixed similarity type. (So we say $\mathscr{A}$ satisfies $x \approx m(x,y,y) \approx m(y,y,x)$ if *there is* a $m \in \mathscr{A}$ such that for all $x, y \in A$: $x = m(x,y,y) = m(y,y,x)$). If $\mathscr{A} = \mathsf{Clo}(\mathbf{A})$ is the term clone of some algebra, these 'clone identities' are also called *Maltsev conditions on* $\mathbf{A}$ (for instance it is a Maltsev condition on $\mathbf{A}$ to have a Maltsev term) [3].

Modulo the equivalence $\sim$, the relation $\leq$ is a partial order. In fact, it is even a lattice order, which is called the *interpretability lattice*. The minimal element in this lattice is the equivalence class given by clone of projections on a two element set (see Exercise 4.11). We write Proj for this particular clone.

There are some natural examples of clone homomorphisms:

Example 4.14.

1. If $\mathscr{A} \subseteq \mathscr{B}$, then the inclusion map $\xi(f) = f$ is a clone homomorphism. We then write $\mathscr{B} \in E(\mathscr{A})$.
2. Let $\mathbf{A}$ be an algebra, and $\mathscr{A} = \mathrm{Clo}(\mathbf{A})$ its term clone. Let $\mathbf{B} \leq \mathbf{A}$, and $\mathscr{B} = \mathrm{Clo}(\mathbf{B})$. Then the map $\xi \colon \mathscr{A} \to \mathscr{B}$ which maps $t \in \mathscr{A}$ to its restriction $t|_B \in \mathscr{B}$ is a clone homomorphism. In this case we write $\mathscr{B} \in S(\mathscr{A})$.
3. Let $\mathbf{A}$ be an algebra, $\mathscr{A} = \mathrm{Clo}(\mathbf{A})$ and $\mathscr{B} = \mathrm{Clo}(\mathbf{A}^n)$ for some n. Then the map $\xi \colon t^{\mathbf{A}} \to t^{\mathbf{A}^n}$ is a clone homomorphism. In this case we write $\mathscr{B} \in P(\mathscr{A})$.
4. Let $\mathbf{A}$ be an algebra, and $\mathbf{B}$ be a homomorphic image of $\mathbf{A}$. Then the map $\xi \colon t^{\mathbf{A}} \to t^{\mathbf{B}}$ is a clone homomorphism from $\mathscr{A} = \mathsf{Clo}(\mathbf{A})$ to $\mathscr{B} = \mathsf{Clo}(\mathbf{B})$. In this case we write $\mathscr{B} \in H(\mathscr{A})$.

Take a moment to think why the maps in Example 4.14 are indeed clone homomorphisms. The clone homomorphisms described in 2,3,4 correspond to closing algebras under HSP. This is not a coincidence! In fact, every clone homomorphism can be constructed as a composition of the clone homomorphisms described in Example 4.14, by the following version of Birkhoff's theorem:

[3] Unfortunately mathematicians are sometimes a bit uncreative in naming things

Theorem 4.15 (Birkhoff's HSP theorem for clones)**.** *Let $\mathscr{A}, \mathscr{B}$ be two clones. Then the following are equivalent:*

1. $\mathscr{A} \leq \mathscr{B}$
2. $\mathscr{B} \in EHSP(\mathscr{A})$

Proof. We already saw that (2) implies (1) in Example 4.14.

For the opposite direction, assume that $\xi \colon \mathscr{A} \to \mathscr{B}$ is a clone homomorphism. Note that $\mathscr{B}$ is an extension of $\xi(\mathscr{A})$. So without loss of generality we can assume that $\mathscr{B} = \xi(\mathscr{A})$; we are then going to prove that $\mathscr{B} \in \mathsf{HSP}(\mathscr{A})$.

