A Geometric Consequence of Residual Smallness *

Keith A. Kearnes † Emil W. Kiss ‡
Matthew A. Valeriote §

Abstract

We describe a new way to construct large subdirectly irreducibles within an equational class of algebras. We use this construction to show that there are forbidden geometries of multitraces for finite algebras in residually small equational classes. The construction is first applied to show that minimal equational classes generated by simple algebras of types **2**, **3** or **4** are residually small if and only if they are congruence modular. As a second application of the construction we characterize residually small locally finite abelian equational classes.

1 Introduction

The Cartesian plane, considered as a geometry, is a system of points and lines related by incidence. Or, one may define it without lines, as a set of points with a given betweenness relation. This same object may be viewed from an algebraic standpoint by taking the set of points as the universe and equipping this set with all operations which preserve betweenness. The resulting algebra can be identified with $\mathbf{R}^2$ considered as an $M_2(\mathbf{R})$-module, expanded by the constant operations.

*1991 *Mathematical Subject Classification* Primary 08A05; Secondary 03C13.

†Research supported by a fellowship from the Alexander von Humboldt Stiftung

‡Research supported by the Hungarian National Foundation for Scientific Research, grant nos. 7442 and 16432.

§Research supported by the NSERC of Canada.

Having both geometric and algebraic standpoints provides multiple ways to view basic geometric concepts, such as the concept of parallelism. For example, a purely geometric way to say that lines ℓ and k are parallel is to say that "ℓ and k do not intersect". An equivalent way to say this which mixes both algebra and geometry is to say that "k is a translate of ℓ". A purely algebraic way to say this is that "any unary algebraic operation which is constant on ℓ is constant on k".

This paper is about geometric properties of finite algebras. Our primary interest is in "multitraces of type **2**", which are special definable subsets of a finite algebra on which the induced structure is that of a vector space $\mathbf{F}^n$ enriched to include the $M_n(\mathbf{F})$-module and constant operations. Here $\mathbf{F}$ can be any finite field. Multitraces of type **2** abound in algebras with abelian properties, and they "patch together" in ways that lead to complex behavior.

Within a single multitrace of type **2** all definitions of parallelness coincide, since within a multitrace the situation is essentially the same as the one discussed in the first paragraph. However, it makes sense to ask if a line from one multitrace is parallel to a line in a different multitrace. Here we discover that different definitions of parallelness describe different concepts. If ℓ is a line in one multitrace and k is a line in any other multitrace, we will call ℓ and k *quasi-parallel* if any unary polynomial function which is constant on ℓ is also constant on k, and conversely. We reserve the word *parallel* for the situation where k can be obtained from ℓ through a sequence of translations, possibly travelling through many multitraces en route. It turns out that under the right centrality hypothesis, which will hold in all applications of this paper, parallelism implies quasi-parallelism. One of the more significant facts proved in this paper (Theorem 5.6) is that in a residually small equational class the converse implication holds.

One consequence of the fact that residual smallness forces parallelism and quasi-parallelism to coincide is that there are forbidden geometries of multitraces in residually small equational classes. One such forbidden geometry appears in Figure 1. Here we assume that the algebra is simple of type **2**. It has seven elements, depicted as points of the geometry, and it has four multitraces, each depicted as a three-element line. There are many inequivalent simple algebras of type **2** which have the geometry pictured in Figure 1, but all generate residually large equational classes. The reason for this is that the geometry of Figure 1 forces the lines N and N' to be quasi-parallel, but not parallel. Please consult the end of Section 5 for an explanation of this fact.

With some effort, one can find such an algebra which in addition generates an abelian equational class.

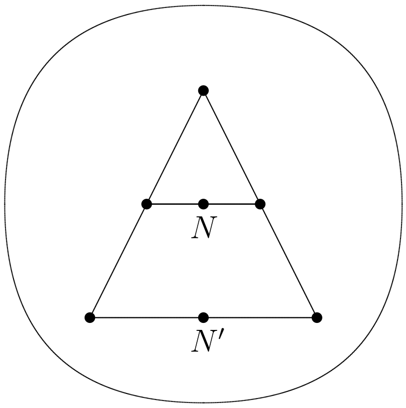

FIGURE 1: A forbidden geometry.

The paper is structured as follows. Section 2 explains the construction of large irreducibles on which the applications are based. Section 3 contains an immediate application: minimal equational classes generated by simple algebras of types **2**, **3** or **4** are residually small if and only if they are congruence modular. Section 4 develops the basic properties of multitraces. Although our eventual interest will be in multitraces of type **2**, the arguments here work equally well for multitraces of type **3** so we include them. Section 5 develops the notions of quasi-parallelism and parallelism for multitraces. Finally, in Section 6 we use the machinery developed to characterize residually small abelian equational classes.

Throughout the paper we use tame congruence theory. The reader can find the necessary background in [2] and [6]. One point of divergence between this paper and those works is that we use the term *irreducible* in place of (the more usual) *subdirectly irreducible*. We will make use of the following notation throughout the paper. Let A be a set and κ a cardinal. If $u \in A$, then $\hat{u}$ denotes the function in A^κ which is constant with value u. The set of these elements, where u runs over A is called the *diagonal* of A^κ. Now let $\mathbf{x}$ be a k-tuple in A (the i-th coordinate of such a tuple will usually be denoted by x_i). For a k-ary operation $f(\mathbf{x})$ on A we let $\hat{f}$ denote the k-ary operation on A^κ which acts coordinatewise like f. Note that when f happens to be a polynomial of an algebra $\mathbf{A}$, then the operation $\hat{f}$ will be a polynomial operation of any subalgebra of $\mathbf{A}^\kappa$ which contains the diagonal.

Acknowledgements. The first and second authors are greatly indebted to Matthew Valeriote for inviting them to Hamilton to work on the topic of this paper. Some of the results of this paper were completed during the Algebraic Model Theory Program at The Fields Institute in 1996/97. All three authors would like to express our appreciation to the staff and directorate of the Fields Institute for providing an excellent and stimulating working environment.

2 Constructing Large Irreducibles

Rather than directly constructing large irreducibles, it is easier to construct algebras which have a homomorphism onto a large irreducible. The following lemma tells us when an algebra has a homomorphism onto an irreducible of cardinality $\geq \kappa$.

LEMMA 2.1 ([3], Lemma 2.1) *An algebra* $\mathbf{B}$ *has an irreducible homomorphic image of cardinality* $\geq \kappa$ *if and only if there is a 4-tuple* (a, b, X, γ) *satisfying the following conditions.*

(1) $a, b \in B$, $X \subseteq B$,

(2) $\gamma \in \operatorname{Con} \mathbf{B}$ *and* $(a, b) \notin \gamma$,

(3) *for every* $\psi \in \operatorname{Con} \mathbf{B}$ *with* $\psi \geq \gamma$ *the following implication holds:*

$$|X/(\psi|_X)| < \kappa \quad \Longrightarrow \quad (a, b) \in \psi\,. \qquad \blacksquare$$

We put Lemma 2.1 to use in the proof of the next theorem, which describes the construction on which all later results in this paper depend.

THEOREM 2.2 *Let* $\mathbf{A}$ *be any algebra and assume that*

(1) $(1, 0)$ *is a* $\mathbf{1}$*-snag of* $\mathbf{A}$*, that is,* $0, 1 \in A$ *and there exists a binary polynomial* s *of* $\mathbf{A}$ *satisfying* $s(0, 0) = 0$, $s(0, 1) = s(1, 0) = 1$;

(2) *f is a unary polynomial of $\mathbf{A}$ such that $f(0) = 0$ and $1' \stackrel{\text{def}}{=} f(1) \neq 1$;*

(3) *in the subalgebra $\mathbf{T}$ of $\mathbf{A}^2$ generated by the diagonal and $\{0,1\}^2$, $(1,1)$ is in a singleton block of the congruence $\tau = \mathrm{Cg}^{\mathbf{T}}((0,1),(0,1'))$.*

Then $\mathsf{V}(\mathbf{A})$ is residually large.

PROOF. Let $\kappa \geq 2$ be a cardinal and denote $\{0,1\}$ by N. Let $\mathbf{B}$ be the subalgebra of $\mathbf{A}^\kappa$ generated by the diagonal and the set $X = N^\kappa$. Let

$$G = \left\{(\mathbf{x}, \hat{f}(\mathbf{x})) \in B^2 \mid \mathbf{x} \in N^\kappa - \{\hat{1}\}\right\} ,$$

and denote by γ the congruence of $\mathbf{B}$ generated by the set G. We now show that $(\hat{1}, \hat{1}', X, \gamma)$ is a 4-tuple witnessing the fact that $\mathbf{B}$ has an irreducible homomorphic image of cardinality $\geq \kappa$.

It is clear that $\hat{1}, \hat{1}' \in B$, $X \subseteq B$ and $\gamma \in \mathrm{Con}\,\mathbf{B}$. First we prove that $(\hat{1}, \hat{1}') \notin \gamma$. We will actually show that $\hat{1}$ is in a singleton block of γ. Since γ is the join of the congruences $\gamma^{\mathbf{x}} = \mathrm{Cg}^{\mathbf{B}}(\mathbf{x}, \hat{f}(\mathbf{x}))$, where $\mathbf{x}$ runs over $N^\kappa - \{\hat{1}\}$, it is sufficient to prove that $\hat{1}$ is in a singleton block for each of these congruences.

For $i, j < \kappa$, $i \neq j$ define τ_{ij} by $\mathbf{u}\ \tau_{ij}\ \mathbf{v}$ iff $(u_i, u_j)\ \tau\ (v_i, v_j)$. This is a congruence of $\mathbf{B}$, because $\mathbf{u} \mapsto (u_i, u_j)$ is a (surjective) homomorphism from $\mathbf{B}$ to $\mathbf{T}$. Now let $\mathbf{x} \in N^\kappa - \{\hat{1}\}$ and set $\mathbf{x}' = \hat{f}(\mathbf{x})$. Then the set $I = \{i < \kappa \mid x_i = 0\}$ is non-empty. Let J be the complement of I in κ. If J is empty, then $\mathbf{x} = \mathbf{x}'$, so $\gamma^{\mathbf{x}} = 0_B$ and we are done. Otherwise, let $i \in I$ and $j \in J$ be arbitrary. Then $\mathbf{x}\ \tau_{ij}\ \mathbf{x}'$, because $(x_i, x_j) = (0,1)$ and $(x'_i, x'_j) = (0, 1')$. Thus $\gamma^{\mathbf{x}} \subseteq \tau_{ij}$ for every such i and j. By (3) therefore we have that if $\hat{1}\ \gamma^{\mathbf{x}}\ \mathbf{z}$ for some $\mathbf{z}$, then $z_i = z_j = 1$. This implies that $\mathbf{z} = \hat{1}$. Thus we have proved that the pair $(\hat{1}, \hat{1}')$ is not in γ.

Finally we must show that if $\psi \geq \gamma$ and $|X/(\psi|_X)| < \kappa$, then $(\hat{1}, \hat{1}') \in \psi$. For any ordinal $\lambda < \kappa$ define $\mathbf{x}^\lambda \in N^\kappa$ by $x_i^\lambda = 1$ if $i < \lambda$ and $x_i^\lambda = 0$ otherwise. $Y = \{\mathbf{x}^\lambda | \lambda < \kappa\}$ is a subset of X and $|Y| = \kappa$. Assume that $\psi \geq \gamma$ and that $|X/(\psi|_X)| < \kappa$. Then we can find ordinals $\mu < \nu < \kappa$ such that $(\mathbf{x}^\mu, \mathbf{x}^\nu) \in \psi$. Let $\mathbf{y}^\nu$ be the complement of $\mathbf{x}^\nu$ in N^κ, and let s be the polynomial in (1). Then $\hat{s}(\mathbf{x}^\nu, \mathbf{y}^\nu) = \hat{1}$, and the element $\mathbf{x} = \hat{s}(\mathbf{x}^\mu, \mathbf{y}^\nu)$ satisfies $x_i = 1$ if $i < \mu$ or $i \geq \nu$ while $x_i = 0$ if $\mu \leq i < \nu$. In particular, $\mathbf{x} \in N^\kappa - \{\hat{1}\}$. Further,

$$\hat{1} = \hat{s}(\mathbf{x}^\nu, \mathbf{y}^\nu)\ \psi\ \hat{s}(\mathbf{x}^\mu, \mathbf{y}^\nu) = \mathbf{x}\,.$$

Setting $\mathbf{x}' = \hat{f}(\mathbf{x})$ we obtain

$$\hat{1}' = \hat{f}(\hat{1}) \;\psi\; \hat{f}(\mathbf{x}) = \mathbf{x}' .$$

Finally, by the definition of γ, we have $(\mathbf{x}, \mathbf{x}') \in \gamma \leq \psi$. Hence

$$\hat{1} \;\psi\; \mathbf{x} \;\psi\; \mathbf{x}' \;\psi\; \hat{1}' .$$

This shows that Lemma 2.1 (3) holds proving that $\mathsf{V}(\mathbf{A})$ is indeed residually large. ∎

Definition 2.3 A pair $(1, 0)$ along with a polynomial $f(x)$ in an algebra $\mathbf{A}$ are said to constitute a *Residually Large Configuration* (or just RL configuration for short) if they satisfy the conditions of Theorem 2.2. In this case we say that the RL configuration occurs in $\mathbf{A}$.

We now provide ways to ensure that condition (3) of Theorem 2.2 is satisfied. First we translate it to the language of polynomials.

LEMMA 2.4 *Let $\mathbf{A}$ be any algebra and $0, 1, 1' \in A$. Let $\mathbf{T}$ be the subalgebra of $\mathbf{A}^2$ generated by the diagonal and $\{0, 1\}^2$. Suppose that $(0, 1') \in T$. Then condition* (3) *of Theorem 2.2 is equivalent to the following:*

(3′) *For every ternary polynomial p of $\mathbf{A}$ satisfying $p(0, 0, 1) = 1$ we have*

$$p(1, 1, 0) = 1 \quad \Longleftrightarrow \quad p(1', 1, 0) = 1 .$$

Assume further that $\mathbf{A}$ has unary polynomials e and c satisfying $e(0) = 0$, $e(1') = e(1) = 1$, $c(0) = 1$ and $c(1) = 0$. Then (3′) *is equivalent to*

(3″) *For every unary polynomial p of $\mathbf{A}$ satisfying $p(0) = 1$ we have*

$$p(1) = 1 \quad \Longleftrightarrow \quad p(1') = 1 .$$

PROOF. As $\mathbf{T}$ is generated by the diagonal of $\mathbf{A}^2$ and $\{0, 1\}^2$, every unary polynomial of $\mathbf{T}$ is of the form $\hat{p}(x, (0, 1), (1, 0))$, where p is a ternary polynomial of $\mathbf{A}$. By Mal'tsev's Lemma, $(1, 1)$ is in a singleton τ-class iff for every unary polynomial r of $\mathbf{T}$ we have $r(0, 1) = (1, 1) \Longleftrightarrow r(0, 1') = (1, 1)$.

Condition (3′) formulates this property, taking into consideration the form of the unary polynomials of $\mathbf{T}$. Thus, (3) $\Longleftrightarrow$ (3′) is proved.

Clearly, (3″) is weaker than (3′), since we can deem p to be a ternary polynomial which does not depend on its second and third variables. If $\mathbf{A}$ satisfies the additional hypotheses, then (3″) implies (3′). Indeed, let p be a ternary polynomial of $\mathbf{A}$ and let $q(x) = p(x, e(x), c(e(x)))$. Clearly, (3″) applied to q implies (3′) for p. ∎

There is another way to ensure that (3) of Theorem 2.2 is satisfied. Both of the main applications in the paper will use the following corollary. It suggests that if the congruence lattice of an algebra contains a pentagon then there is some chance of applying Theorem 2.2.

