VAUGHT'S CONJECTURE FOR VARIETIES

BRADD HART, SERGEI STARCHENKO, AND MATTHEW VALERIOTE

ABSTRACT. We prove that if $\mathcal{V}$ is a superstable variety or one with few countable models then $\mathcal{V}$ is the varietal product of an affine variety and a combinatorial variety. Vaught's conjecture for varieties is an immediate consequence.

Much of the work of modern model theory arises out of the work of Shelah on the classification program for first order theories (see [23]). One of the crowning achievements is an abstract decomposition theorem for models of well-behaved theories. More precisely, if T is a complete classifiable¹ theory in a countable language then every model of T is prime over the union of an independent tree of countable models (see [4 or 12]). For any particular theory T , it is fair to ask if T has a more concrete decomposition theorem.

A variety is a class of algebras in a language $\mathcal{L}$ containing only function symbols and constants which is closed under homomorphic images, submodels and products. In this paper, T will be an equational theory in $\mathcal{L}$. The class of models of T is a variety. In general one can ask whether there is an algebraic decomposition theorem if all completions of T are classifiable.

Vaught's conjecture is the statement that if there are fewer than $2^{\aleph_0}$ many countable models for some countable theory T then there are only countably many countable models. As is the situation in many cases, the structural information gained by considering the uncountable models of a variety leads to an understanding of the countable models as well. To describe some of the early work in this area we need some definitions.

Definition 0.1. 1. If $\mathcal{A}$ is an algebra in the language $\mathcal{L}$ then a polynomial of $\mathcal{A}$ is the interpretation in $\mathcal{A}$ of a term in the language $\mathcal{L}$ together with constants for every $a \in A$.

2. Two algebras in possibly two different languages are said to be polynomially equivalent if they have the same underlying set and the same polynomials.

3. An algebra is affine if it is polynomially equivalent to some left R -module for some ring R .

4. An algebra $\mathcal{A}$ is called abelian if for every term $\tau(\bar{x}, \bar{y})$ and every $\bar{a}, \bar{b}, \bar{c}$ and $\bar{d} \in A$, if $\tau(\bar{a}, \bar{c}) = \tau(\bar{b}, \bar{c})$ then $\tau(\bar{a}, \bar{d}) = \tau(\bar{b}, \bar{d})$.

5. A term $\tau(x, y, z)$ is called a Malcev term for an algebra $\mathcal{A}$ if $\tau(x, x, y) = y$ and $\tau(x, y, y) = x$ hold in $\mathcal{A}$.

Received by the editors October 15, 1991.

1991 *Mathematics Subject Classification.* Primary 03C45; Secondary 03C05, 03C60.

All authors were supported by the NSERC.

¹Superstable, NDOP, NOTOP.

6. A variety is said to be affine or abelian if all its algebras are.
An important connection between these ideas is found in [9].

Theorem 0.2. *An algebra $\mathcal{A}$ is affine iff it is abelian and has a Malcev term.*

Another definition which will be important here is a strengthening of abelian.

Definition 0.3. An algebra $\mathcal{A}$ is called combinatorial if for every term $\tau(\bar{x}, \bar{y})$ and every $\bar{a}, \bar{b}, \bar{c}, \bar{d}$ and $\bar{e} \in A$, if $\tau(\bar{a}, \bar{c}) = \tau(\bar{b}, \bar{d})$ then $\tau(\bar{a}, \bar{e}) = \tau(\bar{b}, \bar{e})$. As usual, a variety is said to be combinatorial if all the algebras in it are.

Remark. This terminology differs from the literature where a combinatorial algebra is usually called strongly abelian or is said to satisfy the strong term condition.

A simple example of a combinatorial algebra is one which is essentially unary. More examples of combinatorial algebras and some discussion can be found in the introduction of [11]. We will need two specific examples of combinatorial varieties for what we are about to say.

Definition 0.4. Suppose G is a group. Let $\mathcal{L}_G$ be the language with a unary function symbol for each $g \in G$. We will write g for this symbol. The variety of G -sets is axiomatized by $(gh)(x) = g(h(x))$ for all $g, h \in G$ and $e(x) = x$ where e is the identity of G . This is just the variety of sets with the group G acting on them.

