Finite Algebra

M. Clasen and M. Valeriote
McMaster University

Contents

CHAPTER 1

The Structure of Finite Algebras

1 Pálfy's theorem

An *algebra* is a pair

$$\mathcal{A} = \langle A, F \rangle,$$

where A is a nonempty set and F is a collection of finitary operations on A. A is the *universe* of $\mathcal{A}$ and the operations in F are sometimes called the *basic operations* of $\mathcal{A}$. Often, F will be presented as an indexed set. This leads to the usual first-order language of $\mathcal{A}$. If F is not indexed, $\mathcal{A}$ will be called a *nonindexed algebra*.

A *clone* of functions on A is a set of functions from A^n to A for varying n's which contains all projections and is closed under composition. The *clone of term operations* of $\mathcal{A}$, denoted by $\mathrm{Clo}(\mathcal{A})$, is the smallest clone on A which contains F. So, a term operation of $\mathcal{A}$ is an operation which is obtained by compositions of basic operations of $\mathcal{A}$ and projections on A.

A *polynomial* of an algebra $\mathcal{A}$ is an operation $p(\bar{x})$ of the form $t(\bar{x}, \bar{a})$, where t is a term operation of $\mathcal{A}$ and $\bar{a}$ is a finite sequence of elements from A. The *clone of polynomials* of $\mathcal{A}$ is the set of all polynomials of $\mathcal{A}$ and is denoted by $\mathrm{Pol}(\mathcal{A})$. It is not hard to see that $\mathrm{Pol}(\mathcal{A})$ is a clone on A, and is in fact the smallest clone on A which contains F and all of the constant operations on A. The set of all n-ary polynomials of $\mathcal{A}$ is denoted by $\mathrm{Pol}_n(\mathcal{A})$.

Two algebras $\mathcal{A}$ and $\mathcal{B}$ with the same universe are *polynomially equivalent* if $\mathrm{Pol}(\mathcal{A}) = \mathrm{Pol}(\mathcal{B})$.

1.1. Exercise Show that the algebra $\mathcal{A} = \langle \{0,1\}; d(x,y,z) \rangle$ with

$$d(x,y,z) = \begin{cases} z, & \text{if } x = y \\ x, & \text{otherwise} \end{cases}$$

is polynomially equivalent to the two-element boolean algebra.

An algebra $\mathcal{A}$ is *minimal* if it is finite and each of its unary polynomials is either constant or a permutation.

1.2. Example

1. Any finite vector space is minimal. The unary polynomials are all of the form $\lambda x + c$, where c is a vector and λ is a scalar.
2. Let G be a group. A G-set is a set A together with a homomorphism from G into the symmetric group on A. Such a set is considered as a unary structure by taking the functions induced by the group elements as basic operations. Any finite G-set is minimal.
3. Any two-element algebra is minimal.

1.3. Exercise Let $\langle G, \cdot \rangle$ be a finite semigroup. Show that there is a $k > 0$ such that $m^k \cdot m^k = m^k$ for all $m \in G$.

A function $f\colon A^n \longrightarrow A$ is called idempotent in the i-th variable if the equation

$$f(x_1,\ldots,x_{i-1},f(x_1,\ldots,x_n),x_{i+1},\ldots,x_n)=f(x_1,\ldots,x_n)$$

holds in A. From the previous exercise, one easily derives the following key fact.

Fact *If A is finite then there exists a natural number $k>0$ such that for all functions $f\colon A\longrightarrow A$, f^k is idempotent.*

If $f(x_1,\ldots,x_n)$ is an operation on A, $i\le n$ and $k\ge 0$, then we define $f^k_{(i)}(x_1,\ldots,x_n)$ inductively by

$$\begin{aligned} f^0_{(i)}(x_1,\ldots,x_n)&=x_i,\\ f^{k+1}_{(i)}(x_1,\ldots,x_n)&=f(x_1,\ldots,x_{i-1},f^k_{(i)}(x_1,\ldots,x_n),x_{i+1},\ldots,x_n). \end{aligned}$$

1.4. Corollary *If A is a finite set, then there exists a $k>0$ such that for all maps $f\colon A^n\longrightarrow A$ and $i\le n$, $f^k_{(i)}(a_1,\ldots,a_{i-1},x,a_{i+1},\ldots,a_n)$ is idempotent for all $a_1,\ldots,a_n\in A$.*

An operation $t(x,y,z)$ on a set A is a *Mal'cev operation*, if it satisfies $t(x,x,y)=y=t(y,x,x)$. An algebra is *Mal'cev*, if it has a term operation which is a Mal'cev operation.

1.5. Example $xy^{-1}z$ is a Mal'cev term on any group.

An algebra $\langle G,\cdot\rangle$ with a binary operation $\cdot$ is a *quasigroup*, if for all $a\in G$ the operations $x\cdot a$ and $a\cdot x$ are permutations of G.

1.6. Lemma *Any finite quasigroup is Mal'cev.*

Proof Let $\mathcal{A}=\langle A,f(x,y)\rangle$ be a finite quasigroup. Choose $k\ge 0$ such that the equation $f^k_{(1)}(x,y)=f^k_{(1)}(f^k_{(1)}(x,y),y)$ holds in $\mathcal{A}$. For all $a\in A$, the operation $f^k_{(1)}(x,a)$ is idempotent and is a permutation of A. Since the only idempotent permutation is the identity, $f^k_{(1)}(x,a)=x$. Thus $f^k_{(1)}(x,y)=x$ holds for all x and y. If we define $d_1(x,y)$ to be $f^{k-1}_{(1)}(x,y)$, we have $f(d_1(x,y),y)=x$. Repeat the above construction to get a term $d_2(x,y)$ such that $f(x,d_2(x,y))=y$. We claim that

$$q(x,y,z)=f(d_1(x,d_2(y,y)),d_2(y,z))$$

is a Mal'cev operation.

First, we have $f(d_1(y,d_2(y,y)),d_2(y,y))=y=f(y,d_2(y,y))$. Since $\mathcal{A}$ is a quasigroup, this implies $d_1(y,d_2(y,y))=y$. So $q(y,y,z)=f(y,d_2(y,z))$. By the choice of d_2, this is z.

On the other hand, $q(z,y,y)=f(d_1(z,d_2(y,y)),d_2(y,y))=z$ by the choice of d_1. □

An algebra is *abelian*, if for all term operations $t(x,\bar{y})$ the algebra satisfies

$$\forall u\forall v\forall\bar{y}\bar{z}(t(u,\bar{y})=t(u,\bar{z})\rightarrow t(v,\bar{y})=t(v,\bar{z})).$$

1.7. Example Any R-module is abelian. The term operations all have the form $t(x,\bar{y})=rx+\sum s_iy_i$ for $r,\ s_i\in R$. Suppose that $t(a,\bar{c})=t(a,\bar{d})$. Then $ra+\sum s_ic_i=ra+\sum s_id_i$, so $rb+\sum s_ic_i=rb+\sum s_id_i$ (just add $rb-ra$ to the first equation).

1.8. Exercise A group is abelian iff it is commutative.

1.9. Theorem (Smith, Gumm) *If $\mathcal{A}$ has a Mal'cev polynomial and satisfies the following weakening of abelianness*

$$\forall \bar{u} \forall \bar{v} \forall yz (t(\bar{u}, y) = t(\bar{u}, z) \to t(\bar{v}, y) = t(\bar{v}, z))$$

for all term operations $t(\bar{x}, y)$, then $\mathcal{A}$ is polynomially equivalent to a module.

Proof Let $d(x,y,z)$ be a Mal'cev polynomial. Let 0 be an arbitrary element of A. Define

$$a + b = d(a, 0, b), \qquad -a = d(0, a, 0).$$

We claim that these operations define an abelian group structure on A.

Let $d_1(x,y,z,u) = d(d(x,0,u), 0, d(y,u,z))$. Note that $d_1(0,b,0,b) = b = d_1(0,b,0,0)$. By abelianness, this implies

$$(a + b) + c = d_1(a,b,c,b) = d_1(a,b,c,0) = a + (b + c).$$

The commutativity is proved similarly.

Claim: We have $f(\bar{x}) = \sum_{1 \leq i \leq n} f_i(x_i) - (n-1) f(0, \ldots, 0)$ for all n-ary polynomials $f(\bar{x})$ of $\mathcal{A}$, where $f_i(x) = f(0, \ldots, 0, x, 0, \ldots, 0)$.

The proof is by induction on n. For $n = 1$, there is nothing to do.

For $n = 2$, consider $f(0,y) - f(0,y) = f(0,0) - f(0,0)$. By abelianness, $f(x,y) - f(0,y) = f(x,0) - f(0,0)$, thus $f(x,y) = f(x,0) + f(0,y) - f(0,0)$ as required.

For $n > 2$, treat $f(x, y_2, \ldots, y_n)$ as a polynomial in $y_2, \ldots, y_n$. The induction hypothesis yields

$$f(x, y_2, \ldots, y_n) =$$
$$f(x, y_2, 0, \ldots, 0) + f(x, 0, y_3, 0, \ldots, 0) + \cdots + f(x, 0, \ldots, 0, y_n) - (n-2) f(x, 0, \ldots, 0).$$

Now apply the case $n = 2$ to each summand to get the result for $f(x, y_2, \ldots, y_n)$.

Define $R = \{p(x) \in \mathrm{Pol}_1(\mathcal{A}) \mid p(0) = 0\}$. Clearly $\langle R, \circ, +, -, 0, id \rangle$ is a ring. If $p(x) \in R$, then $p(x + y) = p(x + 0) + p(0 + y) - p(0 + 0) = p(x) + p(y)$. This shows that $\langle A, +, -, 0 \rangle$ is an R-module M. Since all operations of M can be defined by polynomials of $\mathcal{A}$, we have $\mathrm{Pol}(M) \subseteq \mathrm{Pol}(\mathcal{A})$. For the other inclusion, let $f(\bar{x}) \in \mathrm{Pol}(\mathcal{A})$. If we choose $\alpha_i(x) = f_i(x) - f_i(0) \in R$, an easy calculation shows that $f(\bar{x}) = \sum \alpha_i x_i + f(0, \ldots, 0)$. Thus, $f(\bar{x})$ is a module polynomial. □

1.10. Exercise An abelian algebra which has a Mal'cev polynomial is Mal'cev.

1.11. Lemma (Twin lemma) *Let $\mathcal{A}$ be a minimal algebra with at least three elements and $f(x, \bar{y}) \in \mathrm{Pol}(\mathcal{A})$. If $\bar{c}, \bar{d} \in A$, then $f(x, \bar{c})$ is a permutation of A iff $f(x, \bar{d})$ is.*

Proof Without loss of generality, we may assume that f is binary. (If $f(x, \bar{c})$ is a permutation and $f(x, \bar{d})$ not, we can exchange $\bar{c}$ and $\bar{d}$ elementwise to find single entries c_i and d_i such that $f(x, d_1, \ldots, d_{i-1}, c_i, c_{i+1} \ldots, c_n)$ is a permutation and $f(x, d_1, \ldots, d_{i-1}, d_i, c_{i+1} \ldots, c_n)$ not.)

So assume we have elements a, b such that $f(b, y)$ is a permutation, while $f(a, y)$ is not. Then there are elements c, d such that $f(a, c) = f(a, d)$ and $f(b, c) \neq f(b, d)$. Choose k such that $g(x, y) = f^k_{(2)}(x, y)$ satisfies $g(x, g(x, y)) = g(x, y)$.

Since $f(b, y)$ is a permutation, $g(b, y)$ is an idempotent permutation, so $g(b, y) = y$. Similarly, $g(a, y)$ is constant, since $f(a, y)$ is. In fact, $g(e, y)$ is either constant or the identity for any $e \in A$.

	e	w		y
a	e	e	$\cdots$	e
b	e	w		y
u	e	?		
	$\vdots$			

Let $g(a, y) = e$ for all $y \in A$. Choose elements $w \neq e$ and $u \notin \{a, b\}$. We have $g(b, w) = w$ and $g(b, e) = e$, so $g(x, e)$ is constant, in particular $g(u, e) = e$. There are two possibilities for $g(u, w)$. It can either be e or w. In both cases, the polynomial $g(x, w)$ is neither constant nor a permutation. This contradicts the minimality of $\mathcal{A}$. $\square$

An algebra is called *essentially unary* if every basic operation depends on at most one variable. If this is the case, every polynomial also depends on at most one variable.

1.12. Exercise If the polynomial $g(x_1, \ldots, x_n)$ depends on all variables, there are indices $i < j \leq n$ and elements $c_1, \ldots, c_n$ such that

$$g(c_1, \ldots, c_{i-1}, x, c_{i+1}, \ldots, c_{j-1}, y, c_{j+1}, \ldots, c_n)$$

depends on x and y.

1.13. Theorem (Pálfy) *If $\mathcal{A}$ is a minimal algebra with at least three elements, then $\mathcal{A}$ is essentially unary or polynomially equivalent to a vectorspace.*

Proof Assume that $f(\bar{x})$ is a polynomial of $\mathcal{A}$ which depends on more than one variable. By Exercise 1.12, we can assume that f is binary.

Claim: $\langle A, f \rangle$ is a quasigroup.

Since $f(x, y)$ depends on x, there is a c such that $f(x, c)$ is not constant, hence is a permutation. By the twin lemma, $f(x, c)$ is a permutation for all $c \in A$. Similarly for the second variable. This proves the claim.

By Lemma 1.6, we find a Mal'cev term for $\mathcal{A}$. To apply the theorem of Smith–Gumm, we need to establish

$$f(\bar{a}, c) = f(\bar{a}, d) \quad \Rightarrow \quad f(\bar{b}, c) = f(\bar{b}, d)$$

for all polynomials $f(x, \bar{y})$ and all $\bar{a}, \bar{b}, c, d \in A$.

So let $f(\bar{a}, c) = f(\bar{a}, d)$. If $c = d$, trivially $f(\bar{b}, c) = t(\bar{b}, d)$. If $c \neq d$, then $f(\bar{a}, y)$ is constant, so by the twin lemma, $f(\bar{b}, y)$ is also constant. Again $f(\bar{b}, c) = t(\bar{b}, d)$.

By the theorem of Smith–Gumm, $\mathcal{A}$ is polynomially equivalent to a module M over a finite ring R. The minimality of $\mathcal{A}$ forces each ring element to be invertible, thus R is a division ring. But finite division rings are fields, so M is a vectorspace. $\square$

1.14. Corollary *Let $\mathcal{A}$ be minimal. Then $\mathcal{A}$ is polynomially equivalent to one of the following*

1. *a unary permutational algebra (a G-set),*
2. *a vectorspace over a finite field,*
3. *the two-element boolean algebra,*
4. *the two-element lattice,*
5. *a two-element semilattice.*

Proof The one-element algebra is trivially unary permutational and the minimal algebras with at least three elements are treated in Pálfy's theorem, so it remains to analyze the two-element algebras up to polynomial equivalence. This is done in the next exercise. □

1.15. Exercise Let $\mathcal{A}$ be a two-element algebra. Then $\mathcal{A}$ is polynomially equivalent to one of the following

1. a pure set, or a set with a transposition,
2. a vectorspace,
3. a boolean algebra,
4. a lattice,
5. a semilattice.

For a minimal algebra $\mathcal{A}$ we define the type of $\mathcal{A}$, denoted by $\mathrm{typ}(\mathcal{A})$, to be the number in the enumeration in Pálfy's theorem which applies to it.

2 Localization and Relativization

A *neighbourhood* of an algebra $\mathcal{A}$ is a subset $U \subseteq A$ of the form $e(A)$ for some idempotent polynomial $e(x)$ of $\mathcal{A}$. The algebra *induced by* $\mathcal{A}$ on U, denoted by $\mathcal{A}|_U$, is the algebra

$$\langle U, \{f(\bar{x})|_U \mid f \in \mathrm{Pol}_n(\mathcal{A}), f(U^n) \subseteq U, n < \omega\}\rangle .$$

2.1. Exercise Determine the types of the two-element neighbourhoods of the ring $\mathbb{Z}_4$.

A relation α on a set A is *compatible* with a function $f\colon A^n \longrightarrow A$ if

$$a_1 \;\alpha\; b_1, \dots, a_n \;\alpha\; b_n \quad \Rightarrow \quad f(a_1, \dots, a_n) \;\alpha\; f(b_1, \dots, b_n).$$

A *congruence* on an algebra $\mathcal{A}$ is an equivalence relation that is compatible with the basic operations of $\mathcal{A}$. The set $\mathrm{Con}(\mathcal{A})$ of all congruences of $\mathcal{A}$ is a lattice with respect to $\subseteq$. Its smallest element is $0_{\mathcal{A}} = \{\langle a, a\rangle \mid a \in A\}$ and its largest element is $1_{\mathcal{A}} = A^2$.

2.2. Exercise Let θ be an equivalence relation on the algebra $\mathcal{A}$. Then θ is a congruence iff θ is compatible with all unary polynomials of $\mathcal{A}$.

2.3. Lemma *Let U be a neighbourhood of $\mathcal{A}$.*

(a) If $\bar{a}, c \in U$, $f \in \mathrm{Pol}(\mathcal{A})$ such that $f(\bar{a}) = c$, then there is a polynomial $g \in \mathrm{Pol}(\mathcal{A}|_U)$ such that $g(\bar{a}) = c$.

(b) The restriction map from $\mathrm{Con}(\mathcal{A})$ to $Con(\mathcal{A}|_U)$ is a surjective lattice homomorphism.

Proof (a) Let $U = e(A)$ and $e^2 = e$. Then $g = ef|_U$ is as required.
(b) $\theta \mapsto \theta|_U$ maps θ to an equivalence relation on U. Since θ is compatible with all unary polynomials on $\mathcal{A}$, its restriction is compatible with all restrictions of unary polynomials on $\mathcal{A}$. Thus $\theta|_U$ is a congruence.

It is easy to see that $(\theta_1 \cap \theta_2)|_U = \theta_1|_U \cap \theta_2|_U$. For $\alpha \in \mathrm{Con}(\mathcal{A}|_U)$ define

$$\widehat{\alpha} = \{\langle x, y\rangle \in A^2 \mid \langle ef(x), ef(y)\rangle \in \alpha \text{ for all polynomials } f\}.$$

Claim: $\widehat{\alpha}$ is the largest congruence on $\mathcal{A}$ whose restriction to U is α.

Clearly, $\widehat{\alpha}$ is an equivalence relation on A. By Exercise 2.2, it is enough to show that $\widehat{\alpha}$ is compatible with all unary polynomials. So let $\langle a, b\rangle \in \widehat{\alpha}$ and $f \in \mathrm{Pol}_1(\mathcal{A})$. If g is another unary polynomial, $\langle egf(a), egf(b)\rangle \in \alpha$ by the definition of $\widehat{\alpha}$. But this implies $\langle f(a), f(b)\rangle \in \widehat{\alpha}$. Thus, $\widehat{\alpha}$ is a congruence. $\alpha \le \widehat{\alpha}|_U$ is clear. If $\langle a, b\rangle \in \widehat{\alpha}|_U$, then $\langle e^2(a), e^2(b)\rangle = \langle a, b\rangle \in \alpha$ and so $\widehat{\alpha}|_U = \alpha$.

Finally let θ be another congruence with $\theta|_U = \alpha$. Let $\langle a, b\rangle \in \theta$. Let f be an arbitrary unary polynomial. Then $\langle ef(a), ef(b)\rangle \in \theta$, since θ is a congruence, and $\langle ef(a), ef(b)\rangle \in U^2$ by the choice of e. So $\langle ef(a), ef(b)\rangle \in \theta|_U = \alpha$. The definition of $\widehat{\alpha}$ now implies $\langle a, b\rangle \in \widehat{\alpha}$. This shows that $\theta \subseteq \widehat{\alpha}$ and finishes the proof of the claim.

By the claim, the restriction map is surjective. Let now θ_1 and θ_2 be two congruences of $\mathcal{A}$. Let $\beta = \theta_1 \vee \theta_2$ and $\alpha = \theta_1|_U \vee \theta_2|_U$. Since $\theta_1, \theta_2 \le \beta$, we have $\theta_1|_U, \theta_2|_U \le \beta|_U$ and thus $\alpha \le \beta|_U$. Conversely $\theta_1|_U, \theta_2|_U \le \alpha$, which implies by the claim $\theta_1, \theta_2 \le \widehat{\alpha}$, so $\beta \le \widehat{\alpha}$ and $\beta|_U \le \widehat{\alpha}|_U = \alpha$.

□

For congruences $\alpha \le \beta$ of $\mathcal{A}$, we define the (α, β)-*separating polynomials* to be the elements of the set

$$\mathrm{Sep}(\alpha, \beta) = \{f \in \mathrm{Pol}_1(\mathcal{A}) \mid f(\beta) \not\subseteq \alpha\}\ .$$

Let $M_{\mathcal{A}}(\alpha, \beta)$ be the set of $\subseteq$-minimal elements of $\{f(A) \mid f \in \mathrm{Sep}(\alpha, \beta)\}$. The elements of $M_{\mathcal{A}}(\alpha, \beta)$ are called the (α, β)-*minimal sets* of $\mathcal{A}$.

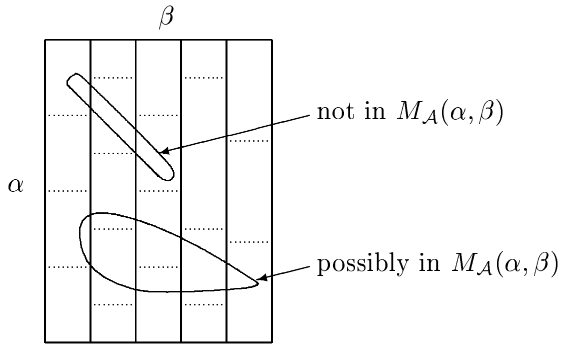

2.4. Exercise $\mathcal{A}$ is minimal iff $M_{\mathcal{A}}(0_{\mathcal{A}}, 1_{\mathcal{A}}) = \{A\}$.

We write $\alpha \prec \beta$ to indicate that there is no congruence between α and β.

2.5. Lemma *Let $\mathcal{A}$ be a finite algebra and $\alpha \prec \beta$ two congruences. The (α, β)-minimal sets are neighbourhoods of $\mathcal{A}$.*

Proof Let U be an (α, β)-minimal set. Let $K = \{f \in \mathrm{Pol}_1(\mathcal{A}) \mid f(A) \subseteq U\}$. Note that K is a right ideal in the semigroup $\mathrm{Pol}_1(\mathcal{A})$. Consider the relation

$$\mu = \{\langle x, y\rangle \in \beta \mid \langle f(x), f(y)\rangle \in \alpha \text{ for all } f \in K\}.$$

Claim: $\mu \in \mathrm{Con}(\mathcal{A})$ and $\alpha \le \mu \le \beta$.