COROLLARY 2.5 *Let* $\mathbf{C}$ *be an algebra,* ρ, σ *congruences of* $\mathbf{C}$ *such that* $\rho < \sigma$ *and* R *a tolerance of* $\mathbf{C}$ *such that* $\sigma \cap R \subseteq \rho$*. Suppose that* $\mathbf{C}$ *has elements* $\tilde{0}$, $\tilde{1}$, $\tilde{0}'$, $\tilde{1}'$ *such that* $\tilde{0}\ R\ \tilde{1}$, $\tilde{0}'\ R\ \tilde{1}'$, $\tilde{0}\ \rho\ \tilde{0}'$, $\tilde{1}\ \sigma\ \tilde{1}'$*:*

$$\begin{array}{ccc} \tilde{0} & \rho & \tilde{0}' \\ R & & R \\ \tilde{1} & \sigma & \tilde{1}' \end{array} .$$

Then the elements $0 = \tilde{0}/\rho = \tilde{0}'/\rho$, $1 = \tilde{1}/\rho$ *and* $1' = \tilde{1}'/\rho$ *of* $\mathbf{A} = \mathbf{C}/\rho$ *satisfy* (3) *of Theorem 2.2 for the subalgebra* $\mathbf{T}$ *generated by the diagonal,* $\{0,1\}^2$ *and* $(0, 1')$*.*

Proof. Let $\mathbf{T}'$ be the subalgebra $\{(x/\rho, y/\rho) \mid x\ R\ y\}$ of $\mathbf{A}^2$, and let τ' be the restriction of the congruence $0_A \times (\sigma/\rho)$ to T'. Since T' contains the diagonal and the pairs $(0,1)$, $(1,0)$, and $(0,1')$ (since $\tilde{0}\ R\ \tilde{1}$ and $\tilde{0}'\ R\ \tilde{1}'$), and τ' contains $\mathrm{Cg}^{\mathbf{T}'}((0,1),(0,1'))$ (since $\tilde{0}\ \rho\ \tilde{0}'$ and $\tilde{1}\ \sigma\ \tilde{1}'$), it is sufficient to prove that $(1,1)$ is in a singleton τ'-class.

So suppose that $(1,1)\ \tau'\ (a,b)$ for some elements $a, b \in A$. Then $1 = a$, $1\ \sigma/\rho\ b$, and $(a,b) \in T'$. This latter condition implies that $a = x/\rho$ and $b = y/\rho$ for some $x, y \in C$, where $x\ R\ y$. Now $x\ \rho\ \tilde{1}$ and $y\ \sigma\ \tilde{1}$, so by transitivity and $\rho \le \sigma$ we have $x\ \sigma\ y$. Thus $R \cap \sigma \subseteq \rho$ implies that $x\ \rho\ y$. Hence $b = a = 1$, proving the statement. The reader is encouraged to give another proof, based on (3′) instead of (3). ∎

As an aid to the reader, we describe why it is 'natural' to want to prove a result like Theorem 2.2. When trying to apply tame congruence theory to the study of residually small equational classes (or to any other problem concerning finite algebras), one eventually wants to investigate the interaction between minimal sets for different quotients. In particular, it is common to need to compare the relationship between $\langle 0, \alpha\rangle$-minimal sets and $\langle \alpha, \beta\rangle$-minimal sets where $0 \prec \alpha \prec \beta$ in $\operatorname{Con} \mathbf{A}$, $I[0, \beta] = \{0, \alpha, \beta\}$, and $\operatorname{typ}(0, \alpha) \neq \operatorname{typ}(\alpha, \beta)$.

Choose $U \in \mathrm{M}_{\mathbf{A}}(0, \alpha)$ and $V \in \mathrm{M}_{\mathbf{A}}(\alpha, \beta)$. If $(a, b) \in \alpha|_U - 0_A$ and $(0, 1) \in \beta|_V - \alpha$, then $(a, b) \in \operatorname{Cg}(0, 1)$, so there is a Mal'tsev chain connecting a to b with polynomial images of $\{0, 1\}$. The polynomials involved can be assumed to have range in U. If the Mal'tsev chain has no trivial links and some polynomial $p(x)$ used in the creation of this chain satisfies $p(\beta|_V) \subseteq \alpha$, then p maps the trace of V which contains $\{0, 1\}$ into a trace of U without identifying 0 and 1.

This is a somewhat strange situation; for then p is a nonconstant function from the minimal algebra $\mathbf{A}|_N$, where $N \subseteq V$ is the $\langle \alpha, \beta\rangle$-trace containing $\{0, 1\}$, to the minimal algebra $\mathbf{A}|_T$, where T is a trace of U. What makes this a little strange is that $\mathbf{A}|_N$ and $\mathbf{A}|_T$ have different types. This does not immediately force the algebra $\mathbf{A}$ to generate a residually large equational class, but it often does. Therefore, when $\mathbf{A}$ generates a residually small equational class, it should not be uncommon in the scenario described above for each polynomial p which was used to create the Mal'tsev chain to satisfy $p(\beta|_V) \nsubseteq \alpha$. Hence, each polynomial maps V to a polynomially equivalent set contained in U. Since a is connected to b by polynomial images of $\{0, 1\}$, this implies that any two elements of a $\langle 0, \alpha\rangle$-trace of U are connected by a chain of overlapping $\langle \alpha, \beta\rangle$-traces contained in U.

A simple way that this could occur is depicted in Figure 2. A different possibility is depicted in Figure 3. In these figures, U has four elements, two of which reside in the unique $\langle 0, \alpha\rangle$-trace. The lines drawn between these four elements indicate the $\langle \alpha, \beta\rangle$-traces contained in U.

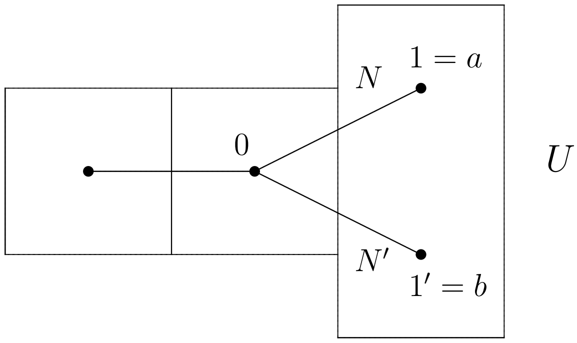

FIGURE 2: N and N' are equal modulo α.

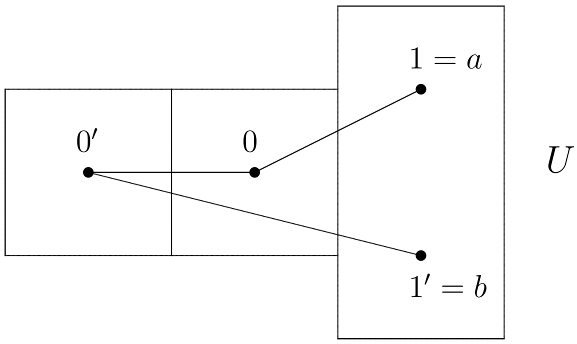

FIGURE 3: No $\langle \alpha, \beta \rangle$-traces are equal modulo α.

In Figure 2, N and N' are distinct $\langle \alpha, \beta \rangle$-traces which are "equal modulo α", which means that they project onto the same set in the quotient modulo α. In the situation we have been describing, it is not necessary that U contain any $\langle \alpha, \beta \rangle$-traces that are equal modulo α, as one can see by Figure 3, but avoiding a pair of traces which are equal modulo α is not so easy since any two α-related elements of U are connected by overlapping $\langle \alpha, \beta \rangle$-traces. (In particular, Figure 3 can only occur when $\mathrm{typ}(0, \alpha) = \mathbf{1}$.) So, it is natural to ask what happens when we have a pair of traces in U which are equal modulo α. Theorem 2.2 proves that if this situation occurs when $\mathrm{typ}(\alpha, \beta) \in \{\mathbf{2}, \mathbf{3}\}$ and the traces share a point in common, then $\mathbf{A}$ generates a residually large equational class. Here is why. When $\mathrm{typ}(\alpha, \beta) \in \{\mathbf{2}, \mathbf{3}\}$, then both $(0, 1)$ and $(1, 0)$ are $\mathbf{1}$-snags. The fact that N and N' are equal modulo α makes it easy to construct the polynomial f having the properties described in Theorem 2.2. Since the type is **2** or **3**, there is a polynomial $c(x)$ which switches 0 and 1. By Lemma 2.4, all that we need to do is show that for any unary polynomial p of $\mathbf{A}$ satisfying $p(0) = 1$ we have

$$p(1) = 1 \quad \Longleftrightarrow \quad p(1') = 1\,.$$

A glance at Figure 2 shows that this is obvious, since either $p(1) = 1 = p(0)$ or $p(1') = 1 = p(0)$ means that p is not 1-1 on U, so $p(\alpha|_U) \subseteq 0_A$. Thus, p identifies any two subsets of U which are equal modulo α; in particular, $p(1) = p(1')$ holds.

It is useful for one to keep in mind Figure 2 when deciding whether to apply Theorem 2.2.

3 Residually Small Minimal Equational Classes

Let $\mathcal{K}$ be a finite set of finite similar algebras. We will say that $\mathsf{HS}(\mathcal{K})$ is **semisimple** if all irreducible algebras in $\mathsf{HS}(\mathcal{K})$ are simple. In this section we will investigate residually small equational classes of the form $\mathcal{V} = \mathsf{HSP}(\mathcal{K})$ where $\mathsf{HS}(\mathcal{K})$ is semisimple. Our main result is the following.

THEOREM 3.1 *Let $\mathcal{K}$ be a finite set of finite similar algebras where $\mathsf{HS}(\mathcal{K})$ is semisimple. Assume that $\mathcal{V} = \mathsf{HSP}(\mathcal{K})$ contains no simple algebras of types* **1** *or* **5**. *Then $\mathcal{V}$ is residually small if and only if $\mathcal{V}$ is congruence modular.*

Only one implication of Theorem 3.1 is hard: it is the proof that residual smallness implies congruence modularity. The argument for the reverse implication goes as follows. If $\mathcal{V}$ is congruence modular and finitely generated, it follows from Theorem 10.15 of [1] that $\mathcal{V}$ is residually small if and only if $\mathcal{V}$ satisfies the commutator equation

$$\text{C1}: \qquad x \wedge [y, y] = [x \wedge y, y]\,.$$

The commutator equation C1 holds for any simple algebra and, as is proved in Theorem 8.1 of [1], C1 is inherited by finite subdirect products. Hence $\mathsf{HS}(\mathcal{K})$ satisfies C1. But now Theorem 8.1 of [1] can be invoked again, together with Remark 8.7 of [1], to conclude that C1 "goes up" from $\mathcal{K}$ to $\mathcal{V}$. Hence $\mathcal{V}$ is residually small if it is congruence modular.

Theorem 3.1 may be compared with Theorem 10.4 of [2] where it is proven that if $\mathcal{V}$ is a locally finite equational class with $\{\mathbf{1},\mathbf{5}\} \cap \operatorname{typ}\{\mathcal{V}\} = \emptyset$ and $\mathcal{V}$ is residually small, then $\mathcal{V}$ is congruence modular. Although the statements of these two theorems are similar, neither implies the other. Theorem 3.1 assumes that $\mathcal{V}$ is generated by simple algebras, while Theorem 10.4 of [2] does not. On the other hand, Theorem 10.4 of [2] assumes that no member of $\mathcal{V}$ has a type **1** or **5** quotient, while Theorem 3.1 only assumes that there are no type **1** or **5** simple algebras in $\mathcal{V}$.

Theorem 3.1 applies to minimal equational classes generated by simple algebras of types **2**, **3** or **4**. Any locally finite minimal equational class is generated by a strictly simple algebra and Theorem 14.8 of [2] proves that if two finite simple algebras generate the same equational class, then they have the same type. Hence, a minimal equational class generated by a strictly simple algebra $\mathbf{A}$ of type **2**, **3** or **4** will satisfy the hypotheses of Theorem 3.1 for $\mathcal{K} = \{\mathbf{A}\}$. It follows that such an equational class is residually small if and only if it is congruence modular. This explains the following corollary to Theorem 3.1.

COROLLARY 3.2 *Assume that* $\mathbf{A}$ *is a strictly simple algebra and that* $\mathcal{V} = \mathsf{HSP}(\mathbf{A})$ *is a residually small minimal equational class.*

(1) *If* $\operatorname{typ}\{\mathbf{A}\} = \{\mathbf{2}\}$, *then* $\mathcal{V}$ *is affine.*

(2) *If* $\operatorname{typ}\{\mathbf{A}\} \in \{\mathbf{3},\mathbf{4}\}$, *then* $\mathcal{V}$ *is congruence distributive.*

The only words of further explanation that we need to add are that when $\mathbf{A}$ is strictly simple and $\mathcal{V} = \mathsf{HSP}(\mathbf{A})$ is congruence modular, then it is a consequence of modular commutator theory that $\mathcal{V}$ is affine when $\mathbf{A}$ is abelian and congruence distributive when $\mathbf{A}$ is nonabelian. (See Theorem 12.1 (1) of [1].)

The abelian strictly simple algebras which generate minimal equational classes are classified in [4], and the classification shows that every such algebra generates a residually small equational class. Hence, the residual smallness hypothesis in Corollary 3.2 is redundant for conclusion (1). It is not redundant for conclusion (2) as is shown by the next result.

COROLLARY 3.3 (See [8]) *Let* $\mathbf{P}$ *be a finite bounded partial order and let* $\mathbf{A}$ *be order primal with respect to* $\mathbf{P}$. *Then* $\mathbf{A}$ *generates a congruence* *distributive equational class if and only if it generates a residually small equational class.*

PROOF. It is not hard to see that $\mathbf{A}$ is strictly simple of type $\mathbf{4}$. To see that the equational class $\mathcal{V} = \mathsf{HSP}(\mathbf{A})$ is minimal, let $0, 1 \in P$ be the bounds. Then the operation $b(x, y)$ defined by $b(x, y) = 0$ for $x \neq 1$ and $b(1, y) = y$ is a basic operation of $\mathbf{A}$ since it is compatible with the order on $\mathbf{P}$. These equations imply that 0 and 1 are terms with different interpretations in each nontrivial member of $\mathcal{V}$. Hence any nontrivial algebra in $\mathcal{V}$ has a nontrivial subuniverse consisting of the interpretations of constant terms. But $\mathbf{F}_{\mathcal{V}}(0) \cong \mathbf{A}$, which is simple, so the subalgebra of constants in any nontrivial algebra of $\mathcal{V}$ is isomorphic to $\mathbf{A}$. Hence $\mathbf{A}$ is embeddable into every nontrivial member of $\mathcal{V}$. This proves that $\mathcal{V}$ contains no proper nontrivial equational class. The rest follows from Corollary 3.2. ∎

Exercise 10.5 of [2] describes an eight-element order primal algebra which generates a residually large (non-congruence distributive) equational class.