For this, let $\mathbf{A}$ be an algebra such that $\mathscr{A} = \mathsf{Clo}(\mathbf{A})$; we can find such an $\mathbf{A}$, by simply fixing an enumeration $f_1, f_2, \ldots$ of all elements of $\mathscr{A}$ and setting $\mathbf{A} = (A; f_1, f_2, \ldots)$. We further define the algebra $\mathbf{B} = (B; \xi(f_1), \xi(f_2), \ldots)$, its term clone is equal to $\mathscr{B}$. We saw in Observation 4.13 that clone homomorphisms preserve identities. Therefore $\mathbf{B}$ satisfies the same identities as $\mathbf{A}$. By Birkhoff's theorem $\mathbf{B} \in \mathsf{HSP}(\mathbf{A})$; and therefore $\mathscr{B} \in \mathsf{HSP}(\mathscr{A})$. □

How is this relevant to the study of finite CSPs? Corollary 4.9 states that $\mathrm{Pol}(\mathbb{B}) \in E(\mathrm{Pol}(\mathbb{A}))$ implies that $\mathbb{B}$ is pp-definable in $\mathbb{A}$, and therefore $\mathrm{CSP}(\mathbb{B}) \leq \mathrm{CSP}(\mathbb{A})$. More generally it is possible to show that $\mathrm{Pol}(\mathbb{B}) \in EHSP(\mathrm{Pol}(\mathbb{A}))$, if and only if there is a so-called *pp-interpretation* of $\mathbb{B}$ in $\mathbb{A}$, which again implies that $\mathrm{CSP}(\mathbb{B})$ reduces to $\mathrm{CSP}(\mathbb{A})$ [4]. We summarize:

Proposition 4.16. *Let* $\mathbb{A}$*,* $\mathbb{B}$ *be finite structures. Assume that there is a clone homomorphism from* $\mathrm{Pol}(\mathbb{B})$ *to* $\mathrm{Pol}(\mathbb{A})$*. Then* $\mathrm{CSP}(\mathbb{A})$ *reduces to* $\mathrm{CSP}(\mathbb{B})$ *in polynomial time.*

In particular, it follows that

- $\mathrm{CSP}(\mathbb{A})$ is NP-complete if $\mathrm{Pol}(\mathbb{A}) \leq \mathrm{Proj}$
- the complexity of $\mathrm{CSP}(\mathbb{A})$ only depends on the identities holding in $\mathrm{Pol}(\mathbb{A})$.

We already saw an Example of this principle in Exercise 4.3. The statement there can in fact be generalized to the following: If $\mathrm{Pol}(\mathbb{A})$ satisfies the identities $g(x, g(y, z)) \approx g(g(x, y), z)$, $g(x, y) \approx g(y, x)$ and $g(x, x) \approx x$, then $\mathrm{CSP}(\mathbb{A}) \in \mathsf{P}$-

We give another example: Let $\mathbb{A}$ be a relational structure, such that there is a constant polymorphism $f \in \mathrm{Pol}(\mathbb{A})$. Note that f is constant if and only if it satisfies $f(x) \approx f(y)$. Let c be the image of f. Note that every non-empty relation in $\mathbb{A}$ contains the tuple $(c, c, \ldots, c)$. This implies that every instance of $\mathrm{CSP}(\mathbb{A})$, which does not contain empty relations, has $(c, c, \ldots, c)$ as a solution. Therefore $\mathrm{CSP}(\mathbb{A}) \in \mathsf{P}$.

Identities that are satisfied by the projections are called *trivial*, since they hold in every clone (every clone contains the projections). Both, the semilattice identities and the identities $f(x) \approx f(y)$ are not satisfied by projections. We therefore call them *non-trivial.*

So all CSPs such that $\mathrm{Pol}(\mathbb{A})$ satisfies only trivial identities are NP-complete. And one might conjecture that non-trivial identities implied that $\mathrm{CSP}(\mathbb{A})$ is in P. However we need to be careful about this, since actually some other natural reductions do not preserve all identities (see Exercise 4.6, 4.7, 4.8).

There is an even more general concept, called a *minion homomorphism*, that preserves the complexity of CPSs. Minion homomorphisms are those maps that preserve *identities of*

[4] think of pp-interpretations as pp-definitions boosted by HSP. We don't need to properly define it here and refer to [8] for the interested reader

height 1. That is, identities of two expressions that contain exactly one operation symbol on each side.

Example 4.17.

- $t(x,y) \approx s(x,y,x)$ is a height 1 identity.
- $t(t(x,y),z) \approx t(x,t(y,z))$ is not of height 1.
- $x \approx m(x,y,y) \approx m(y,y,x)$ is not of height 1 (because x contains no operation symbol).
- $m(x,x,x) \approx m(x,y,y) \approx m(y,y,x)$ is of height 1.

