CONGRUENCE MODULAR VARIETIES WITH SMALL FREE SPECTRA

KEITH A. KEARNES

ABSTRACT. Let $\mathbf{A}$ be a finite algebra that generates a congruence modular variety. We show that the free spectrum of $\mathcal{V}(\mathbf{A})$ fails to have a doubly exponentially lower bound if and only if $\mathbf{A}$ has a finitely generated clone and $\mathbf{A}$ is a direct product of nilpotent algebras of prime power cardinality.

1. INTRODUCTION

Let $\mathbf{A}$ be a finite algebra, and let $\mathcal{V}(\mathbf{A})$ be the variety it generates. If $\mathbf{F}_{\mathcal{V}(\mathbf{A})}(k)$ is the k-generated free algebra in $\mathcal{V}(\mathbf{A})$, then the function

$$\mathrm{Spec}_{\mathbf{A}}(k) := |F_{\mathcal{V}(\mathbf{A})}(k)|$$

is called the *free spectrum of* $\mathcal{V}(\mathbf{A})$ (or the *free spectrum of* $\mathbf{A}$).

We will compare functions with the relation $\preceq$, which is defined by the rule that $f \preceq g$ if $f(k) \leq g(k)$ for all sufficiently large k. In words we say that "g is an upper bound for f" or "f is a lower bound for g". We use $g \succeq f$ to mean the same thing. We will write $\mathrm{Spec}_{\mathbf{A}}(k) \preceq 2^{2^{ck}}$ or $\mathrm{Spec}_{\mathbf{A}}(k) \succeq 2^{2^{ck}}$ to mean that *there exists some* $c > 0$ such that the functions $\mathrm{Spec}_{\mathbf{A}}(k)$ and $2^{2^{ck}}$ are $\preceq$-comparable. Since the number of k-ary operations on a set of size $|A|$ is at most $|A|^{|A|^k}$, and since elements of $\mathbf{F}_{\mathcal{V}(\mathbf{A})}(k)$ may be identified with k-ary (term) operations of $\mathbf{A}$, it is always the case that $\mathrm{Spec}_{\mathbf{A}}(k) \preceq 2^{2^{ck}}$ when $\mathbf{A}$ is finite. We will write $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$ and say that "$\mathrm{Spec}_{\mathbf{A}}(k)$ does not have a doubly exponential lower bound" to mean that $2^{2^{ck}} \not\preceq \mathrm{Spec}_{\mathbf{A}}(k)$.

In this paper we prove that if $\mathbf{A}$ generates a congruence modular variety, then $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$ if and only if $\mathbf{A}$ has a finitely generated clone and $\mathbf{A}$ is a direct product of nilpotent algebras of prime power cardinality.

1991 *Mathematics Subject Classification.* Primary 08B20, Secondary 08A05.

Key words and phrases. free spectrum, nilpotent algebra, twin monoid.

This material is based upon work supported by the National Science Foundation under Grant No. DMS 9802922.

The results of this paper are related to results of Vaughan-Lee [10], Freese and McKenzie [3], and Berman and Blok [1]. To understand the relationship, fix $\mathbf{A}$ to be a finite algebra of finite type that generates a congruence modular variety. By modifying the arguments in [10] it is shown in Chapter 14 of [3] that if $\mathbf{A}$ nilpotent, then $\mathbf{A}$ is finitely based provided that it factors as a direct product of algebras of prime power cardinality. The proof revolves around establishing a bound on the rank of 'commutator terms', and the hypothesis that $\mathbf{A}$ factors into a direct product of prime power algebras is used in a nontrivial way to establish the bound. Later, in [1], it is shown that if there is a finite bound on the rank of commutator terms, then $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$. What we show here is that for $\mathbf{A}$ (as above) the following are equivalent: (i) $\mathbf{A}$ factors as a direct product of nilpotent algebras of prime power cardinality, (ii) $\mathbf{A}$ has a finite bound on the rank of commutator terms, (iii) $\mathrm{Spec}_{\mathbf{A}}(k)$ does not have a doubly exponential lower bound. The key idea behind the proof is to connect these properties with a fourth equivalent property: (iv) $\mathbf{A}$ is nilpotent and its twin monoid is a nilpotent group.

2. The Twin Monoid

Let $\mathbf{A}$ be an algebra, and $f(x)$ and $g(x)$ be unary polynomials of $\mathbf{A}$. We call f and g *twins* if for some n there is an $(n+1)$-ary term operation $t(x, \mathbf{y}) =: t_{\mathbf{y}}(x)$ of $\mathbf{A}$ and tuples $\mathbf{a}, \mathbf{b} \in A^n$ such that $f(x) = t_{\mathbf{a}}(x)$ and $g(x) = t_{\mathbf{b}}(x)$.

Lemma 2.1. *The twin relation $\tau = \{(f, g) \mid f$ and g are twins$\}$ is a tolerance relation of the monoid $\langle \mathrm{Pol}_1(\mathbf{A}); \circ, id \rangle$.*

Proof. Recall that a tolerance relation is a reflexive, symmetric, compatible binary relation. It is clear that the relation τ defined in the lemma is a reflexive, symmetric, binary relation. To see that it is compatible with composition, assume that $(f, g), (f', g') \in \tau$. Then we can find terms t and t' and tuples $\mathbf{a}, \mathbf{b}, \mathbf{a}'$ and $\mathbf{b}'$ such that $f(x) = t_{\mathbf{a}}(x)$ and $g(x) = t_{\mathbf{b}}(x)$ while $f'(x) = t'_{\mathbf{a}'}(x)$ and $g'(x) = t'_{\mathbf{b}'}(x)$. Therefore the term $t_{\mathbf{y}}(t'_{\mathbf{y}'}(x))$ and the tuples $\mathbf{a}\mathbf{a}'$ and $\mathbf{b}\mathbf{b}'$ witness that composition $f(f'(x)) = t_{\mathbf{a}}(t'_{\mathbf{a}'}(x))$ is a twin of $g(g'(x)) = t_{\mathbf{b}}(t'_{\mathbf{b}'}(x))$. □

For any tolerance relation τ on any monoid $\mathbf{M}$, the set of elements τ-related to $1 \in M$ is a submonoid of $\mathbf{M}$. We call the submonoid of $\mathrm{Pol}_1(\mathbf{A})$ that consists of twins of the identity the *twin monoid*. It will be denoted $\mathrm{Tw}(\mathbf{A})$.

Lemma 2.2. *Let $\mathbf{A}$ be a finite algebra. There is a single term $s_{\mathbf{y}}(x)$ such that each element of $\mathrm{Tw}(\mathbf{A})$ can be represented as $s_{\mathbf{a}}(x)$ for some $\mathbf{a}$.*

Proof. Suppose that f and f' are twins of the identity. Then there are terms $t_{\mathbf{y}}(x)$ and $t'_{\mathbf{y}'}(x)$ and tuples $\mathbf{a}, \mathbf{b}, \mathbf{a}'$ and $\mathbf{b}'$ such that $id(x) = t_{\mathbf{a}}(x)$, $f(x) = t_{\mathbf{b}}(x)$, $id(x) = t'_{\mathbf{a}'}(x)$, and $f'(x) = t'_{\mathbf{b}'}(x)$. Now let $T_{\mathbf{y}\mathbf{y}'}(x) = t_{\mathbf{y}}(t'_{\mathbf{y}'}(x))$. Then for the tuple $\mathbf{a}\mathbf{a}'$ we have $T_{\mathbf{a}\mathbf{a}'}(x) = id(id(x)) = id(x)$, while for the tuples $\mathbf{b}\mathbf{a}'$ and $\mathbf{a}\mathbf{b}'$ we have $T_{\mathbf{b}\mathbf{a}'}(x) = f(id(x)) = f(x)$ and $T_{\mathbf{a}\mathbf{b}'}(x) = id(f'(x)) = f'(x)$. Therefore $T_{\mathbf{y}\mathbf{y}'}(x)$ is a term that for different choices of the parameters represents $id(x), f(x)$ and $f'(x)$.