Sets is the variety in the empty language with no axioms.

One final definition is

Definition 0.5. If $\mathcal{V}_1$ and $\mathcal{V}_2$ are subvarieties of $\mathcal{V}$ then $\mathcal{V} = \mathcal{V}_1 \otimes \mathcal{V}_2$ means that $\mathcal{V}$ is the variety generated by $\mathcal{V}_1$ and $\mathcal{V}_2$ and moreover there is a term $d(x, y)$ so that $d(x, y) = x$ holds in $\mathcal{V}_1$ and $d(x, y) = y$ holds in $\mathcal{V}_2$. $\mathcal{V}$ is called the varietal product of $\mathcal{V}_1$ and $\mathcal{V}_2$ and d is called a diagonal term.

Fact 0.6. If $\mathcal{V}_1$ and $\mathcal{V}_2$ are subvarieties of $\mathcal{V}$ and $\mathcal{V} = \mathcal{V}_1 \otimes \mathcal{V}_2$ then for every $\mathcal{A} \in \mathcal{V}$ there is $\mathcal{A}_1 \in \mathcal{V}_1$ and $\mathcal{A}_2 \in \mathcal{V}_2$ so that $\mathcal{A} \cong \mathcal{A}_1 \times \mathcal{A}_2$ and moreover the $\mathcal{A}_i$'s are unique up to isomorphism.

Discussion 0.7. If $\mathcal{A}$ is an $\mathcal{L}$ -algebra then its n th matrix power is an algebra $\mathcal{B}$ in the language $\mathcal{L}' = \mathcal{L} \cup \{d, \rho\}$ where d is n -ary and ρ is unary. The universe of $\mathcal{B}$ is A^n and the constants and functions from $\mathcal{L}$ are interpreted as in $\mathcal{A}^n$. d is interpreted in $\mathcal{B}$ by

$$d((a_1^1, \dots, a_n^1), \dots, (a_1^n, \dots, a_n^n)) = (a_1^1, \dots, a_n^n)$$

and ρ by

$$\rho((a_1, \dots, a_n)) = (a_2, \dots, a_n, a_1).$$

If $\mathcal{V}$ is a variety then the n th varietal power of $\mathcal{V}$ is the variety in $\mathcal{L}'$ generated by the n th matrix power of the free algebra on countably many generators.

The terminology is justified by considering modules. If $\mathcal{V}$ is the variety of all left R -modules for some ring R then the n th varietal power of $\mathcal{V}$ is polynomially equivalent to the variety of left $M_n(R)$ -modules.

The following fact explains why when one is considering the number of models in certain cardinalities, one gets results only up to varietal power. Let $I(\mathcal{V}, \lambda)$ be the number of nonisomorphic models in $\mathcal{V}$ of cardinality λ .

Fact 0.8. If $\mathcal{V}$ is any variety in a language $\mathcal{L}$ and $\mathcal{V}'$ is any varietal power of $\mathcal{V}$ then $I(\mathcal{V}, \lambda) = I(\mathcal{V}', \lambda)$ for all $\lambda \geq |\mathcal{L}|$.

In [1], Baldwin and Lachlan proved, using purely model theoretic means, that if a variety is $\aleph_0$ -categorical then it is $\aleph_1$ -categorical. Their proof did not give an algebraic characterization of $\aleph_1$ -categorical varieties. Givant [6, 7] and Palyutin [19] gave an algebraic characterization of these varieties, and in the $\aleph_0$ -categorical case, McKenzie [15] further refined their results. What is shown is that if $\mathcal{V}$ is an $\aleph_0$ -categorical variety then it is polynomially equivalent to a varietal power of either the variety of sets or a variety of vector spaces over a division ring. McKenzie's proof differs from those of Givant and Palyutin in that his does not rely on any sophisticated model theory, but rather uses tame congruence theory. We will say more about this shortly.

