NEW DIRECTIONS

RALPH MCKENZIE

A variety is called *finitely decidable* (FD for short) iff the class of all its finite models has decidable first-order theory. A variety is called ω-structured iff the class of its finite models does not semantically embed the class of all finite graphs. Finite decidability implies ω-structured. P. Idziak [1] found a small collection of structural properties of algebras that hold in the finite algebras of an ω-structured variety. He proved that for finitely generated varieties omitting type 1, these properties together with a further property concerning a derived variety of modules over a finite ring, form a necessary and sufficient set of conditions for finite decidability. We seek to obtain a similar result that applies to all finitely generated varieties. In this paper, we are seeking new structural implications of ω-structuredness. This has been an ongoing project, pursued by a number of researchers, since about 1992.

1. An Example (the little diamond)

Before beginning the main work of this paper, we present here a three-element algebra that is simple of type 3 (polynomially, it is primal), generates an FD variety, and such that every two-element subalgebra is trivial (all operations are projections). This solves a question we were asking at CRAW16 in Nashville: Can a finite algebra omitting type 1, belonging to an ω-structured variety, possess a subalgebra that has type 1? Answer: Yes. Thus the example rules out what we hoped would be a further new consequence of ω-structuredness.

We put $\mathbf{A} = \langle A, f\rangle = \langle \{0,1,2\}, f(x_1,\ldots,x_6)\rangle$ where

$$\begin{aligned} f(x_1,x_2,x_3,x_4,x_5,x_6) &= x_1 \text{ if } |\{x_1,x_2,x_3\}| < 3 \\ &= x_6 \text{ if } \{x_1,x_2,x_3\} = \{0,1,2\} \text{ and } x_4 = x_5 \\ &= x_4 \text{ if } \{x_1,x_2,x_3\} = \{0,1,2\} \text{ and } x_4 \neq x_5\,. \end{aligned}$$

I claim that every finite $\mathbf{B} \in \mathcal{V} = V(\mathbf{A})$ is isomorphic to some $\mathbf{S} \times \mathbf{A}^k$ where $\mathbf{S} \models f(\bar{x}) = x_1$. Finite decidability of $\mathcal{V}$ easily follows. Clearly, $\mathbf{A}$ has the ternary discriminator as a polynomial, so is functionally complete, simple, and of type 3. Every two elements of $\mathbf{A}$ constitute a trivial two-element algebra.

2. Summary of results proved below

Our readers need to be very familiar with tame congruence theory. There is a large literature on that subject, but everything we use can be found in D. Hobby, R. McKenzie [1]. All algebras are finite and belong to an ω-structured variety $\mathcal{V}$. Given a finite algebra $\mathbf{A}$ in this variety, we have the congruences θ_1, θ_2, θ_3 which are, respectively, the largest strongly solvable congruence (or the largest congruence

Date: February 6, 2017.

α such that $\text{typ}[0_A, \alpha] \subseteq \{1\}$), the largest congruence α such that $\text{typ}[0_A, \alpha] \subseteq \{2\}$, and the largest congruence α such that $\text{typ}[0_A, \alpha] \subseteq \{3\}$. The existence of θ_1 is proved for every finite algebra in [1]. The existence of θ_2 and θ_3 for finite algebras in an ω-structured variety will be proved in this paper. As we shall see, the algebra $\mathbf{B} = \mathbf{A}/\theta_1^{\mathbf{A}}$ has a Maltsev polynomial, and $\theta_1^{\mathbf{B}} = 0_B$. Thus congruences of $\mathbf{A}$ above θ_1 permute in relation product. The chief feature of this paper is to reveal that there are many pairs of congruences lower down in the congruence lattice that are forced to permute if ω-structuredness holds.

We shall show that if α, β are congruences with $\alpha \leq \theta_i$ and $\beta \leq \theta_j$ and $3 \in \{i, j\}$, then α and β permute—$\alpha \circ \beta = \beta \circ \alpha = \alpha \vee \beta$. We conjecture that this is true for $\{i, j\} = \{1, 2\}$, and know that it is not always true for $\{i, j\} = \{1, 1\}$. Using Theorem 6.1 and some other results proved in this paper, it can be proved that if this permutability holds for $\{i, j\} = \{1, 2\}$ then it holds for $\{i, j\} = \{2, 2\}$.

Rather amazingly, our results imply that $\{\theta_1, \theta_2, \theta_3\}$ generate a Boolean sublattice of $\text{Con}(\mathbf{A})$ whose atoms are the non-zero members of $\{\theta_1, \theta_2, \theta_3\}$.

Besides known results in the published literature on FD varieties, we shall be using results from an unpublished manuscript of R. McKenzie and M. Smedberg [3]. Here are the results from all sources, that we shall use; they were all known by 2010. Assume that $\mathcal{V}$ is an ω-structured variety. The only tame congruence types of coverings in the congruence lattices of finite algebras in $\mathcal{V}$ are types 1, 2, 3, and if $\alpha \prec \beta$ in $\text{Con}(\mathbf{A})$ and the type of the covering is 2 or 3 then the (α, β)-minimal sets have no tails. The $(1, 2)$, $(2, 1)$, $(3, 1)$ and $(3, 2)$ transfer principles hold. We have that on any finite $\mathbf{A} \in \mathcal{V}$, $\theta_1 \vee \theta_2$ is the largest solvable congruence of $\mathbf{A}$ and it is an Abelian congruence, while θ_1 is strongly Abelian. If $\mathcal{V}$ is generated by one finite algebra then there are only finitely many non-isomorphic subdirectly irreducible algebras in $\mathcal{V}$, all of them finite. The type of a finite SI algebra in $\mathcal{V}$ is defined to be the type of its monolith. If $\mathbf{S}$ is such an algebra and its type is 3 then the congruence lattice is a chain of type 3 covers and $\theta_1 = \theta_2 = 0_A$. If the type of $\mathbf{S}$ is 2 then in $\mathbf{S}$, θ_2 is the centralizer of the monolith; it is an Abelian congruence comparable to all congruences and the congruences not below θ_2 form a chain of type 3 covers. In this case, $\theta_1 = \theta_3 = 0_A$. The SI algebras of types 2 and 3 have Maltsev polynomials. If $\mathbf{S}$ is an SI of type 1 then θ_1 is the centralizer of the monolith and it is a strongly Abelian congruence. In this type, $\theta_2 = \theta_3 = 0_A$, and again θ_1 is comparable to all congruences and the congruences not below θ_1 form a chain of type 3 covers. Every solvable congruence in any finite algebra of $\mathcal{V}$ is Abelian, and strongly solvable congruences are strongly Abelian.

The stated facts about type 1 SI's and the fact that strongly solvable congruences are strongly Abelian were proved apparently for the first time in McKenzie-Smedberg, and as noted there, all the other results mentioned (with the possible exception of the existence of a Maltsev polynomial for $\mathbf{A}/\theta_1$ for every finite $\mathbf{A}$ in $\mathcal{V}$) had been proved earlier by some subset of R. Willard, M. Valeriote, K. Kearnes, J. Jeong, P. Idziak and D. Delic.

3. Some preliminary results

Theorem 3.1. *Let $\mathbf{A}$ be a finite algebra in an ω-structured locally finite variety. $\mathbf{A}$ has a polynomial operation $p(x, y, z)$ such that*

$$p(x, x, y) \equiv_{\theta_1} y \equiv_{\theta_1} p(y, x, x)$$

for all $x, y \in A$–i.e., p becomes a Maltsev polynomial in $\mathbf{A}/\theta_1$.

Proof. Using what we know, it is easy to see that the algebra $\mathbf{A}/\theta_1$ augmented by all constants generates an ω-structured variety that omits type 1. (This uses the fact that θ_1 is a first order definable relation in $\mathbf{A}$.) Hence it has a Maltsev term. $\bullet$

Theorem 3.2. *Let $\mathbf{A}$ be a finite algebra in an ω-structured locally finite variety. Then $\mathbf{A}$ has a polynomial function J satisfying $J(x) \equiv_{\theta_1} x$ and $x \equiv_{\theta_1} y$ implies $J(x) = J(y)$, for all x, y. Thus J is the projection onto a transversal of the θ_1 equivalence classes. We have $x \equiv_{\theta_1} y$ iff $J(x) = J(y)$.*

Proof. We use the polynomial $p(x, y, z)$ provided by Theorem 3.1. Let $f(x)$ be a polynomial satisfying $f(x) \equiv_{\theta_1} x$ for all x having minimal range for such a polynomial. We can assume that $f(f(x)) = f(x)$ for all x. Consider the operation $f(p(f(x), f(y), f(z)) = q(x, y, z)$. We can iterate $q(x, x, z)$ as a function of z to get a polynomial $b(x, z)$ satisfying $b(x, b(x, z)) = b(x, z)$ for all x, z and such that where $q'(x, y, z) = q(x, x, q(x, x, \cdots, q(x, x, q(x, y, z))$ with the appropriate number of iterations, we have that $q'(x, y, z) \equiv_{\theta_1} q(x, y, z)$ and $q'(x, x, z) = b(x, z)$, and so

$$q'(x, x, q'(x, x, z)) = q'(x, x, z) \equiv_{\theta_1} z$$

. We claim that for all $x, z \in f(A)$, $q'(x, x, z) = z$. If this fails, for some $(x, z) = (a, b)$ then b and $q'(a, a, b) = c$ are distinct elements of $f(A)$ such that where $g(t) = q'(a, a, t)$, gf is a function with range properly contained in the range of f, and satisfying $gf(x) \equiv_{\theta_1} x$ for all x. This contradicts our choice of f. Starting with q' in place of q, we get $q''(x, y, z) \equiv_{\theta_1} p(x, y, z)$ such that $q''(y, x, x) = q''(q''(y, x, x), x, x)$, and $q''(x, x, y) = y$ for all $x, y \in f(A)$. We conclude that q'' is Maltsev acting on $f(A)$. The Maltsev equations imply that $f(A)$ contains no two distinct θ_1 congruent elements. Thus this $f(x)$ has the desired properties of $J(x)$. $\bullet$

Lemma 3.3. *The algebra $\mathbf{A}|_{J(A)}$ is polynomially equivalent to an algebra of the same signature as $\mathbf{A}$ that is isomorphic to $\mathbf{A}/\theta_1$.*

Proof. For each basic operation $F(\bar{x})$ of $\mathbf{A}$, let $F^J(\bar{x})$ be the operation $J(F(\bar{x}))$ restricted to $J(A)$ Write $\mathbf{J}(\mathbf{A})$ for the algebra $\langle J(A), F^J(\text{F basic for } \mathbf{A})\rangle$. It is easy to see that the polynomial operations of $\mathbf{J}(\mathbf{A})$ are precisely the operations of the form G^J where G is any polynomial operation of $\mathbf{A}$. Thus $\mathbf{J}(\mathbf{A})$ is an algebra in the signature of $\mathbf{A}$ that is polynomially equivalent to $\mathbf{A}|_{J(A)}$ (which in Hobby-McKenzie is defined to be the set $J(A)$ supplied with the restrictions of all polynomials of $\mathbf{A}$ under which $J(A)$ is closed). The reader can verify that mapping a/θ_1 to $J(a)$ establishes the desired isomorphism. $\bullet$

Lemma 3.4. *Let $\mathbf{A}$ be a finite algebra in an ω-structured variety. Suppose that $0 \prec_1 \alpha$ and $0 \prec_3 \beta$ are minimal congruences. If M is an $(0, \alpha)$-minimal set, then*

$$(\alpha \vee \beta)|_M = \alpha|_M \cup \beta|_M$$

.