The argument we have just given to construct a single term that witnesses membership in $\mathrm{Tw}(\mathbf{A})$ for any two given polynomials $f, f' \in \mathrm{Tw}(\mathbf{A})$ extends to show that any finite subset of $\mathrm{Tw}(\mathbf{A})$ can be represented by a single term. Since our hypothesis that $\mathbf{A}$ is a finite algebra implies that $\mathrm{Tw}(\mathbf{A})$ is finite, there exists a single term $s_{\mathbf{y}}(x)$ and a tuple $\mathbf{e}$ such that $s_{\mathbf{e}}(x) = id(x)$ while the polynomials of the form $s_{\mathbf{a}}(x)$ represent all elements of $\mathrm{Tw}(\mathbf{A})$. □

Applications of the twin monoid to free spectra are based on the following result.

Theorem 2.3. *Let* $\mathbf{A}$ *be a finite algebra and let* $\mathbf{M} = \mathrm{Tw}(\mathbf{A})$. *For some fixed* n *it is the case that* $\mathrm{Spec}_{\mathbf{M}}(k) \preceq \mathrm{Spec}_{\mathbf{A}}(nk+1)$. *In particular, if* $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$, *then we also have* $\mathrm{Spec}_{\mathbf{M}}(k) \ll 2^{2^{ck}}$.

Proof. Suppose, for example, that the elements $xy^2, yx^2 \in F_{\mathcal{V}(\mathbf{M})}(x, y)$ are distinct. Then there exist $s_{\mathbf{a}}, s_{\mathbf{b}} \in M$ such that the homomorphism from $F_{\mathcal{V}(\mathbf{M})}(x, y)$ to $\mathbf{M}$ induced by

$$\langle x, y \rangle \mapsto \langle s_{\mathbf{a}}, s_{\mathbf{b}} \rangle$$

fails to identify xy^2 and yx^2. This says precisely that there is an element $u \in A$ such that the functions $s_{\mathbf{a}} \circ s_{\mathbf{b}}^2$ and $s_{\mathbf{b}} \circ s_{\mathbf{a}}^2$ disagree at u. It follows that the homomorphism from $\mathbf{F}_{\mathcal{V}(\mathbf{A})}(\mathbf{x}, \mathbf{y}, z)$ to $\mathbf{A}$ induced by

$$\langle \mathbf{x}, \mathbf{y}, z \rangle \mapsto \langle \mathbf{a}, \mathbf{b}, u \rangle$$

fails to identify the elements $s_{\mathbf{x}}(s_{\mathbf{y}}(s_{\mathbf{y}}(z)))$ and $s_{\mathbf{y}}(s_{\mathbf{x}}(s_{\mathbf{x}}(z)))$.

More generally, this type of reasoning shows that if n is the length of the tuple $\mathbf{y}$ that appears in $s_{\mathbf{y}}(x)$ then the assignment

$$\omega(y_1, \ldots, y_k) \mapsto \omega(s_{\mathbf{y}_1}, \ldots, s_{\mathbf{y}_k})$$

is an injective function from $\mathbf{F}_{\mathcal{V}(\mathbf{M})}(k)$ to $\mathbf{F}_{\mathcal{V}(\mathbf{A})}(nk+1)$ for each k.

For the last statement of the theorem, we prove the contrapositive. Assume that $\mathrm{Spec}_{\mathbf{M}}(k) \succeq 2^{2^{ck}}$. Then $\mathrm{Spec}_{\mathbf{M}}(k) \geq 2^{2^{ck}}$ for some fixed $c > 0$ and all large k. From the first part of the theorem we have $\mathrm{Spec}_{\mathbf{A}}(nk+1) \geq 2^{2^{ck}}$ for all large k. Since $\mathrm{Spec}_{\mathbf{A}}(k)$ is an increasing function, this is enough to guarantee that $\mathrm{Spec}_{\mathbf{A}}(k) \geq 2^{2^{c'k}}$ for $c' =$

$c/2n$ and all large k. Thus $\mathrm{Spec}_{\mathbf{A}}(k) \succeq 2^{2^{ck}}$, which concludes the proof. □

Theorem 2.3 indicates that a detailed understanding of free spectra of monoids would be useful for more general free spectra questions. We know very little about free spectra of monoids, but it is not hard to verify the following two facts.

(1) Any nontrivial monoid $\mathbf{M}$ has $\mathrm{Spec}_{\mathbf{M}}(k) \succeq 2^k$.
(2) If $\mathbf{M}$ is a finite monoid, then $\mathrm{Spec}_{\mathbf{M}}(k) \ll 2^{ck\log(k)}$ iff $\mathrm{Spec}_{\mathbf{M}}(k) \preceq 2^{ck}$ iff $\mathbf{M}$ is commutative.

One can find in [6] a complete characterization of those finite algebras $\mathbf{A}$ for which $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{ck}$. By (1) and Theorem 2.3, it is necessary that each such $\mathbf{A}$ have a trivial twin monoid. The result in [6] is that $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{ck}$ if and only if $\mathbf{A}$ has a finitely generated clone and all *local* twin monoids are trivial. (See [6] for the meaning of this.)

In addition to the easy results about free spectra of monoids that we have just listed, the following less obvious result has been known for a long time.

Theorem 2.4. ([2], [9]) *Let* $\mathbf{G}$ *be a finite group. If* $\mathrm{Spec}_{\mathbf{G}}(k) \ll 2^{2^{ck}}$, *then* $\mathbf{G}$ *is nilpotent.*

Fortunately for us, Theorem 2.4 contains everything we will have to know about the free spectra of twin monoids for the results of this paper. This is a consequence of the next theorem and Theorem 12.5 of [4] (which proves that if $\mathbf{A}$ is a finite algebra in a congruence modular variety and $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$, then $\mathbf{A}$ is nilpotent).

Theorem 2.5. *If* $\mathbf{A}$ *is a finite nilpotent algebra in a congruence modular variety, then* $\mathrm{Tw}(\mathbf{A})$ *is a group.*

Proof. The assumption that $\mathbf{A}$ generates a congruence modular variety implies that $\mathbf{1} \notin \mathrm{typ}\{\mathbf{A}\}$, according to Theorem 8.5 of [4]. Therefore, Lemma 4.2 and Theorem 4.3 of [5] prove that idempotent twin polynomials have ranges of the same cardinality. So, if $s_{\mathbf{e}}(x) = id(x)$ and $s_{\mathbf{a}}(x) = f(x) \in \mathrm{Tw}(\mathbf{A})$, then any idempotent iterate f^k of f is a twin of the idempotent polynomial $id^k = id$. Idempotent twins have ranges of the same cardinality, so f^k is a surjective mapping from A to A. This forces f to be a permutation of A. □

Corollary 2.6. *If* $\mathbf{A}$ *is a finite algebra in a congruence modular variety and* $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$, *then* $\mathbf{A}$ *is a nilpotent algebra whose twin monoid is a nilpotent group.*

This corollary makes it clear that in this paper we will be dealing primarily with nilpotent algebras that generate congruence modular varieties. The following fact, which is Exercise 7.6 of [3], will be used throughout the rest of the paper.

Theorem 2.7. *A congruence modular variety generated by a nilpotent algebra is congruence permutable.*

The following result helps to understand twin monoids of algebras in congruence permutable varieties.

If $\mathcal{K}$ is a class of algebras, let $\mathrm{Tw}(\mathcal{K})$ denote $\{\mathrm{Tw}(\mathbf{K}) \mid \mathbf{K} \in \mathcal{K}\}$.

Theorem 2.8. *If $\mathcal{K}$ is a class of algebras that generates a congruence permutable variety, then* $\mathrm{Tw}(\mathsf{HSP}(\mathcal{K})) \subseteq \mathsf{HSP}(\mathrm{Tw}(\mathcal{K}))$.

Proof. We prove the theorem through a sequence of claims.