LEMMA 3.4 *Let $\mathbf{C}$ be a finite algebra and δ, θ, α congruences of $\mathbf{C}$ such that $\delta < \theta$ and $\theta \wedge \alpha = 0_C$. Assume further that $C(\alpha, \theta; \delta)$ fails, and α is a minimal congruence of type $\mathbf{2}$, $\mathbf{3}$, or $\mathbf{4}$. Then the RL configuration occurs in $\mathbf{C}/\delta$.*

PROOF. We shall establish the conditions of Theorem 2.2 in $\mathbf{C}/\delta$, using Corollary 2.5. Let M be a $\langle 0_C, \alpha\rangle$-trace. We have $\neg C(\alpha, \theta; \delta)$, so in fact we must have $\neg C(M^2, \theta; \delta)$, since M^2 generates α. This means that there is a polynomial $p(x, \mathbf{y})$ of $\mathbf{C}$ such that

$$\begin{array}{ccc} \tilde{0} = p(a, \mathbf{u}) & \delta & p(a, \mathbf{v}) = \tilde{0}' \\ \alpha & & \alpha \\ \tilde{1} = p(b, \mathbf{u}) & \theta - \delta & p(b, \mathbf{v}) = \tilde{1}' \end{array} .$$

where $(a, b) \in M^2$ and $\mathbf{u}\ \theta\ \mathbf{v}$. Since $\theta \cap \alpha \subseteq \delta$, the conditions of Corollary 2.5 are satisfied.

We must have $\tilde{0} \neq \tilde{1}$. Indeed, assume otherwise. Then we get $\tilde{0}'\ \theta\ \tilde{1}'$ by transitivity, and thus $\alpha \wedge \theta = 0$ implies that $\tilde{0}' = \tilde{1}'$. But this shows that $\tilde{1}\ \delta\ \tilde{1}'$, which is a contradiction.

As M is a $\langle 0, \alpha\rangle$-trace, $\tilde{0} \neq \tilde{1}$ implies that $M' = p(M, \mathbf{u})$ is also a $\langle 0, \alpha\rangle$-trace. Since the type of $\langle 0, \alpha\rangle$ is $\mathbf{2}$, $\mathbf{3}$, or $\mathbf{4}$, we get that $(\tilde{1}, \tilde{0})$ is a $\mathbf{1}$-snag (because in these types any pair of different elements of a trace is a **1**-snag). By composing a polynomial inverse of $p(x, \mathbf{u}) : M \to M'$ with $p(x, \mathbf{v})$ we obtain a unary polynomial f of $\mathbf{C}$ mapping $\tilde{0}$ to $\tilde{0}'$ and $\tilde{1}$ to $\tilde{1}'$. These remarks show that conditions (1) and (2) of Theorem 2.2 are satisfied in $\mathbf{C}/\delta$ for the elements $0 = \tilde{0}/\delta = \tilde{0}'/\delta$, $1 = \tilde{1}/\delta$ and $1' = \tilde{1}'/\delta$. Corollary 2.5 shows that condition (3) is also satisfied. ∎

The rest of this section is devoted to the proof of Theorem 3.1. Let $\mathcal{K}$ be a fixed finite set of finite similar algebras where $\mathsf{HS}(\mathcal{K})$ is semisimple, and let $\mathcal{V}$ denote $\mathsf{HSP}(\mathcal{K})$. We assume that $\mathcal{V}$ contains no simple algebras of types **1** or **5**.

LEMMA 3.5 *The simple abelian algebras which are in $\mathsf{HS}(\mathcal{K})$ generate a congruence permutable equational class.*

PROOF. We have to show that the type **2** simple algebras in $\mathsf{HS}(\mathcal{K})$ generate a congruence permutable equational class. This follows from Corollary 6.9 of [4] which proves that a locally finite equational class generated by left nilpotent algebras is congruence permutable iff it contains no simple algebra of type **1**. Since $\mathcal{V}$ contains no simple algebra of type **1**, it follows that the equational class generated by the simple abelian algebras in $\mathsf{HS}(\mathcal{K})$ is congruence permutable. ∎

Our aim is to track down the irreducible algebras of $\mathcal{V}$.

LEMMA 3.6 *Let $\mathbf{C} \in \mathcal{V}$ be a finite subdirect product of members of $\mathsf{HS}(\mathcal{K})$ and let δ be a meet-irreducible congruence of $\mathbf{C}$ with upper cover θ. Suppose that $\operatorname{typ}(\delta, \theta) \in \{\mathbf{2}, \mathbf{3}, \mathbf{4}\}$ and there exists a $U \in \mathrm{M}_{\mathbf{C}}(\delta, \theta)$ which has empty tail. Then $\theta = 1_C$ and there exists an atom α in $\operatorname{Con} \mathbf{C}$ which is not below δ.*

PROOF. The algebra $\mathbf{C}$ is a subdirect product of simple algebras, let η_i $(i \in I)$ denote the projection kernels. We may assume that I is minimal for the property that $\bigwedge_{i \in I} \eta_i = 0_C$.

As all congruences of $\mathbf{C}|_U$ can be extended to $\mathbf{C}$, each $\eta_i|_U$ is either a coatom of $\operatorname{Con} \mathbf{C}|_U$ or else $\eta_i|_U = 1_U$. We have

$$\bigwedge_{i \in I} (\eta_i|_U) = (\bigwedge_{i \in I} \eta_i)|_U = 0_U\,.$$

In particular, the coatoms of $\mathrm{Con}\,\mathbf{C}|_U$ meet to 0_U.

Choose a subset $J \subseteq I$ which is minimal under inclusion for the property that $\bigwedge_{j\in J} \eta_j|_U = 0_U$. Clearly, every element j of J satisfies $\eta_j|_U < 1_U$. We claim that

$$\sigma = \bigwedge_{j\in J} \eta_j \le \delta\,.$$

Indeed, if this is not the case, then from the fact that δ is meet-irreducible we get that

$$\theta \le \delta \vee \bigwedge_{j\in J} \eta_j\,,$$

hence

$$\theta|_U \le \delta|_U \vee \bigwedge_{j\in J} (\eta_j|_U) = \delta|_U \vee 0_U = \delta|_U\,,$$

which is a contradiction.

By the results in Section 4 of [2] we know that in the type **3** and **4** cases U has two elements, while in the type **2** case $\mathbf{C}|_U$ is Mal'tsev and nilpotent. Thus, when the type is **3** or **4** every $j \in J$ satisfies $\eta_j|_U = 0_U$, so by the minimality of J we get that $J = \{j\}$ for some j. Thus $\sigma = \eta_j \le \delta$, so δ is indeed a maximal congruence. In fact, $\delta = \eta_j$ is a projection kernel.

If the type is **2**, then the set U contains no **2**-snags, and this implies, using elementary tame congruence theory, that the quotient $\langle \eta_j, 1_C\rangle$ is of abelian type for every $j \in J$. Thus $\mathbf{C}/\sigma$ is a subdirect product of the simple abelian algebras $\mathbf{C}/\eta_j$ for $j \in J$. But, as we proved in Lemma 3.5, the simple abelian algebras in $\mathsf{HS}(\mathcal{K})$ generate a congruence permutable equational class. This class contains $\mathbf{C}/\sigma$ now, and so the interval $I[\sigma, 1_C]$ in $\mathrm{Con}\,\mathbf{C}$ is a modular lattice, whose coatoms intersect to zero. Thus this interval is a relatively complemented lattice. Therefore θ has a complement in the interval $I[\delta, 1_C]$. As δ is meet-irreducible we see that δ must be a coatom, and $\theta = 1_C$.

Now we look for atoms that are not below δ. Suppose first that the type of $\langle \delta, \theta\rangle$ is nonabelian. Then, as we have seen, $\delta = \eta_j$ for some j. It is not possible that $\bigwedge_{i\in I-\{j\}} \eta_i = 0_C$ by the minimality of I. Thus there exists an atom α in $\mathrm{Con}\,\mathbf{C}$ below $\bigwedge_{i\in I-\{j\}} \eta_j$. Then α is not below η_j (since $\bigwedge_{i\in I} \eta_i = 0_C$). So we are done in this case.

In the abelian case recall that the coatoms of $\mathrm{Con}\,\mathbf{C}|_U$ intersect to zero, and as $\mathbf{C}|_U$ is Mal'tsev, $\mathrm{Con}\,\mathbf{C}|_U$ is a complemented modular lattice. The minimality hypothesis on J now implies that the set $\mathcal{B} = \{(\eta_j|_U) \mid j \in J\}$

is a maximal independent set of coatoms of $\mathrm{Con}\,\mathbf{C}|_U$. (Saying that $\mathcal{B}$ is an independent set of coatoms means that no two distinct subsets of $\mathcal{B}$ have the same meet.) The sublattice of $\mathrm{Con}\,\mathbf{C}|_U$ generated by $\mathcal{B}$ is a Boolean lattice which has the same height as $\mathrm{Con}\,\mathbf{C}|_U$. In particular, atoms in the Boolean sublattice generated by $\mathcal{B}$ are atoms in $\mathrm{Con}\,\mathbf{C}|_U$. The atoms of this Boolean sublattice are the elements

$$\alpha_j = \bigwedge_{i \in J-\{j\}} \eta_i|_U \,.$$

The atoms of a finite Boolean algebra join to the top, so there exists a $j \in J$ such that α_j is not below $\delta|_U$. Again, we can find an atom α in $\mathrm{Con}\,\mathbf{C}$ below $\bigwedge_{i \in I-\{j\}} \eta_j$, and as α is not below η_j, we get that $\alpha \vee \eta_j = 1_C$. Therefore $\alpha|_U \vee \eta_j|_U = 1_U$, so $\eta_j|_U \neq 1_U$ implies that $\alpha|_U > 0_U$. On the other hand, $\alpha|_U$ is below the atom α_j, so we have $\alpha|_U = \alpha_j$. Therefore α is not below δ and the proof is complete. ∎

This lemma has several important consequences.

COROLLARY 3.7 *If $\mathbf{C} \in \mathcal{V}$ is a finite subdirect product of members of $\mathsf{HS}(\mathcal{K})$, then the join of the atoms of* $\mathrm{Con}\,\mathbf{C}$ *is* 1_C.

PROOF. If the join of atoms is not 1_C, then there exists a maximal congruence δ of $\mathbf{C}$ which contains all atoms. This contradicts Lemma 3.6 (since the minimal sets for all quotients at the top have empty tail). ∎

COROLLARY 3.8 *Let $\mathbf{S}$ be a finite irreducible algebra in $\mathcal{V}$ with monolith μ. If $C(1_S, \mu; 0_S)$ holds, then $\mathbf{S}$ is simple.*

PROOF. We have $C(\mu, \mu; 0)$, so the monolith has abelian type. If this type is $\mathbf{1}$, then Lemma 6.1 of [4] shows that $\mathcal{V}$ contains a simple algebra of type $\mathbf{1}$, contradicting our assumption on $\mathcal{V}$. Thus the type of μ is $\mathbf{2}$.

Let $\mathbf{S} = \mathbf{C}/\delta$, where $\mathbf{C} \in \mathcal{V}$ is a finite subdirect product of elements of $\mathsf{HS}(\mathcal{K})$ and define θ by $\theta/\delta = \mu$. Then we have $C(1_C, \theta; \delta)$ in $\mathbf{C}$. But by Lemma 2.11 (6) of [4], the body of any $\langle \delta, \theta \rangle$-minimal set U is equal to the intersection with U of a single $(\delta : \theta)$-class. Therefore U has empty tail, and Lemma 3.6 implies that $\mathbf{S}$ is simple. ∎

COROLLARY 3.9 *Let* $\mathbf{C} \in \mathcal{V}$ *be a finite subdirect product of members of* $\mathsf{HS}(\mathcal{K})$ *and* δ *a meet-irreducible congruence of* $\mathbf{C}$ *with upper cover* θ. *Then for every atom* α *of* $\mathrm{Con}\,\mathbf{C}$ *we have* $\mathrm{typ}(0_C, \alpha) \in \{\mathbf{2}, \mathbf{3}, \mathbf{4}\}$. *If* $\theta < 1_C$, *then* α *is either below* δ *or not below* θ.

PROOF. Let α be any atom. There exists a projection kernel η_i which does not contain α, which implies that we have a perspectivity $\alpha/0_C \nearrow 1_C/\eta_i$. Because of our assumption on simple algebras in $\mathcal{V}$, the type of $\langle \eta_i, 1_C\rangle$ is **2**, **3** or **4**. We know that perspective prime quotients have the same type, so the first assertion is proved.

Suppose that an atom α is below θ, but not below δ. Then we have a perspectivity $\alpha/0_C \nearrow \theta/\delta$ as well. Thus θ/δ and $1_C/\eta_i$ are projective prime quotients, and $\mathrm{typ}(\theta, \delta) \in \{\mathbf{2}, \mathbf{3}, \mathbf{4}\}$.

The minimal sets for the quotient $1_C/\eta_i$ have empty tail. For types **1** and **5**, a minimal set may have empty tail with respect to one quotient, but not with respect to a perspective quotient. However, this cannot happen for types **2**, **3** and **4**. Indeed, if U is of type **3** or **4**, then U has empty tail if and only if $|U| = 2$. If U is of type **2**, then U has empty tail if and only if $\mathbf{C}|_U$ has a Mal'tsev polynomial. These properties are independent of the quotient for which U is minimal. Therefore the $\langle \delta, \theta\rangle$-minimal sets have empty tail in our case, and Lemma 3.6 implies that δ is a maximal congruence of $\mathbf{C}$, which contradicts the assumption that $\theta < 1_C$. ∎

COROLLARY 3.10 *If the equational class* $\mathcal{V}$ *is residually small, then every finite irreducible algebra of* $\mathcal{V}$ *is simple.*

PROOF. Assume not. Then there exists an algebra $\mathbf{C} \in \mathcal{V}$, which is a finite subdirect product of members of $\mathsf{HS}(\mathcal{K})$ and has a meet-irreducible congruence δ whose unique upper cover θ is not 1_C. We show that the conditions of Lemma 3.4 are satisfied for this δ, θ, and some atom α.

From Corollary 3.8 we see that $C(1_C, \theta; \delta)$ fails in $\mathbf{C}$. Corollary 3.7 shows that the join of all atoms is 1_C. Therefore there is an atom α such that $C(\alpha; \theta; \delta)$ fails. Of course α is not below δ (since then $C(\alpha; \theta; \delta)$ holds), and so Corollary 3.9 implies that α is not below θ, that is, $\theta \wedge \alpha = 0_C$. By the same Corollary, $\langle 0_C, \alpha\rangle$ has type **2**, **3**, or **4**. Thus all conditions of Lemma 3.4 are satisfied, and so $\mathcal{V}$ is residually large. This contradiction proves the statement. ∎

We can now show that if $\mathcal{V}$ is residually small, then it is congruence modular. In the light of Corollary 3.10, the following lemma finishes the proof of Theorem 3.1.

LEMMA 3.11 *If $\mathcal{V}$ is not congruence modular, then it contains a finite irreducible algebra which is not simple.*

PROOF. Suppose that $\mathcal{V}$ is not congruence modular. Then we can find a finite $\mathbf{C} \in \mathcal{V}$ whose congruence lattice contains a pentagon. That is, there are congruences $\alpha, \beta, \gamma \in \operatorname{Con} \mathbf{C}$ such that $\alpha \prec \beta$, $\alpha \vee \gamma = \beta \vee \gamma$, and $\alpha \wedge \gamma = \beta \wedge \gamma$. Choose a congruence $\delta \geq \alpha$ which is maximal for $\delta \not\geq \beta$. Necessarily, the congruence δ is meet irreducible, so $\mathbf{C}/\delta$ is irreducible. We claim that $\mathbf{C}/\delta$ is not simple. We shall prove this by contradiction.