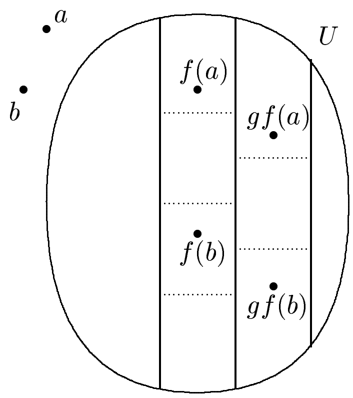

It is easy to see that μ is an equivalence relation between α and β. So it suffices to show that μ is compatible with the unary polynomials. If $\langle a,b\rangle \in \mu$ and $g \in \mathrm{Pol}_1(\mathcal{A})$, then $fg \in K$, so $\langle fg(a), fg(b)\rangle \in \alpha$ for all $f \in K$. But this implies $\langle g(a), g(b)\rangle \in \mu$.

So μ is either α or β. There is an (α,β)-separating polynomial f with range U. For this f, we find a β-related pair $\langle a,b\rangle$ such that $f(a)$ and $f(b)$ are not α-related and so $\langle a,b\rangle \notin \mu$. It follows that $\mu = \alpha$. By the definition of μ, there is a $g \in K$ such that $\langle gf(a), gf(b)\rangle \notin \alpha$.

So $g(A)$ and $gf(A)$ are both subsets of U and gf does not collapse β into α. By the minimality of U, $gf(A) = U$. Since $f(A) = U$, this implies $g(U) = U$, i.e. $g|_U$ is a permutation. By taking a suitable iterate of g, we get an idempotent polynomial $e(x)$ of $\mathcal{A}$ with range U. □

2.6. Exercise If σ is any binary relation on A which is compatible with all unary polynomials, then the equivalence relation generated by σ is a congruence.

Two sets $X, Y \subseteq A$ are *polynomially isomorphic*, denoted $X \simeq Y$, if there are polynomials $f, g \in \mathrm{Pol}_1(\mathcal{A})$ such that $f(X) = Y$, $g(Y) = X$, $fg|_Y = \mathrm{id}_Y$ and $gf|_X = \mathrm{id}_X$. In this event we write $f\colon X \simeq Y$.

2.7. Exercise

(a) If $f\colon X \simeq Y$, then $\mathcal{A}|_X$ and $\mathcal{A}|_Y$ are isomorphic algebras.
(b) A set that is polynomially isomorphic to an (α,β)-minimal set is (α,β)-minimal.

2.8. Theorem *Let $\mathcal{A}$ be a finite algebra and let $\alpha \prec \beta$ be two congruences.*

(a) (Uniformity) Any two (α,β)-minimal sets are polynomially isomorphic.
(b) Every (α,β)-minimal set is a neighbourhood.
(c) (Incompressibility) If U is (α,β)-minimal and $f \in \mathrm{Pol}_1(\mathcal{A})$ is such that $f(\beta|_U) \nsubseteq \alpha$, then $f(U)$ is (α,β)-minimal and $f|_U$ is a polynomial isomorphism between U and $f(U)$.
(d) (Separation) If $\langle a,b\rangle \in \beta \setminus \alpha$ and U is (α,β)-minimal, then there is $f \in \mathrm{Pol}_1(\mathcal{A})$ such that $f(A) = U$ and $\langle f(a), f(b)\rangle \notin \alpha$.
(e) (Connectedness) If U is (α,β)-minimal, then β is the transitive closure of

$$\alpha \cup \{\langle g(x), g(y)\rangle \mid \langle x,y\rangle \in \beta|_U \text{ and } g \in \mathrm{Pol}_1(\mathcal{A})\}.$$

(f) If f is an (α,β)-separating polynomial, then there is an (α,β)-minimal set U with $U \simeq f(U)$.

A pair (α,β) of congruences satisfying (a)–(f) of the previous theorem is called *tame* by McKenzie. This is the origin of the name *tame congruence theory.* McKenzie showed that not only pairs (α,β) with $\alpha \prec \beta$ are tame, but also the pairs such that the interval $[\alpha,\beta]$ is a simple lattice whose largest element is the join of its atoms. Examples of such lattices are subspace lattices of finite-dimensional vectorspaces.

2.9. Example Let $\mathcal{A}$ be a *simple* algebra, i.e. $0_{\mathcal{A}} \prec 1_{\mathcal{A}}$. Then $\mathrm{Sep}(0_{\mathcal{A}}, 1_{\mathcal{A}})$ is the set of all nonconstant unary polynomials and $M_{\mathcal{A}}(0_{\mathcal{A}}, 1_{\mathcal{A}})$ is the set of all minimal ranges of such polynomials.

If $\mathcal{A}$ is, for example, a finite simple nonabelian group, then the minimal sets are two-element boolean algebras. In this case, even more is true: $M_{\mathcal{A}}(0_{\mathcal{A}}, 1_{\mathcal{A}})$ contains all two-element subsets of A. This follows from the fact that any function on $\mathcal{A}$ is a polynomial (see [BS81] for a proof).

The meaning of (e) is best seen in a picture. Here $a \ \beta \ b$.

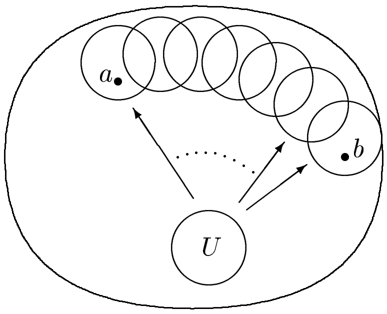

Proof of Theorem 2.8 (b) is already done.
(d) Let U be an (α, β)-minimal set and let e be an idempotent polynomial with range U. Consider the relation

$$\theta = \{\langle x, y\rangle \in \beta \mid \langle ef(x), ef(y)\rangle \in \alpha \text{ for all } f \in \mathrm{Pol}_1(\mathcal{A})\}.$$

This is a congruence relation (as in the proof of Lemma 2.5, since $\{ef \mid f \in \mathrm{Pol}_1(\mathcal{A})\}$ is a right ideal) that lies between α and β. In fact, $\theta = \alpha$, since U is the range of an (α, β)-separating polynomial. Let $\langle a, b\rangle \in \beta \setminus \alpha$. Since $\langle a, b\rangle \notin \theta$, there is a polynomial g such that $\langle eg(a), eg(b)\rangle \notin \alpha$. But $\langle eg(a), eg(b)\rangle \in \beta$ and $eg(A) \subseteq U$. By the minimality of U, $eg(A) = U$. Thus $f(x) = eg(x)$ is the required polynomial.
(e) Let θ be the transitive closure of the relation

$$\gamma = \alpha \cup \{\langle g(x), g(y)\rangle \mid \langle x, y\rangle \in \beta|_U \text{ and } g \in \mathrm{Pol}_1(\mathcal{A})\}.$$

It is easy to see that γ is reflexive, symmetric and compatible with all unary polynomials. By Exercise 2.6, this implies that θ is a congruence. Since $\alpha \subset \gamma \subseteq \theta \leq \beta$, we have $\theta = \beta$.
(a) Choose idempotent polynomials e_0 and e_1 such that $e_0(A) = U$ and $e_1(A) = V$. Furthermore, choose a pair $\langle a, b\rangle \in \beta|_U \setminus \alpha$. By (d), there is a polynomial f with $f(U) = V$ such that $\langle f(a), f(b)\rangle \notin \alpha$.

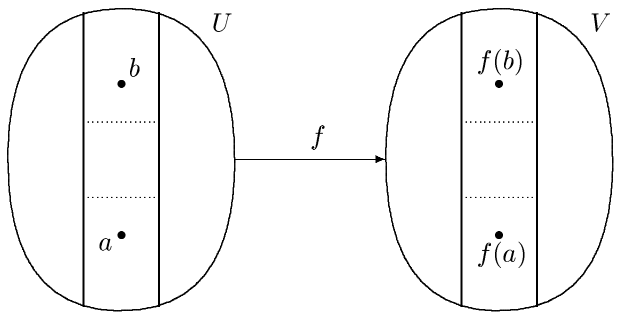

Consider $e_1 f e_0 \in \mathrm{Pol}_1(\mathcal{A})$. The range of this polynomial is contained in V and $e_1 f e_0(\beta) \not\subseteq \alpha$, since $e_1 f e_0(a) = f(a) \not\alpha\, f(b) = e_1 f e_0(b)$. By the minimality of V, we must have $e_1 f e_0(A) = V$. Since $e_0(A) = U$, it follows that $e_1 f(U) = V$. Let $g(x)$ be the polynomial $e_1 f(x)$.

In a similar way we get a polynomial $h(x)$ such that $h(V) = U$. So $hg|_U$ is a permutation of U. Choose a $k > 0$ such that $(hg)^k(x) = x$ for all $x \in U$. Then $l(x) = (hg)^{k-1} h(x)$ is the inverse of $g(x)$ on V, i.e. $l(V) = U$, $lg|_U = \mathrm{id}_U$ and $gl|_U = \mathrm{id}_V$.

(c) Let $U = e(A)$ for a suitable idempotent polynomial e.

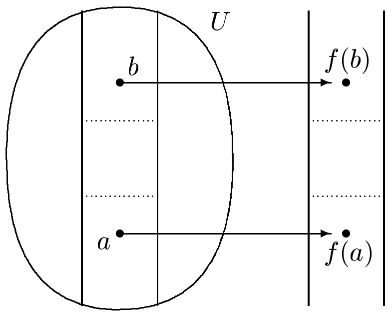

Since $f(\beta|_U) \not\subseteq \alpha$, there is a pair $\langle a, b\rangle \in \beta|_U \setminus \alpha$ such that $f(a)$ and $f(b)$ are not α-related.

Choose by (d) a polynomial $g(x)$ with range U such that $gf(a) \not\alpha\, gf(b)$. Note that $\langle gf(a), gf(b)\rangle = \langle gfe(a), gfe(b)\rangle$. Hence $gfe(x)$ is an (α,β)-separating polynomial whose range is contained in U. By the minimality of U, $gf(U) = gfe(A) = U$. So $g|_{f(U)}$ and $f|_U$ are bijections. By Exercise 2.7, $f(U)$ is (α,β)-minimal.

(f) Let $f \in \mathrm{Sep}(\alpha,\beta)$ and let V be an arbitrary (α,β)-minimal set. There is a pair $\langle a, b\rangle \in \beta \setminus \alpha$ such that $\langle f(a), f(b)\rangle \in \beta \setminus \alpha$. By (e), a and b are connected, modulo α, by a chain of pairs of the form $\langle g_i(c_i), g(d_i)\rangle$ for suitable polynomials g_i and pairs $\langle c_i, d_i\rangle \in \beta|_V \setminus \alpha$ and so the pairs $\langle fg_i(c_i), fg_i(d_i)\rangle$ connect $f(a)$ and $f(b)$ modulo α. There is an index i such that for $g = g_i$ and $\langle c, d\rangle = \langle c_i, d_i\rangle$ the pair $\langle fg(c), fg(d)\rangle$ lies in $\beta \setminus \alpha$.

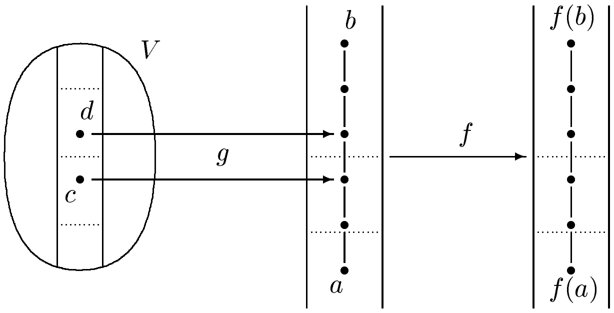

In particular, $g(c)$ and $g(d)$ cannot lie in the same α-class. By (c) we know that $U = g(V)$ is an (α, β)-minimal set. Finally, $f(\beta|_U) \not\subseteq \alpha$, since $fg(c)$ is not α-related to $fg(d)$.

□

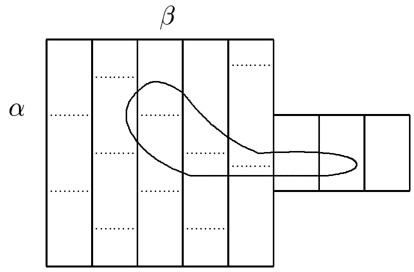

Let U be an arbitrary (α, β)-minimal set. The $\beta|_U$-classes which split into more than one $\alpha|_U$-class are called *(α, β)-traces*. By definition, every (α, β)-minimal set contains at least one trace. The union of all traces of U is called the *body* of U. The remainder is called the *tail* of U.

The congruence lattice of an algebra $\mathcal{A}$ is $\cap$-complete. For a set $X \subseteq A^2$, the *congruence generated by X* is

$$\mathrm{Cg}(X) = \bigcap\{\theta \in \mathrm{Con}(\mathcal{A}) \mid X \subseteq \theta\}.$$

2.10. Exercise

(a) $\mathrm{Cg}(X)$ is the equivalence relation generated by $\{\langle g(x), g(y)\rangle \mid \langle x, y\rangle \in X, g \in \mathrm{Pol}_1(\mathcal{A})\}$. In fact, $\mathrm{Cg}(X)$ is the symmetric transitive closure of this relation.

Stated in another way, $\langle a, b\rangle \in \mathrm{Cg}(X)$ iff there are polynomials $g_i(x)$ and pairs $\langle x_i, y_i\rangle \in X \cup X^{-1}$ such that

$$\begin{aligned} a &= g_0(x_0) \\ g_0(y_0) &= g_1(x_1) \\ g_1(y_1) &= g_2(x_2) \\ &\ddots \\ g_k(y_k) &= b. \end{aligned}$$

(b) If U is a neighbourhood of $\mathcal{A}$, $X \subseteq U^2$ and $a, b \in U$, then $\langle a, b\rangle \in \mathrm{Cg}_{\mathcal{A}}(X)$ iff $\langle a, b\rangle \in \mathrm{Cg}_{\mathcal{A}|_U}(X)$.

2.11. Exercise Let $\mathcal{A}$ be a finite algebra. If $g \in \mathrm{Pol}_1(\mathcal{A})$ is a permutation and $\beta \in \mathrm{Con}(\mathcal{A})$, then g permutes the β-classes of $\mathcal{A}$.

2.12. Proposition *Any two (α, β)-traces of $\mathcal{A}$ are polynomially isomorphic.*

Proof Since any two (α,β)-minimal sets are polynomially isomorphic, it suffices to verify the proposition for traces N and M of the same minimal set U. Since U is a neighbourhood, we may assume $A = U$.

It suffices to find a polynomial $p(x)$ with $p(N) \subseteq M$ and $p(\beta|_N) \not\subseteq \alpha$, since this implies that p is a permutation of A and then Exercise 2.11 gives the result.

Since N and M are traces, there are pairs $\langle a,b\rangle \in \beta|_N \setminus \alpha$ and $\langle c,d\rangle \in \beta|_M \setminus \alpha$. Since α is covered by β,

$$\beta = \mathrm{Cg}(\alpha \cup \{\langle a,b\rangle\}).$$

In particular, $\langle c,d\rangle \in \mathrm{Cg}(\alpha \cup \{\langle a,b\rangle\})$. Choose polynomials g_i and pairs $\langle x_i,y_i\rangle$ according to Exercise 2.10.

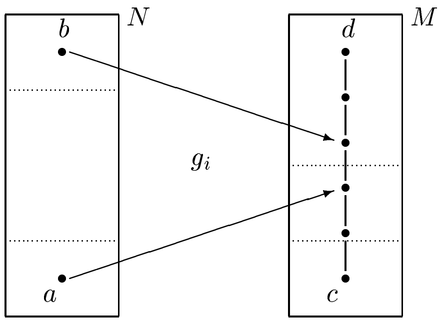

There is an i such that $\langle g_i(x_i), g_i(y_i)\rangle \in \beta \setminus \alpha$. Since $g_i(x_i) \not\alpha\, g_i(y_i)$, we must have $\{x_i,y_i\} = \{a,b\}$. So g_i maps the β-class of a, namely N, into M and $g_i(\beta|_N) \not\subseteq \alpha$. □

2.13. Exercise Let $\alpha \prec \beta$. If U is an (α,β)-minimal set of $\mathcal{A}$, then U/α is an $(0,\beta/\alpha)$-minimal set of $\mathcal{A}/\alpha$.

2.14. Proposition *If $\alpha \prec \beta$ in a finite algebra $\mathcal{A}$ and N is an (α,β)-trace, then $\alpha|_N$ is a congruence on $\mathcal{A}|_N$ and $\mathcal{A}|_N/\alpha|_N$ is a minimal simple algebra.*

Proof By Exercise 2.13, we can assume that $\alpha = 0$. We need to show that every polynomial $p(x) \in \mathrm{Pol}_1(\mathcal{A}|_N)$ is either constant or a permutation.

Let such a p be given. It is of the form $p(x) = q(x)|_N$ for some polynomial $q(x)$ of $\mathcal{A}$ with $q(N) \subseteq N$. Choose a $(0,\beta)$-minimal set U containing N. Since U is the range of some idempotent polynomial e of $\mathcal{A}$, we may assume $q(A) \subseteq U$ by replacing q by eq.

If p is not constant, then $q(\beta|_N) \not\subseteq 0$. By the minimality of U, we must have $q(U) = U$. Hence q is a permutation of U. In particular, $q|_N = p$ is a permutation of N.

The simplicity of $\mathcal{A}|_N$ follows easily from the fact that $0 \prec \beta$ and is left as an exercise. □

The previous proposition allows us to define $\mathrm{typ}(\alpha,\beta)$ to be the type of the minimal algebra $\mathcal{A}|_N/\alpha|_N$ for any (α,β)-trace N. This turns $\mathrm{Con}(\mathcal{A})$ into a labeled lattice. We also define

$$\mathrm{typ}\{\mathcal{A}\} = \{\mathrm{typ}(\alpha,\beta) \mid \alpha \prec \beta \text{ in } \mathrm{Con}(\mathcal{A})\},$$

and for a class $\mathcal{K}$ of algebras, we define

$$\operatorname{typ}(\mathcal{K}) = \bigcup \{ \operatorname{typ}\{\mathcal{A}\} \,:\, \mathcal{A} \in \mathcal{K},\ \mathcal{A} \text{ finite}\}.$$

2.15. Exercise Let $\mathcal{A}$ be either unary or polynomially equivalent to a lattice or semilattice. Then $\mathcal{A}$ is not Mal'cev.

2.16. Lemma *If $\mathcal{A}$ is Mal'cev, then* $\operatorname{typ}\{\mathcal{A}\} \subseteq \{2,3\}$.

Proof Let $\alpha \prec \beta$. We want to determine $\operatorname{typ}(\alpha,\beta)$. Since homomorphic images of Mal'cev algebras are also Mal'cev, we can replace $\mathcal{A}$ by $\mathcal{A}/\alpha$ and assume $\alpha = 0$. Let $p(x,y,z)$ be a Mal'cev polynomial. Choose a $(0,\beta)$-trace N contained in the $(0,\beta)$-minimal set U. Let $e(x)$ be an idempotent polynomial whose range is U.

Claim: $ep(x,y,z)|_N$ is a Mal'cev polynomial of $\mathcal{A}|_N$.

The equations defining a Mal'cev operation are clearly satisfied. It remains to show that $ep(N,N,N) \subseteq N$. So let $a,b,c \in N$. Then $a \;\beta\; b$, so $ep(a,b,c) \;\beta\; ep(b,b,c) = c$ and $ep(a,b,c) \in e(A) = U$. Since N is the intersection of U with the β-class of c, we conclude $ep(a,b,c) \in N$.

By Exercise 2.15 and the claim, $\mathcal{A}|_N$ must be of type 2 or 3. $\square$

3 Centrality

Let α, β and γ be congruences of an algebra $\mathcal{A}$. We say that α *centralizes* β *modulo* γ, denoted by $C(\alpha,\beta,\gamma)$, if for all polynomials $t(x,\bar{y})$ and all elements $a,b,\bar{c},\bar{d} \in A$ such that $a \;\alpha\; b$ and $c_i \;\beta\; d_i$ we have

$$t(a,\bar{c}) \;\gamma\; t(a,\bar{d}) \quad \Rightarrow \quad t(b,\bar{c}) \;\gamma\; t(b,\bar{d}).$$

Note that $\mathcal{A}$ is abelian iff $C(1_{\mathcal{A}},1_{\mathcal{A}},0_{\mathcal{A}})$.

3.1. Proposition

(a) If $\alpha' \leq \alpha$, $\beta' \leq \beta$ and $C(\alpha,\beta,\gamma)$, then $C(\alpha',\beta',\gamma)$.
(b) If $C(\alpha_i,\beta,\gamma)$ for all i, then $C(\bigvee \alpha_i,\beta,\gamma)$.
(c) If $C(\alpha,\beta,\gamma_i)$ for all i, then $C(\alpha,\beta,\bigwedge \gamma_i)$.

Proof (a) $C(\alpha,\beta,\gamma)$ is a condition on all pairs in α and β. If we make these relations smaller, the condition remains true for all the pairs in the smaller relations.
(b) Assume $C(\alpha_i,\beta,\gamma)$ for all i. Let $\theta = \bigvee \alpha_i$. Let $t(x,\bar{y})$ be a term and $a,b,\bar{c},\bar{d} \in A$ such that $a \;\theta\; b$ and so that the corresponding entries of $\bar{c}$ and $\bar{d}$ are β-related. Assume that $t(a,\bar{c}) \;\gamma\; t(a,\bar{d})$.

Since θ is the transitive closure of $\bigcup \alpha_i$, there are elements a_j and congruences α_{i_j} such that

$$a = a_0 \;\alpha_{i_0}\; a_1 \;\alpha_{i_1}\; \ldots \;\alpha_{i_{n-1}}\; a_n = b.$$

Now apply $C(\alpha_{i_j},\beta,\gamma)$ to conclude first $t(a_1,\bar{c}) \;\gamma\; t(a_1,\bar{d})$, then $t(a_2,\bar{c}) \;\gamma\; t(a_2,\bar{d})$ and so on.
(c) Assume $C(\alpha,\beta,\gamma_i)$ for all i. Let $\delta = \bigwedge \gamma_i$. Let $t(x,\bar{y})$ be a polynomial and $a,b,\bar{c},\bar{d} \in A$ such that $a \;\alpha\; b$ and so that the corresponding entries of $\bar{c}$ and $\bar{d}$ are β-related. Assume that $t(a,\bar{c}) \;\delta\; t(a,\bar{d})$. Then $t(a,\bar{c}) \;\gamma_i\; t(a,\bar{d})$ for all i and we can apply the conditions $C(\alpha,\beta,\gamma_i)$ to conclude $t(b,\bar{c}) \;\gamma_i\; t(b,\bar{d})$ for all i. But then $t(b,\bar{c}) \;\delta\; t(b,\bar{d})$.