Claim 2.9. *Let $d(x,y,z)$ be a Mal'tsev term for $\mathcal{V}(\mathcal{K})$. If $\mathbf{B} \in \mathcal{V}(\mathcal{K})$ and e and f are twin polynomials of $\mathbf{B}$, then there is a polynomial $P \in \mathrm{Tw}(\mathbf{B})$ such that $Pe = d(e, e^2, fe)$.*

Since e and f are twins there is a term $t_{\mathbf{y}}(x)$ and tuples $\mathbf{a}$ and $\mathbf{b}$ such that $e(x) = t_{\mathbf{a}}(x)$ and $f(x) = t_{\mathbf{b}}(x)$. Define $T_{\mathbf{yz}}(x) = d(x, t_{\mathbf{y}}(x), t_{\mathbf{z}}(x))$ where $\mathbf{z}$ is a tuple of new variables. Note that $T_{\mathbf{aa}}(x) = id(x)$ and $T_{\mathbf{ab}}(x) = d(x, e(x), f(x))$, so $d(x, e(x), f(x))$ is a twin of the identity. It is this polynomial that we take for P. Clearly $Pe = d(e, e^2, fe)$, as desired.

Claim 2.10. *If $\mathbf{C}, \mathbf{D} \in \mathcal{V}(\mathcal{K})$, then a surjective homomorphism $h\colon \mathbf{C} \to \mathbf{D}$ induces a surjective homomorphism $\widehat{h}\colon \mathrm{Tw}(\mathbf{C}) \to \mathrm{Tw}(\mathbf{D})$.*

The map $\widehat{h}$ is the restriction to $\mathrm{Tw}(\mathbf{C})$ of the map from $\mathrm{Pol}_1(\mathbf{C})$ to $\mathrm{Pol}_1(\mathbf{D})$ that assigns to a polynomial $t(x, \mathbf{a})$ of $\mathbf{C}$ the polynomial $t(x, h(\mathbf{a}))$ of $\mathbf{D}$. If $t_{\mathbf{a}}(x) = id(x)$ on $\mathbf{C}$ and $t_{\mathbf{b}}(x) \in \mathrm{Tw}(\mathbf{C})$, then clearly $t_{h(\mathbf{a})}(x) = id(x)$ on $\mathbf{D}$ and $t_{h(\mathbf{b})}(x) \in \mathrm{Tw}(\mathbf{D})$. Thus h induces a function from $\mathrm{Tw}(\mathbf{C})$ to $\mathrm{Tw}(\mathbf{D})$. It is easy to see that this function preserves composition. The nonobvious part of Claim 2.10 is that this homomorphism is surjective if h is.

Choose $g \in \mathrm{Tw}(\mathbf{D})$, and assume that $t_{\mathbf{y}}(x)$ is a term for which $t_{\mathbf{c}}(x) = id(x)$ and $t_{\mathbf{d}}(x) = g(x)$ for certain tuples $\mathbf{c}, \mathbf{d}$ in $\mathbf{D}$. Let $\mathbf{a}$ and $\mathbf{b}$ be preimages under h for $\mathbf{c}$ and $\mathbf{d}$ respectively. Let $e(x) = t_{\mathbf{a}}(x)$ and $f(x) = t_{\mathbf{b}}(x)$. By construction we have $\widehat{h}(e) = id$ and $\widehat{h}(f) = g$. Claim 2.9 guarantees that there is a $P \in \mathrm{Tw}(\mathbf{C})$ such that $Pe = d(e, e^2, fe)$. Note that, since $\widehat{h}(e) = id$, we have

$$\widehat{h}(P) = \widehat{h}(P)\widehat{h}(e) = \widehat{h}(Pe) = \widehat{h}(d(e, e^2, fe)) = d(id, id^2, g\, id) = g.$$

Thus $P \in \mathrm{Tw}(\mathbf{C})$ is an element that $\widehat{h}$ maps to g.

Claim 2.11. *If* $\mathbf{E}, \mathbf{F} \in \mathcal{V}(\mathcal{K})$*, and* $\mathbf{E}$ *is a subalgebra of* $\mathbf{F}$*, then* $\mathrm{Tw}(\mathbf{E})$ *is a homomorphic image of a submonoid of* $\mathrm{Tw}(\mathbf{F})$*.*

To prove this, choose $g \in \mathrm{Tw}(\mathbf{E})$. There is a term $t_{\mathbf{y}}(x)$ and tuples $\mathbf{c}, \mathbf{d}$ from $\mathbf{E}$ such that $t_{\mathbf{c}}(x) = id(x)$ on $\mathbf{E}$ and $t_{\mathbf{d}}(x) = g(x)$. Since $E \subseteq F$, both $e(x) = t_{\mathbf{c}}(x)$ and $f(x) = t_{\mathbf{d}}(x)$ are polynomials of $\mathbf{F}$, and they satisfy $e|_E = id$ and $f|_E = g$ respectively. Let $P \in \mathrm{Tw}(\mathbf{F})$ be such that $Pe = d(e, e^2, fe)$. In particular, $P|_E = Pe|_E = d(e, e^2, fe)|_E = f|_E = g$. Thus every element $g \in \mathrm{Tw}(\mathbf{E})$ is the restriction to E of some $P \in \mathrm{Tw}(\mathbf{F})$. Let $\mathbf{H}$ be the monoid consisting of all $P \in \mathrm{Tw}(\mathbf{F})$ whose restriction to E agrees with some $g \in \mathrm{Tw}(\mathbf{E})$. $\mathbf{H}$ is a submonoid of $\mathrm{Tw}(\mathbf{F})$ and restriction to E determines a homomorphism from $\mathbf{H}$ onto $\mathrm{Tw}(\mathbf{E})$. This establishes Claim 2.11.

Claim 2.12. *Assume that* $\prod_{i\in I} \mathbf{G}_i \in \mathcal{V}$*. Then* $\mathrm{Tw}(\prod_{i\in I} \mathbf{G}_i)$ *is embeddable in* $\prod_{i\in I} \mathrm{Tw}(\mathbf{G}_i)$*.*

By Claim 2.10, the canonical projections $\pi_j \colon \prod_{i\in I} \mathbf{G}_i \to \mathbf{G}_j$ induce homomorphisms

$$\widehat{\pi}_j \colon \mathrm{Tw}\Big(\prod_{i\in I} \mathbf{G}_i\Big) \to \mathrm{Tw}(\mathbf{G}_j).$$

These homomorphisms determine a natural homomorphism

$$\prod \widehat{\pi}_i \colon \mathrm{Tw}\Big(\prod_{i\in I} \mathbf{G}_i\Big) \to \prod_{i\in I} \mathrm{Tw}(\mathbf{G}_i).$$

Since the kernel of $\widehat{\pi}_j$ consists of those pairs of elements in $\mathrm{Tw}(\prod_{i\in I} \mathbf{G}_i)$ that agree modulo π_j, it follows that the kernel of $\prod \widehat{\pi}_i$ consists of those pairs of elements in $\mathrm{Tw}(\prod_{i\in I} \mathbf{G}_i)$ that agree modulo $\bigwedge \pi_i = 0$. Thus, a pair is in $\ker(\prod \widehat{\pi}_i)$ only if it is a pair of equal polynomials. Hence $\prod \widehat{\pi}_i$ is an embedding.

Through Claims 2.10, 2.11 and 2.12, we have shown that if $\mathcal{K}$ generates a congruence permutable variety, then

- $\mathrm{Tw}(\mathsf{H}(\mathcal{K})) \subseteq \mathsf{H}(\mathrm{Tw}(\mathcal{K}))$,
- $\mathrm{Tw}(\mathsf{S}(\mathcal{K})) \subseteq \mathsf{HS}(\mathrm{Tw}(\mathcal{K}))$, and
- $\mathrm{Tw}(\mathsf{P}(\mathcal{K})) \subseteq \mathsf{SP}(\mathrm{Tw}(\mathcal{K}))$.