Assume that $\mathbf{C}/\delta$ is simple. Then $\delta \prec 1_C$. Let $U \in \mathrm{M}_{\mathbf{C}}(\delta, 1_C)$. As $1_C/\delta$ and β/α are perspective quotients, we have that $\alpha|_U < \beta|_U$. This implies that $\alpha|_U, \beta|_U$ and $\gamma|_U$ generate a pentagon in $\operatorname{Con} \mathbf{C}|_U$. But clearly U is a single trace, since $\delta \prec 1_C$. If $\operatorname{typ}(\delta, 1_C) \in \{\mathbf{3}, \mathbf{4}\}$, then $|U| = 2$, so there aren't enough equivalence relations on U to form a pentagon in $\operatorname{Con} \mathbf{C}|_U$. If $\operatorname{typ}(\delta, 1_C) = \mathbf{2}$, then $\mathbf{C}|_U$ has a Mal'tsev polynomial, so $\operatorname{Con} \mathbf{C}|_U$ cannot contain a pentagon. The remaining possibilities, $\operatorname{typ}(\delta, 1_C) \in \{\mathbf{1}, \mathbf{5}\}$, are ruled out because this would mean that $\mathcal{V}$ contains a simple algebra $\mathbf{C}/\delta$ of type $\mathbf{1}$ or $\mathbf{5}$. This contradiction finishes the proof. ∎

4 Multitraces

Let $\mathbf{A}$ be a finite algebra and $\rho < \gamma$ be congruences of $\mathbf{A}$ with $\langle \rho, \gamma \rangle$ a tame quotient of type $\mathbf{2}$ or $\mathbf{3}$. We are interested in sets of the form $p(N, \ldots, N)$ where p is some polynomial of $\mathbf{A}$ and N is a $\langle \rho, \gamma \rangle$-trace. These sets are called $\langle \rho, \gamma \rangle$-*multitraces* of $\mathbf{A}$. Section 3 of [4] provides us with detailed knowledge of the induced structure on a multitrace in the case where $\rho = 0_A$. We need to consider other situations and so make the following definition.

Definition 4.1 A pair of congruences $\langle \rho, \gamma \rangle$ on a finite algebra $\mathbf{A}$ is called **stiff** if it is tame, and ρ is trivial on all (unary) polynomial images of $\langle \rho, \gamma \rangle$-traces.

Thus, if $\langle \rho, \gamma \rangle$ is stiff, and N is a $\langle \rho, \gamma \rangle$-trace, then for each $p(x) \in \mathrm{Pol}_1\mathbf{A}$ we have that either p is constant on N or $p(N)$ is polynomially isomorphic to N. Of course any tame pair of the form $\langle 0_A, \gamma \rangle$ is stiff; the following proposition provides other examples.

PROPOSITION 4.2 *Let* $\mathbf{A}$ *be a finite algebra and* $\langle \rho, \gamma \rangle$ *a tame pair of congruences of* $\mathbf{A}$.

(1) *Suppose that* $\langle \rho, \gamma \rangle$ *is stiff and that* N *is a* $\langle \rho, \gamma \rangle$*-trace. If the pair* (a, b) *is in the congruence generated by* N^2 *then there is a chain of overlapping* $\langle \rho, \gamma \rangle$*-traces connecting* a *to* b.

(2) *If* $\mathsf{HS}(\mathbf{A}^2)$ *is abelian, and* $\langle \rho, \gamma \rangle$ *is tame of type* **2** *with* ρ *strongly solvable, then* $\langle \rho, \gamma \rangle$ *is stiff.*

PROOF. Part (1) of this proposition is immediate from the definition. Theorem 7.4 of [9] shows that in part (2), the congruence ρ is strongly abelian and hence is trivial on any $\langle \rho, \gamma \rangle$-trace. Arguing as in Lemma 8.2 of [9] we conclude that $\langle \rho, \gamma \rangle$ is indeed stiff. ∎

The following theorem is an extension of some of the results of Section 3 of [4] which will suit our needs in this and subsequent sections.

THEOREM 4.3 *Let* $\mathbf{A}$ *be a finite algebra and* $\langle \rho, \gamma \rangle$ *a stiff pair of* $\mathbf{A}$ *of type* **2** *or* **3**. *Let* N *be a* $\langle \rho, \gamma \rangle$*-trace,* $p(\mathbf{x})$ *a polynomial of* $\mathbf{A}$ *and let* $T = p(N, \ldots, N)$. *Then* T *is an E-trace with respect to* γ *and*

(1) *If* $\langle \rho, \gamma \rangle$ *is of type* **2** *then*

   (a) $\mathbf{A}|_T$ *is term equivalent to* $(\mathbf{A}|_N)^{[k]}$ *for some* k *and hence is polynomially equivalent to a matrix power of a finite vector space. In this case, we say that the* **rank** *of* T *is* k.

   (b) *If* U *is a* $\langle \rho, \gamma \rangle$*-minimal set with body* B *which contains* N *and if* p *has exactly* k *variables then* $C = p(B, \ldots, B)$ *is an E-trace with respect to* $(\rho : \gamma)$ *and the induced structure on this set is isomorphic to* $(\mathbf{A}|_B)^{[k]}$. *Thus* $\mathbf{A}|_C$ *is Mal'tsev.*

   There is a binary polynomial $b(x, y)$ *of* $\mathbf{A}|_C$ *such that for any* $c \in C$, *the mapping* $b(x, c)|_T$ *is a bijection between* T *and the* $\gamma|_C$*-class which contains* c.

(2) *If $\langle \rho, \gamma \rangle$ is of type* **3** *then $\mathbf{A}|_T$ is a primal algebra and ρ is trivial on T. There exists an idempotent polynomial e such that T is the intersection of $e(A)$ with some γ-class, and T is also the union of some $(\rho : \gamma)|_{e(A)}$-classes.*

Proof. If $\langle \rho, \gamma \rangle$ is of type **3** then it is not hard to see that any two elements of T are contained in the image of N under a unary polynomial. Therefore, part (2) of this theorem follows from Lemmas 3.6 and 3.11 of [4]. The last statement follows from a careful examination of the proof of 3.6 part (3), using the fact that the body B of any $\langle \rho, \gamma \rangle$-minimal set U is a union of (two) $(\rho : \gamma)|_U$-classes (see Exercise 4.37 (5) of [2]).

To establish the first claim of (1) it suffices, by Lemma 3.8 of [4], to show that ρ is trivial on T. As ρ is trivial on N then $\mathbf{A}|_N$ is polynomially equivalent to a vector space. Let $0 \in N$ be the additive identity element with respect to this vector space. In this case it is not generally true that any two elements of T are contained in a polynomial image of N (unless N is one-dimensional). First we prove a special case:

Claim 1 *If $v_1, \ldots, v_m$ is a linearly independent set of vectors in N and $t(x_1, \ldots, x_m)$ is a polynomial of $\mathbf{A}$, then $t(0, \ldots, 0)\ \rho\ t(v_1, \ldots, v_m)$ if and only if $t(0, \ldots, 0) = t(v_1, \ldots, v_m)$.*

This can be proved by induction on m. The hypotheses of this theorem handle the case $m = 1$. Suppose that $t(0, \ldots, 0)\ \rho\ t(v_1, \ldots, v_m)$ and that the claim is valid for linearly independent sets with fewer than m elements. Let $\alpha_i = \rho \vee \mathrm{Cg}^{\mathbf{A}}(0, v_i)$ for $i \le m$. As the v_i are linearly independent, we get that $\alpha_m|_N \wedge \bigvee_{i<m} \alpha_i|_N = 0_N$, and using that $\langle \rho, \gamma \rangle$ is tame we see that $\alpha_m \wedge \bigvee_{i<m} \alpha_i = \rho$.

Let $a = t(0, \ldots, 0)$, $b = t(v_1, \ldots, v_{m-1}, 0)$ and $c = t(v_1, \ldots, v_m)$. Then $(b, c) \in \alpha_m \wedge \bigvee_{i<m} \alpha_i = \rho$. On the other hand, $b, c \in t(v_1, \ldots, v_{m-1}, N)$. The stiffness of $\langle \rho, \gamma \rangle$ therefore implies that $b = c$. Thus, $(a, b) \in \rho$, and then by the induction hypothesis we conclude that $a = b = c$, proving the claim.

Now the fact that ρ restricts trivially to T can be proved easily, by applying the claim to polynomials t obtained from p by appropriate linear substitutions.

To prove part (b) of (1), let U be a $\langle \rho, \gamma \rangle$-minimal set which contains N, let B be the body of U and let $e(x)$ be an idempotent polynomial of $\mathbf{A}$

with range U. If p is a k-variable polynomial (with k the rank of the multi-trace T), then by examining the proof of Lemma 3.8 of [4] we see that we may assume that T is coordinatized with respect to p and N via the coordinate polynomials $g_i(x)$, $i \leq k$, i.e.,

$$g_i(p(x_1, \ldots, x_k)) = x_i \text{ for all } x_j \in N \text{ and } i \leq k.$$

Replacing g_i by eg_i and $p(x_1, \ldots, x_k)$ by $p(e(x_1), \ldots, e(x_k))$ we may also assume that $p(e(x_1), \ldots, e(x_k)) = p(\mathbf{x})$ and $eg_i(y) = g_i(y)$ for all $x_j, y \in A$ and $i \leq k$.

Since A is finite then there is some number ℓ such that

$$t^{\ell}_{(i)}(x_1, \ldots, x_{i-1}, t^{\ell}_{(i)}(\mathbf{x}), x_{i+1}, \ldots, x_m) = t^{\ell}_{(i)}(\mathbf{x})$$

holds for all polynomials t of $\mathbf{A}$. We will use this number ℓ to define the following sequence of polynomials of $\mathbf{A}$:

$$\begin{aligned} p_0(\mathbf{x}) &= p(\mathbf{x}); \\ p_{i+1}(\mathbf{x}) &= p_i(x_1, \ldots, x_i, (g_{i+1}p_i)^{\ell-1}_{(i+1)}(\mathbf{x}), x_{i+2}, \ldots, x_k) \end{aligned}$$

for $i < k$.

Claim 2 *For* $0 \leq i \leq k$,

1. $p_i(\mathbf{x}) = p(\mathbf{x})$ *for all* $\mathbf{x} \in N$.

2. *For all* $\mathbf{b}$ *from* B, *and* $1 \leq j \leq k$, *the restriction of the polynomial* $g_j p_i(b_1, \ldots, b_{j-1}, x, b_{j+1}, \ldots, b_k)$ *to* U *is a permutation of* U.

3. *For* $1 \leq j \leq i$, *and* $\mathbf{b}$ *from* B, $g_j p_i(\mathbf{b}) = b_j$.

Note that for a given $i \leq k$, the second part of this claim follows from the first by using Corollary 4.34 of [2] and that $g_j(p(b_1, \ldots, b_{j-1}, x, b_{j+1}, \ldots, b_k))$ is a permutation of U whenever the b_i's are from N.

The remaining two parts of this claim can be proved by induction on i, with the base step being trivial. Suppose that the equalities mentioned in the two parts hold for p_i and consider p_{i+1}. By the induction hypothesis we can easily deduce that $(g_{i+1}p_i)^{\ell-1}_{(i+1)}(\mathbf{x}) = x_{i+1}$ for all $\mathbf{x}$ from N and so conclude that $p_{i+1}(\mathbf{x}) = p_i(\mathbf{x}) = p(\mathbf{x})$ for all $\mathbf{x}$ from N.

If $j \leq i$ and $\mathbf{b}$ is from B then $g_j p_{i+1}(\mathbf{b}) = g_j p_i(b_1, \ldots, b_i, b', b_{i+2}, \ldots, b_k)$ where $b' = (g_{i+1}p_i)^{\ell-1}_{(i+1)}(\mathbf{b})$ is some element from B. Then by induction we have that $g_j p_i(b_1, \ldots, b_i, b', b_{i+2}, \ldots, b_k) = b_j$ as required. For $j = i + 1$ we have that $g_{i+1}p_{i+1}(\mathbf{b})$ is the value of the idempotent polynomial

$$(g_{i+1}p_i)^{\ell}_{(i+1)}(b_1, \ldots, b_i, x, b_{i+2}, \ldots, b_k)$$

applied to b_{i+1}. Since the restriction of this polynomial to U is a permutation of U it follows that this polynomial maps b_{i+1} to b_{i+1} as required. Thus the claim is proved.

If we set $t(\mathbf{x}) = p_k(\mathbf{x})$ then the above claim shows that $C = t(B, \ldots, B)$ is coordinatized with respect to B^k with coordinate maps $g_i(x)$, $i \leq k$ and is contained in $p(B, \ldots, B)$. By Corollary 3.7 of [4] we conclude that C is an E-trace with respect to $(\rho : \gamma)$ since B is and that $\mathbf{A}|_C$ is polynomially isomorphic to $(\mathbf{A}|_B)^{[k]}$. Finally, since $p(B, \ldots, B)$ has size at most $|B|^k$ and the subset C contains exactly this many elements, then it follows that $C = p(B, \ldots, B)$. This also implies that ρ is trivial on C.

Since the type of $\langle \rho, \gamma \rangle$ is **2** then $\mathbf{A}|_B$ has a Mal'tsev polynomial. Using the induced structure on C we therefore see that there exists a polynomial $d(x, y, z)$ of $\mathbf{A}$ which is Mal'tsev on C. Let a be any member of T. Then using the fact that C is contained in a $(\rho : \gamma)$-class, a classical argument shows that the polynomial $b(x, y) = d(x, a, y)$ satisfies the conditions at the end of part (1) of this theorem. ∎

COROLLARY 4.4 *Let* $\mathbf{A}$ *be a finite algebra and* $\langle \rho, \gamma \rangle$ *a stiff pair of congruences of* $\mathbf{A}$ *of type* **2** *or* **3**. *Let* U *and* V *be* $\langle \rho, \gamma \rangle$*-multitraces.*

(1) U *and* V *are polynomially isomorphic if and only if they have the same size, or equivalently, in the type* **2** *case, they have the same rank.*

(2) *Let* $p(x)$ *be a unary polynomial of* $\mathbf{A}$. *There is some multitrace* W *contained in* U *such that* p *is a polynomial isomorphism between* W *and* $p(U)$.

(3) *If* $\rho \prec \beta \leq \gamma$, *then any two* β*-related elements of* U *are contained in a* $\langle \rho, \beta \rangle$*-trace.*

PROOF. Using the coordinate maps one can construct unary polynomials mapping U onto V and vice versa, hence we get (1). To see (2) notice that $p(U)$ is a multitrace of size less than or equal the size of U, and therefore U contains a polynomially isomorphic copy of $p(U)$ by (1). So, it suffices to verify this part, as well as part (3), under the assumption that $A = U$ and $\mathbf{A}$ is polynomially equivalent to a matrix power of a finite vector space or is primal. The latter case is very easy to handle, while the former can be taken care of using some elementary linear algebra. Part (3) follows from the classical observation that in a Mal'tsev algebra the congruence generated by a pair (a, b) is the set of pairs $(p(a), p(b))$, where p is a unary polynomial. ∎

We present an example which shows that even in rather nice situations a set obtained by applying an arbitrary polynomial to the body of a type **2** minimal set need not be coordinatizable with respect to the body. The previous theorem establishes coordinatizability under an additional assumption on the polynomial. Let $\mathbf{A}$ be the algebra with universe $\{0, 1, 2, 3, 4, 5, 6, 7\}$ and having as its only basic operation the following binary function:

$\cdot$	0	1	2	3	4	5	6	7
0	0	1	2	3	2	3	0	1
1	1	0	3	2	3	2	1	0
2	4	5	6	7	6	7	4	5
3	5	4	7	6	7	6	5	4
4	4	5	6	7	6	7	4	5
5	5	4	7	6	7	6	5	4
6	0	1	2	3	2	3	0	1
7	1	0	3	2	3	2	1	0

The congruence γ generated by $\{(0, 1)\}$ is a type **2** cover of 0_A and $\{0, 1, 2, 3\}$ is a $\langle 0_A, \gamma\rangle$-minimal set (having an empty tail). Observe that by applying $x \cdot y$ to this minimal set all 8 elements of $\mathbf{A}$ are obtained. Clearly A itself is not coordinatizable with respect to $\{0, 1, 2, 3\}$ since it does not have size an integral power of 4. We leave as an exercise to check the above details and also that $\mathbf{A}$ is an abelian algebra.