Definition 4.18. A map $\xi\colon \mathscr{A} \to \mathscr{B}$ is called a *minion homomorphism* if

- it preserves arities
- for all operation $f \in \mathscr{A}$, and projections $\pi^n_{i_1}, \dots, \pi^n_{i_m}$:

$$\xi(f \circ (\pi^n_{i_1}, \dots, \pi^n_{i_m})) = \xi(f) \circ (\pi^n_{i_1}, \dots, \pi^n_{i_m})$$

We write $\mathscr{A} \leq_{h1} \mathscr{B}$ if there is a minion homomorphism from $\mathscr{A}$ to $\mathscr{B}$.

So minion homomorphisms are a weakening of clone homomorphisms, in which only h1 identities need to be preserved. Note that because of this, the image of a minion homomorphism does not need to be a clone. Without a proof we state the following result:

Theorem 4.19. *Let $\mathbb{B}$ and $\mathbb{A}$ be finite structures, and assume there is a minion homomorphism from* $\mathrm{Pol}(\mathbb{B})$ *to* $\mathrm{Pol}(\mathbb{A})$*. Then* $\mathrm{CSP}(\mathbb{A})$ *reduces to* $\mathrm{CSP}(\mathbb{B})$*.*

So the complexity of $\mathrm{CSP}(\mathbb{A})$ depends only on the identities of height 1 that hold in $\mathrm{Pol}(\mathbb{A})$. If there is a minion homomorphism $\mathrm{Pol}(\mathbb{A}) \to \mathrm{Proj}$, then $\mathrm{CSP}(\mathbb{A})$ is NP-complete. And, as Bulatov and Zhuk showed, in the other case $\mathrm{CSP}(\mathbb{A})$ is in P.

Next time we are going to give a characterization by Taylor of the non-trivial h1 identities that hold, if $\mathrm{Pol}(\mathbb{A}) \leq_{h1} \mathrm{Proj}$.

4.3 Taylor operations

A clone $\mathscr{A}$ is called *idempotent* if all $t \in \mathscr{A}$ satisfy $t(x,x,\dots,x) \approx x$. In the last section we discussed how clone homomorphism, and more general, minion homomorphisms correspond to typical reductions between CSPs. Using such reductions it can be shown that it is enough to study CSPs with idempotent polymorphism clone.

Lemma 4.20 (without proof)**.** *Let $\mathbb{A}$ be a finite structure. Then there exists another finite structure $\mathbb{A}'$, such that* $\mathrm{Pol}(\mathbb{A}) \leq_{h1} \mathrm{Pol}(\mathbb{A}') \leq_{h1} \mathrm{Pol}(\mathbb{A})$*, and* $\mathrm{Pol}(\mathbb{A}')$ *is idempotent.*

So by Lemma 4.20, every CSP is equivalent to the CSP of a structure with idempotent polymorphism clone. This, together with the reductions from last section, lead to the followings algebraic version of the CSP dichotomy conjecture / now theorem:

Theorem 4.21 (algebraic dichotomy conjecture from ∼2000; proved in [6], [7])**.** *Let $\mathbb{A}$ be a finite structure such that* $\mathrm{Pol}(\mathbb{A})$ *is idempotent. Then either*

1. $\mathrm{Pol}(\mathbb{A}) \leq \mathrm{Proj}$ *and* $\mathrm{CSP}(\mathbb{A})$ *is* NP*-complete, or*

2. $\mathrm{Pol}(\mathbb{A}) \not\leq \mathrm{Proj}$ *and* $\mathrm{CSP}(\mathbb{A})$ *is in* P*.*

As we know $\mathrm{Pol}(\mathbb{A}) \leq \mathrm{Proj}$ implies that $\mathrm{CSP}(\mathbb{A})$ is NP-complete. Thus, the really hard part in proving the conjecture was to show the second case: How can $\mathrm{Pol}(\mathbb{A}) \not\leq \mathrm{Proj}$ be used to find a polynomial time algorithm for $\mathrm{CSP}(\mathbb{A})$?