Proof. Suppose not. Then there are a, b, b' in M such that $a \neq b \neq b'$ and $b \equiv_\alpha b'$ $a \equiv_\beta b$. Obviously, $a \neq b'$ for that would imply $(a, b) \in \alpha \wedge \beta = 0_A$. We choose an idempotent polynomial δ_1 with $\delta_1(A) = M$. There are β-traces, i.e., β-minimal sets, $\{a_0, a_1\}$, $\{a_1, a_2\}$, $\ldots$ $\{a_{n-1}, a_n\}$ such that $a_0 = a$ and $a_n = b$. Now the sets $\{\delta_1(a_0), \delta_1(a_1)\}, \ldots \{\delta_1(a_{n-1}), \delta_1(a_n)\}$ connect a to b and each of these sets is either a singleton or a β-minimal set. Hence there exists $\bar{a} \in M$ such that $\{\bar{a}, b\}$ is a β-minimal set. Since $\beta \cap \alpha = 0_A$, then $\bar{a}, b, b'$ are three distinct elements. Changing notation, we have $a, b, b' \in M$ such that $\{a, b\} = V$ is a β-minimal set and $\{b, b'\}$ is an α-subtrace. We can choose an idempotent polynomial δ_2 so that $\delta_2(A) = \{a, b\}$.

We can now semantically embed into the variety generated by $\mathbf{A}$ a family of finite simple graphs that in turn semantically embeds the class of all finite simple graphs. This will contradict the assumption that $\mathbf{A}$ lies in an ω-structured variety. The family of graphs we use consists of all finite simple graphs without loops, having at least five elements, and such that every vertex has degree at least two. In this construction, and two analogous ones later in the paper, when we have sets A, X, $\{c, d\} \subseteq A$ and $\{x\} \cup U \subseteq X$, we use $[c, d]_U$ to denote the member of A^X that maps x to d if $x \in U$ and maps x to c otherwise. $[c, d]_x$ denotes $[c, d]_{\{x\}}$.

Let $G = (X, E)$ be a finite graph satisfying the above stated restrictions. Let x_0, y_0 be two distinct elements outside of X. Put $Y = X \cup \{x_0, y_0\}$. We take $\mathbf{A}(G)$ to be the subalgebra of $\mathbf{A}^Y$ generated by all the constant functions together with, for all $e = \{x, y\} \in E$, the functions f_e such that $f_e|_X = [a, b']_{\{x,y\}}$, while $f_e(x_0) = a$ and $f_e(y_0) = b$.

We can recover G from $\mathbf{A}(G)$ as follows. Since $\alpha \cap \beta = 0_A$, and $\delta_2(b) = b$, it follows that $\delta_2(b') = b$ (as the other possibility, $\delta_2(b') = a$ would put $(a, b) \in \alpha \cap \beta$). Now $\mathbf{A}|_{\{a,b\}}$ is a two-element Boolean algebra, under operations induced by polynomial operations of $\mathbf{A}$. Then it is easy to see that $\mathbf{A}(G)$ includes $\{a, b\}^Y$. The set $\overline{X}$ of all functions $f_x = [a, b]_{\{x,y_0\}}$ ($x \in X$ is definable in $\mathbf{A}(G)$; and the relation consisting of all triples $(f_x, f_y, f_{x,y})$ where $x \neq y$ in X and $f_{x,y} = [a, b]_{x,y,y_0}$ is definable, too. (In these definitions, we need to use the $|A|$ many constant functions, and the functions $[a, b]_{y_0}$ and $[a, b]_{x_0}$ as parameters.)

We define $\overline{E}$ to be the set of all pairs $\{f_x, f_y\}$, $x \neq y$, such that there exists $h \in A(G)$ with $J(h) = J(f_{x,y})$, $\delta_1(h) = h$, and such that for no β-minimal set $\{u, v\} = q(A)$ ($q = q \circ q$) is $h = q(h)$. It is now up to us to prove that $(\overline{X}, \overline{E})$ is isomorphic to the graph G.

First, if $\{x, y\} = e \in E$ then the generator f_e is a witness h to the fact that $\{f_x, f_y\} \in \overline{E}$. In fact, $J(f_e) = J(f_{x,y})$ since $J(b) = J(b')$. (The polynomial J was introduced in Theorem 3.2.) Moreover, if we had $f_e \in \{u, v\}^Y$ where $\{u, v\}$ is a β-minimal set, it would follow that $(b, b') \in \beta$, which is false.

Now for the hard part, suppose that x, y are two elements of X and $h \in \mathbf{A}(G)$ witnesses that $\{f_x, f_y\} \in \overline{E}$. Letting $f_1, \ldots, f_m$ be a one-to-one list of the non-constant generators f_e ($e \in E$) of $\mathbf{A}(G)$, it follows that there is a polynomial $s(x_1, \ldots, x_m)$ such that $s = \delta_1 s$ and $s(f_1, \ldots, f_m) = h$.

For $e \in E$ let, $g_e(t) = a$ for $t = x_0$, $g_e(t) = f_e(t) = b'$ for $t \in e$, and $g_e(t) = b$ elsewhere in Y. Thus $g_e = f_e$ except at $t \in X \setminus e$ where $g_e(t) = b$ and $f_e(t) = a$. Notice that g_e is congruent modulo β to f_e everywhere and $g_e|_{Y \setminus \{x_0\}} \in \{b, b'\}^{X \cup \{y_0\}}$. For $1 \leq i \leq m$ put $g_i = g_{e_i}$ where $f_i = f_{e_i}$.

We claim that $g = s(g_1, \ldots g_m)$ restricted to $X \cup \{y_0\}$ is not constant (we are not claiming that this function belongs to $A(G)$). To prove the claim, suppose that g is constant on this range, so that the constant value is $\bar{b} = s(b, \ldots, b) = h(y_0) = g(y_0)$. Let $\bar{a} = h(x_0) = s(a, \ldots, a)$. Note that $h \beta g$, i.e., $h(t) \beta g(t)$ for all $t \in Y$. We have that $\{\bar{a}, \bar{b}\}$ is a polynomial image of $\{a, b\}$ and $\bar{a}\, \theta_1\, a$ and $\bar{b}\, \theta_1\, b$ since $J(h) = J(f_{x,y})$. Thus $\bar{a} \neq \bar{b}$ and consequently $\{\bar{a}, \bar{b}\}$ is a β-minimal set.

For $t \in X \setminus \{x,y\}$, we have

$$h(t)\,\beta\,g(t) = \bar{b}\,\beta\,\bar{a}\,,$$

and

$$h(t)\,\theta_1\,a\,\theta_1\,\bar{a}\,;$$

implying $h(t) = \bar{a}$. For $t \in \{x,y\}$, we have that $h(t)\,\beta\,g(t) = \bar{b}$ and $h(t)\theta_1\bar{b}$ so $h(t) = \bar{b}$. We conclude that $j'(h) = h$ where j' is an idempotent polynomial projection onto the β-minimal set $\{\bar{a},\bar{b}\}$. This is a contradiction to one condition in our definition of E.

Thus g restricted to $X \cup \{y_0\}$ is not constant. Since α is strongly Abelian and s restricted to $\{b,b'\}^{X\cup\{x_0)\}}$ is mapping into M, it depends on precisely one variable, and without loss of generality, just on the last variable. Hence g agrees on $X \cup \{y_0\}$ with $\tau(g_m)$ where $\tau(z) = s(z,\dots,z)$ and $\tau(b) = \bar{b}$, $\tau(b') = \bar{b}'$ say, and $\bar{b} \neq \bar{b}'$. Thus τ is a polynomial permutation of M. We have, of course, that $\tau(a) = \bar{a}$. Now say $f_m = f_{e_m}$, so that $g = [\bar{b},\bar{b}']_{e_m}$ on X and takes the values $\bar{a}$, $\bar{b}$ at x_0, y_0.

I claim that $\{x,y\} = e_m$ so that $\{x,y\} \in E$ as we hoped. Suppose that this fails. Then we can suppose that $x \notin e_m$. Then h takes value $\bar{b}$ at x, as $g(x) = \bar{b}$ and the two values at x are in the relation $\beta \cap \theta_1$.

Let us re-order $e_1,\dots,e_{m-1}$ (and $f_1,\dots,f_{m-1}$) so that changing notation, the sequence $f_1(x),\dots,f_m(x)$ is $\bar{r}_1\bar{r}_2a$ where $\bar{r}_1$ is a sequence of k many a's, $\bar{r}_2$ is an iteration of b' ℓ times and $a = f_m(x)$. Define a polynomial $u(z) = s(z,\dots,z,b',\dots,b',z)$ with k initial occurences of x and ℓ adjacent occurences of b. Now u maps M into M. We have that $u(a) = h(x) = g(x) = u(b)$ $(= \bar{b})$. But $u(b') = \bar{b}'$ because this depends only on the last entry and thus equals $\tau(b')$. This polynomial u contradicts the most basic property of an α-minimal set. Namely, u is a polynomial self-map of M that does not collapse α, but is not a permutation. The contradiction establishes that $\{x,y\} = e_m \in E$, as desired. •

4. Towards permutability, under θ_1, θ_3

This whole section constitutes our proof of the theorem below.

Theorem 4.1. *Let* $\mathbf{A}$ *be a finite algebra in an* ω*-structured variety. Suppose that* $0 \prec_1 \alpha$ *and* $0 \prec_3 \beta$ *are minimal congruences. Let* $V = \{a,b\}$ *be a* β*-minimal set,* U *be a* α*-minimal set containing* b*, and* $N = b/\alpha \cap M$*. Finally, suppose that* $b' \in N \setminus \{b\}$ *so that* N *is an* α*-trace. Then there must exist an* α*-minimal set* U' *containing a trace* N' *with* $a \in N'$*, and an element* $a' \in N'$ *such that* $\{a',b'\}$ *is a* β*-minimal set.*

We suppose that $\mathbf{A}$, α, β, U, N and a,b,b' are given as above and that such an element a' does not exist; and we hold these elements, congruences and sets fixed throughout §4. We choose δ_1 and δ_2 to be idempotent polynomials projecting A onto U and onto $V = \{a,b\}$ respectively. We use the construction of Lemma 3.4 to semantically embed into the class of finite subdirect powers of $\mathbf{A}$ the class of all finite graphs with at least five elements, without loops, and in which every vertex has degree larger than 1. This will contradict the ω-structuredness.