The results proved so far show that (in the type **2** or **3** case) every multitrace is a Mal'tsev E-trace. It is interesting to note that the converse also holds:

THEOREM 4.5 *Let $\langle 0_A, \gamma \rangle$ be a tame quotient of type* **2** *or* **3**. *If X is an E-trace with respect to γ and $\mathbf{A}|_X$ has a Mal'tsev polynomial, then X is a $\langle 0_A, \gamma \rangle$-multitrace.*

PROOF. Let U be a multitrace contained in X of maximal size. Since X is an E-trace then it can be connected by overlapping $\langle 0_A, \gamma \rangle$-traces. Thus, if $U \neq X$ then there is a trace N contained in X which contains elements from both U and from $X - U$.

If $d(x, y, z)$ is a polynomial of $\mathbf{A}$ whose restriction to X is Mal'tsev and $a \in N \cap U$ then the set $d(U, a, N)$ is a multitrace which contains both U and N, contradicting the maximality of U. Therefore, X is a multitrace. ∎

An easy corollary of this theorem is the result of Werner on functionally complete algebras.

COROLLARY 4.6 (see [11]) *Let* $\mathbf{A}$ *be a finite algebra in a congruence permutable equational class.* $\mathbf{A}$ *is functionally complete if and only if* Con $\mathbf{A}^2$ *is isomorphic to* $\mathbf{2}^2$, *the square of the two element lattice.*

PROOF. Clearly if Con $\mathbf{A}^2$ is isomorphic to $\mathbf{2}^2$ then $\mathbf{A}$ is a simple algebra. Also, $\mathbf{A}$ cannot be abelian or else the congruence of $\mathbf{A}^2$ generated by the diagonal of A would produce a nontrivial congruence distinct from the kernels of the two projection homomorphisms. Thus $\mathbf{A}$ is a simple nonabelian algebra in a congruence permutable equational class, and hence has type **3**.

From the previous theorem it follows that A itself is a multitrace and then by Theorem 4.3 it follows that $\mathbf{A}|_A$ is a primal algebra, i.e., $\mathbf{A}$ is functionally complete. ∎

For the remainder of this section let $\mathbf{A}$ be a finite algebra and $\langle 0_A, \gamma \rangle$ a tame quotient of $\mathbf{A}$ of type **2** or **3**.

LEMMA 4.7 *Let N be a $\langle 0_A, \gamma \rangle$-trace and let a and b be distinct elements of A contained in the same γ-class as N. Let 0 be any element from N. There is an idempotent unary polynomial of* $\mathbf{A}$ *which maps the γ-class of N onto N and which separates a and b. In particular, if M is any other $\langle 0_A, \gamma \rangle$-trace contained in the same γ-class as N then there is an idempotent polynomial π which maps M onto N.*

If the type of $\langle 0_A, \gamma \rangle$ is **2**, *say the $\langle 0_A, \gamma \rangle$-traces are polynomially equivalent to an* $\mathbf{F}$*-vector space for some finite field* $\mathbf{F}$, *and $f(x)$ is any other* *polynomial which maps M onto N then there is some $v \in N$ and $\lambda \in F$ with $f(x) = \lambda\pi(x) + v$ for all $x \in M$.*

Proof. Since $\langle 0_A, \gamma\rangle$ is of type **2** or **3** then there is a polynomial $+$ of **A** whose restriction to N describes a group operation on N having neutral element 0.

Let $\mathcal{F}$ be the set of unary polynomials of **A** which map the γ-class of N into N. Let $\mathcal{G} \subseteq \mathcal{F}$ be the subset of polynomials which are one-to-one on N and let $\mathcal{H} \subseteq \mathcal{F}$ be the subset of polynomials which separate a and b. By [2], both $\mathcal{G}$ and $\mathcal{H}$ are nonempty. But then $\mathcal{G} \cap \mathcal{H}$ is nonempty, since if $g \in \mathcal{G} - \mathcal{H}$ and $h \in \mathcal{H} - \mathcal{G}$, then $g + h \in \mathcal{G} \cap \mathcal{H}$.

It is clear from the definitions that $\mathcal{G} \cap \mathcal{H}$ is closed under composition, and so by iterating any $f \in \mathcal{G} \cap \mathcal{H}$ one can produce an idempotent polynomial which maps the γ-class of N into N, separates a and b, and is one-to-one on N.

If M is another trace contained in the same γ-class as N, then we may apply the preceding result to distinct elements $a, b \in M$ to obtain an idempotent polynomial π which maps M onto N.

Now, suppose that the type of $\langle 0_A, \gamma\rangle$ is **2** and that $f(x)$ is some polynomial of **A** which maps M into N. Let g be a polynomial inverse of $\pi : M \to N$. Then the map $f(g(x))$ maps N into N and so is of the form $\lambda x + v$ for some $\lambda \in F$ and $v \in N$. From this we see that $f(x) = \lambda\pi(x) + v$ for all $x \in M$ as required. ∎

COROLLARY 4.8 *If M and N are distinct $\langle 0_A, \gamma\rangle$-traces then they have at most one element in common. If M and N have exactly one element in common then there is an idempotent polynomial which maps the γ-class containing N onto N and is constant on M.*

Proof. If M and N are distinct overlapping traces, then we can choose $0 \in M \cap N$ and $a \in M - N$. We shall use Lemma 4.7 to build an idempotent polynomial e which maps the γ-class of a onto N and which maps a (and hence all of M) onto 0. This polynomial is thus both constant and one-to-one on $M \cap N$ and so $|M \cap N| = 1$.

To find such a polynomial, apply Lemma 4.7 to the elements a and 0 to obtain an idempotent polynomial e_1 which maps the γ-class of N onto N and which separates a and 0 (and so $e_1(M) = N$). Let $e_1(a) = b$. Apply the lemma once more to the elements a and b to obtain an idempotent polynomial e_2 which maps the γ-class of N onto N and which separates a and b. Let $c = e_2(a)$ and note that $c \neq b$.

If the type of $\langle 0_A, \gamma\rangle$ is **3** then N has exactly two elements and so $c = 0$ and we are done. On the other hand, if the type is **2** then by the last part of Lemma 4.7 we see that there is some $\lambda \in \mathbf{F}$ with $e_2(x) = \lambda(e_1(x))$ for all $x \in M$. This implies that $c = \lambda b$.

A straightforward calculation shows that the polynomial

$$\frac{\lambda}{\lambda - 1}(e_1(x) - e_2(x)) + e_2(x)$$

maps 0 and a to 0 and is one-to-one on N and so some iterate of it will provide us with the desired idempotent polynomial. ∎

LEMMA 4.9 *Let X be a $\langle 0_A, \gamma\rangle$-multitrace and M a $\langle 0_A, \gamma\rangle$-trace which is not contained in X but which has nonempty intersection with X. Then there is an idempotent polynomial $p(x)$ which maps the γ-class of X onto X and which is constant on M.*

Proof. Let $e(x)$ be an idempotent polynomial such that X is the intersection of a γ-class with the range of e and let 0 be an element in the intersection of M with X. We may assume that e is not constant on M, otherwise we may set $p = e$.

Since M is a trace then so is $e(M)$ and since 0 lies in both M and $e(M)$ then by Corollary 4.8 there is an idempotent polynomial $f(x)$ which maps the γ-class of $e(M)$ onto $e(M)$ and which is constant on M.

Let $+$ be a polynomial whose restriction to X provides an abelian group operation with additive identity element 0 and let $h(x) = e(x) + f(x) - fe(x)$. We claim that h is constant on M and is the identity map on X and so some suitable iterate of h will provide the polynomial p that we are after.

If $m \in M$ then

$$h(m) = e(m) + f(m) - fe(m) = e(m) + 0 - e(m) = 0$$

since f maps all of M to 0 and is the identity on $e(M)$. Also, if $a \in X$ then

$$h(a) = e(a) + f(a) - fe(a) = a + f(a) - f(a) = a$$

since e is the identity map on X. ∎

COROLLARY 4.10 *Let V and W be distinct $\langle 0_A, \gamma\rangle$-multitraces such that V is not contained in W but their intersection is nonempty. Then there is an idempotent polynomial which maps the γ-class of V onto W and which is not one-to-one on V.*

PROOF. All that is needed is to find some $\langle 0_A, \gamma\rangle$-trace contained in V and which contains elements from $V \cap W$ and from $V - W$. Since any two elements in a multitrace can be connected by a chain of overlapping traces then such a trace does indeed exist. The previous corollary provides an idempotent map which maps the γ-class of V onto W and which is constant on this trace. This map cannot be one-to-one on V. ∎

COROLLARY 4.11 *Let V and W be $\langle 0_A, \gamma\rangle$-multitraces. Then $V \cap W$ is either empty or a multitrace.*

PROOF. Let $X = V \cap W$ and assume that X is not empty and not a multitrace. Replacing V and W with smaller multitraces containing X if necessary, we may assume that no proper subset of V or W is a multitrace containing X. We may also assume that $|V| \leq |W|$. By the previous corollary there is an idempotent polynomial $e(x)$ which maps $V \cup W$ onto W and which is not one-to-one on V. This is an impossibility since the multitrace $e(V) \subset W$ is strictly smaller than W and contains X. We conclude that $X = V \cap W$ is a multitrace. ∎

5 Parallel Multitraces

Throughout this section let $\mathbf{A}$ be a finite algebra and $\langle \rho, \gamma\rangle$ a stiff pair of congruences of type **2** or **3**.

Definition 5.1 Two subsets X and Y of an algebra $\mathbf{B}$ are said to be **quasi-parallel** if for all polynomials $p(x)$ of $\mathbf{B}$ we have that p is constant on X if and only if it is constant on Y. We write $X \wr\wr Y$ in this case.

Clearly, $\wr\wr$ is an equivalence relation on the subsets of any algebra. We now list some other important properties of this relation. To formulate this result, let us introduce a new concept.

Recall from Chapter 2 of [2] that subsets X and Y of an algebra $\mathbf{A}$ are polynomially isomorphic if there exist polynomials f, f' of $\mathbf{A}$ such that $f|_X : X \to Y$ and $f'|_Y : Y \to X$ are inverse bijections. It is well known and easy to show that, when $\mathbf{A}$ is finite, to determine if $X, Y \subseteq A$ are polynomially isomorphic it is enough to find f and f' such that $f|_X$ and $f'|_Y$ are bijections, since from these polynomials one can construct inverse (polynomial) bijections between X and Y by composition and iteration. We define subsets X and Y of A to be *E-isomorphic* if they are polynomially isomorphic via idempotent polynomials. That is, there are $e, e' \in \mathrm{E}(\mathbf{A})$ such that $e|_X : X \to Y$ and $e'|_Y : Y \to X$ are inverse bijections. As in the case of ordinary polynomial isomorphism, E-isomorphism can be established from partial information when we are dealing with finite algebras. In particular, if $f|_X : X \to Y$ and $f'|_Y : Y \to X$ are polynomial bijections, and $f(Y) = Y$, then idempotents $e, e' \in \mathrm{E}(\mathbf{A})$ witnessing E-isomorphism can be constructed from f and f' by composition and iteration.

LEMMA 5.2 *Let U and V be $\langle \rho, \gamma \rangle$-multitraces.*

(1) *If X and Y are subsets of A and $p(x)$ is a polynomial of $\mathbf{A}$, then $X \wr\wr Y$ implies $p(X) \wr\wr p(Y)$.*

(2) *If V is properly contained in U, then U is not quasi-parallel to V.*

(3) *Suppose that either*

   (A) *U and V are quasi-parallel and lie in the same γ-class; or*

   (B) *the type of $\langle \rho, \gamma \rangle$ is **3** and U and V are equal modulo $(\rho : \gamma)$.*

   Then U and V are E-isomorphic. If $u_1 \in U$ corresponds to $v_1 \in V$ and $u_2 \in U$ corresponds to $v_2 \in V$ under this E-isomorphism, then for every unary polynomial p of $\mathbf{A}$ we have that $p(u_1) = p(u_2)$ if and only if $p(v_1) = p(v_2)$. In particular, $U \wr\wr V$ holds also in case (B).

(4) *If the type of $\langle \rho, \gamma \rangle$ is **2** and U and V are of rank k and such that their intersection properly contains a multitrace of rank $k - 1$, then $U \wr\wr V$.*

(5) *If U and V are distinct and quasi-parallel, then either they are disjoint or the RL configuration occurs in $\mathbf{A}$. If $\rho = 0_A$ then U and V must be disjoint.*

PROOF. Part (1) follows immediately from the definition of $\wr\wr$.

To prove (2), assume that V is properly contained in U. By Theorem 4.3 we know that both U and V are E-traces with respect to γ. From this and $V \subset U$ we know that there is an idempotent e such that $e(U) = V$ and e is not one-to-one on U. Again by Theorem 4.3 we know that $\mathbf{A}|_U$ is primal or polynomially equivalent to a matrix power of a vector space. In either case there is a binary polynomial $x - y$ of $\mathbf{A}$ whose restriction to U is an abelian group subtraction operation. Subtraction can be used to construct the polynomial $p(x) = x - e(x)$, which is constant on V and not on U. Therefore U is not quasi-parallel to V.

For part (3), suppose first that U and V are quasi-parallel and lie in the same γ-class. As we have mentioned, Theorem 4.3 guarantees an idempotent polynomial e such that U is the intersection of $e(A)$ with some γ-class. Since U and V lie in the same γ-class, then $e(V) \subseteq U$ and so by part (1) we get that $e(V)$ and $e(U) = U$ are quasi-parallel. By part (2), we must have $U = e(V)$. By reversing the roles of U and V we can find an idempotent polynomial e' such that $e'(V) = e'(U) = V$. Thus U and V are E-isomorphic.

Suppose next that the type of $\langle \rho, \gamma \rangle$ is $\mathbf{3}$ and U and V are equal modulo $(\rho : \gamma)$. As $\mathbf{A}|_U$ and $\mathbf{A}|_V$ are both primal, $(\rho : \gamma)$ is trivial on both of these sets. Thus the $(\rho : \gamma)$-classes establish a one-to-one correspondence between U and V. By Theorem 4.3 there is an idempotent polynomial $e(x)$ of $\mathbf{A}$ such that U is the intersection of $e(A)$ with some γ-class and U is the union of $(\rho : \gamma)|_{e(A)}$-classes. From this it follows that if $a \in U$ and $b\ (\rho : \gamma)\ a$ then $e(b) = a$ and so $v\ (\rho : \gamma)\ e(v)$ for all $v \in V$. By reversing the roles of U and V in this argument we can find an idempotent polynomial e' which maps both U and V onto V. Thus U and V are E-isomorphic in this case too.

Let $e' : U \to V$ be an idempotent which maps U onto V. If $p(x)$ is any polynomial of $\mathbf{A}$, then the sets $p(U)$ and $p(V)$ are multitraces which satisfy (A) or (B), respectively, so we can find an idempotent polynomial $f : p(U) \to p(V)$ mapping $p(U)$ onto $p(V)$. Let $x - y$ be a binary polynomial of $\mathbf{A}$ whose restriction to $p(V)$ acts as subtraction with respect to some abelian group structure on $p(V)$. Now consider the polynomial $h(x) = pe'(x) - fp(x)$. For $v \in V$ we have that $h(v) = p(v) - p(v) = 0$ and so h is constant on V. In case (A) we have that h is constant on U since $U \wr\wr V$, and in case (B) we have that if $a, b \in U$ then $h(a)\ (\rho : \gamma)\ h(e'(a)) = h(e'(b))\ (\rho : \gamma)\ h(b)$. As $(\rho : \gamma)$ is trivial on $p(V)$ then we conclude that $h(a) = h(b)$. In both cases, we must have that h is constant on U. Thus for any $u_1, u_2 \in U$ and for

$v_i \stackrel{\text{def}}{=} e'(u_i)$ we have

$$p(v_1) - fp(u_1) = h(u_1) = h(u_2) = p(v_2) - fp(u_2),$$

from which it follows that $p(u_1) = p(u_2)$ if and only if $p(v_1) = p(v_2)$.