We already saw some examples of how non-trivial identities (semi-lattice operations and constant operations) can be used to obtain algorithms for the corresponding CSP. So an important question was: can we characterize all clones in Case 2 of Theorem 4.21 by some concrete identities? The surprising answer is yes; a first such result was obtained by Walter Taylor (already long before the study of CSPs).

Definition 4.22. Let t be an idempotent operation (of some arity n). We then call t a *Taylor operation*, if it satisfies (n many) h1 identities of the form:

$$\begin{aligned} t(x,*,*,\dots,*) &\approx t(y,*,*,\dots,*) \\ t(*,x,*,\dots,*) &\approx t(*,y,*,\dots,*) \\ &\vdots \\ t(*,*,*,\dots,x) &\approx t(*,*,*,\dots,y) \end{aligned}$$

Here, all positions marked by $*$ stand for arbitrary variables.

Example 4.23. 1. Every Maltsev operation m is a Taylor operation, since it is idempotent and satisfies the identities

$$\begin{aligned} m(\mathbf{x},x,y) &\approx m(\mathbf{y},x,x) \\ m(y,\mathbf{x},x) &\approx m(y,\mathbf{y},y) \\ m(y,x,\mathbf{x}) &\approx m(x,x,\mathbf{y}) \end{aligned}$$

2. A (ternary) *majority* operation m satisfies the identities $m(y,x,x) \approx m(x,y,x) \approx m(x,x,y) \approx x$, and is therefore also a Taylor operation.
3. Every operation c satisfy $c(x_1,x_2,\dots,x_n) \approx c(x_2,x_3,\dots,x_n,x_1)$ is Taylor. Such c is called an n-ary *cyclic* operation.
4. A 6-ary *Siggers operation* s satisfies the identity $s(x,y,x,z,y,z) \approx s(y,x,z,x,z,y)$ and is therefore also Taylor.

Note that, whenever a clone has a Taylor term, this can be witnessed by identities in variables x and y alone (since otherwise we can substitute the variables appearing at the $*$-positions by x ans y's). By the following theorem, a clone is in Case 2 of Theorem 4.21 if and only if it has a Taylor operation:

Theorem 4.24 (Walter Taylor '77)**.** *Let $\mathscr{A}$ be an idempotent clone. Then the following are equivalent:*

1. $\mathscr{A} \not\leq \mathrm{Proj}$*,*

2. $\mathscr{A} \not\leq_{h1} \mathrm{Proj}$*,*

3. There is a Taylor operation $t \in \mathscr{A}$.

Note that Taylor's theorem even holds if the universe of the clone is not finite. We are going to prove it in the remaining part of this section.

(2)→(1) holds, since every clone homomorphism is also a minion homomorphism. For (3)→(2) note that, if there was a minion homomorphism $\xi \colon \mathscr{A} \to \mathrm{Proj}$ then it would map the Taylor operation t to a projection $\xi(t)(x_1, \dots, x_n) = x_i$. Since t is Taylor it satisfies an identity $t(*, \dots, x, \dots, *) \approx t(*, \dots, y, \dots, *)$, where x and y are on the i-th position. Since ξ is a minion homomorphism, it preserve this identity, and thus $x = \xi(t)(*, \dots, x, \dots, *) \approx \xi(t)(*, \dots, y, \dots, *) = y$, which is a contradiction!

In order to prove (1)→(2)→(3), we need to do some more work, and use the assumption that $\mathscr{A}$ is idempotent (Theorem 4.24 does not hold for non-idempotent clones, see Exercise 4.9).

We define minors, which help us to discuss identities of height 1:

Definition 4.25. Let g an n-ary operation, and $\alpha \colon [n] \to [k]$ be a map. Then the *minor* g^α is defined as the k-ary operation $g^\alpha(x_1, \dots, x_k) = g(x_{\alpha(1)}, x_{\alpha(2)}, \dots, x_{\alpha(n)})$. For example, if $\alpha \colon [3] \to [4]$ such that $\alpha(1) = \alpha(2) = 1$, and $\alpha(3) = 4$, then $g^\alpha(x_1, x_2, x_3, x_4) = g(x_1, x_1, x_4)$.

Observation 4.26.