Therefore

$$\mathrm{Tw}(\mathsf{HSP}(\mathcal{K})) \subseteq \mathsf{H}(\mathsf{HS}(\mathsf{SP}(\mathrm{Tw}(\mathcal{K})))) = \mathsf{HSP}(\mathrm{Tw}(\mathcal{K})).$$

□

We do not know how general the inclusion

$$\mathrm{Tw}(\mathsf{HSP}(\mathcal{K})) \subseteq \mathsf{HSP}(\mathrm{Tw}(\mathcal{K}))$$

is. It does hold in some situations where $\mathcal{K}$ does not generate a congruence permutable variety. For example, if $\mathcal{K}$ is any class of *bounded*

lattices (meaning that there are equationally definable constants 0 and 1 denoting the bottom and top elements), then the inclusion

$$\mathrm{Tw}\,(\mathsf{HSP}(\mathcal{K})) \subseteq \mathsf{HSP}(\mathrm{Tw}\,(\mathcal{K}))$$

holds. This is because if a lattice $\mathbf{L}$ has a top and a bottom element, then the twin monoid coincides with the monoid of all unary polynomials. (To see this, note that $\mathrm{Tw}\,(\mathbf{L}) \subseteq \mathrm{Pol}_1(\mathbf{L})$ trivially. The reverse inclusion holds because if $p(x) \in \mathrm{Pol}_1(\mathbf{L})$ then $p(x) = (p(x) \wedge 1) \vee (x \wedge 0)$ is a twin of $(p(x) \wedge 0) \vee (x \wedge 1) = id(x)$.) Claims 2.10, 2.11 and 2.12 of Theorem 2.8 hold in any variety whose algebras satisfy $\mathrm{Tw}\,(\mathbf{A}) = \mathrm{Pol}_1(\mathbf{A})$, so the theorem itself holds. (The reason this statement is true is that the proofs of Claims 2.10 and 2.11 involve a pair (e, f) of twins that are modified to a pair $(e, Pe) = (d(e, e^2, e^2), d(e, e^2, fe))$ of twins where P is a twin of the identity. Thus, the pair of twins (id, P) can substitute for (e, f) in any situation where e "acts like the identity". Roughly, this has the effect of modifying f, which may not be a twin of the identity, to a polynomial P that is a twin of the identity. But in a variety where all unary polynomials are twins of the identity, there is no need to make this modification.)

Contrary to the situation for bounded lattices, the theorem does not hold for the variety of all unbounded lattices. (*Unbounded* means *not necessarily bounded*.) The reason for this is that if $\mathbf{F}$ is an infinitely generated free lattice, then it can be shown via Whitman's solution to the word problem for lattices that $\mathrm{Tw}\,(\mathbf{F})$ consists of the identity function alone. Therefore, if $\mathcal{K} = \{\mathbf{F}\}$, then $\mathsf{HSP}(\mathrm{Tw}\,(\mathcal{K}))$ is the variety of trivial monoids. Now, since $\mathsf{HSP}(\mathcal{K})$ is the class of all lattices, to show that $\mathrm{Tw}\,(\mathsf{HSP}(\mathcal{K})) \not\subseteq \mathsf{HSP}(\mathrm{Tw}\,(\mathcal{K}))$ it suffices to exhibit one lattice whose twin monoid is not the trivial monoid. This is easy: any nontrivial lattice with a top and a bottom element will do, since $\mathrm{Tw}\,(\mathbf{L}) = \mathrm{Pol}_1(\mathbf{L})$ in this situation.

3. Prime Power Factorization

We prove our main results in this section. Before getting to them we have to introduce one more concept.

Let $\mathbf{A}$ be a finite algebra that generates a congruence modular variety, and assume that $\delta < \theta$ are congruences on $\mathbf{A}$. We will say that the congruence quotient $\langle \delta, \theta \rangle$ *has characteristic* p, where p is a prime, if the size of each θ/δ-class in $\mathbf{A}/\delta$ is a power of p.

Now suppose that $\delta \prec \theta$ and that $[\theta, \theta] \leq \delta$. (These expressions mean that $\langle \delta, \theta \rangle$ is an abelian prime quotient of $\mathrm{Con}\,(\mathbf{A})$.) Then θ/δ is a minimal abelian congruence of $\mathbf{A}/\delta$. There is a natural way, described in Chapter 9 of [3], of constructing a finite simple module on

the product of the θ/δ-classes. Since every finite simple module has prime power cardinality it follows that all θ/δ-classes have size that is a power of some fixed prime p. This shows that any abelian prime quotient has characteristic p for some prime p.

In this section we deal with finite nilpotent algebras in congruence modular varieties. It is shown in Corollary 7.5 of [3] that congruences on such algebras are *uniform*, which means that all blocks have the same size. For congruence uniform algebras the notion of *index* makes sense for any pair of congruences $\delta < \theta$: the index $[\theta\colon \delta]$ is the number of δ-classes in any θ-class. In the congruence uniform situation, $\langle \delta, \theta \rangle$ has characteristic p precisely when $[\theta\colon \delta]$ is a power of p.

In a congruence uniform algebra, if we have a chain of congruence coverings

$$\delta = \alpha_0 \prec \alpha_1 \prec \cdots \prec \alpha_n = \theta,$$

then

$$[\theta\colon \delta] = [\alpha_n\colon \alpha_{n-1}] \cdots [\alpha_1\colon \alpha_0].$$

Therefore $\langle \delta, \theta \rangle$ has characteristic p if and only if each prime quotient in the chain has characteristic p. In particular, these remarks imply the following theorem.

Theorem 3.1. *A finite nilpotent algebra in a congruence modular variety has cardinality that is a power of the prime p if and only if all its prime quotients have characteristic p.*

We will use tame congruence theory (see [4]) as a tool for detecting the characteristic of a prime quotient.

Lemma 3.2. *Let $\mathbf{A}$ be a finite nilpotent algebra that generates a congruence modular variety. If $\delta \prec \theta$ in $\mathrm{Con}(\mathbf{A})$ and the characteristic of $\langle \delta, \theta \rangle$ is p, then the cardinality of any $\langle \delta, \theta \rangle$-minimal set is a power of p.*

Proof. Let's first argue that no generality is lost by assuming that $\delta = 0$. The assumptions on $\mathbf{A}$ imply that $\mathrm{typ}(\delta, \theta) = \mathbf{2}$, and therefore any minimal set $U \in \mathrm{Min}_{\mathbf{A}}(\delta, \theta)$ has p–power cardinality for some prime p. (These claims can be pieced together from Theorems 4.31, 7.2, 8.5, and 13.9 of [4].) Since $\mathbf{A}|_U$ is nilpotent and generates a congruence modular variety, it is an algebra to which Theorem 3.1 applies: all of its prime quotients have characteristic p. In particular, since $\delta|_U < 1_U$ the index $[1_U\colon \delta|_U]$ is a power of p, which means that the quotient algebra $\mathbf{A}|_U/\delta|_U$ has cardinality which is a power of p. The universe of $\mathbf{A}|_U/\delta|_U$ is $U/\delta|_U$, which is a $\langle 0, \theta/\delta \rangle$-minimal set of $\mathbf{A}/\delta$. Therefore the cardinality of the $\langle 0, \theta/\delta \rangle$-minimal set $U/\delta|_U$ is a power

of same prime p that we started with. (These claims follow from Theorem 2.8 (2), Lemma 2.16 (2) and Lemma 4.36 of [4].) Finally, since the characteristic of $\langle \delta, \theta \rangle$ equals the characteristic of $\langle 0, \theta/\delta \rangle$ by definition, we may replace $\mathbf{A}$ by $\mathbf{A}/\delta$ and $\langle \delta, \theta \rangle$ by $\langle 0, \theta/\delta \rangle$, change notation, and assume henceforth that $\delta = 0$.