Now we establish part (4). Note that if $\rho = 0_A$, then we must have $U = V$ by Corollary 4.11, but not in general. Let $\langle \rho, \gamma \rangle$ be of type **2** and let n be the size of any $\langle \rho, \gamma \rangle$-trace. Let U and V be multitraces of rank k whose intersection properly contains a multitrace M of rank $k-1$. Assume that p is a unary polynomial which is constant on V. Then p is constant on a subset of U which properly contains M. Now U is the union of n disjoint multitraces which are quasi-parallel to M, and so p will be constant on each of these multitraces. Thus $p(U)$, which is a multitrace, can have size 1 or n, but since p is constant on a set which properly contains M, it must be that $p(U)$ has size 1. Thus if p is constant on V it is constant on U. The symmetric fact establishes that U is quasi-parallel to V.

To see why part (5) is true, let U and V be quasi-parallel and suppose that they are distinct but have a nonempty intersection. Let 0 lie in the intersection and let 1 be an element from $U - V$. Since U and V lie in the same γ-class they are E-isomorphic via idempotent polynomials e and e'. In particular, $|U| = |V|$.

If $\rho = 0_A$, then by Corollary 4.10 there is an idempotent polynomial f which maps $U \cup V$ onto V and which is not one-to-one on U. Then by part (1) of this lemma $f(U)$ is a multitrace properly contained in, but quasi-parallel to, the multitrace $f(V) = V$, contradicting part (2) of this lemma. This establishes the second remark of part (5).

If $\rho \neq 0_A$, then we will show that $(1, 0)$ along with e' constitute an RL configuration in $\mathbf{A}$. Since $\mathbf{A}|_U$ supports an abelian group operation with neutral element 0, it follows that $(1, 0)$ is a **1**-snag of $\mathbf{A}$. Since 0 lies in the intersection of U and V then $e'(0) = 0$. Also, for $1' \stackrel{\text{def}}{=} e'(1) \in V - U$ we have $e(1') = e(1) = 1$, and for the polynomial $c(x) = 1 - x$ we have $c(0) = 1$ and $c(1) = 0$. According to Lemma 2.4, to show that our data constitute an RL configuration we need only to verify that for all polynomials $p(x)$, if $p(0) = 1$ then $p(1) = 1$ if and only if $p(1') = 1$. Here we apply part (3) of this lemma to $(u_1, u_2) = (0, 1)$ and $(v_1, v_2) = (0, 1')$. Part (3) tells us that $p(0) = p(1)$ if and only if $p(0) = p(1')$. Hence $p(1) = 1$ if and only if $p(1') = 1$, since $p(0) = 1$. ∎

A different notion of parallelism is given in the next definition.

Definition 5.3 Let T be a tolerance of an algebra $\mathbf{B}$ and let X and Y be two subsets of $\mathbf{B}$. X and Y are said to be T-parallel (denoted: $X \parallel_T Y$) if they are polynomially isomorphic and the pair (X, Y) lies in the transitive closure of the relation

$$\{(t(X,r), t(X,s)) \mid t \in \operatorname{Pol} \mathbf{B} \text{ and } (r,s) \in T\}.$$

When T is the relation B^2 then we write $\parallel$ in place of $\parallel_T$ and use the term "parallel" in place of "T-parallel".

Thus $\parallel_T$ is the transitive closure of the T-twin relation between subsets. Clearly, this is an equivalence relation, and if T is generated (as a tolerance) by the pair (c,d) then we can always assume that $\{r,s\} = \{c,d\}$ in the above definition. We leave it as an exercise to show that if $T \le (0_B : X^2)$, then $X \parallel_T Y$ implies that $X \wr\wr Y$. In particular, if $\mathbf{B}$ is abelian then $X \parallel Y$ implies that $X \wr\wr Y$.

We shall need the following addition to part (3) of Lemma 5.2.

LEMMA 5.4 *If the type of $\langle \rho, \gamma \rangle$ is* $\mathbf{2}$ *and U and V are two quasi-parallel $\langle \rho, \gamma \rangle$-multitraces which are equal modulo some congruence $\varepsilon \le (\rho : \gamma)$, then there is a multitrace W which is ε-parallel to U and E-isomorphic to V.*

PROOF. By Theorem 4.3 there is some set C containing U which is an E-trace with respect to $(\rho : \gamma)$ and such that $\mathbf{A}|_C$ is polynomially isomorphic to $(\mathbf{A}|_B)^{[k]}$ for some $\langle \rho, \gamma \rangle$-body B and natural number k. Let e be an idempotent polynomial such that C is the intersection of $e(A)$ with some $(\rho : \gamma)$-class.

If we apply e to the multitrace V then we obtain a multitrace $e(V)$ contained in C and quasi-parallel to $e(U) = U$. Let W be the $\gamma|_C$-class which contains $e(V)$. By Theorem 4.3 there is some binary polynomial $b(x,y)$ of $\mathbf{A}|_C$ such that for any $c \in C$, the mapping $b(x,c)|_U$ is a bijection between U and the $\gamma|_C$-class containing c. This shows that the set W is a multitrace which is ε-parallel to U and hence is quasi-parallel to it. Then by Lemma 5.2 (2) the multitraces $e(V)$ and W must be equal since $e(V) \subseteq W$ and $e(V) \wr\wr e(U) = U \wr\wr W$.

To prove that V is E-isomorphic to $W = e(V)$ we need to prove that there is a polynomial inverse to $e(x)|_V$ mapping W back onto V. From the remarks following Definition 5.1 we know that it suffices to show that there is some polynomial bijection from W to V, and therefore (in view of Corollary 4.4 (1)) we only need to show that $|V| = |W|$. From the previous paragraph we gather that $|U| = |W| = |e(V)| \leq |V|$. By reversing the role of U and V, we can conclude that in fact U, V and W must have the same size, and hence the multitraces V and W are E-isomorphic. ∎

Next we shall associate congruences with the parallel and quasi-parallel relation. Let W be a $\langle \rho, \gamma \rangle$-multitrace and $\vec{W} = (w_1, \ldots, w_r)$ a listing of the distinct elements of W. Let $\mathbf{B}$ be the subalgebra of $\mathbf{A}^r$ generated by $\vec{W}$ and the diagonal. Then

$$B = \{(t(w_1), \ldots, t(w_r)) \mid t \in \mathrm{Pol}_1 \mathbf{A}\} \ ,$$

hence if a tuple $\mathbf{u}$ in A^r belongs to B then $\{u_i \mid i < r\}$ is a $\langle \rho, \gamma \rangle$-multitrace.

For $\mathbf{u}$ and $\mathbf{v}$ in B and $\varepsilon \in \mathrm{Con}\,\mathbf{A}$, let us define $\mathbf{u}$ and $\mathbf{v}$ to be ε-parallel (and write $\mathbf{u} \parallel_\varepsilon \mathbf{v}$) if this pair lies in the transitive closure of the relation

$$\left\{(\hat{t}(\vec{W}, \hat{c}), \hat{t}(\vec{W}, \hat{d})) \mid t \text{ a polynomial of } \mathbf{A} \text{ and } (c, d) \in \varepsilon\right\}.$$

Further, let us define $\mathbf{u}$ and $\mathbf{v}$ to be quasi-parallel (and write $\mathbf{u} \wr\wr \mathbf{v}$) if for every polynomial $p(x)$ of $\mathbf{B}$, $p(\mathbf{u})$ is a constant sequence if and only if $p(\mathbf{v})$ is.

LEMMA 5.5 *With the notation above, we have the following:*

(1) *The* $\parallel_\varepsilon$ *relation is the congruence of* $\mathbf{B}$ *generated by* $\left\{(\hat{c}, \hat{d}) \mid c \mathrel{\varepsilon} d\right\}$.

(2) *The* $\wr\wr$ *relation is the largest congruence of* $\mathbf{B}$ *for which the diagonal is a union of congruence classes.*

(3) *If* $\varepsilon \leq (\rho : \gamma)$ *then* $\parallel_\varepsilon \subseteq \wr\wr$.

(4) *If* $\mathbf{u} \wr\wr \mathbf{v}$ *then* $\{u_i \mid i < r\}$ *and* $\{v_i \mid i < r\}$ *are quasi-parallel multitraces of* $\mathbf{A}$. *Conversely, if the multitraces* $\{u_i \mid i < r\}$ *and* $\{v_i \mid i < r\}$ *are quasi-parallel and have size* r, *and there are idempotent polynomials* $e_0(x)$ *and* $e_1(x)$ *of* $\mathbf{A}$ *with* $e_0(v_i) = u_i$ *and* $e_1(u_i) = v_i$ *for* $i < r$, *then* $\mathbf{u} \wr\wr \mathbf{v}$. *If* $\varepsilon \leq (\rho : \gamma)$ *and* $\mathbf{u} \parallel_\varepsilon \mathbf{v}$ *in* $\mathbf{B}$ *then* $\{u_i \mid i < r\}$ *and* $\{v_i \mid i < r\}$ *are* ε*-parallel in* $\mathbf{A}$.

(5) *If $\rho = 0_A$ or $\mathbf{A}$ avoids the RL configuration then the congruence $\wr\wr$ intersects trivially with the projection kernels of $\mathbf{B}$.*

PROOF. Using the definitions of $\mathbf{B}$, $\|_\varepsilon$ and $\wr\wr$ it is straightforward to verify (1) and (2).

To show (3) it is sufficient to prove by (1) that if $c\ \varepsilon\ d$ are elements of A, then $\hat{c} \wr\wr \hat{d}$. Let p be a unary polynomial of $\mathbf{B}$. There is some polynomial $r(x, y)$ of $\mathbf{A}$ such that $p(x) = \hat{r}(x, \vec{W})$ for all $x \in B$, and so having p constant on $\hat{c}$ is equivalent to asserting that the polynomial $r(c, x)$ is constant on W. Since W is contained in a γ-class and (c, d) are in $(\rho : \gamma)$ then $r(d, x)$ must map W into some ρ-class. By the stiffness of $\langle \rho, \gamma \rangle$ it follows that this polynomial is actually constant on W. This is equivalent to having $p(\hat{d})$ constant.

The first part of (4) is immediate from the definitions, since if $\mathbf{u} \in B$ then $\{u_i \mid i < r\}$ is a multitrace. For the converse, the first thing to note is that since $\mathbf{u}$ is a vector of r distinct elements in B then there is a polynomial isomorphism in $\mathbf{A}$ which sends w_i to u_i for $i < r$. The inverse to this map can be used to show that $\mathbf{B}$ is also generated by the diagonal and the vector $\mathbf{u}$. Now, suppose that $p(x)$ is a polynomial of $\mathbf{B}$ with $p(\mathbf{u})$ constant. Since $\mathbf{u}$ along with the diagonal generates $\mathbf{B}$ then there is some polynomial $r(x, y)$ of $\mathbf{A}$ such that $p(x) = \hat{r}(x, \mathbf{u})$ for all $x \in B$. We may assume that $r(x, e_0(y)) = r(x, y)$ for all x and y in A. Since $p(\mathbf{u})$ is constant, then the polynomial $d(x) = r(x, x)$ of $\mathbf{A}$ is constant on the multitrace $\{u_i \mid i < r\}$ and thus also on $\{v_i \mid i < r\}$. Then $p(\mathbf{v}) = \hat{r}(\mathbf{v}, \mathbf{u}) = \hat{r}(\mathbf{v}, \mathbf{v}) = \hat{d}(\mathbf{v})$ is a constant vector. Since $\mathbf{B}$ is also generated by the diagonal and $\mathbf{v}$, we get, by symmetry, that if $p(\mathbf{v})$ is constant then so is $p(\mathbf{u})$. Therefore $\mathbf{u} \wr\wr \mathbf{v}$.

It suffices to verify the last part of (4) for pairs in the generating set of $\|_\varepsilon$. So, suppose that $\varepsilon \leq (\rho : \gamma)$ and that $(\mathbf{u}, \mathbf{v}) = (\hat{t}(\vec{W}, \hat{c}), \hat{t}(\vec{W}, \hat{d}))$ for some polynomial t of $\mathbf{A}$ and $(c, d) \in \varepsilon$. By Lemma 4.4 (2) there is some multitrace $U \subseteq W$ with $t(x, c)$ a polynomial isomorphism between U and $\{u_i \mid i < r\}$. From $(c, d) \in (\rho : \gamma)$ we get that the kernels of the mappings $t(x, c)|_W$ and $t(x, d)|_W$ are equal. Hence $\{v_i \mid i < r\} = t(U, d)$ and from this we conclude that the two multitraces are ε-parallel.

To prove (5) suppose that $\mathbf{u} \wr\wr \mathbf{v}$ and that $u_i = v_i$ for some $i < r$. By (4) we know that the multitraces $\{u_i \mid i < r\}$ and $\{v_i \mid i < r\}$ are quasi-parallel and that they share an element in common. By part (5) of Lemma 5.2 we conclude that these multitraces are equal. Call this multitrace U.

Let $x-y$ be some polynomial of $\mathbf{A}$ whose restriction to U acts as subtraction with respect to some abelian group operation on U and let $s(x)$ be the polynomial $x-\mathbf{u}$ of $\mathbf{B}$, where $-$ is to be executed componentwise. Clearly $s(\mathbf{u})$ is constant and so $s(\mathbf{v})$ must be as well. Since $u_i = v_i$ then this can only happen if $\mathbf{u} = \mathbf{v}$. ∎

We are now ready to prove the main result of this section.

THEOREM 5.6 *Let W_0 and W_1 be quasi-parallel $\langle\rho,\gamma\rangle$-multitraces of $\mathbf{A}$ which are equal modulo some congruence $\varepsilon \leq (\rho : \gamma)$. If the equational class generated by $\mathbf{A}$ avoids the RL configuration then $W_0 \parallel_\varepsilon W_1$.*

PROOF. Referring to Lemma 5.2 for the case when $\mathrm{typ}(\rho,\gamma) = \mathbf{3}$ or Lemma 5.4 for the case when $\mathrm{typ}(\rho,\gamma) = \mathbf{2}$, we may assume that W_0 and W_1 are E-isomorphic. That is, we may assume that there are idempotent polynomials e_0 and e_1 of $\mathbf{A}$ with $e_i(W_j) = W_i$ and such that $e_i(e_j(x)) = x$ for all $x \in W_i$ and $i,\ j < 2$. Let r be the number of elements in W_0 (and W_1) and let $\vec{W}_0$ be some r-tuple which is a listing of all of the elements of W_0. Assume that W_0 and W_1 are not ε-parallel. We plan to contradict this assumption by showing, via Corollary 2.5, that it leads to an instance of the RL configuration in the equational class generated by $\mathbf{A}$.

Let $\mathbf{B}$ be the diagonal subalgebra of $\mathbf{A}^r$ generated by $\vec{W}_0$ and define $\vec{W}_1$ to be $\hat{e}_1(\vec{W}_0)$. Note that $\vec{W}_1$ is a listing of the elements of W_1 and that $\hat{e}_0(\vec{W}_1) = \vec{W}_0$. The congruences $\parallel_\varepsilon$ and $\wr\wr$ of $\mathbf{B}$ will play the roles of ρ and σ in Corollary 2.5 respectively. Let R be the kernel of the first projection map from B to A. From Lemma 5.5 (5) we see that $\wr\wr \cap R \subseteq \parallel_\varepsilon$.