1. Condition 2 in Definition 4.18 is equivalent to the statement that for all $g \in \mathscr{A}$ and all suitable maps α: $\xi(g^\alpha) = (\xi(g))^\alpha$. So minion homomorphisms are maps that preserve arities of functions and minors.
2. If g is n-ary, $\alpha \colon [n] \to [m]$ and $\beta \colon [m] \to [k]$, then $(g^\alpha)^\beta = g^{\beta \circ \alpha}$
3. The identities of height 1 are exactly the identities of the form $f^\alpha \approx g^\beta$.
4. In particular, an n-ary idempotent operation t is Taylor, if for every $i = 1, \dots, n$ there are maps $\alpha, \beta \colon [n] \to [2]$ such that $t^\alpha = t^\beta$ and $\alpha(i) \neq \beta(i)$.

Another construction that often appears when working with idempotent clones or algebras is the following:

Definition 4.27. Let $f \colon A^n \to A$ and $g \colon A^m \to A$. Then we define the "star product" $f * g \colon A^{nm} \to A$ by

$$f * g(x_1, \dots, x_{nm}) = f(g(x_1, \dots, x_m), g(x_{m+1}, \dots, x_{2m}), \dots, g(x_{(n-1)m+1}, \dots, x_{nm})).$$

In idempotent clones we can use the the star product to encode multiple operations as the minors of a bigger one, by the following lemma:

Lemma 4.28. *Let T be a finite set of idempotent operations of A. Then there is a $t \in \mathsf{Clo}(T)$ such that for every $f \in T$, there is a map α with $f = t^\alpha$.*

Proof. If $T = \{f\}$, the lemma is clearly true with $t = f$.

For a 2-element set $T = \{f, g\}$, we define $t = f * g$. Since f and g are idempotent:

$$f(g(x_1, x_1 \dots, x_1), g(x_2, \dots, x_2), \dots, g(x_n \dots, x_n)) = f(x_1, x_2, \dots, x_n). \tag{4.1}$$

$$f(g(x_1, x_2 \dots, x_m), g(x_1, \dots, x_m), \dots, g(x_1 \dots, x_m)) = g(x_1, x_2, \dots, x_m). \tag{4.2}$$

Clearly both expressions on the left side are minors of t. For a general finite set $T = \{f_1, \dots, f_n\}$, the term $t = f_1 * (f_2 * (f_3 * \cdots f_n))$ has the desired property. This can be shown by induction on n, and the same argument as for $|n| = 2$. □

Using the same construction, we can prove that minion homomorphism from an idempotent clone to the projections need to preserve the star product.

Lemma 4.29. *Let $\mathscr{A}$ be idempotent, and let $\xi\colon \mathscr{A} \to \mathrm{Proj}$ be a minion homomorphism. Then $\xi(f * g) = \xi(f) * \xi(g)$, for all $f, g \in \mathscr{A}$.*

Proof. Since ξ maps $\mathscr{A}$ to the projection clone, there are projections π_i^n, π_j^m, π_k^{mn}, such that $\xi(f) = \pi_i^n$, $\xi(g) = \pi_j^m$ and $\xi(f * g) = \pi_k^{nm}$. By definition of the star product we have $\pi_i^n * \pi_j^m(x_1, x_2, \dots, x_{nm}) = \pi_i^n(x_j, x_{m+j}, \dots, x_{(n-1)m+j}) = x_{(i-1)m+j}$. So we need to show that $k = (i-1)m + j$. We use the fact, that we can obtain f and g as minors of $t = f * g$, as in the proof of Lemma 4.28.

First, let $\alpha : [nm] \to [n]$ be defined by $\alpha(x) = \lceil x/m \rceil$. This corresponds to the minor in (4.1). Then $\xi(f * g)^\alpha = \xi((f * g)^\alpha) = \xi(f) = \pi_i^n$. On the other hand $\xi(f * g)^\alpha = \pi_{\alpha(k)}^n$, so $\lceil k/m \rceil = i$. Next, let $\beta : [nm] \to [m]$ be defined by $\beta(x) = x \mod m$ if $x \mod m \neq 0$ and m else. Then $\xi(f * g)^\beta = \xi((f * g)^\beta) = \xi(g) = \pi_j^m$. But, on the other hand $\xi(f * g)^\beta = \pi_{\beta(k)}^m$, so k must be equal to j modulo m. Both these observations together imply that $k = (i-1)m + j$. □