The case of affine varieties is of course tied closely to the case of varieties of modules. In [3], Baur proved that all modules are stable. In [5], Garavaglia gave algebraic characterizations, via conditions on chains of positive primitive subgroups, of when a module is superstable or ω -stable. Since varieties are closed under products, a variety of modules is superstable (all completions are superstable) exactly when all modules are ω -stable. In [2] it is shown that if the variety of modules over a countable ring has few countable models then every module over the ring is ω -stable. Rings with this property must be left pure-semisimple. Using these and other facts, Baldwin and Mackenzie characterized in [2] the possible spectrum functions for congruence modular varieties defined in a countable language. In particular, Vaught's conjecture for congruence modular (and as a special case, affine) varieties is verified.

Classifiable combinatorial varieties were dealt with in [11]. A complete characterization of such varieties was also given. In [10], it is shown that the assumption of superstability is enough to obtain this result.

Definition 0.9. 1. A variety is superstable (stable) if the complete theory of each of its models is superstable (stable).

2. A variety in a countable language is said to have few countable models if the number of nonisomorphic countable models is less than $2^{\aleph_0}$.

Definition 0.10. A variety $\mathcal{V}$ is structured if there are two subvarieties $\mathcal{A}$ and $\mathcal{S}$ of $\mathcal{V}$ which are respectively affine and combinatorial so that $\mathcal{V} = \mathcal{A} \otimes \mathcal{S}$.

We are now in the position to state the main theorems of this paper.

Theorem 0.11. *If $\mathcal{V}$ is a superstable variety then $\mathcal{V}$ is structured.*

Theorem 0.12. *If $\mathcal{V}$ is a variety with few countable models then $\mathcal{V}$ is structured.*

These theorems say that the case of superstable varieties and varieties with few countable models comes down to understanding affine and combinatorial varieties with the same properties. We have already said that a superstable affine variety is polynomially equivalent to a variety of left R -modules over a ring R which is left pure-semisimple. Let us be equally explicit about the combinatorial case.

The easiest situation is a combinatorial variety with few countable models. Clearly any variety of G -sets when G is a finite group is an example of a variety with few countable models. Any combinatorial variety with few countable models is polynomially equivalent to the varietal product of finitely many varietal

powers of G -sets for certain finite groups G . This description is implicit in [11] but not stated in the same way as here. An immediate corollary to Theorem 0.12 is

Corollary 0.13. *Vaught's conjecture holds for varieties.*

The case of a superstable combinatorial variety needs a little more preparation. Suppose that $\mathcal{L}$ is a multi-sorted language with sorts $U_1, \dots, U_n$ and only unary function symbols and constants. This means that every function symbol is unary with one sort as domain and one sort as range. The constants are also sorted. Since we are only going to describe these varieties up to polynomial equivalence we will assume that there is a constant c_i with sort U_i for every i . Any variety in such a language is called a multi-sorted unary variety.

Now in the language $\mathcal{L}$, we define the notion of a linear variety (see [11 or 17]). $\mathcal{W}$ is called linear if for every pair of terms $\tau(x)$ and $\sigma(x)$ with the same domain, $\mathcal{W}$ satisfies one of

1. τ is constant or σ is constant or
2. there is γ so that $\gamma\tau = \sigma$ or
3. there is γ so that $\gamma\sigma = \tau$.

$\mathcal{W}$ is called linear because if you consider any $\mathcal{A} \in \mathcal{W}$ and $a \in A$ then the collection of one generated subuniverses of the subuniverse generated by a is linearly ordered by inclusion.

Suppose that $\mathcal{V}$ is a linear multi-sorted unary variety. Then for every $\mathcal{A} \in \mathcal{V}$ and $a \in A$ there is a natural quasi-order on the nonconstant elements of $\langle a \rangle$. If $b, c \in \langle a \rangle$ then $b \leq c$ if there is a term g so that $g(c) = b$. We say that $\mathcal{V}$ is well founded if this quasi-order is always a well quasi-order.