Let $G = (X,E)$ be such a graph. Let x_0, y_0 be two distinct elements outside of X. Put $Y = X \cup \{x_0,y_0\}$. We take $\mathbf{A}(G)$ to be the subalgebra of $\mathbf{A}^Y$ generated by all the constant functions together with, for all $e = \{x,y\} \in E$, the functions f_e such that $f_e|_X = [a,b']_{\{x,y\}}$, while $f_e(x_0) = a$ and $f_e(y_0) = b$.

By "definable in $\mathbf{A}(G)$" we shall mean, in this section, first-order definable using the constant functions and the functions $[a,b]_{y_0}$ and $[a,b]_{x_0}$ as parameters. We note that α restricted to $V = \{a,b\}$ is the identity function, hence $\delta_2(b') = \delta_2(b) = b$ while $\delta_2(a) = a$. Thus for $e \in E$, $\delta_2(f_e) = [a,b]_{\{y_0\}\cup e}$. Also, $\delta_1(a) = b = \delta_1(b)$ while $\delta_1(b') = b'$. This follows from Lemma 3.4. Thus $\delta_1(f_e) = [b,b']_e = g_e$. We shall need these functions $g_e \in A(G)$ several times, as in we did in the proof of Lemma 3.4.

Lemma 4.2. (1) *The set of all functions $f \in A^Y$ such that*

$$\text{for all } y,y' \in Y,\ \ f(y)\,\alpha\circ\beta\circ\alpha\, f(y')$$

is an algebra and $A(G)$ is a subalgebra of this algebra. Thus if $f \in A(G)$ and $f(y)\theta_1 f(y')$ then $f(y)\alpha f(y')$.

(2) *If $K = \{c,d\}$ is any β-minimal set then $A(G)\cap K^Y$ is a definable subset of $\mathbf{A}(G)$. Likewise, for any α-minimal set L, the set $A(G)\cap L^Y$ is definable in $\mathbf{A}(G)$.*

(3) *$A(G)\cap(a/\theta_1 \cup b/\theta_1)^Y$ is a definable subset of $\mathbf{A}(G)$.*

Proof. For (1), note that $(\alpha\circ\beta\circ\alpha)\cap\theta_1 = \alpha$, since if

$$x_0\,\alpha\,x_1\,\beta\,x_2\,\alpha\,x_3\,\theta_1\,x_0\,,$$

then $\alpha \subseteq \theta_1$ implies $x_1\theta_1 x_2$ yielding $x_1 = x_2$ because $\beta\cap\theta_1 = 0_A$. But then obviously, $(x_0,x_3)\in\alpha$.

For (2), for either set $Q = K$ or $Q = L$ there is a polynomial τ such that $\tau\circ\tau = \tau$ and $\tau(A) = Q$. Then $A(G)\cap Q^Y$ is the set of all $f \in A(G)$ satisfying $\tau(f) = f$.

For (3), we use that $A(G)\cap(a/\theta_1\cup b/\theta_1)^Y$ is the set of all $f\in A(G)$ satisfying $J(f) \in \{J(a),J(b)\}^Y$, and $\{J(a),J(b)\}$ is a β-minimal set (since $\{a,b\}$ is a β-minimal set, and (consequently) $(a,b)\notin\theta_1$ so that $J(a)\neq J(b)$. •

Now we claim, and prove, that there are two categories of elements in $A(G)$. The **first category** consists of functions all of whose values lie in one β-equivalence class. The **second category** consists of functions $f\in A(G)$ such that for some unary polynomial p, $\delta_1 p(f) = r$ is non-constant while $J(r)$ is constant. If f is of the second category as witnessed by p, then all values of $\delta_1 p(f)$ are θ_1 equivalent, two distinct values cannot be β-equivalent, and so f cannot be of the first category. We note that it follows from Lemma 4.2 (1), that in this case, all values of $\delta_1 p(f)$ lie in one α-trace inside U.

To prove the claim, that every member of $A(G)$ is either category 1 or category 2, let $f\in A(G)$. We can write $f = t(f_1,\ldots,f_m)$ where $f_i = f_{e_i}$ and $e_1,\ldots,e_m$ is a one-to-one list of all the members of E, and t is some polynomial operation. We can consider f as the result of applying t to the rows of an $|Y|\times m$ matrix whose columns are $f_1,\ldots,f_m$. All entries in this matrix are coming from $\{a,b,b'\}$ and so they are all congruent modulo $\alpha\vee\beta$. Now α centralizes α since it is an Abelian congruence, and β contralizes α because $\alpha\wedge\beta = 0_A$. Consequently, $\alpha\vee\beta$ centralizes α. Thus if

$$t(w_1,\ldots,w_{i-1},b,w_{i+1},\ldots,w_m) = t(w_1,\ldots,w_{i-1},b',w_{i+1},\ldots,w_m)$$

for some $\{w_1,\ldots,w_m\}\subseteq\{a,b,b'\}$ then

$$t(z_1,\ldots,z_{i-1},b,z_{i+1},\ldots,z_m) = t(z_1,\ldots,z_{i-1},b',z_{i+1},\ldots,z_m)$$

for all $\{z_1, \ldots, z_m\} \subseteq \{a, b, b'\}$. We say in this case that restricted to $\{a, b, b'\}$, t has no U-dependence at variable i. t does have U-dependence at variable i if and only if

$$t(b, \ldots, b, b', b \ldots, b) \neq t(b, \ldots, b, b, b \ldots, b)$$

where the lone b' sits at place i.

Now if t has U-dependence at no variables then $f = t(\delta_2(f_1), \ldots, \delta_2(f_m))$, and this obviously implies that all values of f are β-equivalent to $f(y_0) = t(b, \ldots, b)$. Thus f has category 1. On the other hand, if t does have U-dependence at some variable i, then there is a polynomial p mapping the two unequal but α-equivalent elements in the last displayed formula above, into $U = \delta_1(A)$, such that $p = \delta_1 p$ maps $t(b, \ldots, b)$ to a member of an α-trace in U. Since the function $pt(z_1, \ldots, z_m)$ restricted to mapping $\{b, b'\}^m$ into U is essentially unary (a consequence of the fact that $0_A \prec_1 \alpha$), it depends on just the i'th variable. Replacing a by b everywhere it occurs in the matrix exhibiting the production of f from $f_1, \ldots, f_m$, and using the essential unary character of pt restricted to $\{b, b'\}^m$, pf is in U and β-equivalent either to $pt(b, \ldots, b)$ or to $pt(b, \ldots, b, b', b, \ldots, b)$. Then it follows by Lemma 3.4 that every value is equal to one of these two. In fact, where, $g_i = \delta_1(f_i)$ again, this shows that $pf = p\hat{t}(g_i)$ where i is the chosen variable of U-dependence, and $\hat{t}(z) = t(z, \ldots, z)$ is a unary polynomial. Clearly, $J(\delta_1 p(f))$ is constant since $\alpha \subseteq \theta_1$. Thus f is of category 2 in this case.

Lemma 4.3. *The subsets of $A(G)$ consisting of category 1 (respectively category 2) functions are definable.*

Proof. Clearly, category 2 is definable. Since, as we have shown, category 1 is the set-complement of category 2 in $A(G)$, it follows that category 1 is definable. •

Definition 4.4. For $f, g \in U^Y$, we write $f \sim g$ iff there is a polynomial q such that $q(U) = U$ and $q(f) = g$. Note that in this case, some power of q maps U onto U inverting the action of q on U, and so $g \sim f$. This binary relation is an equivalence relation on U^Y and restricted to $A(G)$, is a definable equivalence relation on $A(G) \cap U^Y$.

Recall that for an edge $e \in E$ we have $\delta_1(f_e) = [b, b']_e = g_e$. The collection of all functions $f \in A(G)$ such that $f = \delta_1(f)$ and $f \sim g_e$ for some $e \in E$ we denote by Δ_E.

Corollary 4.5. *Δ_E is definable, in fact, it is the collection of all functions $f \in A(G)$ such that $f = \delta_1(f) \sim p(f')$ for some $f' \in A(G)$ and polynomial p witnessing that f' is of the second category.*

Note that if p witnesses that f is of the second category, then the argument before Lemma 4.3 effectively showed that $p(f) \in \Delta_E$.

Observe that the set $\overline{X}$ of all functions $f_x = [a, b]_{\{x, y_0\}}$ $(x \in X)$ is definable in $\mathbf{A}(G)$; and the relation consisting of all triples $(f_x, f_y, f_{x,y})$ where $x \neq y$ in X and $f_{x,y} = [a, b]_{x,y,y_0}$ is definable, too. (In these definitions, we need to use the $|A|$ many constant functions, and the functions $[a, b]_{y_0}$ and $[a, b]_{x_0}$ as parameters.)

Definition 4.6. Consider these possible properties of a function $f \in A(G)$.

(i): $f = \delta_1(f)$;

(ii): f is not a constant function;
(iii): $J(f)$ is constant;
(iv): $\delta_1(f)$ satisfies (i)–(iii), and for any unary polynomial p such that $\delta_1 p(f)$ satisfies (i)–(iii) we have $\delta_1 p(f) = \delta_1 p(\delta_1(f)) \sim \delta_1(f)$.

We now define the set Γ. It is the set consisting of all $f \in A(G)$ such that

$$J(f) = [J(a), J(b)]_{y_0,x,y}$$

for some $x \neq y$ in X, and f satisfies (iv)

We have seen that the functions f of the second category are defined by the existence of a polynomial p so that $\delta_1 p(f)$ satisfies (i)–(iii). Thus Γ is a subset of the set of functions of the second category. Let us note that for $f \in \Gamma$ with $J(f) = [J(a), J(b)]_{y_0,x,y}$, we have that $\delta_2(f) = [a, b]_{y_0,x,y}$. In general, the first of these conditions is stronger than the second.

<u>Outline of the remainder of the proof of Theorem 4.1.</u> We have a bijection $x \mapsto f_x = [a, b]_{x,y_0}$ between X and a definable set $\overline{X}$ in $A(G)$. We need to show that the relation $\overline{E}$ over $\overline{X}$ which consists of all pairs (f_x, f_y) such that $\{x, y\} \in E$, is definable. The ternary relation consisting of all triples $(f_x, f_y, f_{x,y})$ with $\{x, y\}$ a two-element subset of X is definable. Thus it will suffice to show that the set of functions $f_{x,y}$ ($\{x, y\} \in E$) is definable. This is exactly the set of functions $\delta_2(f)$ where $f = f_e$ is one of our generators.