$\square$

By the second claim of the proposition there is a largest congruence θ with $C(\theta, \beta, \gamma)$. θ is called the *annihilator of* β *modulo* γ. By the third claim, there is a smallest congruence δ with $C(\alpha, \beta, \delta)$. δ is called the *commutator of* α *and* β and is denoted by $[\alpha, \beta]$.

In general, the commutator operation $[\ ,\ ]$ is not commutative and does not behave well with respect to quotients. But in congruence modular varieties, $[\ ,\ ]$ behaves similarly to the group-theoretic commutator in the variety of groups (see [FM87]).

3.2. Exercise The annihilator of β modulo γ is the relation

$$\{\langle a, b\rangle \mid t(a, \bar{c})\ \gamma\ t(a, \bar{d}) \Leftrightarrow t(b, \bar{c})\ \gamma\ t(b, \bar{d}) \text{ for all polynomials } t(x, \bar{y}) \text{ and all } c_i\ \beta\ d_i\}.$$

Let $\alpha \leq \beta$. We say that β is *abelian over* α if $C(\beta, \beta, \alpha)$ holds. β is *solvable over* α if there are congruences

$$\alpha = \alpha_0 \leq \alpha_1 \leq \cdots \leq \alpha_n = \beta$$

such that α_{i+1} is abelian over α_i for all $i < n$. If $\alpha = 0$, we say just that β is *abelian* or *solvable*.

3.3. Exercise β is abelian over α iff β/α is abelian.

3.4. Exercise Define the following descending chain of congruences:

$$\begin{aligned}
[\alpha, \alpha]^0 &= \alpha, \\
[\alpha, \alpha]^{n+1} &= [[\alpha, \alpha]^n, [\alpha, \alpha]^n].
\end{aligned}$$

Show that $\mathcal{A}$ is solvable iff $[1_{\mathcal{A}}, 1_{\mathcal{A}}]^n = 0_{\mathcal{A}}$ for some $n < \omega$.

(In a similar way, one has descending chains $\alpha \geq [\alpha, \alpha] \geq [\alpha, [\alpha, \alpha]] \geq \ldots$ and $\alpha \geq [\alpha, \alpha] \geq [[\alpha, \alpha], \alpha] \geq \ldots$ $\mathcal{A}$ is called *left/right nilpotent* if one of these chains reaches $0_{\mathcal{A}}$ after finitely many steps.)

3.5. Exercise If α is a congruence, $S_1, \ldots, S_n$ are α-classes and $f(x_1, \ldots, x_n)$ is a polynomial, then $f(S_1, \ldots, S_n)$ is contained in an α-class.

3.6. Exercise If $\beta \succ \alpha$, then β is the transitive closure of $\alpha \cup \bigcup\{N^2 \mid N$ an (α, β)-trace$\}$.

3.7. Theorem *Let* $\alpha \prec \beta$ *be congruences of the finite algebra* $\mathcal{A}$. β *is abelian over* α *iff* $\operatorname{typ}(\alpha, \beta) \in \{1, 2\}$.

Proof By Exercise 3.3, we can assume that $\alpha = 0$.

$\Leftarrow$ Let N be a $(0, \beta)$-trace. If $\operatorname{typ}(0, \beta) \notin \{1, 2\}$, $\mathcal{A}|_N$ is polynomially equivalent to a boolean algebra, lattice or a semilattice. Thus we can name the elements of N as 0 and 1 and find a polynomial which defines $\wedge$ on N. Then we have $0 \wedge 0 = 0 \wedge 1$ and $0\ \beta\ 1$, but $1 \wedge 0 \neq 1 \wedge 1$, contradicting the abelianness of β.

$\Rightarrow$ We have to show $C(\beta, \beta, 0_{\mathcal{A}})$, i.e. that β is contained in the annihilator of β modulo $0_{\mathcal{A}}$. Since $0 \prec \beta$, it suffices to find a pair $\langle a, b\rangle \in \beta \setminus 0_{\mathcal{A}}$ that also lies in the annihilator.

Claim: Let $N \subseteq U$ be a $(0_{\mathcal{A}}, \beta)$-trace. Then N^2 is contained in the annihilator of β modulo $0_{\mathcal{A}}$.

Choose elements a and b in N. We have to show for any polynomials $t(x, \bar{y})$ and elements $c_i \;\beta\; d_i$ that $t(a, \bar{c}) = t(a, \bar{d})$ implies $t(b, \bar{c}) = t(b, \bar{d})$.

Let U be a $(0_{\mathcal{A}}, \beta)$-minimal set containing U and e an idempotent polynomial with range U. Let $f(x, \bar{y})$ be a polynomial such that $f(a, \bar{c}) = f(a, \bar{d})$ and $f(b, \bar{c}) \neq f(b, \bar{d})$ for some $c_i \;\beta\; d_i$. We know that $\langle f(b, \bar{c}), f(b, \bar{d}) \rangle \in \beta$. By Theorem 2.8 we may assume that the range of f is contained in U. By Proposition 2.12 we may also assume that $\{f(b, \bar{c}), f(b, \bar{d})\} \subseteq N$. Note that $f(a, \bar{c}) \in N$, since $f(a, \bar{c})$ lies in U and is β-related to $f(b, \bar{c})$, an element of N.

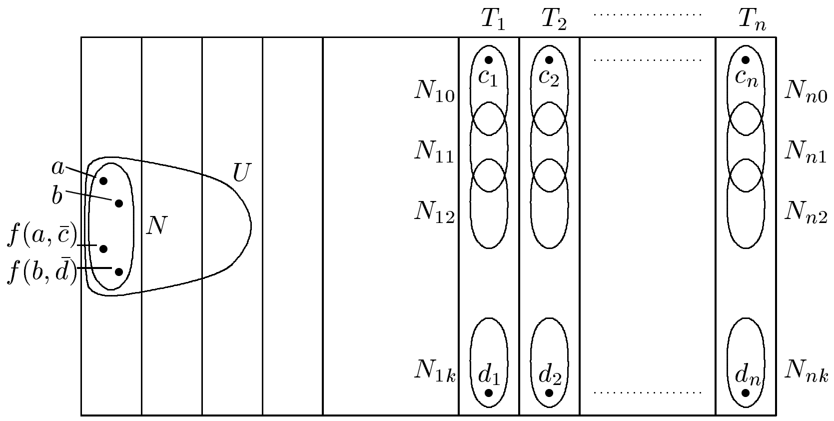

Let T_i be the β-class of c_i and d_i. Then $f(N, T_1, \ldots, T_n) \subseteq N$ by Exercise 3.5, since $f(b, c_1, \ldots, c_n) \in N$. By Exercise 3.6, each c_i can be connected to the corresponding d_i by a chain of overlapping $(0_{\mathcal{A}}, \beta)$-traces. By allowing repetitions we can assume that all chains have the same length. Thus there is a $k < \omega$ and traces N_{ij} for $1 \leq i \leq n$ and $0 \leq j \leq k$ such that

$$c_i \in N_{i0}, \qquad d_i \in N_{ik}, \quad \text{and} \quad N_{ij} \cap N_{ij+1} \neq \emptyset \quad \text{for all } i \text{ and } j.$$

Since each N_{ij} is polynomially equivalent to N, there exist polynomials α_{ij} such that $\alpha_{ij}(N) = N_{ij}$. Define polynomials f_j by

$$f_j(x, y_1, \ldots, y_n) = f(x, \alpha_{1j}(y_1), \ldots, \alpha_{nj}(y_n)).$$

Since $f(N, T_1, \ldots, T_n) \subseteq N$ and $\alpha_{ij}(N) \subseteq T_i$ then $f_j(N, N, \ldots, N) \subseteq N$, so $f_j|_N \in \mathrm{Pol}(\mathcal{A}|_N)$. Choose elments $c'_1, \ldots, c'_n$ and $d'_1, \ldots, d'_n$ in N such that $\alpha_{i0}(c'_i) = c_i$ and $\alpha_{ik}(d'_i) = d_i$.

Now the proof breaks into cases depending on whether $\mathcal{A}|_N$ is essentially unary or polynomially equivalent to a vectorspace.

Case 1. $\mathcal{A}|_N$ is polynomially equivalent to a vectorspace over the finite field F. We can write the $f_j|_N$ in the form

$$f_j|_N(x, y_1, \ldots, y_n) = \mu_j x + \sum \lambda_{ji} y_i + e_j$$

for suitable $\mu_j, \lambda_{j1}, \ldots, \lambda_{jn} \in F$ and $e_j \in N$.

Claim: $\mu_j = \mu_{j+1}$ for all $j < k$.

Let j be fixed. Choose elements $u_i, v_i \in N$ such that $\alpha_{ij}(u_i) = \alpha_{ij+1}(v_i)$. Such elements exist, since $N_{ij} \cap N_{ij+1} \neq \emptyset$. Then we have for all $x \in N$

$$\mu_j x + (\sum \lambda_{ji} u_i + e_j) = f_j(x, \bar{u}) = f(x, \alpha_{1j}(u_1), \dots, \alpha_{nj}(u_n)) =$$
$$f(x, \alpha_{1j+1}(v_1), \dots, \alpha_{nj+1}(v_n)) = f_{j+1}(x, \bar{v}) = \mu_{j+1} x + (\sum \lambda_{j+1i} v_i + e_{j+1}),$$

so $(\mu_j - \mu_{j+1})x$ is constant on N. This forces $\mu_j = \mu_{j+1}$, proving the claim.

By the claim, we have $\mu_0 = \mu_k$. Together with

$$\mu_0 a + \sum \lambda_{0i} c_i' + e_0 = f(a, \bar{c}) = f(a, \bar{d}) = \mu_k a + \sum \lambda_{ki} d_i' + e_d \ ,$$

this implies $\sum \lambda_{0i} c_i' + e_0 = \sum \lambda_{ki} d_i' + e_d$. But then adding $\mu_0 b = \mu_k b$ yields

$$f(b, \bar{c}) = \mu_0 b + \sum \lambda_{0i} c_i' + e_0 = \mu_k b + \sum \lambda_{ki} d_i' + e_d = f(b, \bar{d}).$$

Since this is exactly what was needed to show that $\langle a, b \rangle$ lies in the annihilator of β modulo $0_{\mathcal{A}}$, we are finished.

Case 2. $\mathcal{A}|_N$ is essentially unary. Then the polynomials $f_j|_N$ depend on at most one variable.

Claim: f_j depends on x iff f_{j+1} depends on x.

We show the direction from left to right, the other one is similar. Let $f_j|_N$ depend on x. Then there are $u, u', \bar{v} \in N$ such that $f_j(u, \bar{v}) \neq f_j(u', \bar{v})$. Since $f_j|_N$ does not depend on the variables $\bar{y}$, we can choose the tuple $\bar{v}$ arbitrarily. In particular, we can choose $\bar{v}$ so that $\alpha_{ij}(v_i) \in N_{ij} \cap N_{ij+1}$. But then there is a tuple $\bar{v}' \in N$ such that $\alpha_{ij+1}(v_i') = \alpha_{ij}(v_i)$ and consequently

$$f_{j+1}(u, \bar{v}') = f_j(u, \bar{v}) \neq f_j(u', \bar{v}) = f_{j+1}(u', \bar{v}') \ .$$

Thus $f_{j+1}|_N$ depends on x.

If $f_1|_N$ does not depend on x, then trivially

$$f(b, \bar{c}) = f_1(b, \bar{c}') = f_1(a, \bar{c}') = f(a, \bar{c}) = f(a, \bar{d}) = f_1(a, \bar{d}') = f_1(b, \bar{d}') = f(b, \bar{d}).$$

If $f_1|_N$ depends on x, then by the claim, no f_j depends on the variables $\bar{y}$, so the maps $f_j(b, \bar{y})$ are constant on N^n. By the definition of the f_j's, this just means that $f(b, \bar{y})$ is constant on the sets $N_{1j} \times \cdots \times N_{nj}$. Since these set are overlapping and their union contains $\bar{c}$ and $\bar{d}$, we conclude $f(b, \bar{c}) = f(b, \bar{d})$.

□

Let α, β and γ be congruences. We say that α *strongly centralizes* β modulo γ, denoted $C^*(\alpha, \beta, \gamma)$, if

$$f(a, \bar{c}) \ \gamma \ f(b, \bar{d}) \quad \Rightarrow \quad f(a, \bar{e}) \ \gamma \ f(b, \bar{e})$$

for all polynomials $f(x, \bar{y})$ and all $a, b, \bar{c}, \bar{d}, \bar{e} \in A$ such that $a \ \alpha \ b$ and $c_i \ \beta \ d_i \ \beta \ e_i$. If $\alpha \leq \beta$ and $C^*(\beta, \beta, \alpha)$, we call β *strongly abelian* over α.

An algebra $\mathcal{A}$ is *strongly abelian* if $1_{\mathcal{A}}$ is strongly abelian over $0_{\mathcal{A}}$.

3.8. Example 1. Any essentially unary algebra is strongly abelian. To see this, consider an equation of the form $f(a, \bar{c}) = f(b, \bar{d})$. Either the polynomial $f(x, \bar{y})$ depends on x, in which case we can simply plug in any tuple $\bar{e}$ for the variables $\bar{y}$, or $f(x, \bar{y})$ depends on one of the y's, say the first one.

Then $f(x, e_1, y_2 \dots, y_n)$ does not depend on any of the variables $x, y_2, \dots, y_n$, so in particular $f(a, e_1, e_2 \dots, e_n) = f(b, e_1, e_2, \dots, e_n)$.

2. No nontrivial module is strongly abelian, since we have $0 - 0 = m - m$, but $0 - 0 \neq m - 0$ for any nonzero element m.

3.9. Exercise

(a) Find a finite simple algebra which is of type 1 but which is not essentially unary. Note that this algebra must be strongly abelian.
(b) Find a finite simple algebra which is of type 2 but which is not polynomially equivalent to a vectorspace. Note that this algebra must be abelian.
(c) Find a finite algebra which is strongly abelian but which has a quotient which is not abelian.

3.10. Exercise The following statements are equivalent:

(i) β is strongly abelian over α.
(ii) β/α is strongly abelian.
(iii) $f(\bar{a}, \bar{c})\ \alpha\ f(\bar{b}, \bar{d})$ implies $f(\bar{a}, \bar{e})\ \alpha\ f(\bar{b}, \bar{e})$ for all polynomials $f(\bar{x}, \bar{y})$ and all $\bar{a}, \bar{b}, \bar{c}, \bar{d}, \bar{e} \in A$ such that $a_i\ \beta\ b_i$ and $c_i\ \beta\ d_i\ \beta\ e_i$.

3.11. Exercise Strong abelianness implies abelianness.

3.12. Exercise

(a) The following are equivalent:
  (i) $\mathcal{A}$ is strongly abelian.
  (ii) For each term $t(x_1, \dots, x_n)$ of $\mathcal{A}$ there are equivalence relations $E_1, \dots, E_n$ on A such that for all $\bar{a}$ and $\bar{b}$

$$t(\bar{a}) = t(\bar{b}) \quad \text{iff} \quad a_i\ E_i\ b_i \text{ for all } 1 \leq i \leq n\ .$$

(b) If $\mathcal{A}$ is finite and strongly abelian, then no term can depend on more than $\log_2 |A|$ variables.

3.13. Theorem *Let $\alpha \prec \beta$ be congruences of the finite algebra $\mathcal{A}$. β is strongly abelian over α iff* $\operatorname{typ}(\alpha, \beta) = 1$.

Proof $\Rightarrow$ If $\operatorname{typ}(\alpha, \beta) \in \{3, 4, 5\}$, $\mathcal{A}$ is not even abelian. By Example 3.8.2, $\mathcal{A}$ is also not strongly abelian if $\operatorname{typ}(\alpha, \beta) = 2$.
$\Leftarrow$ By Exercise 3.10 we may assume that $\alpha = 0_{\mathcal{A}}$ and $0_{\mathcal{A}} \prec \beta$. The key to our proof is the following claim:

Claim: If $f(x_1, \dots, x_n)$ is a polynomial, $T_1, \dots, T_n$ are β-classes and $f(T_1, \dots, T_n) \subseteq N$ for some $(0_{\mathcal{A}}, \beta)$-trace N then f depends on at most one variable when restricted to $T = T_1 \times T_2 \times \cdots \times T_n$.

We leave the reduction of this claim to the case when $n = 2$ to the reader, and will assume that f is binary. Suppose that $f|_T$ depends on both of its variables. Then there are elements $a_i \in T_i$, $i = 1, 2$ with $f(a_1, y)$ nonconstant on T_2 and $f(x, a_2)$ nonconstant on T_1. Since the $(0_{\mathcal{A}}, \beta)$-traces cover both T_1 and T_2 then we can find traces $N_i \subseteq T_i$ with $f(a_1, y)$ nonconstant on N_2 and $f(x, a_2)$ nonconstant on N_1.

Since β is abelian (by Theorem 3.7) then we may assume that $a_i \in N_i$ for $i = 1, 2$. Thus, f depends on both of its variables on the set $N_1 \times N_2$. Since N is polynomially isomorphic to N_1 and N_2, we can find unary polynomials $\alpha_1(x)$ and $\alpha_2(x)$ so that $\alpha_i(N) = N_i$, $i = 1, 2$. But then the polynomial $f(\alpha_1(x), \alpha_2(y))$ maps

$N \times N$ into N and depends on both of its variables on this set. This contradicts that the induced structure of $\mathcal{A}$ on N is essentially unary.

To conclude the proof, suppose that a, b, $\bar{c}$, $\bar{d}$, and $\bar{e}$ are in A with $\langle a, b\rangle \in \beta$ and for each $i \le n$ (where n is the length of $\bar{c}$), c_i, d_i and e_i are β-related. If f is a polynomial of $\mathcal{A}$ with $f(a,\bar{c}) = f(b,\bar{d})$ and $f(a,\bar{e}) \neq f(b,\bar{e})$ then by applying a suitable unary polynomial, we may assume that the range of f is contained in some $(0_{\mathcal{A}},\beta)$-minimal set U. By setting T_i to be the β-class which contains c_i, we have that f maps the set $T = T_1 \times T_2 \times \cdots \times T_n$ into some $(0_{\mathcal{A}},\beta)$-trace N contained in U.

Using the claim, we conclude that $f(x,\bar{y})|_T$ depends on at most one variable. By considering the various possibilities, we see that $f(a,\bar{c}) = f(b,\bar{d})$ implies that $f(a,\bar{e}) = f(b,\bar{e})$, contradicting our assumptions.

$\square$

If $\alpha \le \beta$ are congruences of $\mathcal{A}$ and $M_{\mathcal{A}}(\alpha,\beta) = \{A\}$, then $\mathcal{A}$ is said to be an *(α,β)-minimal algebra.*

3.14. Exercise If $\mathcal{A}$ is (α,β)-minimal, then $\mathcal{A}/\alpha$ is $(0_{\mathcal{A}/\alpha},\beta/\alpha)$-minimal.

Two polynomials $f(\bar{x})$ and $g(\bar{x})$ are called *twins* if there is a polynomial $s(\bar{x},\bar{y})$ and two tuples $\bar{a}$ and $\bar{b}$ such that $f(\bar{x}) = s(\bar{x},\bar{a})$ and $g(\bar{x}) = s(\bar{x},\bar{b})$. If the tuples come from the body of a minimal set, we call f and g *body twins.*

3.15. Lemma (Twin lemma) *Let $\mathcal{A}$ be an (α,β)-minimal finite algebra. If $\mathcal{A}$ has two unary body twins of which one is a permutation and the other is not, then the body B of $\mathcal{A}$ consists of a single (α,β)-trace which is a union of two α-classes. In particular, $\alpha \prec \beta$.*

Furthermore, $\mathcal{A}$ has a polynomial which induces a semilattice operation on B/α and $\operatorname{typ}(\alpha,\beta) \in \{3,4,5\}$.

Proof (after E. Kiss) By Exercise 3.14 we can assume that $\alpha = 0_{\mathcal{A}}$. Let $f_0(x) = h_0(x,\bar{c})$ and $g_0(x) = h_0(x,\bar{d})$ be body twins such that $f_0(x)$ is a permutation and $g_0(x)$ is not. By finding the step at which the bijectivity is lost in the following chain

$$h_0(x,\bar{c}) \quad h_0(x,d_1,c_2,\ldots,c_n) \quad h_0(x,d_1,d_2,c_3,\ldots,c_n) \quad \ldots \quad h_0(x,\bar{d}),$$

we can come up with two body twins $f(x) = h(x,c)$ and $g(x) = h(x,d)$ which come from a binary polynomial $h(x,y)$ and with the former a permutation, and the latter not.

By iterating h in the first variable, we may assume that $h(h(x,y),y) = h(x,y)$. Thus any polynomial of the form $h(x,e)$ is idempotent. In particular, any permutation of this form is the identity.

Claim: No twin of f is a permutation.

h	c	b
x	x	x
y	y	y
$\vdots$	$\vdots$	$\vdots$
z	z	z

Suppose that $b \neq c$ is another parameter for which $h(x,b)$ is a permutation. By the choice of h, we have $h(u,b) = u = h(u,c)$ for all $u \in A$. This implies that no row of h is a permutation. So by the minimality of $\mathcal{A}$, we have $h(u,\beta) \subseteq 0_{\mathcal{A}}$ for all $u \in \mathcal{A}$. Put in a different way, if two elements $e, \hat{e} \in A$ are β-related, the functions $h(x,e)$ and $h(x,\hat{e})$ are the same. So c and d cannot be β-related.

Since c and d are from the body, their β-classes are not singletons and we can choose new elements c' and d' such that $c' \mathrel{\beta} c$ and $d' \mathrel{\beta} d$. Consider $h(x,x)$. We have $h(d,d) = h(d',d) = h(d',d')$. By the minimality of $\mathcal{A}$, $h(x,x)$ collapses β to $0_{\mathcal{A}}$, so $c = h(c,c) = h(c',c') = h(c',c) = c'$. This contradiction finishes the proof of the claim.

Claim: B has exactly two elements.

If not, choose $u, v \in B$ such that $u \;\beta\; v$ and $|\{u,v,c\}| = 3$. By the first claim and the minimality of $\mathcal{A}$, the polynomials $h(x,u)$ and $h(x,v)$ are constant on the classes of β. From this it follows that if $\langle a,b\rangle \in \beta$, then $f(a,x)$ and $f(b,x)$ are either both permutations or both not.