Now fix a multi-sorted unary variety $\mathcal{W}$ and its theory T . We will define a one sorted language $\overline{\mathcal{L}}$, a theory $\overline{T}$ and a variety $\overline{\mathcal{W}}$. $\overline{\mathcal{L}}$ will contain a unary function symbol $\overline{f}$ for every $f \in \mathcal{L}$ and a constant $\overline{c}$ for every $c \in \mathcal{L}$. $\overline{\mathcal{L}}$ will also contain a single n -ary function symbol d . We will describe $\overline{T}$ (and $\overline{\mathcal{W}}$) by describing its models up to isomorphism. Fix $\mathcal{A} \models T$. Define an $\overline{\mathcal{L}}$ -structure $\overline{\mathcal{A}}$ as follows:

1. If $\mathcal{A} = \langle A_1, \dots, A_n, \dots \rangle$ where A_i is the interpretation of the sort U_i then let the universe of $\overline{\mathcal{A}}$ be $A_1 \times \dots \times A_n$.
2. Let d be interpreted in $\overline{\mathcal{A}}$ by

$$d(\langle a_1^1, \dots, a_n^1 \rangle, \dots, \langle a_1^n, \dots, a_n^n \rangle) = \langle a_1^1, \dots, a_n^n \rangle.$$

3. If $f \in \mathcal{L}$ has domain U_i and range U_j then define $\overline{f}$ in $\overline{\mathcal{A}}$ by

$$\overline{f}(\langle a_1, \dots, a_n \rangle) = \langle c_1, \dots, f(a_i), \dots, c_n \rangle.$$

$\uparrow$
 j

4. If $c \in \mathcal{L}$ with sort U_i then interpret $\overline{c}$ in $\overline{\mathcal{A}}$ as

$$\langle c_1, \dots, c, \dots, c_n \rangle.$$

$\uparrow$
 i

$\overline{T}$ is the theory of the class of all models $\overline{\mathcal{A}}$ for $\mathcal{A} \models T$ and $\overline{\mathcal{W}}$ is the class of models of $\overline{T}$. It is straightforward to show that $\overline{\mathcal{W}}$ is a variety. Notice that

$\mathcal{A}$ can be essentially recovered from $\overline{\mathcal{A}}$ by letting the sorts be the kernels of d and considering the functions component by component.

A restatement of the main theorem from [11] (and [10]) is

Theorem 0.14. *If $\mathcal{V}$ is a combinatorial variety then $\mathcal{V}$ is superstable iff $\mathcal{V}$ is polynomially equivalent to $\overline{\mathcal{W}}$ for some linear multi-sorted unary variety $\mathcal{W}$ which is well founded.*

The usefulness of this theorem will be demonstrated in §4. Basically, any question about superstable combinatorial varieties reduces to the same question about well-founded linear multi-sorted unary varieties. the models of the latter are just trees of one generated subuniverses and so these questions usually have simple answers.

Before we proceed to the proofs of the main theorems, let us say a word about the connection between this work and tame congruence theory. Tame congruence theory is a beautiful piece of mathematics which deals with the structure of finite algebras. The tools of [13] are used in [16] to give the definitive structure of locally finite decidable varieties. One consequence of this work is

Theorem 0.15. *If $\mathcal{V}$ is a locally finite decidable abelian variety then $\mathcal{V}$ has an affine subvariety $\mathcal{A}$ and a combinatorial subvariety $\mathcal{S}$ so that $\mathcal{V} = \mathcal{A} \otimes \mathcal{S}$.*

In fact, it was this theorem which was the main motivation to undertake this work. In the end, we do not need to use any results from tame congruence theory. However we owe a debt to it nonetheless as a source of what might possibly be proved.

1. PRELIMINARY LEMMAS AND NOTATION

Most of the notation we will use is standard. We draw your attention to a couple of things. When the distinction between a singleton and a tuple does not matter we will not differentiate. If $\mathcal{A}$ is a model, $\varphi(x, y)$ is a formula and $b \in A$ then $\varphi(A, b) = \{a \in A : \mathcal{A} \models \varphi(a, b)\}$. We use $\langle B \rangle$ for the subuniverse generated by B .

Definition 1.1. 1. φ is said to be a product formula if whenever $\mathcal{A}_i$ is an $\mathcal{L}$ -structure for each $i \in I$ then

$$\prod_{i \in I} \mathcal{A}_i \models \varphi(a) \text{ iff } \mathcal{A}_i \models \varphi(a(i)) \text{ for each } i \in I.$$

2. A formula $\varphi(x, y)$ is normal in $\mathcal{A}$ if whenever $\varphi(A, a) \cap \varphi(A, b) \neq \emptyset$ for $a, b \in A$ then $\varphi(A, a) = \varphi(A, b)$.