We will show first that the generators belong to Γ, then define a property, "the bad property". We will define $f \in \Gamma$ to be good if it fails to have the bad property. We will prove that the generators are good; and that if $f \in \Gamma$ is good, then $\delta_2(f) = \delta_2(f')$ where f' is a generator. Thus we will have that the set of functions $[a, b]_{e\cup\{y_0\}}$ ($e \in E$) is precisely the set of $\delta_2(f)$ where f belongs to Γ and is good. This fact will obviously establish that $\overline{E}$ is definable, as desired.

Lemma 4.7. (1) *Every function f_e ($e \in E$) belongs to Γ, and $\delta_1(f_e) \in \Delta_E$.*
(2) *If $f \in \Gamma$, then $\delta_1(f) \sim g_e$ for some $e \in E$.*

Proof. (1) is very easy, using the fact, from Hobby-McKenzie [1], that every polynomial map from U to itself that doesn't identify b and b' is a permutation of U and has a polynomial inverse. (2) is established in Corollary 4.5. •

Now we need to more precisely analyze the form of functions in Γ.

Lemma 4.8. *Let $f \in \Gamma$, $f = t(f_1, \ldots, f_m)$ where $f_1, \ldots, f_m$ are the non-constant generators f_e.*

(1) *$t(z_1, \ldots, z_m)$ has dependence on U at z_i for precisely one i.*
(2) *Defining*
$$f_\beta = t(\delta_2(f_1), \ldots, \delta_2(f_m)) \pmod{\alpha} \text{ and}$$
$$f_\alpha = t(\delta_1(f_1), \ldots, \delta_1(f_m)) \pmod{\beta},$$
we have that $f_\beta \equiv_\alpha f \equiv_\beta f_\alpha$. This means that $f_\beta(z) \equiv f(z) \pmod{\alpha}$ for all $z \in Y$ and that $f_\alpha(z) \equiv f(z) \pmod{\beta}$ for all $z \in Y$.
(3) *Where $\delta_2(f) = [a, b]_{x,y,y_0}$, and $(\bar{a}, \bar{b}) = (f(x_0), f(y_0))$, we have $f_\beta = [\bar{a}, \bar{b}]_{x,y,y_0}$. Moreover, $\delta_2(f_\beta) = \delta_2(f)$; the set $\{\bar{a}, \bar{b}\}$ is a β-minimal set; and $\delta_1(f_\alpha) = \delta_1(f)$.*

(4) *The values of f on $\{x, y, y_0\}$ are all α-equivalent, and likewise, f is constant modulo α on $(X \setminus \{x, y\}) \cup \{y_0\}$.*

Proof. To prove (1), suppose that t has dependence on U at z_i. Then

$$t(b, \ldots, b, b', b, \ldots, b) \neq \hat{t}(b)$$

with the b' being substituted on the left for z_i. By tame congruence theory (Hobby-McKenzie [1]), we get a unary polynomial $p = \delta_1 p$ such that p maps the two unequal but α-equivalent elements to distinct elements of an α-trace in U. Then by Lemma 3.4, $p(f) = p(f_\alpha) \sim g_i$ This cannot happen for two values of i since $g_i \not\sim g_{i'}$ when $i \neq i'$.

For (2), Since $f_i(y_0) = b$ for all i then

$$f(y_0) = \hat{t}(b) = t(b, \ldots, b) = f_\beta(y_0) = \bar{b}$$

and likewise

$$f(y_0) = f_\alpha(y_0) = \bar{a}\,.$$

$\{\bar{a}, \bar{b}\}$, the $\hat{t}$-image of $\{a, b\}$, is a β-minimal set, since $\hat{a}\,\theta_1\, a$ and $\hat{b}\,\theta_1\, b$.

It should be clear that $f_\beta \equiv_\alpha f \equiv_\beta f_\alpha$. Since $V = \{a, b\}$ has no distinct α-related elements, then $\delta_2(f) = \delta_2(f_\beta)$. For $w \in \{x, y, y_0\}$ we have $f_\beta(w)\,\alpha\, f(w)\,\alpha\, f(y_0) = \bar{b}$ (by Lemma 4.2 (1) and the fact that $J(f) = J([a, b]_{x,y,y_0})$. Also, $f_\beta(w) \beta t(b, \ldots, b) = \bar{b}$. Thus $f_\beta(w) = \bar{b}$. For $w \in (X \setminus \{x, y\}) \cup \{x_0\}$, the same argument proves that $f_\beta(w) = \bar{a}$. Thus $f_\beta(w) = [\bar{a}, \bar{b}]_{x,y,y_0}$.

It follows from Lemma 3.4 and Lemma 4.2 (1), and our assumption that $J(\delta_1(f))$ is constant while $\delta_1(f)$ is not constant, and the fact that $f \equiv_\beta f_\alpha$, that $\delta_1(f) = \delta_1(f_\alpha)$. This completes the proof of (3).

For $z \in \{x, y\}$ we have $f(z) \theta_1 \bar{b}$ since $J(f) = J([a, b]_{x,y,y_0})$; so that by Lemma 4.2 (1), we have that $f(z) \equiv_\alpha \bar{b}$. By the same token, for $z \in X \setminus \{x, y\}$ we have that $f(z) \equiv_\alpha \bar{a}$. This is (4). •

Lemma 4.9. (1) *Let $f \in \Gamma$, $J(f) = [J(a), J(b)]_{x,y,y_0}$ with $\{x, y\}$ a two-element subset of X, and $\delta_1(f) \sim [b, b']_e$, $e = \{u, v\} \in E$. We have that f agrees with $f_\beta = [\bar{a}, \bar{b}]_{x,y,y_0}$ except at $w \in e$ where $\{f(w), f_\beta(w)\}$ is an α-subtrace. If $\{x, y\} = e$ then $f(x) = f(y)$. If $\{x, y\} \cap e = \emptyset$ then $f(u) = f(v)$.*

(2) *For $f \in \Gamma$ as above, $f(u)$ is α-equivalent to exactly one of $f(x_0), f(y_0)$, and the α-pair, either $\{f(u), f(x_0)\}$ or $\{f(u), f(y_0)\}$ is an α-subtrace.*

Proof. For (1), let $f \in \Gamma$. We use the same notation as in the previous lemma. We have $f = t(f_1, \ldots, f_m)$ where f_i are the generators f_e and $g_i = \delta_1(f_i)$. And we are writing $\bar{a} = f(x_0) = \hat{t}(a)$ and $\bar{b} = f(y_0) = \hat{t}(b)$.

We have that the m-ary polynomial $\delta_1(t(\bar{x}))$ acting on $\{b, b'\}^m$ and mapping into U depends on at most one variable, since $0 \prec_1 \alpha$; actually depends on exactly one variable since $\delta_1(f)$ is non-constant. We can assume without loss of generality that the variable of dependence is the last. Thus $\delta_1(f) = \delta_1(f_\alpha)) = \delta_1 \hat{t}([b, b']_{e_m})$ where $\delta_1 \hat{t}$ acts as a permutation on U. Hence $\delta_1(f) \sim [b, b']_{e_m}$. Say $e_m = \{u, v\}$.

Now assume that $z \in X \setminus \{x, y, u, v\}$ and $f(z) \neq \bar{a}$. We know that $f(z) \alpha \bar{a}$. By tame congruence theory, there is a polynomial p such that $p = \delta_1 p$ and $p(f(z)) \neq p(f(x_0))$. By Lemma 3.4, $p(f) = p(\delta_1(f))$, and by (iv), $p(f) \sim \delta_1(f) \sim [b, b']_{u,v}$; but this is impossible because $p(f)$ takes distinct values at z and x_0. The contradiction shows that $f(z) = f(x_0)$ for $z \in X \setminus \{x, y, u, v\}$. The same argument shows that if $e_m \cap \{x, y\} = \emptyset$, then $f(x) = f(y) = f(y_0)$ and $f(u) = f(v)$. The argument shows also that $f(z) = f(y_0)$ if $z \in \{x,y\} \setminus \{u,v\}$. If, say, $x \notin \{u,v\}$ but $y = u$, then the argument gives that $f(x) = f(y_0)$ but $f(y) = f(u) \neq f(y_0)$ since $\delta_1(f)$ takes different values at u and y_0.

The first assertion in (2) is obvious. Let us assume, without loss of generatlity, that $f(u)\alpha f(y_0)$. We consider $f(u) = t(u_1,\dots,u_{m-1},f_m(u))$. We know that $f(u) \neq f(y_0)$. Also, we have $f_m(u) = b'$ and we look at

$$t(u_1,\dots,u_{m-1},b) = f_\beta(u) = f_\beta(y_0) = f(y_0) = \bar{b}\,.$$

We have that $\{f(u), f(y_0)\} = \{t(u_1,\dots,u_{m-1},b'), t(u_1,\dots,u_{m-1},b)\}$, a polynomial image of $\{b,b'\}$, and $f(u) \neq f(y_0)$ (else, $(f(u), f_\beta(u) \in \alpha \cap \beta$, which is false). Thus (3) is proved. •

Definition 4.10. For $f \in \Gamma$, $J(f)$ takes the value $J(b)$ at precisely three points, $\{y_0,x,y\}$ where x and y are distinct elements of X. Also $\delta_1(f) = [\bar{b},\bar{b}']_e$ takes the exceptional value $\bar{b}'$ at precisely two points $\{u,v\}$ where $u,v \in X$ and $\{u,v,\} = e \in E$. We call f *aligned* if $\{x,y\} = e$, that is, the two patterns match:

$$e = \{x \in X : J(f)(x) = J(b)\}$$

.

Note that if f is aligned then $\{x,y\} \in E$.

We call $f \in \Gamma$ *un-aligned* if it is not aligned.

We need a definable property that holds for all the generators f_e and fails for all the non-aligned functions in Γ. Call f good if it belongs to Γ and does not have the property described below. If f has the property, call it bad.

The bad property. A function $f \in \Gamma$ has this property iff there is a polynomial p of nine variables and an element $\tau \in \Gamma$ and elements $r_1, s_1, \dots, r_8, s_8 \in \{a,b\}^Y$, all of which take value a at x_0 such that $f = p(\bar{r},\tau)$ and $J(p(\bar{s},\tau))$ is the constant function of value $J(a)$.