$h\|_N$	c	c'
c	c	c''
c'	c'	$\vdots$
c''	c''	c''

Since $h(c,c) = c$, we have $h(N,N) \subseteq N$, where N is the β-class of c. Consider now $h(x,y)$ on N. Of course, $h(x,c)$ is a permutation, while $h(x,c')$ is constant on N, say its value is c''. This gives $h(c'',c) = c'' = h(c'',c')$, i.e. $h(c'',y)$ is not a permutation. But then none of the functions $h(e,y)$ with $e \in N$ is a permutation, and so they are all constant on N. So

$$c = h(c,c) = h(c,c') = c'' = h(c',c') = h(c',c) = c',$$

contradicting the choice of c'.

At this point we know that $B = N = \{c,d\}$. This implies in particular that $0_{\mathcal{A}} \prec \beta$.

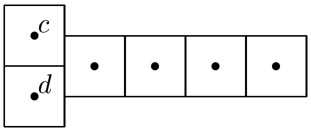

If $h(c,d) = d$ then $h(x,y)$ is a meet operation in the semilattice $\overset{c}{|}d$.
If $h(c,d) = c$, then $h(x,h(d,y))$ is a meet operation in the semilattice $\overset{d}{|}c$.

The presence of a meet operation on N forces $\mathrm{typ}(\alpha,\beta)$ to be 3, 4 or 5. $\square$

3.16. Remark The converse of the lemma is also true. If $\mathcal{A}$ is (α,β)-minimal, $\alpha \prec \beta$ and $\mathrm{typ}(\alpha,\beta) \in \{3,4,5\}$, then there are two body twins of which exactly one is a permutation.

Namely, let $N = \{0,1\}$ be a trace and let p be a polynomial such that $p(x,y)|_N = x \wedge y$. Then $f(x) = p(x,1)|_N = \mathrm{id}_N$, so f does not collapse β to α. By the minimality of $\mathcal{A}$, f is a permutation. It is clear, that $g(x) = p(x,0)$ is not a permutation and that f and g are body twins.

3.17. Lemma *Let $\mathcal{A}$ be finite and (α,β)-minimal for two congruences $\alpha \prec \beta$ with* $\mathrm{typ}(\alpha,\beta) \in \{3,4,5\}$. *Let N be the unique (α,β)-trace of $\mathcal{A}$. Then there is an element $1 \in N$ and a polynomial $p(x,y) \in \mathrm{Pol}_2(\mathcal{A})$ such that*

(a) $N/\alpha = \{O, I\}$, where O and $I = \{1\}$ are the α-classes of N.
(b) N is closed under p and $\langle N, p|_N\rangle/\alpha$ is a semilattice with neutral element I.
(c) $\langle \{a,1\}, p|_{\{a,1\}}\rangle$ is a semilattice for all $a \in A \setminus \{1\}$.
(d) $p(a,u) \;\alpha\; a \;\alpha\; p(u,a)$ for all $a \in A \setminus \{1\}$ and $u \in O$.
(e) $p(x,y)$ is idempotent in the first variable.

Proof Let $g(x,y)$ be a polynomial which induces a semilattice operation on N/α. By renaming O and I, we can assume that g has the following multiplication table.

g	O	I
O	O	O
I	O	I

The polynomial $d(x) = g(x,x)$ doesn't collapse β into α, so it is a permutation of A. Its inverse is of the form d^k, hence it is also a polynomial. Setting $p(x,y) = d^{-1}g(x,y)$, we have $p(x,x) = x$ and p still induces the meet operation on N/α.

By iterating p we may assume that p is idempotent in both variables. Let $z \in I$. If $u \in O$, then $\langle p(z,z), p(u,z)\rangle \in \beta \setminus \alpha$, hence $p(x,z)$ does not collapse β into α and is a permutation of A. Similarly, $p(z,x)$ is a permutation of A. Since these maps are also idempotent, they are the identity and we have

$$p(x,z) = x = p(z,x)$$

for all $z \in I$ and $x \in A$. If z' is another element of I, this gives $z = p(z,z') = z'$, so $|I| = 1$. By now we have established (a), (b) and (e).

(c) $p(1,1) = 1$ is clear since $p(I,I) \subseteq I$. Let $a \neq 1$. By the properties we have already established, $p(a,a) = a$ and $p(a,1) = a = p(1,a)$, so p is indeed a meet operation on $\left[\genfrac{}{}{0pt}{}{1}{a}\right.$.
(d) Let $u \in O$. If $x \in O$, then $p(x,u), p(u,x) \in O$, since $p(O,O) \subseteq O$. This means just $p(x,u)\ \alpha\ x\ \alpha\ p(u,x)$.

If $x \notin N$, then $p(1,x) = x = p(x,1)$, so $p(u,x)\ \beta\ x\ \beta\ p(x,u)$. But since x is in the tail of $\mathcal{A}$, the β-class of x is the same as its α-class.

□

A polynomial satisfying the conditions in Lemma 3.17 is called a *pseudo-meet operation* for $\mathcal{A}$. A *pseudo-join operation* is a pseudo-meet operation where the roles of I and O are exchanged.

3.18. Corollary *If* $\mathrm{typ}(\alpha,\beta) \in \{3,4\}$, *then I and O are both singletons and $\mathcal{A}|_N$ is already polynomially equivalent to a lattice (in the type 4 case) or to a boolean algebra (in the type 3 case). $\mathcal{A}$ has pseudo-meet and pseudo-join operations.*

3.19. Lemma *Let $\mathcal{A}$ be a finite (α,β)-minimal algebra with* $\mathrm{typ}(\alpha,\beta) = 2$. *Let B be the body of $\mathcal{A}$. There is a polynomial $p(x,y)$ such that B is closed under p and $p(x,b)$ and $p(b,x)$ are permutations for all $b \in B$.*

Proof Let $N \subseteq B$ be a trace. There is a polynomial $p(x,y)$ such that N is closed under p and such that p induces a group operation on $\mathcal{A}|_N/\alpha$. Note that $p|_N$ depends on both variables. In fact, $p|_N(n,x)$ and $p|_N(x,n)$ are permutations of N for any $n \in N$. By the twin lemma, any body twin of these polynomials is also a permutation. □

3.20. Remark 1. In the situation of the lemma $\langle B, p|_{B^2}\rangle$ is a quasigroup.
2. The conclusion of the lemma is false if $\mathrm{typ}(\alpha,\beta) = 1$. If $\langle B, p|_{B^2}\rangle$ is a quasigroup, there is a polynomial $d(x,y,z)$ such that $d|_B$ is a Mal'cev operation on B. Let $N \subseteq B$ be a trace of $\mathcal{A}$. Then $d(N,N,N) \subseteq N$, since $d(x,x,x) = x$. Thus $d|_N$ is a Mal'cev operation for N and then $d|_N/\alpha$ is a Mal'cev operation on $\mathcal{A}|_N/\alpha$. So $\mathcal{A}|_N/\alpha$ is not essentially unary.
3. In fact, the Mal'cev polynomial $d(x,y,z)$ gives the addition operations on all traces, via $d|_N(x,0_N,y)$. The polynomial $d|_N(x,0_N,0_M)$ provides a canonical isomorphism between the traces N and M.

3.21. Lemma *Let $\mathcal{A}$, α, β and B be as in Lemma 3.19. There is a polynomial $d(x,y,z)$ such that*

(a) $d(B,B,B) \subseteq B$,
(b) $d(x,x,x) = x$ *for all* $x \in A$.
(c) $d(x,x,y) = y = d(y,x,x)$ *for all* $x \in B$ *and* $y \in A$.
(d) $d(a,b,x)$, $d(a,x,b)$ *and* $d(x,a,b)$ *are permutations of* A *for all* $a,b \in B$.

Proof Let $d_0(x,y,z)$ be a polynomial of $\mathcal{A}$ whose restriction to B is a Mal'cev polynomial of the quasigroup $\langle B, p|_{B^2}\rangle$. Since $q(x) = d_0(x,x,x) = x$ on B, the unary polynomial q does not collapse β into α, thus it is a permutation. Its inverse is of the form $q^{-1} = q^k$ for some $k > 0$ and so it is also a polynomial. Let

$$d_1(x,y,z) = q^{-1}d_0(x,y,z).$$

We then have $d_1(x,x,x) = x$ for all $x \in A$. Since d_0 is a Mal'cev polynomial on B, the maps $d_0(x,a,a)$ and $d_0(a,a,x)$ are permutations of A for each $a \in B$ (using the minimality of $\mathcal{A}$ again). Thus the maps $d_1(x,a,a)$ and $d_1(a,a,x)$ are also permutations.

Let $u(x,y) = d_1(x,x,y)$. Let $u''(x,y)$ be an iterate of $u(x,y)$ in its second variable so that the polynomial $u'(x,y) = u''(x,u(x,y))$ is idempotent in y. Put

$$d_2(x,y,z) = u''(x, d_1(x,y,z)).$$

Note that $u(x,x) = d_1(x,x,x) = x$ and hence $u'(x,x) = u''(x,x) = x$ for all $x \in A$. Thus $d_2(x,x,x) = x$ for all $x \in A$. Note further that we have

$$d_2(a,a,x) = u''(a, d_1(a,a,x)) = u''(a, u(a,x)) = u'(a,x)$$

for $a \in B$. So this map is an idempotent permutation, thus it is the identity. Hence we have $d_2(a,a,x) = x$ for all $a \in B$.

Now let $v(x,y) = d_2(x,y,y)$. Again choose an iterate $v''(x,y)$ of $v(x,y)$ in its first variable so that the polynomial $v'(x,y) = v''(v(x,y),y)$ is idempotent in its first variable. Set

$$d_3(x,y,z) = v''(d_2(x,y,z),z).$$

As with d_2, we can show that $v(x,x) = v'(x,x) = v''(x,x) = x$, $d_3(x,x,x) = x$ and $d_3(x,a,a) = x$ for all $a \in B$ and $x \in A$. Finally, we calculate

$$d_3(a,a,x) = v''(d_2(a,a,x),x) = v''(x,x) = x.$$

By now, we know that d_3 satisfies (b) and (c).

By the twin lemma, the maps $d_3(a,b,x)$ and $d_3(x,a,b)$ are permutations for $a,b \in B$, since the maps $d_3(a,a,x)$ and $d_3(x,a,a)$ are. Assume we had $a,b \in B$ such that $d_3(a,x,b)$ is not a permutation. Choose an element u such that $\langle u,a\rangle \in \beta \setminus \alpha$. By the minimality of $\mathcal{A}$, $d_3(a,x,b)$ collapses β into α and thus we have

$$b = d_3(a,a,b) \; \alpha \; d_3(a,u,b).$$

Since β is abelian over α, we can conclude

$$d_3(u,a,b) \; \alpha \; d_3(u,u,b) = b.$$

But then $d_3(a,a,b) \; \alpha \; d_3(u,a,b)$, i.e. the map $d_3(x,a,b)$ *does* collapse β into α, contradiction. So (d) is satisfied. (a) is an easy consequence of the fact that d_3 maps (α,β)-traces into (α,β)-traces. □

An operation $d(x,y,z)$ with the properties stated in Lemma 3.21 is called a *pseudo-Mal'cev operation*.

4 Labelled congruence lattices

4.1. Exercise Let U be a neighbourhood of $\mathcal{A}$ and let θ be a congruence of $\mathcal{A}$. If N is a $\theta|_U$-class, then the restriction map from $[0_{\mathcal{A}}, \theta]$ to $\mathrm{Con}(\mathcal{A}|_N)$ is a surjective homomorphism.

4.2. Exercise Let $\alpha_1 \wedge \beta_0 = \alpha_0$ and $\beta_0 \leq \beta_1$.

(a) If β_1 is abelian over α_1 then β_0 is abelian over α_0.
(b) If β_1 is strongly abelian over α_1 then β_0 is strongly abelian over α_0.

4.3. Lemma *Let $\mathcal{A}$ be finite and let $\alpha_i \prec \beta_i$ be congruences of $\mathcal{A}$ for $i = 1, 2$. If $\alpha_1 \wedge \beta_0 = \alpha_0$ and $\alpha_1 \vee \beta_0 = \beta_1$, then $M_{\mathcal{A}}(\alpha_0, \beta_0) = M_{\mathcal{A}}(\alpha_1, \beta_1)$ and* $\mathrm{typ}(\alpha_0, \beta_0) = \mathrm{typ}(\alpha_1, \beta_1)$.

Proof If $U \in M_{\mathcal{A}}(\alpha_0, \beta_0)$ is the range of the idempotent polynomial $e(x)$, then $e(\beta_1|_U) \not\subseteq \alpha_1$, since this would imply $e(\beta_0|_U) \subseteq \alpha_1 \wedge \beta_0 = \alpha_0$. So U contains an (α_1, β_1)-minimal set V.

Conversely, if $V \in M_{\mathcal{A}}(\alpha_1, \beta_1)$ is the range of the idempotent polynomial $e(x)$, then $e(\beta_0|_V) \not\subseteq \alpha_0$, since this would imply $e(\beta_1|_V) = e(\beta_0|_V \vee \alpha_1|_V) \subseteq \mathrm{Cg}(e(\beta_0|_V) \cup e(\alpha_1|_V)) \subseteq \alpha_1$. Thus V contains an (α_0, β_0)-minimal set U.

This shows that the two pairs of congruences have the same minimal sets. Let $U \in M_{\mathcal{A}}(\alpha_0, \beta_0)$.

Case 1. $\mathrm{typ}(\alpha_1, \beta_1) \in \{3, 4, 5\}$. Let N_1 be the unique (α_1, β_1)-trace.
Claim: N_1 contains a unique (α_0, β_0)-trace.

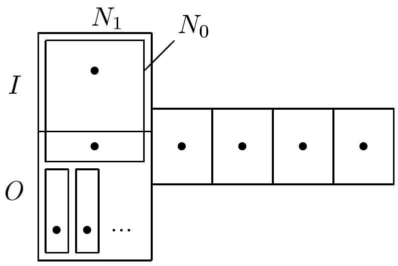

By Exercise 4.1, $\alpha_1|_{N_1} \vee \beta_0|_{N_1} = \beta_1|_{N_1}$. This implies there is a $u \in N_1$ such that $\langle 1, u\rangle \in \beta_0 \setminus \alpha_0$. Therefore N_1 contains at least one (α_0, β_0)-trace N_0, which contains 1. Since $N_1 \setminus N_0 \subseteq O$ and O is an α_1-class, all the other β_0-classes in $N_1 \setminus N_0$ are also α_0-classes, since $\alpha_1|_{N_1} \wedge \beta_0|_{N_1} = \alpha_0|_{N_1}$ by Exercise 4.1. So they belong to the (α_0, β_0)-tail of U.

If $\mathrm{typ}(\alpha_1, \beta_1) \in \{3, 4\}$, then N_1 has two elements. Thus $N_0 = N_1$ and $\alpha_0|_{N_0} = \alpha_1|_{N_1}$. Consequently, $\mathrm{typ}(\alpha_1, \beta_1) = \mathrm{typ}(\alpha_0, \beta_0)$.

If, on the other hand, $\mathrm{typ}(\alpha_0, \beta_0) \in \{3, 4\}$, then the join operation on N_0/α_0 also induces a join operation on N_1/α_1 and so $\mathrm{typ}(\alpha_1, \beta_1) \in \{3, 4\}$.

This shows that, if $\mathrm{typ}(\alpha_1, \beta_1) = 5$, then $\mathrm{typ}(\alpha_0, \beta_0) \notin \{3, 4\}$. It can't be 1 or 2, since the meet operation on N_1/α_1 induces a meet operation on N_0/α_0. So it must also be 5.

Case 2. $\mathrm{typ}(\alpha_1, \beta_1) \in \{1, 2\}$. Then β_1 is abelian or strongly abelian over α_1. By Exercise 4.2, β_0 has the same property over α_0.

So let us assume that $\mathrm{typ}(\alpha_1, \beta_1) = 2$. Let B_1 be the (α_1, β_1)-body and N be an (α_0, β_0)-trace of U. Then $N \subseteq B_1$. Since B_1 has a Mal'cev polynomial, N also has. But then N/α_0 cannot be unary, so must also have type 2.

$\square$

CHAPTER 2

Varieties

1 Subdirectly irreducibles

A *variety* is a class of algebras of the same language defined by a set of equations. When $\mathcal{K}$ is a class of algebras of the same language, we consider the classes

: $\mathsf{H}(\mathcal{K})$ – the class of all homomorphic images of algebras from $\mathcal{K}$.
: $\mathsf{S}(\mathcal{K})$ – the class of all algebras isomorphic to subalgebras of algebras from $\mathcal{K}$.
: $\mathsf{P}(\mathcal{K})$ – the class of all algebras isomorphic to a cartesian product of algebras from $\mathcal{K}$.

The following theorem is fundamental for the theory of varieties.

1.1. Theorem (Birkhoff) *Let $\mathcal{K}$ be a class of algebras of the same language. The following are equivalent:*

(i) $\mathcal{K}$ is closed under taking homomorphic images, subalgebras and cartesian products.
(ii) $\mathcal{K} = \mathsf{HSP}(\mathcal{K})$.
(iii) $\mathcal{K}$ is a variety.

A proof can be found in e. g. [Hod93], [MMT87] or [BS81].

It is a corollary of this theorem, that $V(\mathcal{K}) = \mathsf{HSP}(\mathcal{K})$ is the smallest variety containing $\mathcal{K}$.

Birkhoff's theorem provides a weak sort of structure theorem for the algebras in a variety. A better sort of structure theorem would assert that $V = \mathsf{P}(\mathcal{K})$, i. e. that each member of V is isomorphic to a cartesian product of algebras from some 'well behaved' class $\mathcal{K}$. In some special cases, e. g. if V is locally finite and decidable, V indeed satisfies this sort of structure theorem.

The next best case of structure theorem for a variety V would perhaps assert that $V = \mathsf{SP}(\mathcal{K})$ for some 'nice' class $\mathcal{K}$.

To investigate this possibility further, let us say that an algebra $\mathcal{A}$ is a *subdirect product* of algebras $\mathcal{A}_i$, if $\mathcal{A} \leq \prod_{i\in I} \mathcal{A}_i$ and each projection $\pi_i\colon \mathcal{A} \longrightarrow \mathcal{A}_i$ is surjective.

A homomorphism $\mathcal{A} \longrightarrow \prod_{i\in I} \mathcal{A}_i$ is a *subdirect embedding* if it is injective and its image is a subdirect product of the $\mathcal{A}_i$.

An algebra $\mathcal{A}$ is *subdirectly irreducible* (or just *irreducible* for short) if, whenever $f\colon \mathcal{A} \longrightarrow \prod_{i\in I} \mathcal{A}_i$ is a subdirect embedding, one of the maps $\pi_i \circ f\colon \mathcal{A} \longrightarrow \mathcal{A}_i$ is an isomorphism.

1.2. Theorem (Birkhoff)

(a) The following are equivalent:
(i) $\mathcal{A}$ is irreducible.
(ii) There is a least nonzero congruence $\mu_{\mathcal{A}}$ in $\mathrm{Con}(\mathcal{A})$.

(iii) There are elements $a \neq b$ in $\mathcal{A}$ such that for all $c \neq d$ we have $\langle a, b\rangle \in \mathrm{Cg}_{\mathcal{A}}(\langle c, d\rangle)$.

(b) If $(\theta_i)_{i\in I}$ is a family of nonzero congruences in $\mathrm{Con}(\mathcal{A})$, then the natural map $\mathcal{A} \longrightarrow \prod_{i\in I} \mathcal{A}/\theta_i$ is a subdirect embedding iff $\bigwedge_{i\in I} \theta_i = 0_{\mathcal{A}}$.

(c) Every algebra $\mathcal{A}$ is isomorphic to a subdirect product of irreducible homomorphic images of $\mathcal{A}$.

(d) $V = \mathsf{SP}(V_{SI})$, where V_{SI} is the class of all irreducible members of V.

Proof (b) $\Leftarrow$ It is clear, that $f\colon \mathcal{A} \longrightarrow \prod_{i\in I} \mathcal{A}/\theta_i$ is a homomorphism and that $\pi_i \circ f$ is surjective for each i. It remains to show that f is injective. So let $a \neq b$. Then there is an $i \in I$ such that $\langle a, b\rangle \notin \theta_i$, and thus $a/\theta_i \neq b/\theta_i$. Hence $f(a)$ and $f(b)$ have different i-th coordinates.
(b) $\Rightarrow$ Let $\langle a, b\rangle \in \bigwedge_{i\in I} \theta_i$. Then $a/\theta_i = b/\theta_i$ for all $i \in I$ and thus $f(a) = f(b)$. Since f is an embedding, we must have $a = b$.
(a) It is clear that (ii) and (iii) are equivalent: $\mathrm{Cg}(\langle a, b\rangle) = \mu_{\mathcal{A}}$ is the required congruence and any tuple $\langle a, b\rangle \in \mu_{\mathcal{A}} \setminus 0_{\mathcal{A}}$ will do.
(i)$\Rightarrow$(ii) If there is no such $\mu_{\mathcal{A}}$, then $0_{\mathcal{A}} = \bigwedge_{\theta\neq 0_{\mathcal{A}}} \theta$. By (b), $f\colon \mathcal{A} \longrightarrow \prod_{\theta\neq 0_{\mathcal{A}}} \mathcal{A}/\theta$ is a subdirect embedding. This shows that $\mathcal{A}$ is not irreducible, since we have $\ker(\pi_\theta \circ f) = \theta$ for each θ.
(i)$\Leftarrow$(ii) Let $f\colon \mathcal{A} \longrightarrow \prod_i \mathcal{A}_i$ be a map such that no $\pi_i \circ f$ is an isomorphism. Then $\ker(\pi_i \circ f) \neq 0_{\mathcal{A}}$ for all i and hence $\mu_{\mathcal{A}} \leq \bigwedge_{i\in I} \ker(\pi_i \circ f) = \ker(f)$. Thus f is not an embedding.
(c) Let $a \neq b$. By Zorn's lemma, there is a congruence $\theta_{a,b}$ which is maximal with the property that $\langle a, b\rangle \notin \theta_{a,b}$. By (b), $\mathcal{A} \longrightarrow \prod_{a\neq b} \mathcal{A}/\theta_{a,b}$ is a subdirect embedding, since $\bigwedge_{a\neq b} \theta_{a,b} = 0_{\mathcal{A}}$.

If $\theta \geq \theta_{a,b}$ is any congruence of $\mathcal{A}$, then $\theta \geq \mathrm{Cg}_{\mathcal{A}}(\theta_{a,b} \cup \{\langle a, b\rangle\})$ by the choice of $\theta_{a,b}$. Since the congruence lattice of $\mathcal{A}/\theta_{a,b}$ is isomorphic to the interval $[\theta_{a,b}, 1_{\mathcal{A}}]$, this, together with (a), shows that every algebra $\mathcal{A}/\theta_{a,b}$ is irreducible.
(d) It is clear that $\mathsf{SP}(V_{SI}) \subseteq V$. By (c), the other inclusion is also true, since the algebras $\mathcal{A}/\theta_{a,b}$ are members of V and are irreducible.