We are now ready to finish the proof of Taylor's theorem:

Proof of Theorem 4.24. We first show (1)→(2). So let $\xi\colon \mathscr{A} \to \mathrm{Proj}$ be a minion homomorphism. We want to show that ξ is a clone homomorphism, or in other words, that for $f, g_1, \dots, g_n \in \mathscr{A}$:

$$\xi(f \circ (g_1, \dots, g_n)) = \xi(f) \circ (\xi(g_1), \dots, \xi(g_n)).$$

By Lemma 4.28 there is a $g \in \mathscr{A}$, such that every g_i is equal to a minor $g^{\alpha_i} = g_i$. So

$$(f \circ (g_1, \dots, g_n)(x_1, \dots, x_m) = f(g(x_{\alpha_1(1)}, \dots, x_{\alpha_1(k)}), \dots, g(x_{\alpha_n(1)}, \dots, x_{\alpha_n(k)})),$$

where k denotes the arity of g. We can further write this as a minor $(f * g)^\alpha$, where $\alpha\colon [nk] \to [m]$ is defined by $\alpha((i-1)k + j) = \alpha_i(j)$. So $f \circ (g_1, \dots, g_n) = (f * g)^\alpha$. Since ξ is a minion homomorphism

$$\xi((f * g)^\alpha) = \xi((f * g))^\alpha.$$

By Lemma 4.29 we have

$$\xi((f * g))^\alpha = (\xi(f) * \xi(g))^\alpha = \xi(f) \circ (\xi(g)^{\alpha_1}, \dots, \xi(g)^{\alpha_n})$$

Since ξ is a minion homomorphism, this is equal to

$$\xi(f) \circ (\xi(g^{\alpha_1}), \dots \xi(g^{\alpha_n})) = \xi(f) \circ (\xi(g_1), \dots \xi(g_n)),$$

therefore ξ is a clone homomorphism.

We next show (2) → (3) by an indirect proof. So assume that the clone $\mathscr{A}$ does not have a Taylor operation; we are going to prove that there is a minion homomorphism from $\mathscr{A}$ to Proj, in 3 steps.

Claim 1: For every fixed $t \in \mathscr{A}$ there is a map ξ from the minors of t to Proj such that $\xi(t^\alpha) = \xi(t)^\alpha$.

We know that t is not Taylor. So, by Observation 4.26 (4), there is an index i, such that for all maps α, β either $\alpha(i) = \beta(i)$ or $t^\alpha \neq t^\beta$. So, in some sense the index 'separates' the minors of t (Note that i is not necessarily unique). We define the map by $\xi(t^\alpha) = \pi^k_{\alpha(i)}$ (for $\alpha\colon [n] \to [k]$). This map is well-defined since whenever $t^\alpha = t^\beta$, then $\alpha(i) = \beta(i)$, and therefore $\xi(t^\alpha) = \xi(t^\beta) = \pi^k_{\alpha(i)}$. Moreover it satisfies $\xi(t^\alpha) = \pi^k_{\alpha(i)} = (\pi^n_i)^\alpha = \xi(t)^\alpha$.

Claim 2: For every finite set $T \subseteq \mathscr{A}$, there is a map from T to Proj that preserves minors.

By Lemma 4.29 for every finite subset $T \subseteq \mathscr{A}$, there is an operation $t \in \mathscr{A}$ such that all operations in T are minors of t. Now Claim 2 follows by applying Claim 1 to t.

Claim 3: There is a minion homomorphism $\xi\colon \mathscr{A} \to \text{Proj}$

This follows from a compactness argument; we create a first-order theory Σ, that formally describes minion homomorphisms from $\mathscr{A}$ to Proj. The language of Σ consists of constant symbols π^n_i for all $1 \leq i \leq n$ and $\xi(t)$ for every $t \in \mathscr{A}$. And Σ consist of the sentences

- $\pi^m_i \neq \pi^n_j$ for all indices $i \neq j$ or $m \neq n$
- for all $t \in \mathscr{A}$ of arity n we add the sentence $\xi(t) = \pi^n_1 \vee \xi(t) = \pi^n_2 \vee \cdots \vee \xi(t) = \pi^n_n$.
- for all $t, s \in \mathscr{A}$ such that $s = t^\alpha$ we add the sentences $\xi(t) = \pi^n_1 \to \xi(s) = \pi^k_{\alpha(1)}$, $\xi(t) = \pi^n_2 \to \xi(s) = \pi^k_{\alpha(2)}, \dots \xi(t) = \pi^n_n \to \xi(s) = \pi^k_{\alpha(n)}$.