Lemma 4.11. *If $f \in \Gamma$ is un-aligned, then it has the bad property.*

Proof. We return to the context of Lemmas 4.8 and 4.9 and use the notation employed in their proofs. So we have $f \in \Gamma$,

$$f = t(f_1,\dots,f_{m-1},f_e)\,,$$

where $t(\bar{z})$ depends on U only at the last variable z_m, f_e is one of our generators, and since t is independent of U at z_i for $i < m$, we can replace each f_i by $\delta_2(f_i)$ for $i < m$ and thus write $f = t(h_1,\dots,h_{m-1},f_e)$ with $\{h_1,\dots,h_{m-1}\} \subseteq \{a,b\}^Y$. We assume that f is un-aligned, i.e., $e \neq \{x,y\}$. There are two cases.

Case 1: $e \cap \{x,y\} \neq \emptyset$. We can assume that $e = \{y,v\}$, v,x,y three distinct elements. We have, say $f(v) = \bar{a}' \neq \bar{a} \equiv \bar{a}' \pmod{\alpha}$ and $f(y) = \bar{b}' \neq \bar{b} \equiv \bar{b}' \pmod{\alpha}$.

Case 2: $e \cap \{x,y\} = \emptyset$. $e = \{u,v\}$. In this case, we have by Lemma 4.9 (2) that $f(u) = f(v) = \bar{a}'$ say, with $\bar{a} \neq \bar{a}' \equiv \bar{a} \pmod{\alpha}$.

We can visualize f as the result of applying t to the rows of a $|Y| \times m$ matrix where the final column is f_e and the first $m-1$ columns belong to $\{a,b\}^Y$. We alter this matrix twice, each time by a system of row-by-row changes in the first $m-1$

columns. In all changes, the last position in a row is unchanged, so altogether, we are simply replacing one system of $m-1$ columns in $\{a,b\}^Y$ by another such system. In both changes, each z'th row, where $f(z)=\bar{a}$ and $z \notin e$ gets replaced by a constant row of a's. This does not change the last position in the row, where an a is left standing, and t applied to the new row gives of course the same value $\bar{a}$. Additionally, in the first change, if $e \cap \{x,y\} = \emptyset$, we replace the v'th row by the u'th row and the x'th row by the y'th row. This gives us a system where in the initial $X \times (m-1)$ matrix there are two constant rows and just two other rows, so at most four distinct columns among the first $m-1$. t applied to the new matrix also gives f, because we had that $f(u)=f(v)$ and $f(x)=f(y)=\bar{b}$.

For the second change in this case where $e \cap \{x,y\} = \emptyset$, we replace the y_0'th row by a row that looks like the vth, but with b in last position, not b', and we replace the x'th and y'th rows by the x_0'th row, all a's. By identifying variables among $z_1,\ldots,z_{m-1}$ where both changed matrices have equal columns, we get a five variable polynomial and by adding dummy variables, a nine variable polynomial $p(\bar{z},w)$ with $p(r'_1,\ldots,r'_8,f_e)=f$ and $J(p(s'_1,\ldots,s'_8,f_e))) = \langle J(a)\rangle$, and where all r'_i, s'_i belong to $\{a,b\}^Y$.

In case $e=\{v,y\}$, in the first change we just replace the z'th row for every $z \in X\backslash\{x,y,v\}$ by a constant row of a's, leaving the x'th, y'th, v'th rows unchanged. In the second change, we replace the first $m-1$ entries in the y_0'th row and also those of the y'th row, by the first $m-1$ entries of the v'th row, and all other rows become constant rows of a's. This produces a nine variable polynomial $p(\bar{z},w)$ with $p(r'_1,\ldots,r'_8,f_e)=f$ and $J(p(s'_1,\ldots,s'_8,f_e))) = \langle J(a)\rangle$, and where all r'_i, s'_i belong to $\{a,b\}^Y$.

This ends our proof that if f is un-aligned then it has the bad property. •

Lemma 4.12. *Our chosen generators are good.*

Proof. Let $f=f_e$, $e=\{u,v\}\in E$. We show that f does not have the bad property. We already know that $f\in\Gamma$. Suppose, to get a contradiction, that f has the bad property. Let $f=p(r_1,\ldots,r_8,\lambda)$ and $f^\dagger = p(s_1,\ldots,s_8,\lambda)$ where $\lambda\in\Gamma$ and all the functions $r_1,\ldots,s_8 \in \{a,b\}^Y \subseteq A(G)$ take the value a at x_0, and $J(f^\dagger)$ is constantly $J(a)$. As we noted above, $f^\dagger \in A(G)$. By Lemma 4.2 (1), all values of $f^\dagger$ lie in one α-equivalence class. Then since $f^\dagger(x_0)=f(x_0)=a$, we have that $f^\dagger(u)\,\alpha\, a$. We have that $f(u)=b'\,\alpha\, b$. Thus it is clear that $f(u)$ and $f^\dagger(u)$ are unequal, and congruent modulo β since $r_j(u)$ and $s_j(u)$ are congruent modulo β for $1\le j\le 8$. In fact, we shall show that $\{b', f(u)\}$ is a polynomial image of $\{a,b\}$, hence it is a β-minimal set.

Indeed we have $f(u)=p(\bar{c},w)$ and $f^\dagger(u)=p(\bar{d},w)$ where $w=\lambda(u)$ and $\bar{c},\bar{d}\in\{a,b\}^8$. Permuting the first 8 variabls in p so that those i where $c_i=d_i=a$ come first, and then, those where $c_i=d_i=b$, and then those where $c_i=a$ and $d_i=b$, finally those where $c_i=b$ and $d_i=a$, we get a 5-variable polynomial p' such that $f(u)=p'(a,b,a,b,w)$, $f^\dagger(u)=p'(a,b,b,a,w)$. Let ι be a polynomial such that $\iota(a)=b$ and $\iota(b)=a$. ($\mathbf{A}|_{\{a,b\}}$ is a Boolean algebra). Then put $p''(x)=p'(a,b,x,\iota(x),w)$. One sees that $p''(a)=f(u)$, while $p''(b)=f^\dagger(u)$. Thus $\{b', f^\dagger(u)\}$ is a β-minimal set.

We know that $\{a, f^\dagger(u)\} \subseteq \alpha$. We need to show that this set is actually an α-subtrace. Since $\delta_1(f) = [b, b'_{u,v}$, it follows by Lemma 3.4 that

$$\delta_1(f) = \delta_1 p(b, \ldots, b, \lambda) = \delta_1 p(s_1, \ldots, s_r, \lambda) = \delta_1 f^\dagger) \,.$$

Thus $\delta_1(f^\dagger) = \delta_1(f)$. Then by Lemma 4.9 (2), exactly one of $\{\lambda(u), \lambda(x_0)\}$, $\lambda(u), \lambda(y_0)\}$ is a subtrace. The remainder of our proof is almost the same in the two cases, so we assume that $\{\lambda(u), \lambda(y_0)\}$ is the α-subtrace. Consider the element

$$c = p(s_1(u), \ldots, s_8(u), \lambda(y_0)) \,.$$

We claim that $c = f^\dagger(y_0)$. We know that $(\lambda(u), \lambda(y_0)) \in \alpha$. Thus $(c, f^\dagger(y_0) \in \alpha \cap \beta)$ and the claim follows. Next we claim that $f^\dagger(y_0) = a$. Indeed $f^\dagger(x_0) = f(x_0) = a$ since $r_i(x_0) = a = s_i(x_0)$ for all i. Moreover, $(\lambda(x_0), \lambda(y_0)) \in \beta$ so that $(f^\dagger(y_0), f^\dagger(x_0)) \in \beta \cap \alpha$, establishing this claim. Putting two claims together, we get that

$$\{a, f^\dagger(u)\} = \left\{ p(\overline{s(u)}, \lambda(y_0)), p(\overline{s(u)}, \lambda(u)) \right\} \,,$$

a polynomial image of $\lambda(y_0), \lambda(u))$. Thus by Lemma 4.9 (2), the two elements constitute an α-subtrace unless they are equal. But equality of these elements implies that $(b', a) \in \beta$, which in turn implies $(b', b) \in \beta$, which is definitely false.

Our starting assumption as we began the proof of Theorem 4.1 was that no such element as $f^\dagger(u)$ exists. This contradiction proves that $f = f_e$ is good. It is aligned by definition. •

Now every function f_e is of the form $\delta_2(f)$ for some good $f \in \Gamma$; and it follows that $\{x, y\} \in E$ (where $x \neq y$, $\{x, y\} \subseteq X$, the function $[a, b]_{x,y,y_0}$ is $\delta_2(f)$ for some good $f \in \Gamma$.

The remaining details of demonstrating that we have a semantic embedding of the class $\mathcal{G}$ of graphs into the finite subdirect powers of $\mathbf{A}$ are just the same as in the proof of Lemma 3.4 and will be omitted.

We have now contradicted our assumption that $\mathbf{A}$ belongs to an ω-structured variety, and this completes our proof of Theorem 4.1.

5. Interaction of type 1 and type 2 minimal sets

This section is the beginning of an attempt to prove that α and β permute whenever $\alpha \leq \theta_1$ and $\beta \leq \theta_2$, which would imply also that any two congruences contained in θ_2 permute. It seems it would also imply that any ω-structured Abelian variety is the varietal product of a type 1 variety with a type 2 variety. So far, all I have is the theorem below.

Theorem 5.1. *Let $\mathbf{A}$ be a finite algebra in an ω-structured variety. Suppose that $0 \prec_1 \alpha$ and $0 \prec_2 \beta$ are minimal congruences. If U is an $(0, \alpha)$-minimal set, then*

$$(\alpha \vee \beta)|_U = \alpha|_U \cup \beta|_U$$

.

Proof. We prove this with a modification of a construction that was used in Hobby, McKenzie [2] (Chapter 10) to prove Lemma 8.4 in that book. In Hobby, McKenzie [2] there was the extra assumption that $\mathbf{A}$ is Abelian, which we do not assume.

For the construction, assume that α, β, U are as in the statement of the lemma. Let M be an α-trace contained in U, and let $0, 0', 1$ be elements such that $\{0, 0'\} \subseteq M$, $0 \neq 0'$, and $\{0, 1\} \in \beta$, $0 \neq 1$. Let δ_1 be an idempotent polynomial projecting

A onto U. Since β classes are connected by their traces and $e_1(0) = 0$, $e_1(1) = 1$, there is a chain $0 = a_0 = \delta_1(b_0), a_1 = \delta_1(b_1), \ldots, a_n = \delta_1(b_n) = 1$ where each set $\{b_i, b_{i+1}\}$ is a β-subtrace. Since $1 \neq 0$, there is an i so that $0 = a_i \neq a_{i+1}$. Now δ_1 restricted to a $(0_A, \beta)$-minimal set containing $\{b_i, b_{i+1}\}$ does not collapse β to 0_A in that set, so it follows that e_1 restricted to that set is a polynomial isomorphism, and thus there is a $(0_A, \beta)$-minimal set $V \subseteq U$ having trace N with $0, c \in N$ for some $c \neq 0$. we change notation so that $0, 1 \in N$. Since $\alpha \cap \beta = 0_A$, the elements $0', 0, 1$ are all distinct. Let δ_2 be an idempotent polynomial projecting A onto V.