The following definition of h-formula comes from [18].

Definition 1.2. 1. An h-formula is any formula which is in the smallest set S which contains all the atomic formulas, is closed under quantification and conjunction and satisfies the condition that if $\varphi \in S$ and $\psi \in S$ then $\exists x \varphi \wedge \forall x(\varphi \rightarrow \psi) \in S$.

2. A positive primitive formula (pp formula or just ppf) is any formula which is in the smallest set S containing all the atomic formulas and closed under conjunction and existential quantification.

3. A pp!-formula is either a pp formula or a formula of the form $\exists! x \varphi$ where φ is a ppf and $\exists!$ is short for “there exists unique”.

Remark 1.3. Any pp!-formula is an h-formula.

The following proposition is straightforward. See for example [18].

Proposition 1.4. *h-formulas are product formulas.*

One final definition which is relevant for our work on Vaught's conjecture.

Definition 1.5. We say a possibly incomplete theory in a countable language is small if every one of its completions has only countably many complete types over the empty set. We say an elementary class is small if its theory is small.

There are two crucial properties which will unify all the different cases we intend to handle in this paper.

Definition 1.6. A class K satisfies the normality condition if all pp!-formulas are normal in all $\mathcal{A} \in K$.

Lemma 1.7. *If K is a stable class closed under products then K satisfies the normality condition.*

Proof. In fact, all product formulas are normal in any stable class closed under products. This is basically folklore; for a proof see for example [11] or [18]. $\square$

Lemma 1.8. *If K is a class closed under Boolean powers and has few countable models then K satisfies the normality condition.*

Proof. This is essentially Theorem 4.1 of [2]. There they only consider atomic formulas but the changes required to handle pp!-formulas are minor. $\square$

Corollary 1.9. *If a variety $\mathcal{V}$ is stable or has few countable models then $\mathcal{V}$ satisfies the normality condition.*

For a proof of the following corollary see either [2] or [11].

Corollary 1.10. *If $\mathcal{V}$ satisfies the normality condition then $\mathcal{V}$ is abelian.*

Remark 1.11. In fact, in the previous corollary only the atomic formulas need be normal.

Definition 1.12. A class K satisfies the tree condition if there is no finitely generated $\mathcal{A} \in K$ with product formulas $\varphi_n(x, y_n)$ and tuples $a_\eta \in A$ for $\eta \in 2^{<\omega}$ so that if $\text{len}(\eta) = n$ then

$$\varphi_n(A, a_\eta) \neq \emptyset, \quad \varphi_{n+1}(A, a_{\eta 0}), \quad \varphi_{n+1}(A, a_{\eta 1}) \subseteq \varphi_n(A, a_\eta)$$

and

$$\varphi_{n+1}(A, a_{\eta 0}) \cap \varphi_{n+1}(A, a_{\eta 1}) = \emptyset.$$

Remark 1.13. In other words, K satisfies the tree condition means that there is no finitely generated algebra with a uniformly defined binary tree of product formula defined sets.

Lemma 1.14. *If K is superstable and is closed under products then K satisfies the tree condition.*

Proof. Straightforward. The proof is in the same spirit as the proof of Theorem 1.3 of [11]. In fact, in the definition of the tree condition one does not have to restrict oneself to finitely generated algebras. $\square$

Lemma 1.15. *If K is a small elementary class then K satisfies the tree condition.*

Proof. Clear. $\square$

Corollary 1.16. *If K has a few countable models then K satisfies the tree condition.*

Definition 1.17. A class K is called amenable if it satisfies both the normality condition and the tree condition.

Conclusion 1.18. If $\mathcal{V}$ is a variety which

1. is superstable or
2. has few countable models or
3. is stable and small

then $\mathcal{V}$ is amenable.

The main technical theorem of this paper which will be proved in §3 is

Theorem 1.19. *If $\mathcal{V}$ is an amenable variety then $\mathcal{V}$ is structured.*

Remark 1.20. The combination of Conclusion 1.18 and Theorem 1.19 gives a proof of Theorems 0.11 and 0.12. It follows then by our discussion in the introduction that every variety with few countable models is superstable.