If N is a $\langle 0, \theta \rangle$-trace of $U \in \mathrm{Min}_{\mathbf{A}}(0, \theta)$, then the facts that N is a congruence class of $\mathbf{A}|_U$ and that $\mathbf{A}|_U$ is congruence uniform of prime power cardinality imply that N has prime power cardinality for the same prime. So what we have left to show is this: if $0 \prec \theta$, then the characteristic of $\langle 0, \theta \rangle$ divides the cardinality of any $\langle 0, \theta \rangle$-trace. One way to see this is to note that any θ-class is an E-trace with respect to θ (meaning that it is the intersection of a θ-class with the image of an idempotent polynomial — take $id(x)$ for the polynomial), and that θ is an abelian minimal congruence of an algebra in a congruence permutable variety. These facts together with Theorem 4.5 of [8] show that any θ-class is a *multitrace of type* **2**. According to the structure theorem for such objects, given in Theorem 3.10 of [7], this means that the size of a θ-class is a power of the size of any $\langle 0, \theta \rangle$-trace. Therefore, since θ-classes and $\langle 0, \theta \rangle$-traces each have prime power cardinality, the primes must agree. This concludes the proof. □

Lemma 3.3. *Let* $\mathbf{A}$ *be an algebra that generates a congruence modular variety. Let* α *be a central congruence on* $\mathbf{A}$*, and let* λ *be an element of* $\mathrm{Tw}(\mathbf{A})$*. If* V *is an* α*-class, and* λ *fixes an element of* V*, then* λ *fixes every element of* V*.*

Proof. Assume that $0 \in V$ is fixed by λ, that $s_{\mathbf{e}}(x) = id(x)$, and $s_{\mathbf{a}}(x) = \lambda(x)$. If $(0, b) \in \alpha$, then since α is a central congruence we can change the underlined entries from 0 to b in

$$s_{\mathbf{e}}(\underline{0}) = 0 = s_{\mathbf{a}}(\underline{0})$$

and preserve the equality of the left and right sides. This yields

$$s_{\mathbf{e}}(\underline{b}) = s_{\mathbf{a}}(\underline{b}),$$

so $\lambda(b) = s_{\mathbf{a}}(b) = s_{\mathbf{e}}(b) = b$. □

Let $h \colon \mathbf{A} \to \mathbf{A}/\alpha$ be the natural homomorphism. We will use the notation $\widehat{\alpha}$ to denote the kernel of the induced homomorphism $\widehat{h} \colon \mathrm{Tw}(\mathbf{A}) \to \mathrm{Tw}(\mathbf{A}/\alpha)$ that we described in the proof of Claim 2.10 of Theorem 2.8. To be explicit, $(\kappa, \lambda) \in \widehat{\alpha}$ provided $\kappa(x) \equiv_\alpha \lambda(x)$ for all $x \in A$.

Lemma 3.4. *Let* $\mathbf{A}$ *be a finite nilpotent algebra that generates a congruence modular variety. Assume that*

(1) $\delta \prec \theta$ *in* $\mathrm{Con}(\mathbf{A})$,
(2) $\lambda \in \mathrm{Tw}(\mathbf{A})$ *has prime power order, and*
(3) $(id, \lambda) \in \widehat{\theta} - \widehat{\delta}$.

Then the order of λ *is a power of the characteristic of* $\langle \delta, \theta \rangle$.

Proof. Factoring modulo δ does not affect the hypotheses or conclusion, so there is no loss of generality in assuming that $\delta = 0$.

Let p be the prime that is the characteristic of $\langle 0, \theta \rangle$ and let q be the prime for which $\lambda^{q^k} = id$ for some k. The assumptions imply that λ maps every θ-class into itself, λ permutes each θ-class, and λ is not the identity on some θ-class. Let V be a θ-class on which λ is not the identity. V is a union of λ-orbits, each of which has size q^r for some r. The size of V is a power of p, so if $q \neq p$ then λ must have a fixed point on V. But now Lemma 3.3 implies that λ is the identity on V, contrary to the choice of V. It must be that $q = p$. □

Lemma 3.5. *Let* $\mathbf{A}$ *be a finite nilpotent algebra that generates a congruence modular variety. If* θ *is a minimal congruence on* $\mathbf{A}$*, then* $\widehat{\theta}$ *is a nontrivial abelian congruence of* $\mathrm{Tw}(\mathbf{A})$*. If* θ *is minimal and* $\langle 0, \theta \rangle$ *has characteristic* p*, then so does* $\langle 0, \widehat{\theta} \rangle$.

Proof. Let $d(x, y, z)$ be a Mal'tsev term for $\mathbf{A}$. To see that $\widehat{\theta} > 0$, choose $(u, v) \in \theta - 0$. Then $f(x) = d(x, u, v)$ is a twin of $d(x, u, u) = id(x)$ and $f(x) = d(x, u, v) \equiv_\theta d(x, u, u) = id(x)$ for all $x \in A$. This shows that $(id, f) \in \widehat{\theta}$. Thus, to prove that $\widehat{\theta}$ is nontrivial it is enough to observe that $f \neq id$ (since $f(u) = v \neq u$).

Now we show that $\widehat{\theta}$ is abelian. Since a group congruence is abelian if and only if the elements congruent to the identity element commute with each other, we must prove that if $e, f \in \mathrm{Tw}(\mathbf{A})$ satisfy $(id, e), (id, f) \in \widehat{\theta}$, then $ef = fe$. Note first that since $(id, e), (id, f) \in \widehat{\theta}$ we have $e(x) \equiv_\theta x$ and $f(x) \equiv_\theta x$ for all $x \in A$, so both e and f map every θ-class into itself. Therefore, to prove that they commute it will suffice to prove that they commute on any θ-class.

Select a θ-class V and pick $0 \in V$. Set $1 = e(0) \equiv_\theta 0$. Since $e(x)$ is a twin of $id(x)$, and $d(x, 1, 0)$ is a twin of $d(x, 0, 0) = id(x)$, it follows from Lemma 2.1 that $e'(x) = d(e(x), 1, 0) \in \mathrm{Tw}(\mathbf{A})$. Moreover, $e'(0) = d(1, 1, 0) = 0$, so according to Lemma 3.3 we must have $e'(x) = x$ on V. Hence $d(x, 1, 0)$ is the inverse of $e(x)$ on V. Similarly, if we define $2 = f(0) \equiv_\theta 0$, then we get that $d(x, 2, 0)$ is the inverse of $f(x)$ on V. To prove that e and f commute on V it is enough to prove that their inverses $d(x, 1, 0)$ and $d(x, 2, 0)$ commute on V. This follows trivially from the fact, proved in Chapter 5 of [3], that on V the operation $d(x, y, z)$ interprets as $x - y + z$ with respect to some abelian group

operations on V (since θ is abelian). Thus, on V, the polynomials we are interested in are just $d(x,1,0) = x-1+0$ and $d(x,2,0) = x-2+0$, which are translations with respect to the abelian group structure on V. Since translations commute, we have $ef = fe$ on V.

Finally we must show that the characteristic of $\langle 0, \widehat{\theta} \rangle$ is the same as $\langle 0, \theta \rangle$. Let N be the normal subgroup of $\mathrm{Tw}(\mathbf{A})$ consisting of elements that are $\widehat{\theta}$-related to *id*. We need to show that N is a p-group for the prime p that is the characteristic of $\langle 0, \theta \rangle$. If this is not the case, then N contains an element λ of order q where q is a prime different from p. Lemma 3.4 proves that this is impossible. □

Corollary 3.6. *Let* $\mathbf{A}$ *be a finite nilpotent algebra that generates a congruence modular variety.* $\mathrm{Tw}(\mathbf{A})$ *is solvable. If, moreover,* $\mathbf{A}$ *is a direct product of algebras of prime power cardinality, then* $\mathrm{Tw}(\mathbf{A})$ *is nilpotent.*

Proof. Choose a sequence of congruences

$$0 = \theta_0 \prec \theta_1 \prec \cdots \prec \theta_n = 1.$$

This chain induces a chain

$$0 = \widehat{\theta}_0 < \widehat{\theta}_1 < \cdots < \widehat{\theta}_n = 1$$

of congruences on $\mathrm{Tw}(\mathbf{A})$. Moreover, since $[1, \theta_{i+1}] \leq \theta_i$ for each i, Lemma 3.5 applied to $\mathbf{A}/\theta_i$ shows that $[\widehat{\theta}_{i+1}, \widehat{\theta}_{i+1}] \leq \widehat{\theta}_i$ in $\mathrm{Tw}(\mathbf{A})$. This proves that $\mathrm{Tw}(\mathbf{A})$ is solvable.

Now suppose that $\mathbf{A} = \mathbf{A}_1 \times \cdots \times \mathbf{A}_k$ where each $\mathbf{A}_i$ has prime power cardinality. Then for $\mathcal{K} = \{\mathbf{A}_1, \ldots, \mathbf{A}_k\}$ we have that $\mathbf{A} \in \mathsf{HSP}(\mathcal{K})$. Since we proved in Theorem 2.8 that $\mathrm{Tw}(\mathsf{HSP}(\mathcal{K})) \subseteq \mathsf{HSP}(\mathrm{Tw}(\mathcal{K}))$, to prove that $\mathrm{Tw}(\mathbf{A})$ is nilpotent it will suffice to prove that each $\mathrm{Tw}(\mathbf{A}_i)$ is nilpotent.