After defining the construction, we will show that it serves to semantically embed the class $\mathcal{G}$ of all finite simple graphs without loops and with at least three vertices, into the class of finite subdirect powers of $\mathbf{A}$. This contradicts our assumption that $\mathbf{A}$ lies in an ω-structured variety. The contradiction proves the theorem.

So let $\mathbf{G} \in \mathcal{G}$, $\mathbf{G} = (G, E)$. Let $X = G \cup \{p_1, p_2, p_3, p_4\}$ where p_1, p_2, p_3, p_4 are distinct points outside G. We define some elements of A^X.

For $v \in G$, let $f_v : X \to \{0, 0', 1\}$ be defined as

$$f_v(x) = \left\{ \begin{array}{ll} 0' & \text{if } x = p_2 \\ 1 & \text{if } x = v \\ 0 & \text{otherwise,} \end{array} \right\}$$

and for $e \in E$ and $i = 1, 2$ let $f_e^i : X \to \{0, 0', 1\}$ be defined as

$$f_e^i(x) = \left\{ \begin{array}{ll} 0' & \text{if } x = p_i \\ 1 & \text{if } x \in e \\ 0 & \text{otherwise.} \end{array} \right\}$$

Let

$$G^\star = \{f_v : v \in G\},$$

$$E^\star = \{f_e^i : e \in E, i = 1, 2\},$$

$$N^\star = \{f \in N^X : f(p_1) = f(p_2), f(p_3) = 0 \text{ and } f(p_4) = 1\}.$$

We define

D is the subalgebra of A^X
generated by the set $G^\star \cup E^\star \cup N^\star$
together with all the constants.

Notice that for $\mu \in D$ we have $\{\mu(p_1), \mu(p_2)\} \subseteq \alpha$ since the generators of $\mathbf{D}$ satisfy this.

Lemma 5.2. *Let $t(x_1, \ldots, x_k)$ be a k-ary polynomial of $\mathbf{A}$ such that $t(A^k) \subseteq U$ and for all $a_1, \ldots, a_k \in M \cup N$ and all $i \leq k$ the unary polynomial*

$$t(a_1, \ldots, a_{i-1}, x, a_{i+1}, \ldots, a_k)$$

is non-consant on $M \cup N$. Then either all these functions are constant on M, or else $k = 1$ and $t(x_1)$ maps U bijectively onto U.

Proof. If this lemma is false, by substituting constants for some of the variables in t, we can find a binary polynomial $t(x,y)$ such that $t(x,a)$ is non-constant on M for some $a \in M \cup N$ while $t(c,x)$ is non-constant on $M \cup N$ for all $c \in M \cup N$. Since $M \cup N$ is contained in one equivalence class of the Abelian congruence $\alpha \vee \beta$ then $t(x,a)$ induces a permutation of U for all $a \in M \cup N$. Since α is a type 1 atom, then $t|_{M^2}$ depends only on its first variable. For $a,b,c \in M$, $t(a,b) = t(a,c)$. Since $\alpha \vee \beta$ is Abelian, we have $t(a,b) = t(a,c)$ whenever $a \in M \cup N$ and $b,c \in M$, in fact because U is α-minimal, $t(a,b) = t(a,c)$ whenever $a \in M \cup N$ and $(b,c) \in \alpha|_U$.

For all $c \in M \cup N$, $t(c,x)$ is not constant on N, since this function is not constant on $M \cup N$. Hence $t(c,V)$ is a β-minimal set inside U, and $t(c,N)$ is a β-trace.

Since $t(U,a) = U$, for $a \in M \cup N$, then $t(x,a)$ permutes the β-equivalence classes. Thus

$$t(0/\beta, 1) = (t(0,1)/\beta \cap U = (t(0,0)/\beta) \cap U = t(0/\beta) \cap U, 0)\,,$$

since $0,1 \in \beta$. Thus we can find $c \in U$ such that $(0,c) \in \beta$ and $t(c,1) = t(0,0)$. Now, $(0,c) \notin \alpha$ for if not then

$$(t(c,1), t(0,1)) \in \alpha\,, \text{ implying } (t(0,0), t(0,1)) \in \alpha \cap \beta\,,$$

and so $t(0,0) = t(0,1)$. This contradicts the fact noticed above that $t(c,x)$ does not collapse N.

Since 0 and c are distinct β-equivalent elements, we may choose a polynomial g with $g(A) \subseteq V$ and $g(0) \neq g(c)$. Since all β-traces in V are polynomially isomorphic, and $\mathbf{A}|_N$ is polynomially equivalent to a one-dimensional vector space, we may choose h so that $h(A) \subseteq V$ and $h(0) = 0$, $h(c) = 1$. Let $p(x) = t(x, h(x))$. Then $p(0) = t(0,0)$, $p(c) = t(c,1) = t(0,0) = p(0)$, but $p(0') = t(0', h(0')) = t(0', h(0)) \neq p(0)$. The last inequality implies that $p|_U$ is a permutation, and contradicts the equality $p(0) = p(c)$. This concludes our proof of this lemma. •

Now put

$$S = \{g \in \mathrm{Pol}_1 \mathbf{A} : g(A) \subseteq U = g(U)\}\,.$$
$$D(U) = D \cap U^X\,.$$
$$D(V) = D \cap V^X\,.$$

Note that $D(U), D(V)$ are definable in $\mathbf{D}$. Where δ_1, δ_2 are the polynomial projections of A onto U, V respectively, $D(U)$ consists of all $\mu \in D$ with $\delta_1(\mu) = \mu$ and likewise for $D(V)$. For $\mu, \nu \in D(U)$, we put $\mu \sim \nu$ iff there is some $g \in S$ with $\nu = g(f)$.

For $\mu \in D(U)$ we can write $\mu = t(\nu_1, \dots, \nu_k)$ for a term t and generating elements $\nu_1, \dots, \nu_k$. We can always arrange that t satisfies the hypothesis of Lemma 5.2. If $k = 1$ and $t(x_1)$ acts as a permutation on U, then $\mu \sim \nu$ for some generator ν. Else t is a collapsing function, i.e., the one-variable specializations of t all collapse $\alpha|_M$. In this case, obviously, we can take $\nu_1, \dots, \nu_k \in N^X$ and $\mu \in N(t)^X$, where we write $N(t) = t(N^k)$.

Since A is finite we can choose finitely many collapsing functions $t_1, \dots, t_m$ such that for every collapsing function t we have $N(T) = N(t_i)$ for some $1 \leq i \leq m$. This choice of the t_i is fixed and independent now of choice of $\mathbf{G}$. The arity of t_i will be denoted by n_i.

Suppose that $\mu = t(\nu_1, \dots, \nu_k)$ for collapsing function t and generators $\nu_j \in N^\star$, as above, and choose i with $N(t) = N(t_i)$. Then the range of f is contained in $N(t_i)$, hence for $x \in X$ we can choose $\nu_{x,1}, \dots, \nu_{x,n_i} \in N$ with $f(x) = t_i(\nu_{x,1}, \dots, \nu_{x,n_i})$.

For $1 \leq i \leq n_i$ let $\nu_i' = (\nu_{x,i} : x \in X)$. Then clearly, $\mu = t_i(\nu_1', \ldots, \nu_{n_i'})$ and $\nu_j' \in D(V)$.

We take $\mathrm{Coll}(x)$ as a first order formula (with some parameters from D) equivalent to x is not constant, $x \in D(U)$, and $x = t_i(y_1, \ldots, y_{m_i})$ for some $\{y_1, \ldots, y_{n_i}\} \subseteq D(V)$.

We take $\mathrm{Gen}(x)$ as a first order formula equivalent to $x \in D(U)$ and x is not constant and $\neg\mathrm{Coll}(x)$, and for all $y \in D(U)$ with $\neg\mathrm{Coll}(y)$ and polynomial g with $g(A) \subseteq U$, if $g(y) = x$ then $y \sim x$ and in fact $g(U) = U$.

Claim 1: For $\mu \in D$, there is a generator $\nu \in G^\star \cup E^\star$ with $\mu \sim \nu$ iff $\mathbf{D} \models \mathrm{Gen}(\mu)$.

To prove this, suppose first that $\mu \sim \nu \in G^\star \cup E^\star$. Then any member of $D(U)$ equivalent to μ takes two distinct values that are α-equivalent. This rules out $\mu \sim \nu \in D(V)$. If $\mathrm{Coll}(\mu)$ then all values of μ are β-equivalent; again this is false. The prove of the converse is very easy, using the same facts, and Lemma 5.2.

Clearly, if $\mathrm{Gen}(\mu$ then $\mu \sim \nu$ for a unique $\nu \in G^\star \cup E^\star$. We put $G^\dagger$ equal to those μ with $\mathrm{Gen}(\mu)$ such that whenever $\mathrm{Gen}(\mu')$ and $\mu \sim \lambda$, $\mu' \sim \lambda'$, $\lambda, \lambda' \in D(U)$ and $\delta_2(\lambda) = \delta_2(\lambda')$ then $\mu \sim \mu'$. We put $E^\dagger$ equal to the set of those μ such that $\mathrm{Gen}(\mu)$ and $\mu \notin G^\dagger$. It is easy to see that $G^\dagger = G^\star/\sim$ and $E^\dagger = E^\star/\sim$. We can choose a polynomial operation $x + y$ so that restricted to N, $x + y$ is the addition of vectors in the vector space $\mathbf{A}|_N$ with 0 chosen as the zero element. Thus for an edge $e = \{x, y\} \in E$, $\delta_2(f_x) + \delta_2(f_y) = \delta_2(f_e^i)$ $(i = 1, 2)$. Where $m(x, y, z)$ is a polynomial giving a Maltsev operation on V, we must have $x + y = m(x, 0, z)$. The operation m has the property that whenever $x, y, z, x', y', z' \in V$ and exactly one of the pairs $(x, x'), (y, y'), (z, z')$ is an unequal pair, then $m(x, y, z) \neq m(x', y', z')$. That is, each one-variable specialization of m acts as a permutation of V. Using this, we can easily show that if $\mu \sim f_x$ and $\mu' \sim f_y$ then $\{x, y\} \in E$ iff $\mu \not\sim \mu'$ and there is $\lambda \in E^\dagger$ $\bar{\mu} \sim \mu$, $\bar{\mu}' \sim \mu'$ $(\bar{\mu}, \bar{\mu}' \in D(U))$, and $\delta_2(\lambda) = \delta_2(\bar{\mu}) + \delta_2(\bar{\mu}')$.