2. THE AFFINE SUBVARIETY

The goal of this section is to prove that the affine algebras and the combinatorial algebras in an amenable variety form subvarieties.

Definition 2.1. 1. If $g(\bar{x}, \bar{z})$ is a term then $\ker_{\bar{x}}(g)(\bar{u}; \bar{v})$ is the following formula: $\forall \bar{z}(g(\bar{u}, \bar{z}) = g(\bar{v}, \bar{z}))$.

2. If $\alpha \in \text{Con}(\mathcal{A})$ then α is called co-affine if $\mathcal{A}/\alpha$ is an affine algebra.

3. If $\theta \in \text{Con}(\mathcal{A})$ then θ is combinatorial if for every $a, b, \bar{c}, \bar{d}, e \in A$ and every term τ , if $a\theta b\theta e$, $\bar{c}\theta\bar{d}$ and $\tau(a, \bar{c}) = \tau(b, \bar{d})$ then $\tau(e, \bar{c}) = \tau(e, \bar{d})$.

Remark 2.2. 1. Notice that $\ker_{\bar{x}}(g)$ is a product formula and defines an equivalence relation on tuples of elements in any algebra $\mathcal{A}$. If in addition $\mathcal{A}$ is abelian then $\ker_{\bar{x}}(g)$ is equivalent in $\mathcal{A}$ to $\exists \bar{y}(g(\bar{u}, \bar{y}) = g(\bar{v}, \bar{y}))$. For any algebra $\mathcal{A}$, we will not distinguish between the formula $\ker_{\bar{x}}(g)$ and the equivalence relation it defines on $\mathcal{A}$.

2. If $\mathcal{A} \in \mathcal{V}$ and $\mathcal{V}$ is abelian then $\alpha \in \text{Con}(\mathcal{A})$ is co-affine iff $\mathcal{A}/\alpha$ has a Malcev term.

3. The standard terminology for the term combinatorial is strongly abelian.

The main theorem we will prove in this section is

Theorem 2.3. *If $\mathcal{V}$ is an amenable variety then there is a term $\rho(x, y)$ so that $\rho(\rho(x, y), z) = \rho(x, z)$ holds in $\mathcal{V}$ and if $\mathcal{A} \in \mathcal{V}$ and $\theta = \ker_x(\rho)$ then*

1. θ is the minimal co-affine congruence on $\mathcal{A}$ and
2. θ is the maximal combinatorial congruence on $\mathcal{A}$.

Corollary 2.4. *If $\mathcal{V}$ is amenable then the affine algebras in $\mathcal{V}$ form a subvariety.*

Proof. In any variety the affine algebras are closed under submodels and homomorphisms. It suffices to show that if $\mathcal{A}_i \in \mathcal{V}$ for $i \in I$ are affine algebras then $\prod_{i \in I} \mathcal{A}_i$ is affine. Let ρ be the term guaranteed by Theorem 2.3. $\ker_x(\rho) = 0_{\mathcal{A}_i}$

for each $i \in I$ and so $\ker_x(\rho) = 0$ in $\prod_{i \in I} \mathcal{A}_i$ which means, using Theorem 2.3, that $\prod_{i \in I} \mathcal{A}_i$ is affine. $\square$

Corollary 2.5. *If $\mathcal{V}$ is amenable then the combinatorial algebras in $\mathcal{V}$ form a subvariety.*

Proof. By the previous theorem, the class of combinatorial algebras in $\mathcal{V}$ is axiomatized by the equation $\rho(x, z) = \rho(y, z)$ and so this class forms a subvariety of $\mathcal{V}$. $\square$

We will now prove some technical lemmas before the proof of Theorem 2.3

Notation 2.6. If $g(x, \bar{y})$ is a term then we define $g_x^n(x, \bar{y})$ inductively by

1. $g_x^0(x, \bar{y}) = g(x, \bar{y})$ and
2. $g_x^{n+1}(x, \bar{y}) = g_x^n(g(x, \bar{y}), \bar{y})$.