Now clearly, if Σ has a model, then $t \mapsto \xi(t)$ is a minion homomorphism. By Claim 2, every finite subset of Σ has a model. But, by the compactness theorem of first-order logic, also Σ has a model, which concludes the proof.

□

Having a Taylor operation t is still a very general condition, since we don't know anything about the arity of t, or the specific form. But it can be show that for idempotent clone $\mathscr{A}$ on a *finite* set, the following are equivalent:

- $\mathscr{A}$ has a Taylor term
- $\mathscr{A}$ has a 6-ary Siggers operation $s(x, y, x, z, y, z) \approx s(y, x, z, x, z, y)$
- $\mathscr{A}$ has a 4-ary Siggers operation $s(a, r, e, a) \approx s(r, a, r, e)$
- $\mathscr{A}$ has a WNU (weak near unanimity) operation, so an operation satisfying $w(y, x, \dots, x) \approx w(x, y, x, \dots, x) \approx \cdots \approx w(x, x, \dots, y)$.
- $\mathscr{A}$ has a cyclic operation $c(x_1, \dots, x_p) \approx c(x_2, \dots, x_p, x_1)$, for every prime $p > |A|$

Due to our time restrictions we were not able to discuss any of these results in more details; the proofs of the CSP dichotomy conjecture both used the existence of WNU polymorphisms. Surprisingly, even if we drop the finiteness condition, there is a uniform characterization of idempotent Taylor clones by the identities

$$t(x, y, y, y, x, x) \approx t(y, x, y, x, y, x) \approx t(y, y, x, x, x, y)$$

This result was proved by Mirek Olšák in 2016.

4.4 Exercises

Exercise 4.1. We claimed that for every $n \geq 2$ there is structure $\mathbb{A}$ with $|A| = n$ and a NP-complete CSP($\mathbb{A}$). For $n = 2, 3$ we saw this in Example 4.3 and 4.2. For general $n > 3$ prove it as follows:

- Show that 3-COLOR reduces to n-COLOR (the problem coloring a graph with n colors).
- Thus n-COLOR is NP-complete.
- Give a description of n-COLOR as CSP of a structure on the set $\{1, 2, \ldots, n\}$.

Exercise 4.2. We say in Example 4.4 that solving systems of linear equations over $\mathbb{Z}_p$ can be modeled as a CSP. Prove that for every finite algebra $\mathbf{A}$ (so finite universe and finite type) there is a structure $\mathbb{A}$, such that solving systems of polynomial equations over $\mathbf{A}$ is equivalent to CSP($\mathbb{A}$). (Hint: use the function graphs of the basic operations as relations of $\mathbb{A}$).

Exercise 4.3. Let A be finite and $\emptyset \neq R \subseteq A^2$ be a binary relation, and assume that R is preserved by a semilattice operation.

1. Show that, if R is subdirect, every primitive positive sentence $\exists x_1, \ldots, x_n R(x_{i_1}, x_{j_1}) \wedge \cdots \wedge R(x_{i_k}, x_{j_k})$ is true (and hence CSP$(A; R) \in \mathsf{P}$).
2. Show that (1) is not true if R is not subdirect.
3. (voluntary) For general such R, can you still prove, that CSP$(A; R) \in \mathsf{P}$?

Exercise 4.4. Have a look at Section 1 and 2 of [8] if you want more background / examples of CSPs.

Exercise 4.5. What do you think: Does Proposition 4.8 also hold for structures $\mathbb{A}$

- with finite domain A, but infinitely many relations?
- with infinite domain A, but finitely many relations?