Fix one $\mathbf{A}_i$ and pick a chain of congruences as above:

$$0 = \theta_0 \prec \theta_1 \prec \cdots \prec \theta_n = 1.$$

Since $\mathbf{A}_i$ has cardinality that is a power of some prime p, therefore each $\langle \theta_i, \theta_{i+1} \rangle$ has characteristic equal to this p. From what we have proved already, this chain induces a chain

$$0 = \widehat{\theta}_0 < \widehat{\theta}_1 < \cdots < \widehat{\theta}_n = 1$$

in $\mathrm{Con}(\mathrm{Tw}(\mathbf{A}_i))$ where each $\langle \widehat{\theta}_i, \widehat{\theta}_{i+1} \rangle$ has characteristic equal to p. Hence $\mathrm{Tw}(\mathbf{A}_i)$ is a p-group. Since p-groups are nilpotent, the proof is complete. □

We now prove a partial converse to the second claim of Corollary 3.6. The full converse is proved in Theorem 3.12

Theorem 3.7. *Let* $\mathbf{A}$ *be a finite, subdirectly irreducible, nilpotent algebra that generates a congruence modular variety. If* $\mathrm{Tw}(\mathbf{A})$ *is nilpotent, then* $\mathbf{A}$ *has prime power cardinality.*

Proof. Let μ denote the monolith of $\mathbf{A}$ and let p be the characteristic of $\langle 0, \mu\rangle$. If every prime quotient of $\mathbf{A}$ has characteristic p, then the cardinality of $\mathbf{A}$ is a power of p. Therefore, to establish the theorem, we must prove that if $\mathbf{A}$ has a prime quotient of characteristic $q \neq p$, then $\mathrm{Tw}(\mathbf{A})$ is not nilpotent. Assume that $\langle \delta, \theta\rangle$ is a prime quotient of $\mathbf{A}$ with characteristic $q \neq p$.

Claim 3.8. $\mathrm{Tw}(\mathbf{A})$ *has an element* λ *of order* q*.*

Since the characteristic of $\langle \delta, \theta\rangle$ is q, the characteristic of $\langle \widehat{\delta}, \widehat{\theta}\rangle$ is also q, as one can deduce by applying Lemma 3.5 to $\mathbf{A}/\delta$. Therefore q divides the order of $\mathrm{Tw}(\mathbf{A})$. Cauchy's Theorem implies that there is a $\lambda \in \mathrm{Tw}(\mathbf{A})$ of order q.

Claim 3.9. *If* $\lambda \in \mathrm{Tw}(\mathbf{A})$ *has order* q*, then there exist* $u, v, w \in A$ *and* $g \in \mathrm{Pol}_1(\mathbf{A})$ *such that*

(1) $\lambda(u) = v$,
(2) $(u, v) \notin \mu$,
(3) $(g(u), g(v)) = (u, w) \in \mu - 0$, *and*
(4) $\forall x \neq y \in A \left((x, y) \notin \mathrm{Cg}^{\mathbf{A}}(g(x), g(y))\right)$.

Choose θ minimal so that $\lambda(x) \equiv_\theta x$. (Equivalently, θ is a congruence minimal for the property that $\widehat{\theta}$ contains (id, λ).) Since $\lambda \neq id$, we do not have $\theta = 0$. Thus there is a $\delta \prec \theta$. By the minimality of θ, there is an element $u \in A$ that has the property that $\lambda(u) \not\equiv_\delta u$. This is the element we take for u, and $\lambda(u)$ is the element we take for v. Already we have that (1) holds, and that $(u, v) \in \theta - \delta$.

The characteristic of $\langle \delta, \theta\rangle$ must be $q \neq p$. This follows from Lemma 3.4 and the facts that the order of λ is q and $(id, \lambda) \in \widehat{\theta} - \widehat{\delta}$. In particular, we cannot have $\langle \delta, \theta\rangle = \langle 0, \mu\rangle$ since the characteristics differ. So, if we had $(u, v) \in \mu \leq \delta$, then we would contradict $(u, v) \notin \delta$. We conclude that (2) holds.

Congruence uniformity allows us to choose an element $w \in A$ for which $(u, w) \in \mu - 0$. We need to locate a polynomial g for which (3) and (4) hold. Let U be a $\langle \delta, \theta\rangle$-minimal set, and let h be a unary polynomial of $\mathbf{A}$ for which $h(A) = U$ and $(h(u), h(v)) \in \theta|_U - \delta$. Such an h exists by Theorem 2.8 (4) of [4]. As is the case for any algebra, if $a, b \in A$, then the set of all pairs of the form $(k(a), k(b))$, where k runs over all unary polynomials of $\mathbf{A}$, is equal to the diagonal subalgebra of $\mathbf{A}^2$ generated by (a, b). In a congruence permutable variety the

diagonal subalgebras of $\mathbf{A}^2$ are precisely the congruences, so the set of all $(k(a), k(b))$ is precisely $\mathrm{Cg}^{\mathbf{A}}(a, b)$. Thus, in the particular case where $(a, b) = (h(u), h(v))$, the fact that $\mathrm{Cg}^{\mathbf{A}}(h(u), h(v)) \geq \mu = \mathrm{Cg}^{\mathbf{A}}(u, w)$ implies that there is a unary polynomial k such that $(kh(u), kh(v)) = (u, w)$. We take $g = kh$. It is automatic that (3) holds.

To see that (4) holds, assume to the contrary that there exist distinct elements $x, y \in A$ such that $(x, y) \in \mathrm{Cg}^{\mathbf{A}}(g(x), g(y))$. Fix such x and y and let $\beta = \mathrm{Cg}^{\mathbf{A}}(g(x), g(y)) = \mathrm{Cg}^{\mathbf{A}}(x, y)$. Since $x \neq y$ it is possible to choose $\alpha \prec \beta$. Then, since $(x, y) \in \beta$ and $(g(x), g(y)) \notin \alpha$ we get that $g(\beta) = kh(\beta) \not\subseteq \alpha$. In particular, $h(\beta) \not\subseteq \alpha$. From this, and Definition 2.5 of [4], the set $h(A) = U \in \mathrm{Min}_{\mathbf{A}}(\delta, \theta)$ contains an $\langle \alpha, \beta \rangle$-minimal set. But algebras induced on minimal sets of $\mathbf{A}$ are E-minimal, by Theorems 4.31, 7.2, and 8.5 of [4], and there are no proper containments between induced E-minimal algebras. Thus U is itself an $\langle \alpha, \beta \rangle$-minimal set, and moreover by Theorem 2.8 (1) of [4] we conclude that $\mathrm{Min}_{\mathbf{A}}(\delta, \theta) = \mathrm{Min}_{\mathbf{A}}(\alpha, \beta)$. Next, since $g(\beta) = kh(\beta) \not\subseteq \alpha$ and $U = h(A)$, it follows that $k(\beta|_U) \not\subseteq \alpha$. Hence, by Theorem 2.8 (3) of [4] we have $k(U) \in \mathrm{Min}_{\mathbf{A}}(\delta, \theta)$. But we chose k so that it contains distinct μ-related elements u and w in its range. Thus, $\mu|_{k(U)} > 0$, and it follows that μ restricts nontrivially to any $\langle \delta, \theta \rangle$-minimal set. But this implies that each $\langle \delta, \theta \rangle$-minimal set contains a $\langle 0, \mu \rangle$-minimal set. As already noted, such a containment cannot be proper, so we are led to

$$\mathrm{Min}_{\mathbf{A}}(\alpha, \beta) = \mathrm{Min}_{\mathbf{A}}(\delta, \theta) = \mathrm{Min}_{\mathbf{A}}(0, \mu).$$

Now we have a characteristic problem: Lemma 3.2 proves that, since the characteristics of $\langle 0, \mu \rangle$ and $\langle \delta, \theta \rangle$ differ, their minimal sets do not have the same size. This contradiction establishes (4), and completes the proof of Claim 3.9.