□

The least nonzero congruence of an irreducible algebra is often called its *monolith*.

The equivalence of (i) and (ii) in (a) shows that subdirect irreducibility is an internal property of algebras, despite its external definition.

$V = \mathsf{SP}(\mathcal{K})$ for a *set* $\mathcal{K}$ iff up to isomorphism V_{SI} is a set iff there is a cardinal bound κ on the size of the irreducibles in V. When this happens, V is said to be *residually small*, in particular *residually* $< \kappa$. Otherwise, V is *residually large*.

The main focus of the remainder of this text will be on residually small varieties generated by a finite algebra. In the end we will consider abelian varieties and give a criterion for such a variety to be residually small.

We start off with some old results on residual smallness.

1.3. Theorem (Taylor) *If κ is the cardinality of the language of V and if V has an irreducible of size $> 2^\kappa$, then V is residually large.*

A proof of this can be found in [BS81].

1.4. Theorem (McKenzie, Shelah) *If the language of V is countable and V has an infinite irreducible, then V has irreducibles of size κ for every $\aleph_0 \leq \kappa \leq 2^{\aleph_0}$.*

A variety V is *finitely generated* if $V = \mathsf{HSP}(\mathcal{A})$ for some finite algebra $\mathcal{A}$. A class $\mathcal{K}$ is said to be *abelian* if every algebra in it is abelian. $\mathcal{K}$ is *congruence modular* if every algebra in it has a modular congruence lattice.

1.5. Theorem (Quackenbush) *If V is finitely generated and has an infinite irreducible, then it has arbitrarily large finite ones.*

Question: If V is finitely generated and has arbitrarily large finite irreducibles, must it have an infinite one?

The answer is in general **No** (McKenzie), but for 'nice', e. g. congruence modular, varieties, it is **Yes** (Freeze–McKenzie). Other 'nice' varieties for which the answer is **Yes** are varieties omitting the types 1 and 5 (McKenzie–Hobby). We will be concerned for the rest of these notes with showing that abelian varieties are also among the 'nice' ones. This is the content of the following theorem.

1.6. Theorem (Kiss, Kearnes, Valeriote) *If V is a finitely generated residually small abelian variety, then V is residually $< n$ for some $n < \omega$.*

If the language of V is finite, then there is a recursive function f such that $V = \mathsf{HSP}(\mathcal{A})$ is residually small iff V is residually $< f(|\mathcal{A}|)$. There is an algorithm to determine if $\mathcal{A}$ generates a residually small variety.

The proof involves finding bounds on the size of irreducibles in V which have congruence lattices of the following sorts.

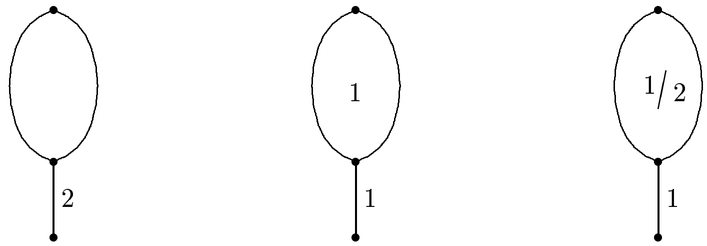

$typ(0_S, \mu_S) = 2$ $\quad$ $typ(0_S, \mu_S) = 1$, $typ(S) = \{1\}$ $\quad$ $typ(0_S, \mu_S) = 1$, $typ(S) = \{1, 2\}$

We will treat the three cases in Sections 3, 4 and 6. But first we need some more facts about the abelian condition.

1.7. Exercise Let V be a locally finite abelian variety and let $\mathcal{K}_1$ be the class of all finite algebras $\mathcal{A}$ in V such that $\text{typ}(\mathcal{A}) = \{1\}$. Then $V_1 = \mathsf{HSP}(\mathcal{K}_1)$ contains all of the strongly abelian algebras in V. We'll see later on that V_1 contains only strongly abelian algebras.

1.8. Exercise Let $\mathcal{K}$ be a class of finite algebras. If $\mathcal{A}$ is a finite member of $\mathsf{HSP}(\mathcal{K})$, then there are $\mathcal{A}_1, \ldots, \mathcal{A}_n$ in $\mathcal{K}$ such that $\mathcal{A}$ is already contained in $\mathsf{HS}(\mathcal{A}_1 \times \cdots \times \mathcal{A}_n)$.

1.9. Exercise Let a and b belong to $\mathcal{A}$ and let $\mathcal{C}$ be the subalgebra of $\mathcal{A}^2$ generated by $0_{\mathcal{A}} \cup \{\langle a, b\rangle\}$. Then $\langle c, d\rangle \in C$ iff $\langle c, d\rangle = \langle p(a), p(b)\rangle$ for some unary polynomial p of $\mathcal{A}$.

The function $q(\langle x, y\rangle)$ is a unary polynomial of $\mathcal{C}$ iff there is a binary polynomial $r(x, y)$ of $\mathcal{A}$ such that $q(\langle x, y\rangle) = \langle r(x, a), r(y, b)\rangle$.

1.10. Exercise (Kearnes) If R is a finite ring and M a subdirectly irreducible R-module, then $|M| \leq |R|$.

1.11. Exercise A locally finite variety with an infinite irreducible member has arbitrarily large finite irreducible members. A variety with language L containing an irreducible of size λ contains irreducibles of every cardinality $|L| \leq \kappa \leq \lambda$.

1.12. Exercise Let $\mathcal{A}$ be an (α, β)-minimal algebra with $\mathrm{typ}(\alpha, \beta) = 2$. Show that the following relation is a congruence on $\mathcal{A}$ (called the *twin congruence*):

$$\theta = \{\langle a, b\rangle \mid \text{ for all binary polynomials } p\text{: } p(x, a) \text{ is a permutation iff } p(x, b) \text{ is}\}$$

Show that the body of $\mathcal{A}$ is a θ-class.

2 Facts about the abelian condition

An algebra $\mathcal{A}$ is abelian iff it satisfies the formulas

$$\forall x x' \forall \bar{y} \bar{y}' \, (t(x, \bar{y}) = t(x, \bar{y}') \to t(x', \bar{y}) = t(x', \bar{y}'))$$

for all terms $t(x, \bar{y})$. Thus abelianness is determined by a set of universal Horn conditions. Since these are preserved by cartesian products and subalgebras, then

$$\mathcal{A} \text{ abelian} \quad \Longrightarrow \quad \mathsf{SP}(\mathcal{A}) \text{ abelian.}$$

But there are finite abelian algebras with nonabelian quotients, so abelianness is in general not preserved by homomorphic images. If $\mathcal{A}$ is abelian and lies in a congruence modular variety, then every homomorphic image of $\mathcal{A}$ is abelian ([HM88]).

2.1. Proposition

(a) Let V be an abelian variety. Then $\mathrm{typ}(V) \subseteq \{1, 2\}$, i. e. $\mathrm{typ}(\alpha, \beta) \in \{1, 2\}$ for every finite algebra $\mathcal{A}$ in V and all congruences $\alpha \prec \beta$ in $\mathrm{Con}(\mathcal{A})$.

(b) If $\mathcal{A}$ is finite and abelian and $\alpha \prec \beta$ in $\mathrm{Con}(\mathcal{A})$ such that $\mathrm{typ}(\alpha, \beta) = 2$, then every (α, β)-minimal set has empty tail, hence has a Mal'cev polynomial.

Proof (a) Since V is closed under homomorphic images then it suffices to show that if $\mathcal{A} \in V$ is finite and β is a congruence of $\mathcal{A}$ which covers $0|_A$ then $\mathrm{typ}(0|_A, \beta)$ is 1 or 2. This follows immediately from Theorem 3.7.
(b) We can assume $\alpha = 0_{\mathcal{A}}$. Now assume that $U \in M_{\mathcal{A}}(0_{\mathcal{A}}, \beta)$ has nonempty tail. Choose a polynomial $d(x, y, z)$ such that $d|_U$ is a pseudo-Mal'cev operation on U. If b is an element of the body and t an element of the tail of U, then

$$d(t, t, t) = t = d(b, b, t)\ .$$

By abelianness, we can replace the last occurrence of t in both terms by b to get

$$d(t, t, b) = d(b, b, b) = b.$$

Consider the polynomial $h(x) = d(x, d(t, d(t, x, b), b), b)$. Let N be the trace which contains b. For all $u, v \in N$, we have $d(t, u, v) \mathrel{\beta} d(t, u, u) = t$, hence $d(t, u, v) = t$. Now this gives

$$h(u) = d(u, d(t, d(t, u, b), b), b) = d(u, d(t, t, b), b) = d(u, b, b) = u,$$

so h is the identity on N. By the minimality of U, h is a permutation on U. But on the other hand,

$$h(t) = d(t, d(t, d(t, t, b), b), b) = d(t, d(t, b, b), b) = d(t, t, b) = b = h(b).$$

This contradiction finishes the proof.

□

An algebra $\mathcal{A}$ is said to be Hamiltonian if every nonempty subuniverse of $\mathcal{A}$ is a congruence class of some congruence on $\mathcal{A}$. A variety is Hamiltonian if each member is.

2.2. Example Every module and in particular every abelian group is Hamiltonian. The 8-element quaternion group is an example of a nonabelian Hamiltonian group.

2.3. Exercise

(a) Show that a nonempty subuniverse B of $\mathcal{A}$ is a block of a congruence if and only if for all polynomials $p(x)$ of $\mathcal{A}$, and b, c from B, $p(b) \in B$ if and only if $p(c) \in B$.

(b) Show that V is Hamiltonian if for each term $t(x, \bar{y})$ there is a ternary term $k_t(x, y, z)$ ('k' for Klukovits) such that

$$V \models k_t(t(x, \bar{y}), x, z) \approx t(z, \bar{y}) \ .$$

(c) Prove the converse to (b). Hint: For $t(x, \bar{y})$ a term, consider the free algebra of V generated by x, z and $\bar{y}$ and apply the Hamiltonian property to the subalgebra generated by x, z and $t(x, \bar{y})$.

(d) Show that every Hamiltonian variety is abelian. The converse is true for locally finite varieties. It is proved in Theorem 2.4.

(e) Use the Hamiltonian property to show that if $\mathcal{A}$ belongs to a locally finite abelian variety and $p(x)$ is a polynomial then there is some term $t(x, y, z)$ such that for all $c \in A$, $p(x) = t(x, c, p(c))$ for all $x \in A$.

Not every abelian variety is Hamiltonian (E. Kiss), but the following is true.

2.4. Theorem (Kiss, Valeriote) *If V is locally finite and abelian, then it is Hamiltonian.*

Proof It will be enough to establish the following claim.

Claim: If $\mathcal{A}$ is finite and $\mathsf{HS}(\mathcal{A}^{|\mathcal{A}|})$ is abelian, then $\mathcal{A}$ is Hamiltonian.

Let $\mathcal{B}$ be a subalgebra of $\mathcal{A}$. B is a block of a congruence iff it is a block of $\beta = \mathrm{Cg}_{\mathcal{A}}(B^2)$.

B is not a block of β iff there is a pair $\langle b_1, b_2 \rangle \in B^2$ and a polynomial $p(x)$ of $\mathcal{A}$ with $p(b_1) \in B$ and $p(b_2) \notin B$. Thus, B fails to be a block of a congruence of $\mathcal{A}$ if and only if there is a polynomial $p(x)$ such that $p(B) \cap B \neq \emptyset$ and $p(B) \not\subseteq B$. Let

$$\mathcal{P} = \{s(B, \ldots, B) \mid s(\bar{x}) \in \mathrm{Pol}(\mathcal{A})\} \ .$$

Note that every set $S \in \mathcal{P}$ is contained in some β-class.

Assume that B is not a block of some congruence of $\mathcal{A}$ and let $T \in \mathcal{P}$ be a $\subseteq$-maximal member of $\mathcal{P}$ with the property that $T \cap B \neq \emptyset$ and $T \not\subseteq B$, say $T = t(B, \ldots, B, \bar{a})$ for some term t and parameters $\bar{a} \in A$. Let $\bar{b} \in B$ and $T' = t(B, \ldots, B, \bar{b})$. Since $\mathcal{B}$ is a subalgebra, $T' \subseteq B$. By abelianness, $|T| = |T'|$. We'll show that our assumptions on $\mathcal{A}$ imply that in fact $T = T'$. This, of course, contradicts the fact that T is not contained in B.

Claim: Let $s(\bar{x}, \bar{y})$ be a polynomial and $\bar{a} \in A$ be such that $S = s(B, \ldots, B, \bar{a})$ is maximal in $\mathcal{P}$. If $a_i \ \beta \ b_i$ for all entries a_i of $\bar{a}$, then $S = s(B, \ldots, B, \bar{b})$.

Just consider the case when $\bar{y}$ is a single variable y. Since β is generated by B^2, it is enough to verify this for a, b of the form $p(u), p(v)$ for some polynomial p and $u, v \in B$. By the maximality of S,

$$S = s(B, \ldots, B, p(u)) \subseteq s(B, \ldots, B, p(B)) \subseteq S\ .$$

But then also

$$s(B, \ldots, B, p(v)) \subseteq s(B, \ldots, B, p(B)) \subseteq S = s(B, \ldots, B, p(u))\ .$$

Since the two sets have the same size, they must be equal.

Claim: Let $s(B, \ldots, B, \bar{a}) \in \mathcal{P}$ be maximal. If $S_{\bar{c}} = s(B, \ldots, B, \bar{c})$ lies in the same β-class as S, then $S = S_{\bar{c}}$.

If not, then there is a $\bar{v} \in B$ such that

$$0_{\bar{a}} = s(\bar{v}, \bar{a}) \neq s(\bar{v}, \bar{c}) = 0_{\bar{c}}\ .$$

Since S and $S_{\bar{c}}$ lie in the same β-class, then $0_{\bar{a}}\ \beta\ 0_{\bar{c}}$. Let $\bar{S}$ be an enumeration of S, say of length k. Then there are elements s_i^j from B with

$$\bar{S} = \begin{pmatrix} s_1 \\ \vdots \\ s_k \end{pmatrix} = \begin{pmatrix} s(s_1^1, \ldots, s_n^1, \bar{a}) \\ \vdots \\ s(s_1^k, \ldots, s_n^k, \bar{a}) \end{pmatrix}$$

Let $\bar{S}_j$ be the j-th column of the matrix (s_i^j). Let $\mathcal{C}$ be the subalgebra of $\mathcal{A}^k$ generated by the vectors $\bar{S}_j$ and the constants. (For $a \in A$, $\hat{a}$ denotes the k-tuple $\langle a, \ldots, a\rangle$.) We have $s(\bar{S}_1, \ldots, \bar{S}_n, \widehat{a_1}, \ldots, \widehat{a_l}) = \bar{S}$ and $s(\bar{S}_1, \ldots, \bar{S}_n, \widehat{c_1}, \ldots, \widehat{c_l}) = \bar{S}_{\bar{c}}$, an enumeration of the elements in $S_{\bar{c}}$. Note that $s(\widehat{v_1}, \ldots, \widehat{v_n}, \widehat{a_1}, \ldots, \widehat{a_l}) = \widehat{0_{\bar{a}}}$ and $s(\widehat{v_1}, \ldots, \widehat{v_n}, \widehat{c_1}, \ldots, \widehat{c_l}) = \widehat{0_{\bar{c}}}$.

Let $\theta = \mathrm{Cg}_{\mathcal{C}}(\langle \widehat{0_{\bar{a}}}, \widehat{0_{\bar{c}}}\rangle)$. If we can prove that $\langle \bar{S}, \bar{S}_{\bar{c}}\rangle \notin \theta$, then we reach the contradiction that $\mathcal{C}/\theta$ is not abelian and are finished.

It will be enough to show, that if $p(x)$ is a polynomial of $\mathcal{C}$ and $p(\widehat{0_{\bar{a}}})$ is an enumeration of S, then so is $p(\widehat{0_{\bar{c}}})$. But if p is such a polynomial, then $p(x) = r(x, \bar{S}_1, \ldots, \bar{S}_n)$ for some polynomial r of $\mathcal{A}$. Since $r(\widehat{0_{\bar{a}}}, \bar{S}_1, \ldots, \bar{S}_n)$ is an enumeration of S, then $r(0_{\bar{a}}, B, \ldots, B) \in \mathcal{P}$ and $S \subseteq r(0_{\bar{a}}, B, \ldots, B)$, so $S = r(0_{\bar{a}}, B, \ldots, B)$. By the first claim, $S = r(0_{\bar{c}}, B, \ldots, B)$, so $r(\widehat{0_{\bar{c}}}, \bar{S}_1, \ldots, \bar{S}_n)$ is another enumeration of S.

To conclude the proof, we apply the claim to the sets T and T'. Since T is maximal and T' lies in the same β-class as T, it follows that $T = T'$. □

3 The case $\mathrm{typ}(0, \mu) = 2$

In this section we will consider finite irreducible algebras with $\mathrm{typ}(0, \mu) = 2$, where μ is the monolith of the algebra. Our goal is to show that in a locally finite abelian variety, there is an upper bound to the size of such algebras.

3.1. Lemma *Let K and L be abelian groups and let $a \in L \setminus \{0\}$ and $R \subseteq \mathrm{Hom}(K, L)$ be finite such that for each $c \in K \setminus \{0\}$ there is an $r \in R$ with $r(c) = a$. Then $|K| \leq (|R| + 1)!$.*

Proof by induction on $|R|$. If $R = \emptyset$, we have nothing to do, since $K = \{0\}$ in this case. So let $|R| > 0$. Choose $r \in R$ so that $|r^{-1}(a)| = |\ker(r)|$ is maximal. Then $|K| \leq 1 + |R| \cdot |r^{-1}(a)|$.

For each $c \in \ker(r)\setminus\{0\}$ there is an $s \in R$ such that $s(c)=a$. Since $r(c)=0\neq a$, we know $s \in R\setminus\{r\}$. By the induction hypothesis, $|\ker(r)| \leq |R|!$. We conclude

$$|K| \leq 1+|R|\cdot|\ker(r)| \leq 1+|R|\cdot|R|! \leq (|R|+1)! \tag{1}$$

□

For V a variety and λ a cardinal, let $F_\lambda(V)$ be the *free algebra on* λ *generators* in V and let $M_\lambda = |F_\lambda(V)|$.

3.2. Theorem *Let* $\mathcal{A}$ *be a finite irreducible abelian algebra and let* V *be the variety generated by* $\mathcal{A}$. *If* $\operatorname{typ}(0_{\mathcal{A}},\mu_{\mathcal{A}})=2$ *Then* $|\mathcal{A}| \leq M_2^{M_2}$.

Proof Let $U \in M_{\mathcal{A}}(0_{\mathcal{A}},\mu)$ and let $e(x)$ be an idempotent polynomial with range U. Note that $\mu|_U$ is the smallest nontrivial congruence in $\operatorname{Con}(\mathcal{A}|_U)$, hence $\mathcal{A}|_U$ is polynomially equivalent to an irreducible module. This follows from the fact that U has an empty tail and so has a Mal'cev polynomial.

Choose elements $0\neq 1$ such that $\langle 0,1\rangle \in \mu|_U$. For all $u \neq v$, we have $\mu \leq \operatorname{Cg}(\langle u,v\rangle)$ and in particular $\langle 0,1\rangle \in \operatorname{Cg}(\langle u,v\rangle)$. From this we get a polynomial $p(x)$ of $\mathcal{A}$ with $ep(u)=0\neq ep(v)$ or $ep(u)\neq 0 = ep(v)$. In both cases, $ep(x)$ has range contained in U and $ep(u)\neq ep(v)$.

Let $p(x)=q(x,\bar{a})$ for some term q and $\bar{a}$ from A. By the abelianness of $\mathcal{A}$, we conclude $eq(u,\bar{0})\neq eq(v,\bar{0})$. So $er(u,0)\neq er(v,0)$ for the binary term $r(x,y)=q(x,y,\ldots,y)$. Thus the map

$$\gamma\colon A \longrightarrow U^{F_2(V)},\ \gamma(a)(t(x,y)) = et(a,0)$$

is injective and hence $|\mathcal{A}| \leq |U|^{M_2}$.

The following argument shows that $|U| \leq M_2!$ and is due to Freese and McKenzie. A more complicated argument due to Kearnes shows that in fact $|U| \leq M_2$.

Claim: $\mathcal{A}|_U$ is polynomially equivalent to an irreducible R-module with $|R| \leq M_2$.

Recall that $\mathcal{A}|_U$ is polynomially equivalent to a module over the ring $R=\{r(x)\in \operatorname{Pol}_1(\mathcal{A}|_U) \mid r(0)=0\}$, where 0 is a fixed but arbitrary element of the body. By Lemma 3.21, there is a pseudo-Mal'cev polynomial $d(x,y,z)$ for U. Let $r(x)\in R$. Then $r(x)=es(x,\bar{a})|_U$ for some term $s(x,\bar{y})$ and some tuple $\bar{a}\in A$. We have $es(0,\bar{a})=r(0)=0$, so

$$d(es(0,\bar{a}),es(0,\bar{a}),es(x,\bar{a})) = es(x,\bar{a}) = d(es(x,\bar{a}),es(x,\bar{a}),es(x,\bar{a})).$$

By abelianness, we conclude that for $x\in U$,

$$d(\underbrace{es(0,\bar{a})}_{=0},es(0,\bar{0}),es(x,\bar{0})) = d(es(x,\bar{a}),es(x,\bar{0}),es(x,\bar{0})) = es(x,\bar{a}) = r(x).$$

If we set $\hat{r}(x,y)=s(x,y,\ldots,y)$, we finally get $r(x)=d(0,e\hat{r}(0,0),e\hat{r}(x,0))$ for $x\in U$. Thus each $r(x)\in R$ is determined by an element of $F_2(V)$. This implies $|R|\leq M_2$.

Claim: If R is a finite ring and M an irreducible R-module then $|M|\leq |R|!$.

Let N be the least nonzero submodule of M, say $N=\langle a\rangle$. Then for any $c\in M\setminus\{0\}$, we have $N\subseteq Rc$. In particular, there is an $r\in R$ with $rc=a$. Now the claim follows from Lemma 3.1. □

3.3. Corollary *Let* V *be locally finite and abelian. If* $\mathcal{S}\in V$ *is a finite irreducible algebra such that* $\operatorname{typ}(0_{\mathcal{S}},\mu_{\mathcal{S}})=2$, *then* $|\mathcal{S}|\leq M_2^{M_2}$.