The only tricky piece of the proof: Suppose that $\mu, \mu', \lambda \in G^\dagger$, $\lambda \in E^\dagger$, and $\delta_2(\mu) + \delta_2(\mu') = \delta_2(\lambda)$. Say $\mu \sim f_x$, $\mu' \sim f_y$, $\lambda \sim f_e^i$, $e \in E$, $x \neq y$. Then say $\delta_2(\mu)$ is constantly equal to a except at x where the value is $b \neq a$, and $\delta_2(\mu')$ is constantly equal to a' except at y where the value is $b' \neq a'$. Then $\mu + \mu'$ is constantly equal to $a + a'$ except at x, y where the respectively values are $a' + b$ and $a + b'$. Since λ constantly equal to some a'' except at exactly two places where it takes equal values different from a'', it follows that $e = \{x, y\}$ (and that $a + a' = a''$ and $b = b'$).

So we have a bijection between X and $G^\dagger/\sim$ where both $G^\dagger$ and the equivalence relation $\sim$ have been shown to be definable. And the edge relation in $\mathbf{G}$ corresponds to a relation defined in essentially first order language at the end of the last paragraph.

This concludes our proof that the construction $\mathbf{D}(G)$ semantically embeds $\mathcal{G}$ if Lemma 5.1 fails. Thus Theorem 5.1 is proved.

6. Permutability with pure type 3 congruences

Theorem 6.1. *Let $\mathbf{A}$ be a finite algebra in an ω-structured variety. If α, β are congruences such that $0 \prec_1 \alpha$ and $0 \prec_3 \beta$, then $\alpha \circ \beta = \beta \circ \alpha$ and $\alpha|_M \vee \beta|_M = \alpha_M \cup \beta|_M$ for every $(0_A, \alpha)$-minimal set M.*

Proof. From Hobby-McKenzie [1], we know that α is the transitive closure of Tr_α where Tr_α is the set of all pairs (x, y) such that $\{x, y\}$ is an α-subtrace. Likewise, β is the transitive closure of Tr_β. Theorem 4.1 states that

$$\mathrm{Tr}_\alpha \circ \mathrm{Tr}_\beta \subseteq \mathrm{Tr}_\beta \circ \mathrm{Tr}_\alpha\,.$$

Since these relations are symmetric, we have that Tr_α permutes with Tr_β. It follows that α permutes with β. •

Theorem 6.2. *Suppose that $\alpha \leq \theta_1$ and typ$[0_A, \beta] \subseteq \{3\}$. Then $\alpha \circ \beta = \beta \circ \alpha$.*

Proof. We first prove this in the case that $0_A \prec \beta$. There is a chain

$$0_A \prec \alpha_1 \prec \cdots \prec \alpha_k = \alpha$$

. We shall show that if $i < k$ and α_i permutes with β, then α_{i+1} permutes with β. So suppose that $\alpha_i \circ \beta = \alpha_i \vee \beta$. Now in the quotient algebra $\mathbf{B} = \mathbf{A}/\alpha_i$, we have that $0_B \prec \alpha_{i+1}/\alpha_i$. We show that $0_B \prec (\alpha_i \vee \beta)/\alpha_i$ also. Indeed, suppose that in $\mathbf{A}$, congruence θ is a congruence with $\alpha_i < \theta \leq \alpha_i \vee \beta$. Choose $(c, d) \in \theta$ outside of α_i. We have $(c, e) \in \alpha_i$ and $(e, d) \in \beta$ for some e. Clearly, $e \neq d$. Thus (e, d) generates β and there is a polynomial f so that $\{f(e), f(d)\}$ is a $(0_A, \beta)$-minimal set. We have $f(e) = f(c)$ since the only other possibility is $f(c) = f(d)$ but then $(f(e), f(d)) \in \alpha_i \cap \beta$ which is impossible—$\alpha_i \cap \beta = 0_A$. Thus $(f(c), f(d))$ generates β, implying that $\beta \leq \theta$ and so $\theta = \alpha_i \vee \beta$.

Thus Theorem 6.1 applied to $\mathbf{B}$ yields that the congruences α_{i+1}/α_i and $(\alpha_i \vee \beta)/\alpha_i$ permute. In $\mathbf{A}$, this means that α_{i+1} permutes with $\alpha_i \vee \beta$. Consequently,

$$\alpha_{i+1} \circ \beta \subseteq (\alpha_{i+1} \circ (\alpha_i \vee \beta) = (\beta \circ \alpha_i) \circ \alpha_{i+1} = \beta \circ \alpha_{i+1}\,.$$

Now given α and β as in the statement of the theorem, we string a chain

$$0_A \prec \beta_1 \prec \cdots \prec \beta_k = \beta\,.$$

We showed above that α permutes with β_1. We show inductively that α permutes with β_i implies α permutes with β_{i+1}. Then it will follow that α permutes with β.

So assume that α permutes with β_i. Now β_i is strongly solvably equivalent with $\alpha \vee \beta_i$, which means that in $\mathbf{B} = \mathbf{A}/\beta_i$ we have that $(\alpha \vee \beta_i)/\beta_i \leq \theta_i^{\mathbf{B}}$. Also, in $\mathbf{B}$, $0_B \prec_3 \beta_{i+1}/\beta_i$. Thus the case proved above implies that $(\alpha \vee \beta_i)/\beta_i$ permutes with β_{i+1}/β_i. In $\mathbf{A}$, this means that $\alpha \vee \beta_i$ permutes with β_{i+1}. Thus we have the calculation

$$\alpha \circ \beta_{i+1} \subseteq \beta_{i+1} \circ (\alpha \vee \beta_i) = \beta_{i+1} \circ (\beta_i \circ \alpha) = \beta_{i+1} \circ \alpha\,.$$

This concludes our proof of Theorem 6.2. •

Lemma 6.3. *Let $0_A \prec_2 \alpha$ and M be a $(0_A, \alpha)$-minimal set. Suppose that N is a (β, β')-minimal set where $\beta \prec_3 \beta'$. Then N is a two-element set not included in M.*

Proof. Since M has no tail, then the induced algebra $\mathbf{M} = \mathbf{A}|_M$ is a nilpotent Maltsev algebra. On the other hand, since N has no tail, $\mathbf{A}|_N$ is polynomially equivalent to a 2-element Boolean algebra. Since both M and N are the projected images of A under idempotent polynomials, then if $N \subseteq M$ we would have that $\mathbf{A}|_N = \mathbf{M}|_N$. Thus N would be a 2-snag in $\mathbf{A}|_M$, contradicting that $\mathbf{M}$ is nilpotent. •

Lemma 6.4. *Let $0_A \prec_2 \alpha$ and M be a $(0_A, \alpha)$-minimal set. Suppose typ$[0_A, \beta] \subseteq \{3\}$. then $\beta|_M = 0_M$.*

Proof. We can suppose that this fails for β, but there is $\beta' \prec \beta$ such that $\beta'|_M \subseteq 0_M$. By the assumption, $\beta' \prec_3 \beta$. Tame congruence theory tells us that for any two β-congruent elements, there is a chain connecting one to the other in which every pair of consecutive elements either are β'-equivalent or constitute a (β', β)-minimal set. Let $M = e(A)$ where $e = e \circ e$ is a polynomial. Suppose that we do have $(b, b') \in \beta$, $b \neq b'$, and $b, b' \in M$. Connect b to b', say $b = a_0, \dots, a_n = b'$ where each a_i, a_{i+1} are β'-congruent or constitute a minimal set for the cover. Now we have $b = e(a_0), e(a_1), \dots, e(a_n) = b'$. Since $e(a_i), e(a_{i+1}\}$ cannot be a (β', β)-minimal set, then tame congruence theory says that we have $e(a_i)$ congruent to $e(a_{i+1})$ modulo β', holding for all i. It follows that $(b, b') \in \beta'$, contradiction. •

Theorem 6.5. *Let $0_A \prec_1 \alpha$ and suppose that typ$[0_A, \beta] \subseteq \{3\}$. Let M be an α-minimal set. Then $(\alpha \vee \beta)|_M = \alpha|_M \cup \beta|_M$.*

Proof. Let $0_A = \mu_0 \prec \cdots \prec \mu_n = \beta$ where each cover $\mu_i \prec \mu_{i+1}$ is of type 3. Let a, a' be two distinct elements of an α-trace in M. We show by induction on i that $a\mu_i b \in M$ implies $b = a$.

Suppose this is true for μ_i but not for μ_{i+1}, say $a \neq b \in M$, $(a, b) \in \mu_{i+1}$. Let $\mathbf{B} = \mathbf{A}/\mu_i$. By Theorem 6.2, $\alpha \vee \mu_i = \alpha \circ \mu_i$. Thus it easily follows that in $\mathbf{B}$ we have $\mu_i \prec_1 (\alpha \vee \mu_i)/\mu_i = \alpha'$ and in $\mathbf{A}$ the intervals $[0_A, \alpha]$ and $[\mu_i, \alpha \vee \mu_i]$ are projective. By Hobby-McKenzie [1], M/μ_i is an α'-minimal set and a/μ_i, $a'.\mu_i$ belong to an α'-trace in this set. Moreover, $\mu_{i+1}/\mu_i \succ_3 \mu_i$ in $\mathbf{B}$. The element $b/\mu_i \in B$ then contradicts Lemma 3.4. This shows that in fact a is μ_{i+1} congruent to no element of M but itself. •

Theorem 6.6. *Let β, δ be congruences of $\mathbf{A}$ such that typ$[0_A, \beta] \cup$ typ$[0_A, \delta] \subseteq \{3\}$. Then typ$[0_A, \beta \vee \delta] \subseteq \{3\}$ and $\beta \circ \delta = \delta \circ \beta$.*

Proof. Suppose that the first statement fails. Then from the $(3, 2)$ and $(3, 1)$ exchange properties, which both hold in $\mathbf{A}$, we have some α such that $0 \prec_1 \alpha \leq \beta \vee \delta$ or $0 \prec_2 \alpha \leq \beta \vee \delta$. In either case, let M be an α-minimal set. Choose $e = e^2$ with $e(A) = M$.