Claim 3.10. *If u, v, w and g have the properties listed in Claim 3.9, and $d(x, y, z)$ is a Mal'tsev term for $\mathbf{A}$, then $\Sigma(x) := d(x, g(x), u)$ is a permutation of A that fixes all elements of the μ-class of u, maps the μ-class of v into itself, and moves all elements of the μ-class of v.*

Suppose that $\Sigma(a) = \Sigma(b)$. Let $\theta = \mathrm{Cg}^{\mathbf{A}}(g(a), g(b))$. We have

$$a' := d(a, g(a), u) = \Sigma(a) = \Sigma(b) = d(b, g(b), u) \equiv_\theta d(b, g(a), u) =: b'.$$

By Corollary 7.4 of [3] the mapping $d(x, g(a), u)$ is a polynomial permutation that has a polynomial inverse. If the inverse is $p(x)$, then

$$a = p(a') \equiv_\theta p(b') = b.$$

Hence $(a, b) \in \theta = \mathrm{Cg}^{\mathbf{A}}(g(a), g(b))$. By property (4) of Claim 3.9 we conclude that $a = b$. This shows that Σ is 1-1, and therefore is a permutation of $\mathbf{A}$.

Arbitrarily choose u' from the μ-class of u. Then, by the properties of g, we have that $g(u') = g(u) = u$. Thus $\Sigma(u') = d(u', u, u) = u'$, and so Σ fixes every element of u/μ.

Finally, choose v' from the μ-class of v. By the properties of g we have $g(v') = g(v) = w \equiv_\mu u$, so $\Sigma(v') = d(v', w, u) \equiv_\mu d(v, u, u) = v$. This proves that Σ maps the μ-class of v into itself. If $\Sigma(v') = v'$, then

$$d(\underline{v'}, w, u) = \Sigma(v') = v' = d(\underline{v'}, u, u).$$

Applying the $1, \mu$-term condition to the underlined position we get that

$$u = d(\underline{w}, w, u) = d(\underline{w}, u, u) = w.$$

But $u \neq w$, so we cannot have $\Sigma(v') = v'$. Thus, Σ fixes no element of v/μ.

Claim 3.11. *No q-Sylow subgroup of* $\mathrm{Tw}(\mathbf{A})$ *is normal.*

The polynomial permutation $\Sigma^{-1} \circ \lambda \circ \Sigma$ has order q, since it is a conjugate of λ and λ has order q. Moreover, λ is a twin of id, so $\Sigma^{-1} \circ \lambda \circ \Sigma$ is a twin of $\Sigma^{-1} \circ id \circ \Sigma = id$ by Lemma 2.1. This shows that both λ and $\Sigma^{-1} \circ \lambda \circ \Sigma$ are elements of $\mathrm{Tw}(\mathbf{A})$, and both have order q.

If a q-Sylow subgroup was normal it would be the unique q-Sylow subgroup, and this would force it to contain all elements of $\mathrm{Tw}(\mathbf{A})$ whose order is a power of q. In particular, it would contain both λ and $\Sigma^{-1} \circ \lambda \circ \Sigma$, and therefore it would contain the element $\Gamma := \lambda^{-1} \circ \Sigma^{-1} \circ \lambda \circ \Sigma$. If so, the order of Γ would be a power of q. We show that this is not the case.

Let's evaluate Γ at u:

	Evaluation:	Justification:
$\Gamma(u)$	$= \lambda^{-1} \circ \Sigma^{-1} \circ \lambda \circ \Sigma(u)$	Defn. of Γ
	$= \lambda^{-1} \circ \Sigma^{-1} \circ \lambda(u)$	$\Sigma(u) = u$
	$= \lambda^{-1} \circ \Sigma^{-1}(v)$	$\lambda(u) = v$
	$\equiv_\mu \lambda^{-1}(v)$	$\Sigma^{-1}(v) \equiv_\mu v$
	$= u$	$\lambda^{-1}(v) = u$

This proves two things. First, since Γ is a polynomial that maps u back into its μ-class, therefore Γ maps the entire μ-class of u into itself. Second, Γ moves u. For if not, then in the above derivation we would have equality at the beginning and end. This would force equality on the fourth line: $\Sigma^{-1}(v) = v$. But this contradicts the part of Claim 3.10 that asserts that Σ has no fixed points μ-related to v.

Thus Γ acts on u/μ in a way that moves u. If Γ has order that is a power of q, then the fact that u/μ has cardinality that is a power of p, and $q \neq p$, means that Γ must have a fixed point on u/μ. Now we have a contradiction to Lemma 3.3: $\Gamma \in \mathrm{Tw}(\mathbf{A})$ has a fixed point on u/μ, but does not fix all elements of u/μ. This proves the claim.

Claim 3.11 finishes the proof of the theorem, because all Sylow subgroups of a nilpotent group are normal. □

Two varieties $\mathcal{V}_1$ and $\mathcal{V}_2$ in the same language are *independent* if there is a binary term $t(x, y)$ for which

$$\mathcal{V}_1 \models t(x, y) = x \quad \text{and} \quad \mathcal{V}_2 \models t(x, y) = y.$$

When this is so, then $\mathcal{V}_1$ intersects $\mathcal{V}_2$ trivially, and any algebra in the join $\mathcal{V}_1 \vee \mathcal{V}_2$ factors as a direct product of an algebra in $\mathcal{V}_1$ and an algebra in $\mathcal{V}_2$; moreover, all homomorphisms between algebras in $\mathcal{V}_1 \vee \mathcal{V}_2$ respect these direct factorizations. We write $\mathcal{V}_1 \times \mathcal{V}_2$ to denote the join of $\mathcal{V}_1$ and $\mathcal{V}_2$ when they are independent.

It is not difficult to prove that if $\mathcal{V}_1$ and $\mathcal{V}_2$ are subvarieties of a congruence permutable variety and $\mathcal{V}_1 \cap \mathcal{V}_2$ is trivial, then $\mathcal{V}_1$ and $\mathcal{V}_2$ are independent. We will use this fact in the proof of the next theorem. (We only use the fact in the situation when $\mathcal{V}_1$ and $\mathcal{V}_2$ are subvarieties of a *nilpotent* congruence permutable variety. In this situation the fact is a special case of Theorem 11.3 of [3]. However, nilpotence is not a necessary hypothesis.)

Theorem 3.12. *Let* $\mathbf{A}$ *be a finite nilpotent algebra that generates a congruence modular variety.* $\mathbf{A}$ *factors into a direct product of algebras of prime power cardinality if and only if* $\mathrm{Tw}(\mathbf{A})$ *is nilpotent.*

Proof. We proved in Corollary 3.6 that if $\mathbf{A}$ factors into a direct product of algebras of prime power cardinality, then $\mathrm{Tw}(\mathbf{A})$ is nilpotent. Here we prove the reverse direction, so assume that $\mathrm{Tw}(\mathbf{A})$ is nilpotent.

For each prime p, let $\mathcal{K}_p$ denote the class of subdirectly irreducible homomorphic images of $\mathbf{A}$ whose monolith has characteristic p. Let $\mathcal{V}_p$ denote the subvariety of $\mathcal{V}(\mathbf{A})$ generated by $\mathcal{K}_p$.

Claim 3.13. *For each prime* p*, each finite algebra in* $\mathcal{V}_p$ *has order that is a power of* p*.*

To show this, first note that by Claim 2.10 of Theorem 2.8 the twin group of each member of $\mathcal{K}_p$ is a homomorphic image of $\mathrm{Tw}(\mathbf{A})$, which we assumed to be a nilpotent group. Therefore the twin groups of members of $\mathcal{K}_p$ are nilpotent. Since $\mathcal{K}_p$ consists of subdirectly irreducible algebras whose monolith has characteristic p, it is a consequence of

Theorem 3.7 that all members of $\mathcal{K}_p$ have cardinality that is a power of p. For each $\mathbf{S} \in \mathcal{K}_p$ we can apply Lemma 3.5 repeatedly to successive quotients to obtain that the twin group $\mathrm{Tw}(\mathbf{S})$ is a p-group for the same p. As proved in Theorem 2.8, $\mathrm{Tw}(\mathsf{HSP}(\mathcal{K}_p)) \subseteq \mathsf{HSP}(\mathrm{Tw}(\mathcal{K}_p))$, therefore any finite algebra in the variety $\mathcal{V}_p$ has a twin group that is a p-group for this p. Now it cannot be that some finite $\mathbf{C} \in \mathcal{V}_p$ has cardinality divisible by a prime $q \neq p$, for if this happened then q would appear as the characteristic of some prime quotient of $\mathbf{C}$, and thus it would appear as the characteristic of some congruence quotient of the p-group $\mathrm{Tw}(\mathbf{C})$. This proves the claim.