Let $0 \in W_0$ denote the first component of $\vec{W}_0$ and let $0' = e_1(0) \in W_1$. The elements $\hat{0}, \hat{0}' = \hat{e}_1(\hat{0}), \vec{W}_0, \vec{W}_1 \in B$ will play the roles of $\tilde{0}, \tilde{0}', \tilde{1}$ and $\tilde{1}'$ from Corollary 2.5. Clearly we have $\hat{0}\ R\ \vec{W}_0$ and $\hat{0}'\ R\ \vec{W}_1$. To verify that

$$\begin{array}{ccc} \hat{0} & \parallel_\varepsilon & \hat{0}' \\ R & & R \\ \vec{W}_0 & \wr\wr & \vec{W}_1 \end{array},$$

as Corollary 2.5 requires, we need to show that $\hat{0} \parallel_\varepsilon \hat{0}'$ and $\vec{W}_0 \wr\wr \vec{W}_1$. The latter follows from Lemma 5.5 (4). For the former, it is enough to show that $(0,0') \in \varepsilon$. As W_0 and W_1 are equal modulo ε, then there is some $c \in W_1$ such that $0\ \varepsilon\ c$. Then $0' = e_1(0)\ \varepsilon\ e_1(c) = c$, which yields that $(0,0') \in \varepsilon$.

We have now proved that the conditions of Corollary 2.5 are satisfied. By Lemma 5.5 (4) it follows that $(\vec{W}_0, \vec{W}_1) \notin \ \|_\varepsilon$, since W_0 and W_1 were assumed not to be ε-parallel. Thus, in order to prove that $\{\hat{0}/ \|_\varepsilon , \vec{W}_0/ \|_\varepsilon \}$ along with the polynomial $\hat{e}_1/ \| \ \varepsilon$ constitute an RL configuration in $\mathbf{B}/ \|_\varepsilon$ it is sufficient to show that $(\vec{W}_0, \hat{0})$ is a $\mathbf{1}$-snag of $\mathbf{B}$. This can be done using the fact that $\mathbf{A}|_{W_0}$ has an abelian group operation. ∎

We can now explain why the configuration in Figure 1 leads to residual largeness. If $\mathbf{A}$ is any seven-element simple algebra of type $\mathbf{2}$ whose collection of minimal sets consists of the three-element subsets from Figure 1, then we claim that the minimal sets N and N' are quasi-parallel but not parallel.

To see that $N \wr\wr N'$, suppose that $p(x)$ is some polynomial of $\mathbf{A}$ which maps N to a single element. It follows that p maps the two remaining minimal sets distinct from N and N' to the same set (by Corollary 4.8) and this set is either a minimal set or a singleton. In either case, this forces p to map all of N' to a point. Similarly, any polynomial which maps N' to a point also maps N to one. Thus N and N' are quasi-parallel and are in fact the only pair of distinct quasi-parallel minimal sets in the algebra since every other pair of minimal sets have nonempty intersection.

The scarcity of quasi-parallel traces in $\mathbf{A}$ implies that if N and N' are parallel, then they must lie in some multitrace. This is not possible since A has only 7 elements and any multitrace which contains both N and N' must have rank at least 2 and hence contain at least 9 elements. We gather from Theorem 5.6 that $\mathbf{A}$ generates a residually large equational class.

6 Residually Small Abelian Equational Classes

In this section we will show that a locally finite abelian equational class is residually small if and only if it avoids the RL configuration. We have already shown in Section 2 that the presence of the RL configuration leads to residual largeness. To show the converse we need to locate an instance of the RL configuration in any locally finite abelian equational class which has a proper class of irreducibles. We will do more than this. We will locate an instance of the RL configuration in any locally finite abelian equational class $\mathcal{V}$ which has a sufficiently large irreducible, where 'sufficiently large' means 'exceeding some finite bound determined by the free spectrum of $\mathcal{V}$'.

Let $\mathcal{V}$ be a locally finite abelian equational class. For i a nonnegative integer, let M_i be the size of the i-generated free algebra in $\mathcal{V}$. The sequence $(M_0, M_1, M_2, \ldots)$ is known as the *free spectrum* of $\mathcal{V}$. The main result of this section is that if $\mathcal{V}$ contains an irreducible whose cardinality exceeds the finite number $2^{M_3} M_2^{M_2^4}$, then the RL configuration occurs in $\mathcal{V}$. To make reading the proof easier, we now give a rough sketch of the path the argument will take.

Assume that the RL configuration does not occur in $\mathcal{V}$ and that $\mathbf{S}$ is an irreducible member of $\mathcal{V}$. Let ρ denote the strongly solvable radical of $\mathbf{S}$. We will show that the cardinality of $\mathbf{S}$ is no more than $2^{M_3} M_2^{M_2^4}$ by showing that ρ-classes are no bigger than 2^{M_3} and that the index of ρ is no more than $M_2^{M_2^4}$. The first part of the argument (bounding ρ-classes) is a modification of Shapiro's proof for bounding the size of irreducibles in strongly abelian equational classes (Theorem 6.3). The second part of the argument (bounding the index of ρ) is accomplished by showing that the index of ρ is at most $M_2^{KM_2}$, where K denotes the number of covers of ρ in $\operatorname{Con} \mathbf{S}$ (Theorem 6.4). Then we argue that in the absence of RL configurations we must have $K \leq M_2^3$ (Corollary 6.11). All of the machinery of multitraces is applied to establish this bound on K.

We begin by listing a few facts about abelian equational classes.

PROPOSITION 6.1 *Let* $\mathbf{A}$ *be a member of* $\mathcal{V}$, *a locally finite abelian equational class.*

(1) *If* ρ *is a locally strongly solvable congruence of* $\mathbf{A}$ *then* ρ *is strongly abelian.*

(2) *The typeset of* $\mathcal{V}$ *is contained in* $\{\mathbf{1}, \mathbf{2}\}$ *and all type* $\mathbf{2}$ *minimal sets of finite algebras in* $\mathcal{V}$ *have empty tails and hence are Mal'tsev.*

(3) *If* $\mathbf{A} \in \mathcal{V}$ *and* $X, Y \subseteq A$, *then* $X \parallel Y$ *implies* $X \wr\wr Y$.

(4) *If* $p(\mathbf{x}) \in \mathrm{Pol}_{\mathrm{n}}\mathbf{A}$ *then there is an* $(n+2)$*-ary term* $t(\mathbf{x}, u, v)$ *such that for any* $0 \in A$ *we have* $p(\mathbf{x}) = t(\mathbf{x}, 0, p(0, \ldots, 0))$ *for all* $\mathbf{x} \in A$.

(5) *If* U *is the range of an idempotent polynomial of* $\mathbf{A}$, *then there is a binary term* $t(x, y)$ *satisfying* $t(t(x, y), y) = t(x, y)$ *such that all sets of the form* $t(A, a)$, *for* $a \in A$ *are E-isomorphic and parallel to* U.

(6) *If* $\mathbf{A}$ *is finite and* $\langle \rho, \gamma \rangle$ *is a stiff pair of congruences of* $\mathbf{A}$ *of type* $\mathbf{2}$ *then modulo the parallel (or quasi-parallel) relation there are at most* M_2 $\langle \rho, \gamma \rangle$*-multitraces.*

Proof. Parts (1) and (2) repeat some basic facts about abelian equational classes which can be found in [9, 2]. Part (3) follows from the remark made after Definition 5.3.

For part (4) we use the fact proved in [5] that locally finite abelian equational classes are hamiltonian. The following argument is similar to that given by Klukovits [7] in his study of hamiltonian equational classes. Now if $p(\mathbf{x})$ is an n-ary polynomial, then $p(\mathbf{x}) = s(\mathbf{x}, \mathbf{a})$ for some $(n+k)$-ary term s and some $\mathbf{a} \in A^k$. Let $\mathbf{F}$ be the $\mathcal{V}$-free algebra generated by $2(n+k)$ distinct elements x_i, x_i', y_j, y_j' where $1 \leq i \leq n$ and $1 \leq j \leq k$. Let $\mathbf{x} = (x_1, \ldots, x_n)$, $\mathbf{x}' = (x_1', \ldots, x_n')$, $\mathbf{y} = (y_1, \ldots, y_k)$, and $\mathbf{y}' = (y_1', \ldots, y_k')$. Let S be the subuniverse of $\mathbf{F}$ generated by the three elements $s(\mathbf{x}, \mathbf{y}')$, $s(\mathbf{x}', \mathbf{y}')$ and $s(\mathbf{x}', \mathbf{y})$. Since $\mathcal{V}$ is hamiltonian, S is a block of a congruence which we denote σ. In $\mathbf{F}/\sigma$ we have

$$s(\mathbf{x}', \mathbf{y}') = s(\mathbf{x}', \mathbf{y}),$$

so the term condition guarantees that

$$s(\mathbf{x}, \mathbf{y}') = s(\mathbf{x}, \mathbf{y}).$$

Therefore it must be that $s(\mathbf{x}, \mathbf{y}) \in S$. There must be a ternary term r such that

$$\mathcal{V} \models r(s(\mathbf{x}, \mathbf{y}'), s(\mathbf{x}', \mathbf{y}'), s(\mathbf{x}', \mathbf{y})) = s(\mathbf{x}, \mathbf{y}).$$

Using this equation, and $p(\mathbf{x}) = s(\mathbf{x}, \mathbf{a})$, we can construct the desired term t. Let $\mathbf{u}_i$ denote the tuple of length i whose entries are all equal to u. The term $t(\mathbf{x}, u, v) \stackrel{\text{def}}{=} r(s(\mathbf{x}, \mathbf{u}_k), s(\mathbf{u}_n, \mathbf{u}_k), v)$ has the property claimed in (4), since

$$\begin{aligned} t(\mathbf{x}, 0, p(0, \ldots, 0)) &= t(\mathbf{x}, 0, s(\mathbf{0}_n, \mathbf{a})) \\ &= r(s(\mathbf{x}, \mathbf{0}_k), s(\mathbf{0}_n, \mathbf{0}_k), s(\mathbf{0}_n, \mathbf{a})) \\ &= s(\mathbf{x}, \mathbf{a}) = p(\mathbf{x}). \end{aligned}$$

To prove (5), suppose that $U = e(A)$ for e an idempotent polynomial. Then there is a term $r(x, \mathbf{y})$ and elements $\mathbf{u}$ of A such that $e(x) = r(x, \mathbf{u})$ for all $x \in A$. Since $\mathcal{V}$ is locally finite and e is idempotent we may assume that the equation $r(x, \mathbf{y}) = r(r(x, \mathbf{y}), \mathbf{y})$ holds in $\mathcal{V}$. We claim that any set of the form $r(A, \mathbf{v})$ is polynomially isomorphic to U via the idempotent polynomials $r(x, \mathbf{u})$ and $r(x, \mathbf{v})$. It is an elementary exercise, using the abelian property and the fact that r is idempotent in x to show this. Finally, setting $t(x, y) = r(x, y, y, \ldots, y)$ produces the sought-after term.

For part (6), let T be a $\langle\rho, \gamma\rangle$-multitrace. By Theorem 4.3 (1)(b), there is an idempotent $e(x)$ such that T is a γ-class of $U = e(A)$, and each other γ-class of U is a $\langle\rho, \gamma\rangle$-multitrace parallel to T. Using part (5) of this proposition there is a term $t(x, y)$ such that $t(A, a)$ is E-isomorphic and parallel to U for any $a \in A$. Now fix $0 \in A$ and let $T_0 = t(T, 0)$. Then T_0 is parallel and E-isomorphic to T. Note that, up to parallelism, T_0 (hence T) is determined by the choice of t. This is because T_0 is a γ-class of $t(A, 0)$ and all such classes are parallel to one another. Since there are at most M_2 choices for t, we are done. ∎

The first step in our proof will be to show that the "large" finite irreducibles in $\mathcal{V}$ must have a strongly abelian monolith but cannot themselves be strongly abelian. The next theorem is a special case of Theorem 5.1 of [3].

THEOREM 6.2 *If* $\mathbf{S}$ *is a finite abelian irreducible with monolith* μ *such that* $\operatorname{typ}(0_S, \mu) = \mathbf{2}$ *then* S *has at most* $M_2^{M_2}$ *elements.*

For the remainder of this section let $\mathbf{S}$ be a finite irreducible member of $\mathcal{V}$, let μ be the monolith of $\mathbf{S}$ and let ρ be its strongly solvable radical (the largest strongly solvable congruence of $\mathbf{S}$). Note that under the assumption that $\mathcal{V}$ is abelian, ρ must be strongly abelian.

The following is a modification of an argument found in [10].

THEOREM 6.3 *If* σ *is a strongly solvable congruence of* $\mathbf{S}$, *then each* σ*-class has no more than* 2^{M_3} *elements. Consequently, each strongly abelian irreducible in* $\mathcal{V}$ *has size at most* 2^{M_3}.

Proof. We may assume that σ contains μ, the least nontrivial congruence of $\mathbf{S}$. Let C be a σ-class and let $(a, b) \in \mu$ with $a \neq b$. For each $c \in C$ let T_c denote the set of ternary terms $t(x, y, z)$ (modulo $\mathcal{V}$-equivalence) for which $a = t(c, c, a)$. Since there are M_3 ternary terms up to equivalence, there are at most 2^{M_3} different sets T_c, $c \in C$. To prove the theorem it will suffice to show that $T_c \neq T_d$ whenever c and d are distinct elements of C.

Since (a, b) is in the congruence generated by (c, d), there is a unary polynomial $p(x)$ such that $p(c) = a \neq p(d)$ (or the same with c and d switched, in which case the argument is the same). From Proposition 6.1 (4) there is a ternary term t such that $p(x) = t(x, c, p(c)) = t(x, c, a)$. Thus, since $a = p(c) = t(c, c, a)$ we get that $t \in T_c$. We now show that $t \notin T_d$.

If $a = t(d, d, a) = t(c, c, a)$, then the strong term condition implies that $t(d, d, a) = t(d, c, a)$. Therefore $a = t(d, d, a) = t(d, c, a) = p(d) \neq a$, which is a contradiction. This completes the proof. ∎

Using the bound established in the previous theorem we can now bound the size of a finite irreducible in $\mathcal{V}$ in terms of the number of covers of its strongly solvable radical ρ. Let $\mathcal{C}$ be the set of covers of ρ and let $K = |\mathcal{C}|$. Note that for each $\alpha \in \mathcal{C}$ we have $\mathrm{typ}(\rho, \alpha) = \mathbf{2}$.

THEOREM 6.4 $|S/\rho| \leq M_2^{KM_2}$ *and hence* $|S| \leq 2^{M_3} M_2^{KM_2}$.

Proof. For each $\alpha \in \mathcal{C}$, let α_* be a congruence maximal amongst those whose intersection with α is ρ. It is clear that α_* is meet irreducible and that the type of the pair $\langle \alpha_*, \alpha^* \rangle$ is $\mathbf{2}$, where α^* is the unique cover of α_*, since $\langle \rho, \alpha \rangle$ is perspective with this pair. The intersection of $\{\alpha_* \mid \alpha \in \mathcal{C}\}$ contains ρ and contains no cover of ρ, and so it is ρ. Thus, $\mathbf{S}/\rho$ can be embedded in $\prod_{\alpha \in \mathcal{C}} \mathbf{S}/\alpha_*$. By Theorem 6.2 we know that each factor of this product has size at most $M_2^{M_2}$ and so S/ρ can have at most $M_2^{KM_2}$ elements. The second statement follows from Theorem 6.3. ∎

We now set out to show that if $\mathcal{V}$ avoids the RL configuration then the strongly solvable radical of $\mathbf{S}$ must have no more than M_2^3 covers (so $K \leq M_2^3$). To avoid a trivial situation, assume that the type of $(0_S, \mu)$ is $\mathbf{1}$.