4 The case $\mathrm{typ}(\mathcal{S}) = \{1\}$

Our next task is to examine irreducible algebras $\mathcal{S}$ where $\mathrm{typ}(0_{\mathcal{S}}, \mu_{\mathcal{S}}) = 1$. This task divides in the (easy) case where $\mathrm{typ}(\mathcal{S}) = \{1\}$ and the (hard) case where $\mathrm{typ}(\mathcal{S}) = \{1, 2\}$. We will treat the easy case in this section.

4.1. Lemma *Let $\mathcal{B}$ be an algebra and θ a strongly abelian congruence of $\mathcal{B}$. If there is a polynomial $t(x, \bar{y})$ of $\mathcal{B}$ and elements $a, b, \bar{c}, \bar{d}$ such that $t(a, \bar{c}) = t(b, \bar{d})$, but $\langle t(a, \bar{c}), t(b, \bar{c})\rangle \in \theta \setminus 0_{\mathcal{B}}$, then $\mathsf{HS}(\mathcal{B}^2)$ contains a nonabelian algebra.*

Proof If possible, choose a, b as in the statement, with $\langle a, b\rangle \in \theta$. Assume that $\mathsf{H}(\mathcal{B})$ is abelian and let

$$\begin{aligned} 0 &= t(a, \bar{c}) = t(b, \bar{d}), \\ 0' &= t(b, \bar{c}), \\ 0'' &= t(a, \bar{d}). \end{aligned}$$

Then $\langle 0, 0'\rangle \in \theta \setminus 0_{\mathcal{B}}$ and since $\mathcal{B}/\theta$ is abelian and $t(a, \bar{c})\ \theta\ t(b, \bar{c})$, we also have $\langle 0, 0''\rangle \in \theta$.

Claim: For all polynomials $p(x)$, $p(0) = p(0')$ if and only if $p(0) = p(0'')$.

The abelianness of $\mathcal{B}$ implies that for any polynomial $p(x)$, $pt(a, \bar{c}) = pt(b, \bar{c})$ iff $pt(a, \bar{d}) = pt(b, \bar{d})$. So we have $p(0) = p(0')$ iff $p(0) = p(0'')$ for all polynomials $p(x)$ of $\mathcal{B}$.

Let $\mathcal{C}$ be the subalgebra of $\mathcal{B}^2$ generated by $0_{\mathcal{B}} \cup \{\langle a, b\rangle\}$ and let γ be the congruence on $\mathcal{C}$ generated by $\{\langle\langle 0, 0'\rangle, \langle 0'', 0\rangle\rangle\}$ (The two pairs $\langle 0, 0'\rangle$ and $\langle 0'', 0\rangle$ are in $\mathcal{C}$ by Exercise 1.9)

Claim: $\mathcal{C}/\gamma$ is nonabelian.

Since

$$\begin{aligned} t(\langle a, b\rangle, \langle c_1, c_1\rangle, \ldots, \langle c_n, c_n\rangle) &= \langle 0, 0'\rangle\ \gamma\ \langle 0'', 0\rangle = t(\langle a, b\rangle, \langle d_1, d_1\rangle, \ldots, \langle d_n, d_n\rangle), \\ t(\langle b, b\rangle, \langle c_1, c_1\rangle, \ldots, \langle c_n, c_n\rangle) &= \langle 0', 0'\rangle, \\ t(\langle b, b\rangle, \langle d_1, d_1\rangle, \ldots, \langle d_n, d_n\rangle) &= \langle 0, 0\rangle, \end{aligned}$$

the above claim follows from

Claim: $\langle\langle 0, 0\rangle, \langle 0', 0'\rangle\rangle \notin \gamma$.

If $\langle 0, 0\rangle\ \gamma\ \langle 0', 0'\rangle$, there is a polynomial $g(\langle x, y\rangle)$ of $\mathcal{C}$ with $g(\langle 0, 0'\rangle) \neq g(\langle 0'', 0\rangle)$ and $\langle 0, 0\rangle$ equal to one of them.

Without loss of generality, suppose that $g(\langle 0, 0'\rangle) = \langle 0, 0\rangle$. There is a polynomial $s(x, y)$ of $\mathcal{B}$ such that $g(\langle x, y\rangle) = \langle s(x, a), s(y, b)\rangle$ and so we have $0 = s(0, a) = s(0', b)$. By showing that $0 = s(0'', a) = s(0, b)$ we will obtain the contradiction $g(\langle 0'', 0\rangle) = \langle 0, 0\rangle = g(\langle 0, 0'\rangle)$.

Case 1. $\langle a, b\rangle \in \theta$. Since θ is strongly abelian, we have $0 = s(0, a) = s(0, b)$. The abelianness of $\mathcal{B}$ then implies $s(x, a) = s(x, b)$ for all x. Hence we have $s(0', a) = s(0', b) = 0 = s(0, a)$. By the first claim, this implies that $s(0'', a) = 0 = s(0, a)$.

Case 2. $\langle a, b\rangle \notin \theta$. Since $0\ \theta\ 0'$ and $s(0, a) = s(0', b)$, then $\langle s(0, a), s(0', a)\rangle \in \theta \setminus 0_{\mathcal{B}}$ is impossible. Otherwise, we could have selected 0 and $0'$ in place of a and b and s in place of t, putting us in Case 1. Thus, $s(0, a) = s(0', a)$.

By the first claim, $s(0,a) = 0 = s(0'',a)$ and by abelianness, $s(0',a) = s(0',b)$ implies $s(0,b) = s(0,a) = 0$.

□

4.2. Theorem *Let V be a locally finite abelian variety. Let V_1 be the subvariety generated by all finite members of V with type set $\{1\}$. Then $V_1 = \{\mathcal{A} \in V \mid \mathcal{A}$ strongly abelian$\}$.*

Proof If $\mathcal{A}$ is a strongly abelian member of V then, being locally finite, it lies in the variety generated by its finite subalgebras. All of these subalgebras are members of V_1 and so $\mathcal{A}$ lies in V_1 as well.

By Theorem 7.2 and Corollary 7.6 of [HM88] it follows that $\mathrm{typ}(V_1) = \{1\}$. If V_1 is not strongly abelian, then V_1 contains a finite algebra $\mathcal{A}$ which is not strongly abelian, since V is locally finite. Let us assume that $\mathcal{A}$ is as small as possible.

Claim: $\mathcal{A}$ is irreducible.

Assume not. By Theorem 1.2, $\mathcal{A}$ is isomorphic to a subdirect product of irreducible proper quotients. Since these quotients are all smaller than $\mathcal{A}$ and also in V_1, they are strongly abelian by the minimal choice of $\mathcal{A}$. Since strong abelianness is a universal Horn property, it is preserved by SP. Thus $\mathcal{A}$ is strongly abelian, contradiction.

By the minimality of $\mathcal{A}$, $\mathcal{A}/\mu$ is strongly abelian, where μ is the monolith of $\mathcal{A}$. Furthermore, since the typeset of $\mathcal{A}$ is $\{1\}$, then $\mathrm{typ}(0|_A, \mu) = 1$. Since $\mathcal{A}$ is not strongly abelian then there are a, b, $\bar{c}$, and $\bar{d}$ and a polynomial t with $t(a,\bar{c}) = t(b,\bar{d})$ and $t(a,\bar{c}) \neq t(b,\bar{c})$. We at least have that $\langle t(a,\bar{c}), t(b,\bar{c})\rangle \in \mu$, since $\mathcal{A}/\mu$ is strongly abelian. Then by the Lemma, we conclude that $\mathsf{HS}(\mathcal{A}^2)$ is not abelian, contradicting our assumption that V is abelian. □

4.3. Corollary *If $\mathcal{A}$ is a finite member of an abelian variety and* $\mathrm{typ}(\mathcal{A}) = \{1\}$ *then $\mathcal{A}$ is strongly abelian.*

4.4. Theorem *Let V be a locally finite abelian variety and let $\mathcal{S}$ be a finite irreducible member of V with a strongly abelian monolith μ. If ρ is a strongly abelian congruence of $\mathcal{S}$ then each ρ-class has size bounded by 2^{M_3}.*

Proof Let C be a ρ-class and let $\langle a,b\rangle \in \mu$ with $a \neq b$. For each $c \in C$, let T_c be the set of all term operations $t(x,y,z)$ on $\mathcal{S}$ with $a = t(c,c,a)$. Since there are at most M_3 ternary term operations on $\mathcal{S}$ then there are at most 2^{M_3} sets of the form T_c, $c \in C$. We will establish the bound by showing that $T_c \neq T_d$ when c and d are distinct members of C.

Let c, $d \in C$ with $c \neq d$. Since $\langle a,b\rangle$ is in the congruence generated by $\langle c,d\rangle$ then there is a polynomial $p(x)$ of $\mathcal{S}$ with $p(c) \neq p(d)$ and $a \in \{p(c), p(d)\}$. Without loss of generality, assume that $p(c) = a$. It follows from exercise 2.3 that there is a term $t(x,y,z)$ such that $p(x) = t(x,c,p(c)) = t(x,c,a)$. In fact, we can let $t(x,y,z)$ be $k_s(z,y,x)$, where k_s is a term as in part b) of exercise 2.3 and $s(x,\bar{u}) = p(x)$ for some term s and tuple $\bar{u}$. Since $p(c) = a$ then $t(c,c,a) = a$ and so $t \in T_c$.

To conclude the proof, we need only show that $t \notin T_d$. If $a = t(d,d,a) = t(c,c,a)$ then since ρ is strongly abelian, it follows that $t(d,d,a) = t(d,c,a)$ and so $a = t(d,c,a) = p(d) \neq a$, a contradiction. Thus, $t(d,d,a) \neq a$ and so $t \notin T_d$. □

4.5. Corollary (Shapiro) *Let V be a locally finite abelian variety. If $\mathcal{S}$ is a finite irreducible member of V such that* $\mathrm{typ}(\mathcal{S}) = \{1\}$*, then* $|\mathcal{S}| \leq 2^{M_3}$.

4.6. Corollary *If V is abelian and generated by the finite algebra $\mathcal{A}$ and $\mathcal{S}$ is an irreducible member of V with* $\operatorname{typ}(\mathcal{S}) = \{1\}$, *then*

$$|\mathcal{S}| \leq 2^{|A|^{|A|^3}} .$$

5 The residually large configuration

5.1. Example (A residually large abelian variety) Let $V = \mathsf{HSP}(\mathcal{A})$, where

$$\mathcal{A} = \langle \{0, 1, 1'\}, +, ', 0 \rangle$$

and the nonconstant operations are given by

$+$	0	1	$1'$
0	0	1	1
1	1	0	0
$1'$	1	0	0

	$'$
0	0
1	$1'$
$1'$	$1'$

Claim: V is abelian.

To prove abelianness, we find normal forms for the terms of V. There are a number of 'obvious' identities:

$$x + x = 0 \qquad x + y = y + x \qquad x + (y + z) = (x + y) + z$$
$$x + y' = x + y \qquad x'' = x' \qquad (x + y)' + 0 = x + y$$

From these equations it is easy to prove by induction on the complexity of the terms that each term $t(x_1, \ldots, x_n)$ which depends on all of its variables is equal in V to one of the following:

$$0, \quad x_1, \quad x_1 + 0, \quad x_1', \quad x_1 + \cdots + x_n, \quad (x_1 + \cdots + x_n)' .$$

Let now $\mathcal{B}$ be an algebra in V and $t(x_1, \ldots, x_n)$ a term depending on all variables and assume that $t(a, \bar{c}) = t(a, \bar{d})$. We have to show that $t(b, \bar{c}) = t(b, \bar{d})$.

In the case $n \leq 1$ there is nothing to do, so we can assume that t is of the form $x_1 + \cdots + x_n$ or $(x_1 + \cdots + x_n)'$. In the first case, we derive the following sequence of equations:

$$\begin{aligned}
t(a, \bar{c}) &= t(a, \bar{d}) \\
a + c_2 + \cdots + c_n &= a + d_2 + \cdots + d_n \\
a + a + c_2 + \cdots + c_n &= a + a + d_2 + \cdots + d_n \\
0 + c_2 + \cdots + c_n &= 0 + d_2 + \cdots + d_n \\
b + 0 + c_2 + \cdots + c_n &= b + 0 + d_2 + \cdots + d_n \\
b + c_2 + \cdots + c_n &= b + d_2 + \cdots + d_n \\
t(b, \bar{c}) &= t(b, \bar{d}) .
\end{aligned}$$

The other case is similar.

We leave it as an exercise to prove that $\mu_{\mathcal{A}} = 0_{\mathcal{A}} \cup \{\langle 1, 1' \rangle, \langle 1', 1 \rangle\}$ is the only nontrivial congruence on $\mathcal{A}$ and that

$$\begin{aligned}
&\operatorname{typ}(0_{\mathcal{A}}, \mu_{\mathcal{A}}) = 1, \\
&\operatorname{typ}(\mu_{\mathcal{A}}, 1_{\mathcal{A}}) = 2, \\
&M_{\mathcal{A}}(0_{\mathcal{A}}, \mu_{\mathcal{A}}) = \{A\}, \\
&M_{\mathcal{A}}(\mu_{\mathcal{A}}, 1_{\mathcal{A}}) = \{\{0, 1\}, \{0, 1'\}\}.
\end{aligned}$$

Claim: V is residually large.

Let κ be any cardinal. Consider $\mathcal{A}^\kappa$. Let $\mathcal{B}$ be $\{0,1\}^\kappa \cup \{0,1'\}^\kappa$, a proper subalgebra of $\mathcal{A}^\kappa$. Let θ be the congruence of $\mathcal{B}$ generated by

$$\{\langle f, f'\rangle \mid f \in \{0,1\}^\kappa \setminus \{\hat{1}\}\} ,$$

where for $a \in A$, $\hat{a} = \langle a, a, \ldots\rangle$.

Since for any nontrivial polynomial p of $\mathcal{B}$, we have $p(f) = p(f')$, it follows that $\{f, g\} = \{h, h'\}$ for some $h \in \{0,1\}^\kappa \setminus \{\hat{1}\}$, whenever $f\ \theta\ g$ and $f \neq g$. From this it follows that $\langle \hat{1}, \hat{1}'\rangle \notin \theta$ and that $|\mathcal{B}/\theta| \geq 2^\kappa$. It remains to show that $\mathcal{B}/\theta$ is irreducible. For this it is enough to show that

$$\langle \hat{1}, \hat{1}'\rangle \in \theta \vee \mathrm{Cg}(\langle f, g\rangle), \quad \text{for all } \langle f, g\rangle \notin \theta.$$

So let f, $g \in B$ with $\langle f, g\rangle \notin \theta$ and let $\alpha = \theta \vee \mathrm{Cg}(\langle f, g\rangle)$. We have

$$\hat{1} = (\hat{1} + f + f)\ \alpha\ (\hat{1} + f + g) \overset{(*)}{\theta} (\hat{1} + f + g)'\ \alpha\ (\hat{1} + f + f)' = \hat{1}' ,$$

where at (*) we used the fact that $f + g \neq \hat{0}$ and so $\hat{1} + f + g \neq \hat{1}$.

The following lemma unravels the general construction behind this example.

5.2. Lemma (The residually large configuration) *Let $\mathcal{A}$ be an algebra with*

1. *subsets $N = \{0,1\}$ and $N' = \{0,1'\}$, where $1 \neq 1'$,*
2. *a 'subtraction polynomial', i. e. a polynomial $s(x,y)$ of $\mathcal{A}$ such that* $\begin{array}{c|cc} s & 0 & 1 \\ \hline 0 & 0 & * \\ 1 & 1 & 0 \end{array}$,
3. *polynomials $e(x)$ and $e'(x)$ such that*

$$\begin{aligned} e(0) &= e'(0) = 0, \\ e(1) &= e(1') = 1, \\ e'(1) &= e'(1') = 1', \end{aligned}$$

such that for all polynomials $p(x)$ of $\mathcal{A}$ the following holds:

$$\textit{if } p(0) = p(1) \textit{ or } p(0) = p(1'),\textit{ then } p(1) = p(1') . \tag{*}$$

Then $\mathsf{HSP}(\mathcal{A})$ is residually large.

In Example 5.1, we can choose $e(x) = x + 0$ and $e'(x) = x'$.

Proof Let κ be a cardinal. Let $\mathcal{B}$ be the subalgebra of $\mathcal{A}^\kappa$ generated by the set $\{0,1\}^\kappa \cup \{\hat{a} \mid a \in A\}$ and γ the congruence on $\mathcal{B}$ generated by

$$\{\langle f, e'(f)\rangle \mid f \in \{0,1\}^\kappa \setminus \{\hat{1}\}\} .$$

Claim: $\langle \hat{1}, \hat{1}'\rangle \notin \gamma$.

If $\langle \hat{1}, \hat{1}'\rangle \in \gamma$, then there is a polynomial $p(x)$ of $\mathcal{B}$ and an $f \in \{0,1\}^\kappa \setminus \{\hat{1}\}$ such that $p(f) \neq p(e'(f))$ and $\hat{1} \in \{p(f), p(e'(f))\}$.

Since $\mathcal{B}$ is generated by $\{0,1\}^\kappa$ and the constants, there is a polynomial $q(x, \bar{y})$ of $\mathcal{A}$ and $\bar{g}$ a tuple from $\{0,1\}^\kappa$ such that $p(x) = q(x, \bar{g})$. Since $p(f) \neq p(e'(f))$, there is an $i < \kappa$ with $p(f)(i) \neq p(e'(f))(i)$. Let us assume, without loss of generality, that

$$1 = q(f(i), \bar{g}(i)) \neq q(e'(f(i)), \bar{g}(i)) .$$

This implies in particular $q(f,\bar g)=\hat 1$. As $f(i)\neq e'(f(i))$ then we must have $f(i)=1$ and $e'(f(i))=1'$, implying that

$$1=q(1,\bar g(i))\neq q(1',\bar g(i))\ .$$

Since $f\neq\hat 1$, there is a $j<\kappa$ such that $f(j)=0$. Then $q(0,\bar g(j))=1$ and we can conclude that

$$q(0,\bar g(j))=1=q(1,\bar g(i))\neq q(1',\bar g(i))\ .$$

By rearranging the variables of q, we may assume that $\bar g(i)$ and $\bar g(j)$ are of the form

$$\begin{aligned}\bar g(i)&=\overbrace{0\dots0}^{I_1}\overbrace{1\dots1}^{I_2}\overbrace{0\dots0}^{I_3}\overbrace{1\dots1}^{I_4},\\ \bar g(j)&=1\dots1\,0\dots0\,0\dots0\,1\dots1.\end{aligned}$$

Let $c(x)=s(1,x)$ and let

$$r(x)=q(x,\underbrace{e(x),\dots,e(x)}_{I_1},\underbrace{c(e(x)),\dots,c(e(x))}_{I_2},\underbrace{0,\dots,0}_{I_3},\underbrace{1,\dots,1}_{I_4})\ .$$

Then

$$\begin{aligned}r(0)&=q(0,\bar g(j))=1,\\ r(1)&=q(1,\bar g(i))=1,\\ r(1')&=q(1',\bar g(i))\neq 1,\end{aligned}$$

contradicting (*). So $\hat 1\ \gamma\ \hat 1'$, and in fact $\hat 1/\gamma=\{\hat 1\}$.

Let θ be a congruence above γ maximal with $\langle\hat 1,\hat 1'\rangle\notin\theta$. Then $\mathcal B/\theta$ is irreducible. It remains to show that $|\mathcal B/\theta|\geq\kappa$.

For $\lambda\leq\kappa$ let

$$a_\lambda(i)=\begin{cases}1 & \text{if } i<\lambda\\ 0 & \text{otherwise.}\end{cases}$$

Now let $\lambda<\delta<\kappa$ and suppose that $a_\lambda\ \theta\ a_\delta$. Then

$$\hat 1=s(\hat 1,s(a_\delta,a_\delta))\ \theta\ s(\hat 1,s(a_\delta,a_\lambda))\ \gamma\ e'(s(\hat 1,s(a_\delta,a_\lambda)))\ \theta\ e'(s(\hat 1,s(a_\delta,a_\delta)))=e'(\hat 1)=\hat 1'\ ,$$

contradiction. Thus the a_λ for $\lambda\leq\kappa$ give rise to κ many distinct elements of $\mathcal B/\theta$. □

5.3. Example (A residually small abelian variety) Let $V=\mathsf{HSP}(\mathcal A)$, where

$$\mathcal A=\langle\{0,0',1,1'\},+,f,g,h,0,0',1,1'\rangle$$

and the nonconstant operations are given by

$+$	0	$0'$	1	$1'$
0	0	0	1	1
$0'$	0	0	1	1
1	1	1	0	0
$1'$	1	1	0	0

	f	g	h
0	0	1	$0'$
$0'$	0	1	$0'$
1	0	$1'$	$1'$
$1'$	1	$1'$	$1'$

It is left as an exercise to prove that V is abelian and that $\mathcal{A}$ has exactly one nontrivial congruence $\mu_{\mathcal{A}}$ with

$$\begin{aligned} \operatorname{typ}(0_{\mathcal{A}},\mu_{\mathcal{A}}) &= 1, \\ \operatorname{typ}(\mu_{\mathcal{A}},1_{\mathcal{A}}) &= 2, \\ M_{\mathcal{A}}(0_{\mathcal{A}},\mu_{\mathcal{A}}) &= \{A\}, \\ M_{\mathcal{A}}(\mu_{\mathcal{A}},1_{\mathcal{A}}) &= \{\{0,1\},\{1,1'\},\{0,1'\}\} \end{aligned}$$

Claim: V is residually < 6.

Let $\mathcal{S}$ be an irreducible in V. The equations of $\mathcal{A}$ force the range of $+$ to be a 2-group G_0. The ranges G_1 of g and G_2 of h are polynomially isomorphic to G_0 and so are 2-groups as well. Also, every nontrivial polynomial of S has range contained in G_i for some i.

Let $a \in G_0 \cap G_1$. Then $a = g(b)$ for some $b \in S$ and $a + 0 = a$ and so $g(b) + 0 = a$. But $g(x) + 0 = 1$ is an equation of $\mathcal{A}$ and so $a = 1$. We have just shown that $G_0 \cap G_1 = \{1\}$. In a similar way one can show that $G_1 \cap G_2 = \{1'\}$. This gives us the following picture of $\mathcal{S}$:

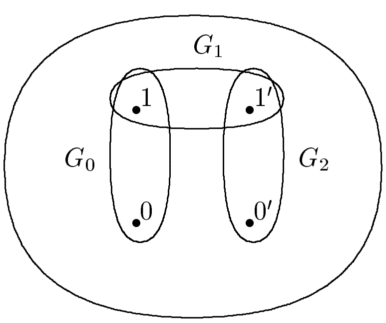

Claim: Disjoint congruences on $\mathcal{S}|_{G_0}$ give rise to disjoint congruences on $\mathcal{S}$.