We call g_x^n an iterate of g .

Lemma 2.7. *If $\mathcal{V}$ is amenable then for every term $g(x, y, \bar{u})$ there is an $n \in \omega$ so that*

$$g_x^n(g(A, a, \bar{c}), d, \bar{c}) = g_x^n(g(A, b, \bar{c}), d, \bar{c})$$

for all $\mathcal{A} \in \mathcal{V}$ and $a, b, \bar{c}, \bar{d} \in A$.

Proof. First of all, it suffices to prove this statement when $a = d$. Secondly, it suffices to prove that for any fixed $\mathcal{A} \in \mathcal{V}$ there is such an n . Otherwise, suppose that for every n there is $\mathcal{A}_n$ so that $g_x^n(g(\mathcal{A}_n, a_n, \bar{c}_n), a_n, \bar{c}_n) \neq g_x^n(g(\mathcal{A}_n, b_n, \bar{c}_n), a_n, \bar{c}_n)$ for some $a_n, b_n, \bar{c}_n \in \mathcal{A}_n$. Then no n will work for $\prod_{n \in \omega} \mathcal{A}_n$.

So fix $\mathcal{A} \in \mathcal{V}$ and $\bar{c} \in A$ and let $\tau(x, y) = g(x, y, \bar{c})$. Also fix $a, b \in A$. We need to show that for some n , $\tau_x^n(\tau(A, a), a) = \tau_x^n(\tau(A, b), a)$. Suppose not. That is, assume that for all n , $\tau_x^n(\tau(A, a), a) \cap \tau_x^n(\tau(A, b), a) = \emptyset$ (use the normality of pp formulas). Define polynomials $\sigma_m(x, y_0, \dots, y_{m-1})$ by induction.

1. $\sigma_0(x, y_0) = \tau(x, y_0)$ and
2. $\sigma_{m+1}(x, y_0, \dots, y_m) = \tau(\sigma_m(x, y_1, \dots, y_m), y_0)$.

Let $\varphi_m(x, y_0, \dots, y_{m-1}) = \exists z(x = \sigma_m(z, y_0, \dots, y_{m-1}))$.

Suppose $\eta: m \rightarrow \{a, b\}$. Then

$$\varphi_{m+1}(A, \eta, a), \varphi_{m+1}(A, \eta, b) \subseteq \varphi_m(A, \eta)$$

and by the fact that $\mathcal{A}$ is abelian,

$$\varphi_{m+1}(A, \eta, a) \cap \varphi_{m+1}(A, \eta, b) = \emptyset.$$

Hence if no n exists as required then we contradict the tree property in the algebra generated by a, b and $\bar{c}$. Since $\mathcal{V}$ is amenable, this cannot happen and so we are done. $\square$

Definition 2.8. An equivalence relation $E(\bar{x}; \bar{y})$ is kernel defined if it is the intersection of finitely many equivalence relations of the form $\ker_{\bar{x}}(g)$ for terms $g(\bar{x}, \bar{z})$.

Discussion 2.9. Suppose $\mathcal{A}$ is any structure and E is a pp-defined equivalence relation on A (in our case, usually a kernel defined equivalence). We will call

$\mathcal{A}/E$ an affine structure if there is a term $\sigma(x, y, z)$ on $\mathcal{A}$ which respects E and for all $a, b \in A$, $\sigma(a, b, b)E\sigma(b, b, a)Ea$. That is, σ is a Malcev term on $\mathcal{A}/E$.

Here is how the affine structures will be used. Fix a structure $\mathcal{A}$ with normal ppfs. Suppose α_1 and α_2 are pp-defined equivalences on $\mathcal{A}$ and α_1 properly refines α_2 . Moreover, suppose that $\mathcal{A}/\alpha_i$ is an affine structure with Malcev term σ_i for each i . (In fact it is enough that only $\mathcal{A}/\alpha_2$ is an affine structure.)

Claim 2.10. σ_1 respects α_2 and every α_2 -class contains at least two α_1 -classes.