If $0_A \prec_2 \alpha$, We can choose b, b' in M, $b \neq b'$, $(b, b') \in \alpha \cap (\beta \vee \delta)$; and applying e to a Maltsev chain, we have that $(b, b') \in \beta|_M \vee \delta|_M$. This certainly implies that there is some $a, a' \in M$, $a \neq a'$, such that $(a, a') \in \beta \cup \delta$. But this contradicts Lemma 6.5

If $0_A \prec_1 \alpha$, M contains a pair $(b, b') \in \alpha \cap (\beta \vee \delta)$. As before, $(b, b') \in \beta|_M \vee \delta|_M$. So there certainly must exist $a \in M$, $a \neq b$, $(a, b) \in \beta|_M \cup \delta_M$. The triple a, a', b contradicts Theorem 6.3

Now let us prove the second statement. We do have that $(\beta \vee \theta_1)$ and $\delta \vee \theta_1$ commute, since $\mathbf{A}/\theta_1$ has a Maltsev polynomial. We also have that β and δ permute with θ_1 by Theorem 6.2. Thus

$$\beta \circ \delta \subseteq (\delta \circ \theta_1) \circ (\beta \circ \theta_1) = \delta \circ \beta \circ \theta_1 .$$

So if this theorem fails, we'll have a pair $(a, a') \in \theta_1 \cap (\beta \vee \delta)$, $a \neq a'$, and choosing some $0_A \prec_1 \alpha \leq \mathrm{Cg}(a, a')$, we'll have a pair $(b, b') \in \alpha$, $b \neq b'$, $\{b, b'\}$

contained in a $(0_A, \alpha)$-minimal set U. This pair (b, b') also belongs to $\beta \vee \delta$, and so to $\beta|_U \vee \delta_U$, and then clearly we get a contradiction of the first statement of this theorem, proved above. This finishes our proof of the theorem. •

Theorem 6.7. (1) *There is a unique largest congruence θ_3 with only type 3 covers below it. Similarly, there is a unique largest conruence θ_2 with only type 2 covers below it.*

(2) *We have that, $\theta_1 \cap (\theta_2 \vee \theta_3) = \theta_2 \cap (\theta_1 \vee \theta_3) = \theta_3 \cap (\theta_1 \vee \theta_2) = 0_A$.*

(3) *$\theta_1 \vee \theta_2$ is the solvable radical, the largest solvable congruence. It is Abelian.*

Proof. The existence of θ_3 follows from Theorem 6.6 For θ_2 we have a more basic argument. Suppose that $\mathrm{typ}[0_A, \delta] \cup \mathrm{typ}[0_A, \theta] \subseteq \{2\}$.

To prove that $\mathrm{typ}[0_A, \delta \vee \theta] \subseteq \{2\}$, suppose not. Let $0_A \leq \rho \prec \lambda \leq \delta \vee \theta$, with the type of this cover different from 2. The type cannot be 3, since $\delta \vee \theta$ is solvable. Thus $\rho \prec_1 \lambda$. By the $(2,1)$-transfer principle, there is $0_A \prec_1 \alpha \leq \delta \vee \theta$. $\alpha \leq \theta$ is impossible, so $\delta \vee \theta \geq \alpha \vee \theta > \theta$. The interval $[\theta, \delta \vee \theta]$ is solvable, so by the $(2,1)$-transfer principle, there is $\theta \leq \tau \prec_1 \delta \vee \theta$. Clearly, $\theta \not\leq \tau$ so we have a non-trivial strongly solvable interval $[\theta \cap \tau, \theta]$. But then $1 \in \mathrm{typ}[0_A, \theta]$, a contradiction

To prove (2), we begin by noting that $\theta_1 \vee \theta_2$ is solvable, so $\theta_3 \cap (\theta_1 \vee \theta_2) = 0_A$. Our arguments for the disjointness of θ_1 with $\theta_2 \vee \theta_3$, and disjointness of θ_2 with $\theta_1 \vee \theta_3$, parallel the proof above that θ_2 exists.

Suppose that $\theta_2 \cap (\theta_1 \vee \theta_3) \geq \lambda \succ_2 0_A$. Now $\theta_3 \cap \lambda = 0_A$, so $\theta_3 \vee \lambda > \theta_3$. Find $\theta_3 \leq \rho \prec \theta_3 \vee \lambda$. We have that the intervals $[0_A \lambda]$ and $[\rho, \theta_3 \vee \lambda]$ are projective so that they have the same type. That type is 2. But θ_3 is strongly solvably equivalent to $\theta_3 \vee \theta_1$. This is a contradiction.

Finally, suppose that $\theta_1 \cap (\theta_2 \vee \theta_3) \geq \lambda \succ_1 0_A$. Then $\lambda \vee \theta_3$ is larger than θ_3 and strongly solvably equivalent to it. Since $\theta_2 \vee \theta_3$ is solvably equivalent to θ_3, by the $(1,2)$-exchange property, there is $\theta_3 \leq \alpha \prec_1 \theta_2 \vee \theta_3$. Then $\alpha \cap \theta_2$ is smaller than θ_2 and strongly solvable equivalent to it. This gives $1 \in \mathrm{typ}[0_A, \theta_2]$, a contradiction.

For (3), note that the join of two solvable congruences is solvable. Suppose that there is a larger solvable congruence that $\theta_1 \vee \theta_2$. Then $\theta_1 \vee \theta_2$ has a cover of type 1 or 2. Suppose first that $\theta_1 \vee \theta_2 \prec_1 \lambda$. By the $(2,1)$-exchange property, the interval from θ_1 to λ has a type 1 atom $\theta_1 \prec_1 \delta$. But then δ is strongly solvable, hence lies below θ_1, a contradiction.

Finally, suppose that $\theta_1 \vee \theta_2 \prec_2 \lambda$. The exchange properties applied in the interval from θ_2 to λ yield $\theta_2 \prec_2 \delta$ for some δ. We get a chain $0_A \beta_1 \prec_2 \cdots \prec_2 \beta_m = \delta$. Suppose that there is a type 1 cover below δ, hence a type 1 atom $0_A \prec_1 \alpha \leq \delta$. Let i be largest with $\alpha \not\leq \beta_i$. Then $i < m$ and $\alpha \leq \beta_{i+1}$. Since $0_A \prec \alpha$ and $\beta_i \prec \beta_{i+1}$, it follows that $[0_A, \alpha]$ and $[\beta_i, \beta_{i+1}]$ are projective intervals. By tame congruence theory, the have the same type. Thus $1 = 2$, which is absurd. •

Theorem 6.8. *Suppose that $\alpha \leq \theta_2$ and $\beta \leq \theta_3$. Then α and β permute.*

Proof. We have

$$\beta \circ \alpha \leq (\alpha \vee \theta_1) \circ (\beta \vee \theta_1) \subseteq (\alpha \vee \theta_1) \circ \beta\,,$$

since $\mathbf{A}/\theta_1$ has permuting congruences and β permutes with θ_1 (by Theorem 6.2). So assuming that $a \,\beta\, b \,\alpha c$ there is $d \in A$ with $(a, d) \in \alpha \vee \theta_1$ and $(d, c) \in \beta$. Thus

$$(a, d) \in (\theta_2 \vee \theta_1) \cap (\beta \,\circ\, \alpha \,\circ \beta) \leq \theta_2 \vee \theta_3\,.$$

Let $\theta = \mathrm{Cg}(a, d)$. This is a solvable congruence below $\theta_2 \vee \theta_3$ so by Theorem 6.7 it contains no type 1 atom, and by the $(2, 1)$-exchange property, has no type 1 congruence covers. Thus $\theta \leq \theta_2$.

Choose a chain $0_A = \theta_0 \prec_2 \theta_1 \cdots \prec_2 \theta_k = \theta$. We shall prove by induction on $i \leq k$ that $\theta_i \leq \alpha$. So suppose that $i < k$ and $\theta_i \leq \alpha$. Choose any (θ_i, θ_{i+1})-minimal set M, say $M = e(A)$, $e = e^2$ a polynomial. Choose $(x, y) \in \theta_{i+1}|_M \setminus \theta_i|_M$. Now $(x, y) \in \alpha \vee \beta$ so by applying e we obtain $(x, y) \in \alpha|_M \vee \beta|_M$. By Lemma 6.3, $\beta|_M = 0_M$. Thus $(x, y) \in \alpha$. Since θ_{i+1} is generated by $\theta_i \cup \{(x, y)\}$, it follows that $\theta_{i+1} \leq \alpha$. So we conclude that $\theta \leq \alpha$, $(a, d) \in \alpha$, and $(a, c) \in \alpha \circ \beta$.

We have shown that $\beta \circ \alpha \subseteq \alpha \circ \beta$. Applying converse to this inclusion, we get $\alpha \circ \beta \subseteq \beta \circ \alpha$. •

7. The remaining results

Theorem 7.1. *Let α, λ be congruences of $\mathbf{A}$ such that typ$[0_A, \alpha] \cup$ typ$[0_A, \lambda] \subseteq \{2, 3\}$. Then typ$[0_A, \alpha \vee \lambda] \subseteq \{2, 3\}$.*

Proof. •

Theorem 7.2. $\theta_1 \vee \theta_2 = \theta_2 \circ \theta_2 \circ \theta_2$.

Proof. Let $(a, b) \in \theta_2 \vee \theta_2$. Then $a\theta_1 J(a)$ and $b\theta_1 J(b)$ and $(J(a), J(b)) \in \theta_1|_{J(A)} \vee \theta_2|_{J(A)}$ We know that $\theta_1|_{J(A)} = 0_{J(A)}$. Thus $J(a)\theta_2 J(b)$. •

Theorem 7.3. 0_A, $\theta_1, \theta_2\theta_3$, $\theta_1 \vee \theta_2$, $\theta_1 \vee \theta_3$, $\theta_2 \vee \theta_3$, $\theta_1 \vee \theta_2 \vee \theta_3$ *constitute a Boolean sublattice of Con* $\mathbf{A}$.

Proof. •

Theorem 7.4. *Every pair of congruences below $\theta_2 \vee \theta_3$ commute.*

Proof. •

Theorem 7.5. *Suppose that $0_A \prec_i \alpha$ and $0_A \prec_j \beta$ where $\{i, j\} \subseteq \{1, 2, 3\}$, $i \neq j$. Let U be an α-minimal set. Then $(\alpha \vee \beta)|_U = \alpha|_U \cup \beta|_U$.*

Proof. •

Theorem 7.6. 0_A, $\theta_1, \theta_2\theta_3$, $\theta_1 \vee \theta_2$, $\theta_1 \vee \theta_3$, $\theta_2 \vee \theta_3$, $\theta_1 \vee \theta_2 \vee \theta_3$ *constitute a Boolean sublattice of Con* $\mathbf{A}$.

Proof. •

References

(1) 1. P. Idziak, A characterization of finitely decidable congruence modular varieties, Transactions of the AMS, Volume 340, Numer 3 (1997), 903–934.
(2) 2. D. Hobby, R. McKenzie, The Structure of Finite Algebras, Birkhäuser, 1989.
(3) 3. R. McKenzie, M. Smedberg, Strong solvability and residual finiteness for finitely decidable varieties (manuscript 2013).

Department of Mathematics, Vanderbilt University, Nashville, U.S.A.
E-mail address: ralph.n.mckenzie@vanderbilt.edu