Let θ be a congruence on $\mathcal{S}|_{G_0}$ and let $\hat{\theta}$ be the congruence on $\mathcal{S}$ generated by θ. Since G_0 is the range of an idempotent polynomial then $\theta \cap \hat{\theta} = \theta$. We claim that if $\langle 0,0'\rangle \in \hat{\theta}$ then $\langle 0,1\rangle \in \theta$. This follows, with some effort, from the facts, noted earlier, that every nontrivial polynomial has range contained in G_i for some i and that the intersection of distinct G_i's have size at most 1.

Now, let θ_0 and θ_1 be congruences of $\mathcal{S}|_{G_0}$ with $\langle x,y\rangle \in \hat{\theta}_0 \cap \hat{\theta}_1$ and with $x \neq y$. If $\{x,y\} = \{0,0'\}$ then we have that $\langle 0,1\rangle \in \theta_0 \cap \theta_1$. Otherwise, x and y must be in the ranges of some nontrivial polynomials and so must belong to $G_0 \cup G_1 \cup G_2$. If they both lie in G_i for some i, then as G_i is polynomially isomorphic to G_0 there are x', $y' \in G_0$ with $\{x,y\}$ polynomially isomorphic to $\{x',y'\}$. It follows that $\langle x',y'\rangle \in \theta_0 \cap \theta_1$.

By symmetry, the other cases to consider are when $x \in G_0$ and $y \in G_1$ or $x \in G_0$ and $y \in G_2$. In both cases, we can use various polynomial projections onto G_0 to find a pair of distinct elements of G_0 which lie in $\theta_0 \cap \theta_1$. We can conclude that if $\theta_0 \cap \theta_1 \neq 0_{G_0}$ then $\hat{\theta}_0 \cap \hat{\theta}_1 \neq 0_S$.

Thus, the irreducibility of $\mathcal{S}$ implies that $\mathcal{S}|_{G_0}$ is, up to polynomial equivalence, an irreducible 2-group and so has size 2.

Our picture now is the following:

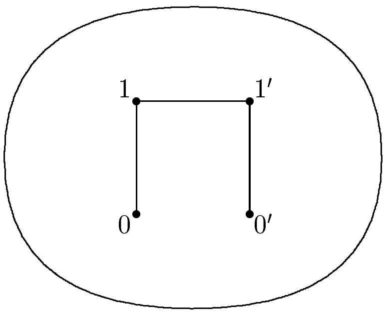

Note that if a and b are elements which behave identically with respect to the basic operations of $\mathcal{S}$ then the only two distinct elements identified by the congruence generated by $\langle a,b\rangle$ are a and b. The elements 0 and $0'$ have this property. If S has more than 5 elements then, by considering the various possibilities, it can be shown that there must be another pair of elements besides $\{0, 0'\}$ which behave identically. From this it follows that, as $\mathcal{S}$ is irreducible, S can contain no more than 5 elements.

5.4. Exercise This exercise continues the discussion of Example 5.3.

(a) Show that V is abelian.
(b) Prove the claims made about congruence lattice and the minimal sets of $\mathcal{A}$.
(c) Find an irreducible in V which contains exactly 5 elements. Hint: Look for a quotient of a subalgebra of $\mathcal{A}^3$.

5.5. Exercise Let $\mathcal{A}$ be a finite algebra and $\alpha \succ \beta$ and $\mu \succ \nu$ in the congruence lattice of $\mathcal{A}$. Show that if some subset U is a minimal set for both (α, β) and (μ, ν) then the set of (α, β)-minimal sets coincides with the set of (μ, ν)-minimal sets.

5.6. Exercise Call two covers α, β of a congruence ρ equivalent if the pairs (ρ, α) and (ρ, β) have the same minimal sets. Show that if $\mathcal{A}$ is finite, $0_{\mathcal{A}} \prec \alpha$ and J is the set of covers of $0_{\mathcal{A}}$ equivalent to α then any $(0, \alpha)$-minimal set U is also a $(0_{\mathcal{A}}, \gamma)$-minimal set, where γ is the join of the congruences in J.

If the type of $(0, \alpha)$ is 2 show that any $\gamma|_U$ class V contained in the body of U is polynomially equivalent to a vector space having $|J|$ one dimensional subspaces. Also, show that if the $(0, \alpha)$-traces are vector spaces over the finite field F then $\mathcal{A}|_V$ is polynomially equivalent to a vector space over F as well.

5.7. Exercise Let V be a locally finite abelian variety and suppose that ρ is a strongly abelian congruence of the finite algebra $\mathcal{A}$ from V. If $t(x_1, \ldots, x_n)$ is a term of V and $C_1, \ldots, C_n$ are ρ-blocks then the number of variables which $t|_{C_1 \times \cdots \times C_n}$ can depend on is at most the product of the size of the 2-generated free algebra in V with the size of the quotient $\mathcal{A}/\rho$.

5.8. Exercise Let $\mathcal{A}$ be a finite algebra. Define two congruences α and β to be strongly solvably related if $\mathrm{typ}(\{\alpha \wedge \beta, \alpha \vee \beta\}) = \{1\}$. Show that the strong solvability relation is a congruence of the congruence lattice of $\mathcal{A}$.

5.9. Exercise Let $\alpha \prec \beta$ in the congruence lattice of the finite algebra $\mathcal{A}$ with $\mathrm{typ}(\alpha, \beta) = 2$. Show that if α is strongly solvable (i.e., $\mathrm{typ}([0, \alpha]) = \{1\}$) and U is an (α, β)-minimal set with body B then $\alpha|_B = 0_B$.

5.10. Exercise Let $\mathcal{A}$ be a finite algebra in an abelian variety and let α be a strongly solvable congruence of $\mathcal{A}$. Prove that α is strongly abelian.

Hint: Assume that we are dealing with a minimal counterexample and then use Lemma 4.1.

6 The case $\mathrm{typ}(\mathcal{S}) = \{1, 2\}$

6.1. Proposition *Let $\mathcal{B}$ be a finite algebra and let α and β be congruences of $\mathcal{B}$.*

(a) $\mathrm{typ}([0_{\mathcal{B}}, \alpha]) = \{1\}$ *iff there is a chain* $0_{\mathcal{B}} = \alpha_0 \prec \alpha_1 \prec \ldots \prec \alpha_n = \alpha$ *such that* $\mathrm{typ}(\alpha_i, \alpha_{i+1}) = 1$ *for all* $i < n$

(If this is the case, α is called strongly solvable.*)*

(b) If $\mathrm{typ}([0_{\mathcal{B}}, \alpha]) = \{1\}$ *and* $\mathrm{typ}([0_{\mathcal{B}}, \beta]) = \{1\}$, *then* $\mathrm{typ}([0_{\mathcal{B}}, \alpha \vee \beta]) = \{1\}$.

Proof (a) $\Rightarrow$ obvious.

$\Leftarrow$ Suppose there is a pair $\mu \prec \nu \leq \alpha$ with $\mathrm{typ}(\mu,\nu) \neq 1$. Let N be a (μ,ν)-trace and $\langle a,b\rangle \in \nu|_N \setminus \mu$. Since $\mathrm{typ}(\mu,\nu) \neq 1$, there is a polynomial $p(x,y)$ under which N is closed and, after possibly interchanging a and b, such that

p	a	b
a	$*$	a
b	a	b

Choose an i such that $\langle a,b\rangle \in \alpha_{i+1} \setminus \alpha_i$. Since α_{i+1} is strongly abelian over α_i and $\langle p(a,b), p(b,a)\rangle \in \alpha_i$, then $\langle p(a,b), p(b,b)\rangle \in \alpha_i$, contradiction.

(b) Since we have already established (a), it suffices to show that $\mathrm{typ}([\beta, \alpha\vee\beta]) = \{1\}$.

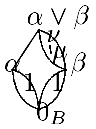

Let $0_{\mathcal{B}} = \alpha_0 \prec \alpha_1 \prec \ldots \prec \alpha_n = \alpha$ be chosen according to (a). Suppose that there is a pair $\mu \prec \nu$ between β and $\alpha \vee \beta$ such that $\mathrm{typ}(\mu,\nu) \neq 1$.

Let U be a (μ,ν)-minimal set and $\langle a,b\rangle \in \nu|_U \setminus \mu|_U$. Since $\nu \leq \alpha \vee \mu$, then $\nu|_U \leq \alpha|_U \vee \mu|_U$ and so $\langle a,b\rangle \in \alpha|_U \vee \mu|_U$. So there is a $c \in U$ with $\langle a,c\rangle \in \alpha \setminus 0_B$ (after possibly adjusting a). Since $\mathrm{typ}(\mu,\nu) \neq 1$, then there is a polynomial $p(x,y)$ with (*) $\begin{array}{c|cc} p & c & a \\ \hline c & * & c \\ a & c & a \end{array}$, after interchanging a and c if necessary. p is derived from either a pseudo-meet, a pseudo-join or a pseudo-Mal'cev operation on U. As in (a), this implies $\mathrm{typ}(\alpha_i, \alpha_{i+1}) \neq 1$ for some i.

□

In [HM88], the configuration (*) which appeared in the last proof is called a *1-snag*.

6.2. Corollary *There is a largest congruence $\rho_{\mathcal{B}}$ of $\mathcal{B}$, such that* $\mathrm{typ}([0_{\mathcal{B}}, \rho_{\mathcal{B}}]) = \{1\}$. *If $\alpha \succ \rho_{\mathcal{B}}$, then* $\mathrm{typ}(\rho_{\mathcal{B}}, \alpha) \neq 1$.

The congruence $\rho_{\mathcal{B}}$ is called the strongly solvable radical *of $\mathcal{B}$.*

From now on we will assume the following situation unless otherwise stated: $\mathcal{A}$ is a finite algebra, $V = \mathsf{HSP}(\mathcal{A})$ is abelian and $\mathcal{S}$ is a finite irreducible with monolith μ and strongly solvable radical ρ such that $\mathrm{typ}(0_{\mathcal{S}}, \mu) = 1$ and $\mathrm{typ}(\mathcal{S}) = \{1,2\}$. From the corollary, we get that every cover of ρ is of type 2 and that $\mathrm{Con}(\mathcal{S})$ looks like:

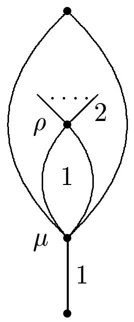

6.3. Proposition *Let $\mathcal{B} \in V$.*

(a) If $\mathcal{B}$ is finite and β is a strongly solvable congruence of $\mathcal{B}$, then β is strongly abelian.

(b) If $p(\bar{x})$ is a polynomial of $\mathcal{B}$, then there is a term $t(\bar{x}, y, z)$ and elements $a, b \in B$ such that $p(\bar{x}) = t(\bar{x}, a, b)$ for all $\bar{x} \in B$.

Proof (a) See Exercise 5.10.

(b) See Exercise 2.3.

□

6.4. Theorem *Each ρ-class has at most 2^{M_3} elements and thus $|\mathcal{S}| \leq |\mathcal{S}/\rho| 2^{M_3}$.*

Proof This is just Theorem 4.4. □

To finish the proof of Theorem 1.6, we have to bound $|\mathcal{S}/\rho|$. To achieve this, we have to analyze the covers of ρ a bit closer.

Let N be the number of covers of ρ.

6.5. Proposition *If $\mathcal{S}/\rho$ is big, then ρ has a lot of covers. In fact, $|\mathcal{S}/\rho| \leq \left(M_2^{M_2}\right)^N$.*

Proof Let I be the set of covers of ρ. For $\alpha \in I$ let α' be a congruence of $\mathcal{S}$ which is maximal with the property that $\alpha \wedge \alpha' = \rho$.

Claim: $\mathcal{S}/\alpha'$ is an irreducible with monolith of type 2.

Since $\alpha \succ \rho$, then $\alpha \vee \alpha' \succ \alpha'$ and so, by Lemma 4.3, $\operatorname{typ}(\alpha', \alpha \vee \alpha') = 2$ since $\operatorname{typ}(\rho, \alpha) = 2$. By the maximality of α', every congruence lying strictly above α' also lies above α and so $\mathcal{S}/\alpha'$ is irreducible.

By Theorem 3.2, $|\mathcal{S}/\alpha'| \leq M_2^{M_2}$. It is easy to see that $\bigwedge_{\alpha \in I} \alpha' = \rho$ and thus $\mathcal{S}/\rho$ embeds naturally into $\prod_{\alpha \in I} \mathcal{S}/\alpha'$. This implies the inequality. □

6.6. Corollary $|S| \leq \left(M_2^{M_2}\right)^N 2^{M_3}$.

We call two covers θ_0 and θ_1 of ρ equivalent if the (ρ, θ_0) and (ρ, θ_1) minimal sets are the same. It follows from the properties of minimal sets given in Theorem 2.8 that this is indeed an equivalence relation on the covers of ρ.

6.7. Lemma *Let $\mathcal{B} \in V$ be finite and $\alpha \prec \beta$ be a pair of congruences of $\mathcal{B}$. Then there exist a term $t(x, y)$ such that $t(t(x, y), y) = t(x, y)$ holds in V and $t(B, a) \in M_{\mathcal{B}}(\alpha, \beta)$ for all $a \in B$.*

In particular, there are at most M_2 equivalence classes of covers of ρ.

Proof Let $U \in M_{\mathcal{B}}(\alpha, \beta)$ and let e be an idempotent polynomial with range U. Choose a term $s(x, \bar{y})$ and a tuple $\bar{c}$ such that $e(x) = s(x, \bar{c})$. Since e is idempotent, then by suitably iterating s in its first variable, we may assume that $s(x, \bar{y})$ is idempotent in x.

Let $a \in B$ be an arbitrary element. Since $\mathcal{B}$ is abelian,

$$|U| = |s(U, \bar{c})| = |s(U, a, \ldots, a)| = |s(B, a, \ldots, a)| .$$

Also, since $\mathcal{B}/\alpha$ is abelian and $s(\beta, \bar{c}) \not\subseteq \alpha$, then $s(\beta|_U, a, \ldots, a) \not\subseteq \alpha$. So $s(B, a, \ldots, a)$ contains an (α, β)-minimal set. But since all such minimal sets have the same cardinality as U and $s(B, a, \ldots, a)$ has size equal to U then it too is an (α, β)-minimal set. Hence the term $t(x, y) = s(x, y, \ldots, y)$ is as required.

Since for each cover θ of ρ there is some binary term operation which determines some (ρ, θ)-minimal set, and there are at most M_2 binary term operations on $\mathcal{S}$, then there are at most M_2 equivalence classes of ρ covers. □

The final, and most difficult, step in proving Theorem 1.6 is to show that under the assumption of residual smallness the number of covers of ρ can be bounded as a function of the size of the generating algebra $\mathcal{A}$. Theorem 1.6 follows immediately from the following result.

6.8. Theorem *Let V be a locally finite abelian variety. The following are equivalent:*

(i) V is residually small.

(ii) V avoids the residually large configuration.

(iii) V is residually $< M_2^{M_2^4} 2^{M_3}$.

Proof (iii)⇒(i) clear.
(i)⇒(ii) This follows from Lemma 5.2.
(ii)⇒(iii) We know that for S a finite irreducible from V, S has size bounded by $(M_2^{M_2})^N 2^{M_3}$, where N is the number of covers of the strongly solvable radical of S. So, what is left is to show how avoiding the residually large configuration leads to $|N| \le M_2^3$.

□

To show that when V avoids the residually large configuration, the number of covers of ρ is at most M_2^3 we need only prove that each equivalence class of covers of ρ has at most M_2^2 elements, since we know that there are at most M_2 equivalence classes.

Let's fix a cover α of ρ and let J be the set of all covers of ρ equivalent to α and

$$\gamma = \bigvee_{\beta \in J} \beta \ .$$

Let $U \in M(\rho, \alpha)$. A modification of Exercise 5.6 can be used to show that U is also a (ρ, γ)-minimal set and that this pair of congruences is tame. A consequence of this is that the (ρ, γ)-minimal sets have most of the properties given in Theorem 2.8. A definition of tameness, as well as a proof of this can be found in [HM88].

A (ρ, γ)-trace of U is a $\gamma|_U$-class which is not also a $\rho|_U$-class. If W is such a trace then $\mathcal{S}|_W$ is a minimal algebra which is polynomially isomorphic to a vector space. The one dimensional subspaces of this vector space correspond to the covers of ρ which lie in J. Another important fact about W is noted in the following Lemma.

6.9. Lemma *If p is any polynomial of $\mathcal{S}$, then either p is constant on W or p induces an isomorphism between W and $p(W)$.*

Proof Suppose p is not injective on W. Then p is not injective on U, so p must collapse γ into ρ.

If p is not constant on W, there are $c, d \in W$ such that $p(c) \neq p(d)$. Consider $t(x, y) = p(x - y + c)$, where $+$ and $-$ are polynomials whose restrictions to W provide vector space addition and subtraction, and note that

$$t(c, c) = p(c) = t(d, d) \ .$$

But we also have

$$\langle t(c, c), t(d, c) \rangle = \langle p(c), p(d) \rangle \in \rho \setminus 0_{\mathcal{S}} \ .$$

By Lemma 4.1, the variety V is not abelian, contradiction. □

It is beyond the scope of these notes to provide a complete proof of Theorem 6.8. Instead, we will conclude by developing two key ideas which are used in the proof and which are of independent interest.

7 Multitraces

We'll see in this section that subsets larger than traces, but which are derived from traces, can carry an affine structure. These larger structures are called multitraces, and are obtained by applying a polynomial (of several variables) to a trace. The behaviour of the multitraces in a finitely generated abelian variety determines the residual character of the variety.

7.1. Example (The matrix power of a module) If M is an R-module, then M^k can be regarded as both an R-module and a module over the ring $\mathsf{M}_k(R)$ of $k \times k$-matrices over R.

As a module over R, the terms of M^k are of the form

$$r_1 I_k x_1 + \cdots + r_l I_k x_l \ ,$$

where the r_i are ring elements, I_k is the $k \times k$ identity matrix and the elements of M^k are treated as column vectors.

As a module over $\mathsf{M}_k(R)$, the terms of M^k are of the form

$$A_1 x_1 + \cdots + A_l x_l$$

for arbitrary $k \times k$-matrices A_i over R.

The latter set of matrices can be produced from the former set using the projection matrices (with a single 1 on the diagonal) and the matrix that cyclically permutes the variables.

For an arbitrary algebra $\mathcal{A}$, we consider the following operations on A^k:

$$\begin{aligned} \pi_i(x) &= (0, \ldots, 0, x_i, 0, \ldots, 0), \\ s(x) &= (x_k, x_1, \ldots, x_{k-1}), \\ d(x_1, \ldots, x_k) &= (x_{11}, \ldots, x_{kk}) \end{aligned}$$

and define the k-th *matrix power* of $\mathcal{A}$ to be the algebra

$$\mathcal{A}^{[k]} = \langle \mathcal{A}^k, d, s \rangle \ .$$

7.2. Proposition

(a) $\mathrm{Con}\mathcal{A} \simeq \mathrm{Con}\mathcal{A}^{[k]}$ *via* $\theta \mapsto \theta^k$.

(b) $^{[k]}$ *is an equivalence between the varieties* $\mathsf{HSP}(\mathcal{A})$ *and* $\mathsf{HSP}(\mathcal{A}^{[k]})$.

(c) If V is a F-vector space, then $V^{[k]}$ is term equivalent to V^k considered as a module over $\mathsf{M}_k(F)$.

Proof The proof of this Proposition is left as an exercise. $\square$

7.3. Theorem *Let $\mathcal{A}$ be finite and $0 \prec \alpha$ with* $\mathrm{typ}(0,\alpha) \in \{1,2\}$ *and let N be a $(0,\alpha)$-trace. If*

$$T = f(N, \ldots, N)$$

for some polynomial f, then $\mathcal{A}|_T$ is polynomially equivalent to a matrix power $(\mathcal{A}|_N)^{[k]}$ for some $k \leq n$.

Proof Assume that $\mathrm{typ}(0,\alpha) = 2$ and let F be the field over which $\mathcal{A}|_N$ is a vector space. Let $+$ be the addition operation for this vector space and let 0 be the zero vector. Let $0' = f(0, \ldots, 0)$ and define

$$G = \{g(x)|_T \mid g \in \mathrm{Pol}(\mathcal{A}), g(T) \subseteq N, g(0') = 0\} \subseteq N^T \ .$$

It is not difficult to see that G is a finite dimensional vector space over F. Let $g_1(x), \ldots, g_k(x)$ be a basis of G.

Since, for all i, $g_i f$ maps N into N, and $\bar{0}$ to 0, then $g_i f|_N$ is a polynomial of the vector space $\mathcal{A}|_N$ which sends 0 to 0 and so there are $\lambda_1^i, \ldots, \lambda_k^i \in F$ such that

$$g_i f(x_1, \ldots, x_n) = \sum_j \lambda_j^i x_j \ .$$

Let M be the $k \times n$-matrix (λ_j^i). It follows that since the g_i's are linearly independent, then so are the rows of M. Thus M has rank k and so $k \leq n$. Let M' be an $n \times k$-matrix M' such that $MM' = I_k$, say $M' = (\mu_j^i)$. For $i \leq n$ let $l_i(x_1, \ldots, x_k)$ be a polynomial of $\mathcal{A}$ whose restriction to N is $\mu_1^i x_1 + \cdots + \mu_k^i x_k$. Let $f'(x_1, \ldots, x_k) = f(l_1(x_1, \ldots, x_k), \ldots, l_n(x_1, \ldots, x_k))$. It is easy to check that $g_i f(\bar{x}) = x_i$ and that $f'(N, \ldots, N) \subseteq T$.

Claim: If $a, b \in T$ are distinct, then there is an i such that $g_i(a) \neq g_i(b)$.

Since T is contained in an α-class, $\langle a, b \rangle \in \alpha \setminus 0$. Since N is an $(0, \alpha)$-trace, a and b can be separated by a polynomial g mapping T into N. Let $g'(x) = g(x) - g(0)$. Then $g'(a) \neq g'(b)$ and $g'(0') = 0$, so $g' \in G$. Since the g_i's form a basis of G, the result follows.

Now let $a \in T$ and $b = f'(g_1(a), \ldots, g_k(a))$. For all i, we have $g_i(b) = g_i(f'(g_1(a), \ldots, g_k(a))) = g_i(a)$ since $MM' = I_k$. By the previous claim, $a = b$ and thus $f'(N, \ldots, N) = T$.