Claim 3.13 implies that if p and q are distinct primes, then no nontrivial algebra in $\mathcal{V}_p$ has cardinality equal to the cardinality of an algebra in $\mathcal{V}_q$. Thus $\mathcal{V}_p$ and $\mathcal{V}_q$ intersect trivially. As observed directly before the statement of the theorem, this means that $\mathcal{V}_p \vee \mathcal{V}_q = \mathcal{V}_p \times \mathcal{V}_q$. Moreover, if r is a prime different from both p and q, then since the order of a finite algebra in $\mathcal{V}_r$ is a power of r and the order of a finite algebra in $\mathcal{V}_p \times \mathcal{V}_q$ is a product of p's and q's, we get that $\mathcal{V}_r$ intersects trivially with $\mathcal{V}_p \times \mathcal{V}_q$. Thus

$$(\mathcal{V}_p \vee \mathcal{V}_q) \vee \mathcal{V}_r = (\mathcal{V}_p \times \mathcal{V}_q) \vee \mathcal{V}_r = (\mathcal{V}_p \times \mathcal{V}_q) \times \mathcal{V}_r.$$

Generalizing this, if $p_1, \ldots, p_k$ is the sequence of primes p for which $\mathcal{K}_p$ is nonempty, then the variety generated by the union of the $\mathcal{K}_{p_i}$'s is

$$\mathcal{V}_{p_1} \times \cdots \times \mathcal{V}_{p_k}.$$

In particular, since $\mathbf{A}$ is a subdirect product of algebras in $\bigcup_{i=1}^{k} \mathcal{K}_{p_i}$, we get that $\mathbf{A}$ is in this variety. Hence $\mathbf{A}$ is a product of algebras of prime power cardinality. □

A *commutator term* for an algebra $\mathbf{A}$ is a term $\omega(x_1, \ldots, x_r, z)$ such that

$$\mathbf{A} \models \omega(x_1, \ldots, \overset{i\text{-th}}{z}, \ldots, x_r, z) = z$$

for each i. We call a commutator term $\omega(x_1, \ldots, x_r, z)$ *nontrivial* if

$$\mathbf{A} \not\models \omega(x_1, \ldots, x_r, z) = z.$$

The *rank* of a nontrivial commutator term is the number r that appears in these equations.

Theorem 3.14. *Let* $\mathbf{A}$ *be a finite nilpotent algebra of finite type that generates a congruence modular variety. The following conditions are equivalent.*

(1) $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$.
(2) $\mathrm{Tw}(\mathbf{A})$ *is nilpotent.*

(3) $\mathbf{A}$ *factors as a direct product of algebras of prime power cardinality.*
(4) $\mathbf{A}$ *has a finite bound on the rank of nontrivial commutator terms.*
(5) $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{p(k)}$ *for some polynomial* $p(k)$.

Proof. Corollary 2.6 proves that (1) ⇒ (2). Theorem 3.12 proves that (2) ⇒ (3). Theorem 14.9 of [3] proves that (3) ⇒ (4). (The implication (3) ⇒ (4) is the only place where we need to assume that $\mathbf{A}$ has finite type.) The implication (4) ⇒ (5) is established in the proof of Theorem 1 of [1]. The implication (5) ⇒ (1) is trivial. □

In fact, it is easy to characterize those algebras in congruence modular varieties that have small free spectrum without assuming finite type, as we did in Theorem 3.14.

Corollary 3.15. *Let* $\mathbf{A}$ *be a finite algebra for which* $\mathrm{typ}\{\mathcal{V}(\mathbf{A})\} \cap \{\mathbf{1}, \mathbf{5}\} = \emptyset$. *Then* $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$ *if and only if*

(1) $\mathbf{A}$ *has a finitely generated clone, and*
(2) $\mathbf{A}$ *factors as a direct product of nilpotent algebras of prime power cardinality.*

Proof. First assume that (1) and (2) hold. Given (1), there is no loss of generality assuming that $\mathbf{A}$ has finite type. By (2), $\mathbf{A}$ is nilpotent; therefore $\mathcal{V}(\mathbf{A})$ is *locally solvable*, a concept introduced in [4]. Since are assuming that $\mathbf{1} \notin \mathrm{typ}\{\mathcal{V}(\mathbf{A})\}$ we conclude from Theorem 7.11 of [4] and the fact that $\mathcal{V}(\mathbf{A})$ is locally solvable that $\mathcal{V}(\mathbf{A})$ is congruence permutable. Hence Theorem 3.14 (3)⇒(1) shows that $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$.

Now assume that $\mathrm{typ}\{\mathcal{V}(\mathbf{A})\} \cap \{\mathbf{1}, \mathbf{5}\} = \emptyset$ and $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$. By Theorem 12.5 of [4], $\mathbf{A}$ is a finite nilpotent algebra that generates a congruence permutable variety. The proof of Theorem 1 of [1] shows that in this situation $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$ if and only if there is a finite bound on the rank of nontrivial commutator terms, *even if the type of* $\mathbf{A}$ *is infinite*. But Lemma 14.6 of [3] proves that if there is a finite bound on the rank of nontrivial commutator terms, then the clone of $\mathbf{A}$ is finitely generated. (What is shown there is that the clone of $\mathbf{A}$ is generated by a fixed Mal'tsev term from the clone, a collection of unary terms representing all unary terms, and a collection of commutator terms representing all nontrivial commutator terms.) Thus $\mathrm{typ}\{\mathcal{V}(\mathbf{A})\} \cap \{\mathbf{1}, \mathbf{5}\} = \emptyset$ and $\mathrm{Spec}_{\mathbf{A}}(k) \ll 2^{2^{ck}}$ imply that (1) holds. Now, given (1) and that $\mathbf{A}$ is nilpotent we can derive (2) from Theorem 3.14 (1) ⇒ (3). □

Acknowledgment. Valuable discussions with Joel Berman led to the discovery of these results.

References

[1] J. Berman and W. Blok, *Free spectra of nilpotent varieties*, Algebra Universalis **24** (1987), 279–282.
[2] G. Higman, *The orders of relatively free groups*, Proc. Intern. Conf. Theory of Groups, Austral. Nat. Univ. Canberra, 1965, 153–165.
[3] R. Freese and R. McKenzie, *Commutator Theory for Congruence Modular Varieties*, LMS Lecture Notes v. 125, Cambridge University Press, 1987.
[4] D. Hobby and R. McKenzie, *The Structure of Finite Algebras*, Contemporary Mathematics v. 76, American Mathematical Society, 1988.
[5] K. A. Kearnes, *An order-theoretic property of the commutator*, Internat. J. Algebra and Comput. **3** (1993), 491–533.
[6] K. A. Kearnes and E. W. Kiss, *Finite algebras of finite complexity*, to appear in Discrete Mathematics.
[7] K. A. Kearnes, E. W. Kiss and M. A. Valeriote, *Minimal sets and varieties*, Trans. Amer. Math. Soc. **350** (1998), 1–41.
[8] K. A. Kearnes, E. W. Kiss and M. A. Valeriote, *A geometric consequence of residual smallness*, to appear in Ann. Pure Appl. Logic.
[9] P. M. Neumann, *Some indecomposable varieties of groups*, Quart. J. Math. Oxford **14** (1963), 46–50.
[10] M. R. Vaughan-Lee, *Nilpotence in permutable varieties*, Universal Algebra and Lattice Theory, 293–308, Lecture Notes in Mathematics v. 1004, Springer Verlag, Berlin-New York, 1983.

(Keith A. Kearnes) Department of Mathematics, University of Louisville, Louisville, KY 40292, USA
E-mail address: kearnes@louisville.edu