Exercise 4.6. Let $\mathbb{A}$, $\mathbb{A}'$ be two relational structures of the same type (Example: two graphs). We say that they are *homomorphically equivalent*, if there are homomorphisms $h \colon \mathbb{A} \to \mathbb{A}'$ and $h' \colon \mathbb{A}' \to \mathbb{A}$ (here a homomorphism is a map that preserves all relations). Show that then CSP($\mathbb{A}$) = CSP($\mathbb{A}'$). (Hint: homomorphism preserve pp-formulas)

Exercise 4.7. Let $\mathbb{A} = (\{0, 1\}; \leq)$, and $\mathbb{A}' = (\{0\}; \leq)$.

- Show that $\mathbb{A}$ and $\mathbb{A}'$ are homomorphically equivalent.
- Show that Pol($\mathbb{A}'$) contains a Maltsev polymorphism, but Pol($\mathbb{A}$) does not. Therefore there is no clone homomorphism $\xi \colon \mathrm{Pol}(\mathbb{A}') \to \mathrm{Pol}(\mathbb{A})$

Exercise 4.8. There is however a minion homomorphism in the example from Exercise 4.7 (and in general for homomorphically equivalent structures): Let $h \colon \mathbb{A} \to \mathbb{A}'$ and $h' \colon \mathbb{A}' \to \mathbb{A}$ be the two homomorphisms. Then show that $\xi \colon \mathrm{Pol}(\mathbb{A}) \to \mathrm{Pol}(\mathbb{A}')$, which is defined by $\xi(f)(x_1, \ldots, x_n) = hf(h'(x_1), \ldots, h'(x_n))$ is a minion homomorphism.

Exercise 4.9. An operation f is called essentially injective, if $f(x_1, \ldots, x_n) = g(x_{i_1}, \ldots, x_{i_k})$, for an injective g. The essentially injective operations on $\mathbb{N}$ form a clone $\mathscr{A}$.

- Show that this clone contains no Taylor term.

- However show that there are $f, u \in \mathscr{A}$ such that $f(x,y) = u \circ f(y,x)$. Conclude that $\mathscr{A} \not\leq \mathrm{Proj}$.
- (*) try to find a *system* of non-trivial height 1 identities in $\mathscr{A}$.

Exercise 4.10. In the proof of (1)→(2) of Theorem 4.24, we actually only checked that ξ preserves compositions. Complete the proof that ξ is actually a clone homomorphism.

Exercise 4.11. Let Proj^A be the clone of the projections on a set A. Note that we defined $\mathrm{Proj} = \mathrm{Proj}^{\{0,1\}}$.

- Show that $\mathrm{Proj}^A \sim \mathrm{Proj}$ for every $|A| > 1$. Conclude that the equivalence class of Proj is the smallest element of the interpretability lattice.
- What is the problem with Proj^A for $|A| = 1$? Prove that its equivalence class under $\sim$ is the maximum of the interpretability lattice.

Bibliography

[1] Jaroslav Ježek. *Universal Algebra* `http://ka.karlin.mff.cuni.cz/jezek/ua.pdf` Script, 2008

[2] Clifford Bergman. *Universal algebra: Fundamentals and selected topics.* Chapman and Hall/CRC, 2011. (Book available in our library)

[3] Ralph Freese, Ralph McKenzie. *Commutator theory for congruence modular varieties.* CUP Archive, 1987.

[4] Andrew Moorhead. *Supernilpotent algebras are nilpotent.* Preprint on `https://arxiv.org/abs/1906.09163` 2019.

[5] Paweł Idziak, Jacek Krzaczkowski. *Satisfiability in multi-valued circuits.* Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science. 2018.

[6] Dmitriy Zhuk. *A proof of CSP dichotomy conjecture.* 58th Annual Symposium on Foundations of Computer Science (FOCS). 2017.

[7] Bulatov, Andrei A. *A dichotomy theorem for nonuniform CSPs.* 58th Annual Symposium on Foundations of Computer Science (FOCS). 2017.

[8] Libor Barto, Andrei Krokhin, Ross Willard. *Polymorphisms, and how to use them.* `https://drops.dagstuhl.de/opus/volltexte/2017/6959/pdf/DFU-Vol7-15301-1.pdf`

Thanks

Thanks to Anna Gajdova, Juraj Hartmann, Štěpán Hudeček, Zuzana Procházková, Svilen Popov, Lucien Šima, Žaneta Semanišinová, Ester Sgallová, and Yifan Zhang for pointing out mistakes and typos in earlier versions.