Using f' and the g_i's, we have bijections

$$F \colon N^k \longrightarrow T, (x_1, \ldots, x_k) \mapsto f'(x_1, \ldots, x_k) \ .$$

and

$$G \colon T \longrightarrow N^k, t \mapsto (g_1(t), \ldots, g_k(t))$$

In a natural way, F and G induce bijections between the functions on T and the functions on N^k: If $h \colon T^m \longrightarrow T$, then

$$G(h)(\bar{x}_1, \ldots, \bar{x}_m) = G(h(F(\bar{x}_1), \ldots, F(\bar{x}_m)))$$

and if $s \colon (N^k)^m \longrightarrow N^k$, then

$$F(s)(t_1, \ldots, t_m) = F(s(G(t_1), \ldots, G(t_m)))$$

.

Claim: If $s(x)$ is a polynomial of $(\mathcal{A}|_N)^{[k]}$, then $F(s) \in \mathrm{Pol}(\mathcal{A}|_T)$.

It suffices to establish this for $x+y$, λx, $\pi_i(x)$ and $s(x)$ since these functions generate the polynomial clone of $(\mathcal{A}|_N)^{[k]}$ (exercise).

$$F(x+y)(t_1, t_2) = f'(G(t_1) + G(t_2)) = f'(g_1(t_1) + g_1(t_2), \ldots, g_k(t_1) + g_k(t_2)) \in \mathrm{Pol}(\mathcal{A}),$$

$$F(\pi_i)(t) = f'(\pi_i(g_1(t), \ldots, g_k(t))) = f'(0, \ldots, 0, g_i(t), 0, \ldots, 0) \in \mathrm{Pol}(\mathcal{A}),$$

$$F(s)(t) = f'(g_k(t), g_1(t), \ldots, g_{k-1}(t)) \in \mathrm{Pol}(\mathcal{A}).$$

Claim: If $h \in \mathrm{Pol}(\mathcal{A}|_T)$, then $G(h) \in \mathrm{Pol}((\mathcal{A}|_N)^{[k]})$.

If $h(x_1, \ldots, x_m) \in \mathrm{Pol}(\mathcal{A}|_T)$, then

$$G(h)(\bar{u}_1, \ldots, \bar{u}_m) = \langle g_1 h(f'(\bar{u}_1), \ldots, f'(\bar{u}_m)), \ldots g_k h(f'(\bar{u}_1), \ldots, f'(\bar{u}_m)) \rangle \ .$$

This is a vector of linear maps on $\mathcal{A}|_N$ and so it is in $\mathrm{Pol}(\mathcal{A}|_N^{[k]})$.

The above claims establish that $\mathcal{A}|_T$ is polynomially equivalent to $\mathcal{A}|_N^{[k]}$ in the case where the type of $(0_A, \alpha)$ is 2. The argument for the type 1 case is similar and is left to the reader to reproduce. □

A set T as in the previous theorem is called a $(0, \alpha)$-*multitrace* of rank k.

8 Parallelism

Consider a field F and the n-dimensional vector space $V = F^n$ for some $n \geq 1$. The lines in V are equal to the cosets of the one-dimensional subspaces of V. Two lines L_1 and L_2, are parallel if they are cosets of the same one-dimensional subspace. Alternatively, they are parallel if and only if one can be obtained from the other via translation (i.e., addition with some fixed vector).

More generally, for β a congruence of a finite algebra $\mathcal{B}$, let us call two polynomially isomorphic subsets X, Y β-*parallel* and write

$$X \parallel_\beta Y$$

if $\langle X, Y\rangle$ lies in the transitive closure of $\{\langle t(X,a), t(X,b)\rangle \mid t(x,y) \in \mathrm{Pol}_2(\mathcal{B}), \langle a, b\rangle \in \beta\}$. If $\beta = 1_{\mathcal{B}}$, we just write $X \parallel Y$. So, if X and Y are parallel then one can be obtained from the other via a sequence of polynomial translations.

8.1. Proposition *Let $\mathcal{B}$ be a finite algebra.*

(a) For each $X \subseteq B$ and each congruence β on $\mathcal{B}$, $\parallel_\beta$ is an equivalence relation on the set of all subsets of B which are polynomially isomorphic to X.

(b) If $\mathcal{B}$ is abelian and $X \parallel Y$, then $|p(X)| = |p(Y)|$ for all polynomials p.

Proof (a) The relation is naturally reflexive. To establish symmetry and transitivity use the fact that the sets in question are all polynomially isomorphic.
(b) Let p be a polynomial. It suffices to consider pairs of the form $\langle t(X,a), t(X,b)\rangle$, where $t(x,y)$ is a polynomial. Since $\mathcal{B}$ is abelian, $p(t(x,a)) = p(t(y,a))$ if and only if $p(t(x,b)) = p(t(y,b))$ for all x, y and so $|p(t(X,a))| = |p(t(X,b))|$.

□

The previous proposition highlights an important feature of parallelism. Namely, that in an abelian algebra, polynomial projections of polynomially isomorphic parallel subsets have the same size. Let's define two subsets X and Y of an algebra $\mathcal{B}$ to be *quasi-parallel* if for all polynomials $p(x)$ of $\mathcal{B}$, $p(X)$ and $p(Y)$ have the same size.

In our vector space example, any two subsets of V which have the same size are quasi-parallel, since all nonconstant polynomials of V are permutations. If we regard V as a module over the ring of $n \times n$-matrices over the field F then it turns out that two lines of V are parallel if and only if they are quasi-parallel. This is because, as a module, there are enough polynomial projections available distinguish between lines which are not parallel.

8.2. Proposition

(a) If $\alpha \prec \beta$ such that $\mathrm{typ}(\alpha, \beta) = 2$, then all (α, β)-traces lying in the same (α, β)-minimal set are parallel.

(b) If $\mathcal{B}$ is a member of a finitely generated abelian variety and $\alpha \prec \beta$ such that $\mathrm{typ}(\alpha, \beta) = 2$, then

$$|M_{\mathcal{B}}(\alpha, \beta)/ \parallel | \leq M_2$$

Proof (a) Let N, M be traces of the same minimal set U. Let $d(x,y,z)$ be a pseudo-Mal'cev operation on U. Choose elements $n \in N$ and $m \in M$. Then $d(N,n,n) = N$ and $d(N,n,m) = M$, which is all we need.
(b) Let U be an (α,β)-minimal set. U is of the form $t(B,\bar{c})$ for some term $t(x,y)$ which is idempotent in the first variable. Choose $a \in B$. Since $\mathcal{B}$ lies in an abelian variety, $t(B,a,\dots,a)$ is also (α,β)-minimal. This shows that every (α,β)-minimal set is parallel to one determined by a binary term.

□

8.3. Remark The previous proposition applies not only to (α,β)-traces where $\alpha \prec \beta$ and $\operatorname{typ}(\alpha,\beta) = 2$, but also to (ρ,γ)-traces in our subdirectly irreducible $\mathcal{S}$.

8.4. Lemma *Let $\mathcal{A}$ be a finite abelian algebra and $0 \prec \alpha$ in $\operatorname{Con}(\mathcal{A})$ with $\operatorname{typ}(0,\alpha) = 2$. If N_0 and N_1 are distinct $(0,\alpha)$-traces which have a nonempty intersection then they are not quasi-parallel.*

Proof We need to show that there is a polynomial which is constant on precisely one of N_1 and N_0. Let U be a minimal set which contains N_0 and let $e(x)$ be an idempotent polynomial with range U. If e is constant on N_1 then we are done, and so we may assume that e is nonconstant on N_1. It follows that $e(N_1) = N_0$.

Let 0 lie in $N_1 \cap N_0$ and let $a \in N_1 \setminus N_0$. Since e is nonconstant on N_1, then $b = e(a) \neq 0$. Let $g(x)$ be a polynomial with range U and with $g(b)$ and $g(a)$ distinct members of N_0. If g is constant on one of N_0 and N_1 and one-to-one on the other, then we are done, and so either g is constant on both N_0 and N_1 or it is one-to-one on both of them. If the former holds, then we would have $g(a) = g(b)$ and so the latter must hold. By suitably iterating g, we may assume that it is idempotent. Then $c = g(a) \neq g(b) = b$ in N_0.

As $\mathcal{A}|_{N_0}$ is polynomially isomorphic to a vector space then there is a binary polynomial $-$ of $\mathcal{A}$ which behaves as subtraction on N_0 with 0 as the neutral element. Then $e(x) - g(x)$ is constant on N_0 since $e(0) - g(0) = 0 - 0 = 0$ and $e(b) - g(b) = b - b = 0$. But, $e(a) - g(a) = b - c \neq 0$, showing that $e(x) - g(x)$ is nonconstant on N_1. We conclude that N_0 is not quasi-parallel to N_1. □

8.5. Theorem *Let $\mathcal{A}$ be a finite abelian algebra and $0 \prec \alpha$ with $\operatorname{typ}(0,\alpha) = 2$. Let N_0 and N_1 be quasi-parallel $(0,\alpha)$-traces which lie in the same α-class. If $\mathsf{HSP}(\mathcal{A})$ omits the residually large configuration, then $N_0 \parallel_\alpha N_1$.*

Proof Assume $N_0 \not\parallel_\alpha N_1$ and let U_i, $i = 0,1$, be $(0,\alpha)$-minimal sets with $N_i \subseteq U_i$. Choose idempotents e_0 and e_1 with ranges U_0 and U_1 respectively. Since $e_i(N_i) = N_i$ and N_0 and N_1 are quasi-parallel and lie in the same α-class, then $e_i(N_0) = e_i(N_1) = N_i$ for $i = 0,1$. By suitably iterating these polynomials we may assume that $e_0(e_1(x)) = x$ for all $x \in N_0$.

Let $\bar{N_0}$ be an enumeration of N_0, say $\bar{N_0} \in A^l$, and let $\bar{N_1} = e_1(\bar{N_0})$. As the type of $(0,\alpha)$ is 2 then $\mathcal{A}|_{N_0}$ is polynomially equivalent to a vector space and so has an additive structure. Let $0 \in N_0$ be the neutral element with respect to this structure and let $0' = e_1(0)$. Let $-$ be a binary polynomial of $\mathcal{A}$ which acts as subtraction on N_0. Let $\mathcal{B}$ be the subalgebra of $\mathcal{A}^l$ generated by $\{\bar{N_0}\} \cup \{\hat{c} \mid c \in A\}$.

Define $\hat{\parallel}_\alpha$ to be the transitive closure in $\mathcal{B}$ of the relation

$$\{\langle t(\bar{N_0},c), t(\bar{N_0},d)\rangle \in B^2 \mid t(x,y) \text{ a polynomial of } \mathcal{A} \text{ and } \langle c,d\rangle \in \alpha\}\ .$$

Claim: $\hat{\parallel}_\alpha$ is a congruence of $\mathcal{B}$.

$\hat{\|}_\alpha$ is naturally symmetric and transitive and it is not difficult to show that it is reflexive (since $\mathcal{B}$ is generated by $\bar{N}_0$ and the constants). We need only show that the generating set of $\hat{\|}_\alpha$ is a subalgebra of $\mathcal{B}^2$. That is, we need to show that if $t(x, y)$ is a polynomial of $\mathcal{A}$, $\langle c, d\rangle \in \alpha$ and $s(x)$ is a polynomial of $\mathcal{B}$ then $\langle st(\bar{N}_0, c), st(\bar{N}_0, d)\rangle \in \hat{\|}_\alpha$. Since $\mathcal{B}$ is generated by $\bar{N}_0$ and the constants, then $s(x) = r(x, \bar{N}_0)$ for some polynomial r of $\mathcal{A}$. If we set $t'(x, y) = r(t(x, y), x)$ then $st(\bar{N}_0, c) = t'(\bar{N}_0, c)$ and $st(\bar{N}_0, d) = t'(\bar{N}_0, d)$ as required.

Claim: $\hat{0} \mathrel{\hat{\|}_\alpha} \hat{0}'$.

The pair $\langle 0, 0'\rangle$ lies in α. With the polynomial $t(x, y) = y$, we get $t(\bar{N}_0, 0) = \hat{0}$ and $t(\bar{N}_0, 0') = \hat{0}'$.

Claim: Let $\bar{C}$ and $\bar{D}$ be nonconstant members of $\mathcal{B}$ with $\bar{C} \mathrel{\hat{\|}_\alpha} \bar{D}$. Then the set of elements of $\bar{C}$ is α-parallel to the set of elements of $\bar{D}$. Consequently, $\bar{N}_0 \mathrel{\hat{\nparallel}_\alpha} \bar{N}_1$.

Let C and D be the set of elements of the tuples $\bar{C}$ and $\bar{D}$, respectively. Then C and D are $(0, \alpha)$-traces. By transitivity it suffices to verify this for pairs of the form $\bar{C} = t(\bar{N}_0, c)$ and $\bar{D} = t(\bar{N}_0, d)$ with $\langle c, d\rangle \in \alpha$.

As C and D are traces, then they are polynomially isomorphic to N_0 and so there is some polynomial $e(x)$ with $e(C) = N_0$. Then $C = t(e(C), c)$ and $D = t(e(C), d)$, showing that C and D are α-parallel.

To complete the proof, we will show that $\mathcal{B}/\hat{\|}_\alpha$ contains the residually large configuration. Let $N = \{\hat{0}, \bar{N}_0\}/\hat{\|}_\alpha$ and $N' = \{\hat{0}, \bar{N}_1\}/\hat{\|}_\alpha$. The polynomials e_0 and e_1 provide the required projections between N and N' in $\mathcal{B}/\hat{\|}_\alpha$ and the polynomial $s(x, y) = x - y$ on N_0 is a subtraction polynomial on N.

So, we need only show that for all polynomials $p(x)$ of $\mathcal{B}$, if $p(\hat{0}) \mathrel{\hat{\|}_\alpha} p(\bar{N}_0)$ or $p(\hat{0}) \mathrel{\hat{\|}_\alpha} p(\bar{N}_1)$ then $p(\bar{N}_0) \mathrel{\hat{\|}_\alpha} p(\bar{N}_1)$. Let's assume that $p(\hat{0}) \mathrel{\hat{\|}_\alpha} p(\bar{N}_0)$. Since p is a polynomial of $\mathcal{B}$ and $\mathcal{B}$ is generated by $\bar{N}_0$ and the constants, there is a polynomial $q(x, y)$ of $\mathcal{A}$ such that $p(x) = q(x, \bar{N}_0)$. Our assumption now reads $q(0, \bar{N}_0) \mathrel{\hat{\|}_\alpha} q(\bar{N}_0, \bar{N}_0)$.

By the previous claim, it follows that the sets $q(0, N_0)$ and $\{q(w, w) \mid w \in N_0\}$ are α-parallel in $\mathcal{A}$ and hence are quasi-parallel. Since $q(0, 0)$ lies in both sets,

$$q(0, N_0) = \{q(w, w) \mid w \in N_0\}$$

by Lemma 8.4.

Claim: If $\bar{Y}, \bar{Z} \in B$ with $\bar{Y} \mathrel{\hat{\|}_\alpha} \bar{Z}$ and for some i, $\bar{Y}_i = \bar{Z}_i$, then $\bar{Y} = \bar{Z}$.

Let Y and Z be the elements of the tuples $\bar{Y}$ and $\bar{Z}$ respectively. If either of $\bar{Y}$ or $\bar{Z}$ is constant, then both are and so $\bar{Y} = \bar{Z}$ in this case. Otherwise, both Y and Z are quasi-parallel traces with a nonempty intersection. By Lemma 8.4 we conclude that $Y = Z$. Let $x - y$ be a polynomial of $\mathcal{A}$ whose restriction to Y is subtraction and let $s(x) = x - \bar{Y}$. Then as $\hat{\|}_\alpha$ is a congruence we have that $s(\bar{Y}) \mathrel{\hat{\|}_\alpha} s(\bar{Z})$. But $s(\bar{Y})$ is constant and so $s(\bar{Z})$ must be as well. As $\bar{Y}$ and $\bar{Z}$ have a common entry, then we conclude that $\bar{Y} = \bar{Z}$.

From the previous claim, we can deduce that $q(0, \bar{N}_0) = q(\bar{N}_0, \bar{N}_0)$, i.e., $q(0, w) = q(w, w)$ for all $w \in N_0$. By abelianness it follows that $q(0, 0) = q(w, 0)$ for all $w \in N_0$. Since $N_0 \parallel N_1$, the polynomial $q(x, 0)$ must be constant on N_1, i.e., $q(0', 0) = q(e_1(w), 0)$ and by abelianness again, $q(0', w) = q(e_1(w), w)$ for all $w \in N_0$.

This means that

$$p(\hat{0}') = p(e_1(\bar{W}_0)) = p(\bar{W}_1) ,$$

and so

$$p(\bar{W}_0) = p(\hat{0}) \mathrel{\hat{\|}_\alpha} p(\hat{0}') = p(\bar{W}_1) . \tag{2}$$

This establishes that $\mathcal{B}/ \hat{\|}_\alpha$ contains a residually large configuration, contrary to our assumptions. Thus N_0 and N_1 must be α-parallel. □

8.6. Corollary *Let $\mathcal{A}$ be a finite simple abelian algebra which generates a residually small variety. Then any two quasi-parallel minimal sets of $\mathcal{A}$ must be parallel.*

The above connection between quasi-parallel and parallel traces in abelian varieties turns out to be crucial in the study of residual smallness. For a complete illustration of this, the reader is encouraged to consult [KKV99]. Also in that paper can be found a description of an algorithm which determines if a finite algebra generates a residually small abelian variety.

8.7. Exercise Let $\mathcal{A}$ be a finite abelian algebra such that for every congruence α which covers 0_A, $\mathrm{typ}(0_A, \alpha) = 2$. Let α_1 and α_2 be covers such that $M_{\mathcal{A}}(0_{\mathcal{A}}, \alpha_1) = M_{\mathcal{A}}(0_{\mathcal{A}}, \alpha_2)$ and let $\gamma = \alpha_1 \vee \alpha_2$. Let U be a $(0_{\mathcal{A}}, \alpha_1)$ minimal set with $e(A) = U$ for some idempotent polynomial $e(x)$.

(a) Show that U is a $(0_{\mathcal{A}}, \gamma)$-minimal set and that if δ is any congruence of $\mathcal{A}$ below γ with $\delta|_U = \gamma|_U$ then $\delta = \gamma$.

(b) Show that if β is a cover of $0_{\mathcal{A}}$ below γ then $\beta|_U \neq 0_U$ and so U is a $(0_{\mathcal{A}}, \beta)$-minimal set. (What is being established here is that $(0_{\mathcal{A}}, \gamma)$ is a *tame interval*.)

Hint: First show that if V is a $(0_{\mathcal{A}}, \beta)$-minimal set then $\alpha_i|_V \neq 0_V$ and that $\beta|_V \leq \alpha_1|_V \vee \alpha_2|_V$. Next, use the fact that $\mathcal{A}|_V$ is Mal'cev to conclude that the interval $[0_{\mathcal{A}|_V}, \alpha_1|_V \vee \alpha_2|_V]$ in the congruence lattice of $\mathcal{A}|_V$ is a height 2 modular lattice, and finally use this to show that the interval $[0_{\mathcal{A}|_V}, \beta|_V]$ is perspective to $[0_{\mathcal{A}|_V}, \alpha_i|_V]$.

(c) Let

$$\theta = \{\langle x, y\rangle \in \gamma \mid ef(x) = ef(y) \text{ for all polynomials } f \text{ of } \mathcal{A}\} .$$

Show that θ is a congruence and in fact $\theta = 0_{\mathcal{A}}$.

(d) Show that if $\langle a, b\rangle \in \gamma \setminus 0_{\mathcal{A}}$ then there is a polynomial of $\mathcal{A}$ with range U which separates a and b. (Note: this fact is needed in order to understand the structure of $(0_{\mathcal{A}}, \gamma)$-multitraces.)

8.8. Exercise Use Corollary 8.6 and Lemma 8.4 to show that any 7 element simple algebra whose minimal sets are all polynomially isomorphic to the 3 element vector space and are arranged according to the following picture (with lines representing the minimal sets) must generate a residually large equational class.

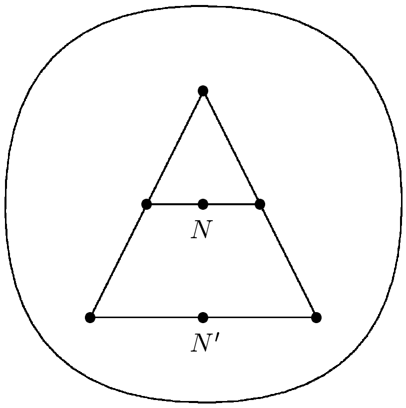

8.9. Exercise Let $\mathcal{A}$ be an algebra which has a Mal'cev polynomial. Show that if the pair $\langle u, v\rangle$ is in the congruence of $\mathcal{A}$ generated by the pair $\langle a, b\rangle$ then there is a polynomial $p(x)$ of $\mathcal{A}$ with $p(a) = u$ and $p(b) = v$.

8.10. Exercise Show that the algebra $\mathcal{A}$ defined as follows generates a residually large equational class. $\mathcal{A}$ is the algebra

$$\langle\{0,1,2,3\}, +, f(x), g(x)\rangle$$

where

+	0	1	2	3
0	0	1	0	1
1	1	0	1	0
2	0	1	0	1
3	1	0	1	0

	f	g
0	1	2
1	2	3
2	1	3
3	2	2

8.11. Exercise Let $\mathcal{A}$ be a finite simple algebra of type 3, i.e., the $(0_{\mathcal{A}}, 1_{\mathcal{A}})$-minimal sets are 2 element boolean algebras, and let U be a minimal set of $\mathcal{A}$. Show that if $p(x_1, \ldots, x_n)$ is a polynomial of $\mathcal{A}$, then the 'multitrace' $T = p(U, \ldots, U)$ is *primal*, i.e., every function on the set T is equal to the restriction of some polynomial of $\mathcal{A}$ to T.

Bibliography

[BS81] S. Burris and H. P. Sankappanavar. *A course in Universal Algebra*. Graduate Texts in Mathematics. Springer, 1981.

[FM87] R. Freese and R. McKenzie. *Commutator Theory for Congruence Modular Varieties*, volume 125 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 1987.

[HM88] D. Hobby and R. McKenzie. *The Structure of Finite Algebras*, volume 76 of *Contemporary Mathematics*. American Mathematical Society, 1988.

[Hod93] Wilfried Hodges. *Model Theory*. Oxford University Press, 1993.

[KKV99] K. Kearnes, E. Kiss, and M. Valeriote. A geometric consequence of residual smallness. *Annals of Pure and Applied Logic*, 99:137–169, 1999.

[MMT87] R. McKenzie, G. McNulty, and W. Taylor. *Algebras, Lattices, Varieties Volume I*. Wadsworth, 1987.