Definition 6.5 Let $\mathbf{A}$ be a finite algebra and $\delta \in \mathrm{Con}\,\mathbf{A}$. Two covers α and β of δ will be called **equivalent** if $\mathrm{M}_{\mathbf{A}}(\delta, \alpha) = \mathrm{M}_{\mathbf{A}}(\delta, \beta)$.

Every minimal set in $\mathbf{S}$ is the image of an idempotent unary polynomial of $\mathbf{S}$. By Proposition 6.1 (5) it is possible to choose a collection of at most M_2 subsets of S representing every idempotent image up to E-isomorphism. In particular, $\mathbf{S}$ has at most M_2 minimal sets up to E-isomorphism. This implies that there are at most M_2 equivalence classes of covers of the congruence ρ of $\mathbf{S}$. The rest of the section is devoted to showing that no equivalence class can contain more than M_2^2 covers if $\mathcal{V}$ avoids the RL configuration.

LEMMA 6.6 *Let $\mathbf{A}$ be a finite algebra and $\delta \in \mathrm{Con}\,\mathbf{A}$. Let $\mathcal{E}$ be a set of equivalent covers of δ such that the type of $\langle \delta, \alpha \rangle$ is $\mathbf{2}$ for each $\alpha \in \mathcal{E}$. If $\gamma = \bigvee \mathcal{E}$ and β is a cover of δ below γ with $\mathrm{typ}(\delta, \beta) = \mathbf{2}$, then β is equivalent to each member of $\mathcal{E}$.*

Proof. Let β be a cover of δ below γ with $\mathrm{typ}(\delta, \beta) = \mathbf{2}$, and U a $\langle \delta, \beta \rangle$-minimal set with body B and tail T. Since $\beta \leq \bigvee \mathcal{E}$ and U is the range of some idempotent polynomial e of $\mathbf{A}$, then $\beta|_U \leq \bigvee_{\alpha \in \mathcal{E}} \alpha|_U$. If $\alpha \in \mathcal{E}$, then α is abelian over δ, hence we get from Lemma 4.27 of [2] that $\alpha|_U \subseteq B^2 \cup T^2$. Therefore $\beta|_B \leq \bigvee_{\alpha \in \mathcal{E}} \alpha|_B$ and so there must be at least one $\alpha \in \mathcal{E}$ with $\alpha|_B > \delta|_B$. Fix such an α, and let $a, b \in B$ with $(a, b) \in \alpha - \delta$. Connect a to b by $\langle \delta, \alpha \rangle$-traces, and pull this chain into U by e. We get that U contains a $\langle \delta, \alpha \rangle$-minimal set V whose intersection with B has more than 1 element. From Lemma 4.30 of [2] we conclude that $U = V$. Thus $\mathrm{M}_{\mathbf{A}}(\delta, \beta) = \mathrm{M}_{\mathbf{A}}(\delta, \alpha)$ for each $\alpha \in \mathcal{E}$ as required. ∎

COROLLARY 6.7 *Let $\mathbf{A}$ be a finite algebra and $\mathcal{E}$ a set of equivalent covers of a congruence δ of $\mathbf{A}$ with $\mathrm{typ}(\delta, \alpha) = \mathbf{2}$ for each $\alpha \in \mathcal{E}$. If every cover of δ below $\gamma = \bigvee \mathcal{E}$ is of type $\mathbf{2}$, then the pair $\langle \delta, \gamma \rangle$ is tame and of type $\mathbf{2}$.*

Proof. If U is a $\langle \delta, \alpha \rangle$-minimal set for some $\alpha \in \mathcal{E}$, then U is also a $\langle \delta, \gamma \rangle$-minimal set and the restriction map from the interval $[\delta, \gamma]$ in $\mathrm{Con}\,\mathbf{A}$ into $\mathrm{Con}\,\mathbf{A}|_U$ is 1-separating. The previous lemma can be used to show that this map is also 0-separating and so it follows that $\langle \delta, \gamma \rangle$ is indeed tame. The type of a tame quotient is determined by the polynomial structure on any minimal set, so the type of $\langle \delta, \gamma \rangle$ is the same as the type of $\langle \delta, \alpha \rangle$: it is $\mathbf{2}$. ∎

Now we resume our efforts to bound the size of the irreducible algebra $\mathbf{S}$. Fix a cover α of the strongly solvable radical ρ in $\mathrm{Con}\,\mathbf{S}$, and let $\mathcal{E}$ be the set of covers of ρ equivalent to α. We have to prove that $|\mathcal{E}| \leq M_2^2$. Let $\gamma = \bigvee \mathcal{E}$. The previous corollary demonstrates that $\langle \rho, \gamma \rangle$ is tame and of type $\mathbf{2}$. By Proposition 4.2 (2) we have that this pair is stiff.

Let W be a $\langle \rho, \gamma \rangle$-minimal set and $V \subseteq W$ a $\langle \rho, \gamma \rangle$-trace. Fix $(a, b) \in \mu - 0_S$. By Proposition 4.2 (2), for any $\beta \in \mathcal{E}$ the quotient $\langle \rho, \beta \rangle$ is stiff, so by Proposition 4.2 (1) there is a chain of overlapping $\langle \rho, \beta \rangle$-traces which connect a to b. Let $N_\beta^0, \ldots, N_\beta^{k_\beta}$ be such a chain, with k_β as small as possible.

Since each $\langle\rho,\beta\rangle$-trace is contained in a $\langle\rho,\gamma\rangle$-trace, we can fix $\langle\rho,\gamma\rangle$-traces V^i_β such that $N^i_\beta \subseteq V^i_\beta$ for all $i \leq k_\beta$. The following lemma records some relevant facts about these chains under the assumption that $\mathcal{V}$ avoids the RL configuration.

LEMMA 6.8 *Assume that $\mathcal{V}$ avoids the RL configuration and that β, $\nu \in \mathcal{E}$.*

(1) *Each k_β is at least 1.*

(2) *If $i < k_\beta$ and M is a $\langle\rho,\gamma\rangle$-multitrace which contains V^{i+1}_β, then $M \cap V^i_\beta$ is a singleton and equals $N^i_\beta \cap N^{i+1}_\beta$.*

(3) *If V^0_β is quasi-parallel to V^0_ν, then they are equal.*

Proof. If $k_\beta = 0$ then we would have a, $b \in N^0_\beta$, contradicting that μ is trivial on any $\langle\rho,\beta\rangle$-multitrace.

To prove (2), let $e(x)$ be an idempotent polynomial of $\mathbf{S}$ such that M is the intersection of $e(S)$ with some γ-class. If $M \cap V^i_\beta$ contains more than one element, then as e is one-to-one on M it must also be one-to-one on V^i_β. Thus, $e(V^i_\beta)$ is polynomially isomorphic to V^i_β and has at least two elements in common with it. Since V^i_β is a multitrace of rank 1, from Lemma 5.2 (4) and (5) we conclude that these two sets must be equal, and so M contains V^i_β. Applied to $M = V^{i+1}_\beta$, this argument shows that $V^i_\beta \cap V^{i+1}_\beta$ contains exactly one element since these two traces are distinct.

To get a contradiction from $V^i_\beta \cup V^{i+1}_\beta \subseteq M$ recall that N^j_β, $j \leq k_\beta$ is a chain of overlapping $\langle\rho,\beta\rangle$-traces *of minimal length* which connect a to b. Let $a_{-1} = a$, $a_{k_\beta} = b$ and for $j < k_\beta$, let a_j be an element in the intersection of N^j_β and N^{j+1}_β. By Corollary 4.4 (3) there exists a $\langle\rho,\beta\rangle$-trace N which contains both a_{i-1} and a_{i+1}. Then the chain $N^0_\beta, \ldots, N^{i-1}_\beta, N, N^{i+2}_\beta, \ldots, N^{k_\beta}_\beta$ is a shorter chain of $\langle\rho,\beta\rangle$-traces which connect a to b. From this contradiction we get that $M \cap V^i_\beta$ is a singleton. Since $N^i_\beta \cap N^{i+1}_\beta \neq \emptyset$, $N^i_\beta \subseteq V^i_\beta$ and $N^{i+1}_\beta \subseteq V^{i+1}_\beta \subseteq M$, we get that $N^i_\beta \cap N^{i+1}_\beta$ is the same singleton.

If V^0_β is quasi-parallel to V^0_ν, then as these two multitraces contain the element a it follows from Lemma 5.2 (5) that they must be equal. ∎

LEMMA 6.9 *If $\mathcal{E}$ has more than M_2^2 elements then there is some subset $\mathcal{I}$ of $\mathcal{E}$ of size greater than M_2 with $V^0_\beta = V^0_\nu$ for all β, $\nu \in \mathcal{I}$. For β, $\nu \in \mathcal{I}$, the singletons $V^0_\beta \cap V^1_\beta$ and $V^0_\nu \cap V^1_\nu$ are distinct.*

PROOF. From Proposition 6.1 (6) we learn that modulo the quasi-parallel relation there are at most M_2 distinct $\langle \rho, \gamma \rangle$-traces. Thus if $\mathcal{E}$ has size greater than M_2^2 then there must be some subset $\mathcal{I}$ of $\mathcal{E}$ of size greater than M_2 such that V_β^0 and V_ν^0 are quasi-parallel for all β, $\nu \in \mathcal{I}$. By part (3) of Lemma 6.8 it follows that for β, $\nu \in \mathcal{I}$, $V_\beta^0 = V_\nu^0$.

By construction, for any $\beta \in \mathcal{E}$, the unique element in $V_\beta^0 \cap V_\beta^1$ is β-related to a. As $\beta \cap \nu = \rho$ and ρ is trivial on V_β^0 and V_ν^0 for any two covers β, ν of ρ it follows that the singletons $V_\beta^0 \cap V_\beta^1$ and $V_\nu^0 \cap V_\nu^1$ are distinct. ∎

We now do something rather unexpected (and having roots in the paper [5]) which will lead quickly to the main result of this section. For each $\beta \in \mathcal{E}$, let M_β be some maximal $\langle \rho, \gamma \rangle$-multitrace which contains V_β^1. Define ε to be the congruence of $\mathbf{S}$ generated by V^2.

LEMMA 6.10 *Suppose that $\mathcal{V}$ avoids the RL configuration. If M is a maximal $\langle \rho, \gamma \rangle$-multitrace and M' is a multitrace which is quasi-parallel to M and lies in the same ε-class as M then $M = M'$.*

PROOF. By Theorem 5.6 it follows that M and M' are ε-parallel and so there are polynomials $p_i(x, y)$ of $\mathbf{S}$, $i \leq n$ for some n, and pairs (u_i, v_i) from V^2 such that $p_0(M, u_0) = M$, $p_i(M, v_i) = p_{i+1}(M, u_{i+1})$ and $p_n(M, v_n) = M'$. As $\mathbf{S}$ is abelian we see that $|p_i(M, u_i)| = |p_i(M, v_i)|$ for each $i \leq n$. Using the maximality of M it follows that $M = p_i(M, V) = p_i(M, v_i)$ for each $i \leq n$. Thus, $M = M'$ as required. ∎

We are now able to derive our final contradiction.

COROLLARY 6.11 *If $\mathcal{V}$ avoids the RL configuration, then $\mathcal{E}$ has no more than M_2^2 elements.*

PROOF. Suppose instead that $\mathcal{V}$ avoids the RL configuration and that $\mathcal{E}$ has more than M_2^2 elements. Let $\mathcal{I}$ be a subset of $\mathcal{E}$ satisfying the conditions of Lemma 6.9. Since $\mathcal{I}$ has more than M_2 elements then by Proposition 6.1 there must be two members β and ν such that M_β and M_ν are parallel. These maximal multitraces lie in the same ε-class as a and hence by Lemma 6.10 must be equal. Then M_β is a multitrace which contains V_β^1 and which contains at least two elements of V_β^0, namely the elements in $V_\beta^0 \cap V_\beta^1$ and $V_\nu^0 \cap V_\nu^1$ (since $V_\beta^0 = V_\nu^0$ and $V_\nu^1 \subseteq M_\nu = M_\beta$). This contradicts part (2) of Lemma 6.8. ∎

Putting all the results together, we get:

THEOREM 6.12 *Let $\mathcal{V}$ be a locally finite abelian equational class. The following are equivalent:*

(1) *$\mathcal{V}$ is residually small,*

(2) *$\mathcal{V}$ is residually bounded by $2^{M_3} M_2^{M_2^4}$,*

(3) *$\mathcal{V}$ avoids the RL configuration.*

The following corollary provides an algorithm to determine whether or not a finite algebra generates a residually small abelian equational class.

COROLLARY 6.13 *The equational class $\mathcal{V}$ generated by a finite algebra $\mathbf{A}$ is abelian and residually small if and only if $\mathsf{HS}(\mathbf{A}^{A^3})$ is hamiltonian and every 2-generated algebra in $\mathcal{V}$ avoids the RL configuration. In this case $\mathcal{V}$ is residually bounded by $2^{n^{n^3}} n^{n^{4n^2+2}}$, where $|A| = n$.*

Proof. By [5], the first condition is necessary and sufficient for a finite algebra to generate an abelian equational class and by Theorem 2.2 the second condition is necessary if $\mathcal{V}$ is to be residually small.

Suppose that $(1, 0)$ along with the polynomial $f(x)$ constitute an RL configuration in the algebra $\mathbf{B}$ from $\mathcal{V}$. Then there is a polynomial $s(x, y)$ of $\mathbf{B}$ such that $s(0, 0) = 0$ and $s(0, 1) = s(1, 0) = 1$. Since $s(0, 0) = 0$ and $f(0) = 0$, then by Proposition 6.1 there are terms $t(x, y, z)$ and $r(x, y)$ such that $s(x, y) = t(x, y, 0)$ and $f(x) = r(x, 0)$ for all x, $y \in B$. From this it is easy to see that the RL configuration occurs in the subalgebra of $\mathbf{B}$ generated by $\{0, 1\}$. ∎

References

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

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

[3] K. Kearnes. Cardinality bounds for subdirectly irreducible algebras. *Journal of Pure and Applied Algebra*, 112:293–312, 1996.

[4] K. Kearnes, E. Kiss, and M. Valeriote. Minimal sets and varieties. *Trans. Amer. Math. Soc.*, 350(1):1–41, 1998.

[5] E. Kiss and M. Valeriote. Abelian algebras and the Hamiltonian property. *The Journal of Pure and Applied Algebra*, 87:37–49, 1993.

[6] E. W. Kiss. An easy way to minimal algebras. *The International Journal of Algebra and Computation*, 7:55–75, 1997.

[7] L. Klukovits. Hamiltonian varieties of universal algebras. *Acta. Sci. Math.*, 37:11–15, 1975.

[8] R. McKenzie. Monotone clones, residual smallness and congruence distributivity. *Bulletin of the Australian Mathematics Society*, 41:283–300, 1990.

[9] R. McKenzie and M. Valeriote. *The Structure of Locally Finite Decidable Varieties*, volume 79 of *Progress in Mathematics*. Birkhäuser Boston, 1989.

[10] Jacob Shapiro. Finite algebras with Abelian properties. *Algebra Universalis*, 25:334–364, 1988.

[11] H. Werner. Congruences on products of algebras and functionally complete algebras. *Algebra Universalis*, 4:99–105, 1974.

Department of Mathematics
University of Louisville
Louisville, KY 40292, USA.

Department of Algebra and Number Theory
Eötvös Lóránd University
1088 Budapest, Múzeum krt. 6–8, Hungary.

Department of Mathematics and Statistics
McMaster University
Hamilton, Ontario, Canada, L8S 4K1.