Proof. We treat one variable at a time. Suppose $a\alpha_2b$ for some $a, b \in A$.

$$\sigma_1(a, b, b)\alpha_1a \quad \text{and} \quad \sigma_1(b, b, b)\alpha_1b$$

so $\sigma_1(a, b, b)\alpha_2\sigma_1(b, b, b)$. Of course,

$$\sigma_1(a, u, v)\alpha_2\sigma_1(a, u, v) \quad \text{and} \quad \sigma_1(a, b, b)\alpha_2\sigma_1(a, b, b)$$

so by the normality of ppfs, $\sigma_1(a, u, v)\alpha_2\sigma_1(b, u, v)$. The other variables are handled similarly.

Now suppose $a\alpha_2b$ but a and b are not α_1 -related. Fix any c . Then $\sigma_1(a, b, c)\alpha_2\sigma_1(a, a, c)$ by the claim and so $\sigma_1(a, b, c)\alpha_2c$. If $\sigma_1(a, b, c)\alpha_1c$ then $\sigma_1(a, a, c)\alpha_1\sigma_1(a, b, c)$ so by the normality of ppfs,

$$b\alpha_1\sigma_1(a, a, b)\alpha_1\sigma_1(a, b, b)\alpha_1a.$$

This is a contradiction. $\square$

If $\mathcal{B}$ is a finitely generated structure and $\langle \alpha_i : i \in \omega \rangle$ is a properly descending chain of kernel defined equivalences so that $\mathcal{B}/\alpha_i$ is an affine structure for each i and ppfs are normal in $\mathcal{B}$ then every α_i -class is refined by an α_{i+1} -class. This means that $\mathcal{B}$ would fail the tree condition. Hence if $\mathcal{V}$ is amenable then no finitely generated algebra can have such a descending chain of equivalences.

Lemma 2.11. *Suppose $\mathcal{V}$ is amenable and $\tau_i(x, y, \bar{z})$ are terms for $i < n$ so that for every $\mathcal{A} \in \mathcal{V}$, $a, b, \bar{c} \in A$ and $i < n$, $\tau_i(A, a, \bar{c}) = \tau_i(A, b, \bar{c})$. Then if $\alpha = \bigcap_{i < n} \ker_y \tau_i(x, y, \bar{z})$,*

1. $\mathcal{B}/\alpha$ is an affine structure for any $\mathcal{B} \in \mathcal{V}$ and,
2. for any finitely generated $\mathcal{B} \in \mathcal{V}$ there is a kernel defined congruence β on $\mathcal{B}$ so that $\beta \subseteq \alpha$ and $\mathcal{B}/\beta$ is affine.

Proof. For the first, it suffices to find a Malcev term on the α -classes. Consider $\mathcal{F}$, the free algebra on the generators x, y and z . By our assumption, there are $v_i(x, y, z) \in F$ for $i < n$ so that $\tau_i(y, y, \bar{z}) = \tau_i(v_i, z, \bar{z})$. By normality, there is $g(x, y, z) \in F$ so that $\tau_i(y, x, \bar{z}) = \tau_i(v_i, g, \bar{z})$ for all $i < n$.

We will show that g respects α and is a Malcev term for the quotient structure F/α . We first show that $g(x, x, z)$ and $g(z, x, x)$ are α -related to z in $\mathcal{F}$. From the above equalities we can deduce the following, using the abelian property and the freeness of $\mathcal{F}$,

$$\tau_i(x, x, \bar{z}) = \tau_i(v_i(x, x, z), z, \bar{z})$$

and

$$\tau_i(x, x, \bar{z}) = \tau_i(v_i(x, x, z), g(x, x, z), \bar{z})$$

for each $i < n$. Thus $g(x, x, z)$ and z are α -related. A similar argument can be used to show that $g(z, x, x)$ is α -related to z . It now easily follows that for any $\mathcal{B} \in \mathcal{V}$, the term g respects α and is a Malcev term on $\mathcal{B}/\alpha$.

For the second, fix any term $\sigma(y, \bar{w})$. Let

$$\alpha_\sigma = \bigcap_{i < n} \ker_y \tau_i(x, \sigma(y, \bar{w}), \bar{z}).$$

For any finite set of terms Σ , $\alpha_\Sigma = \bigcap_{\sigma \in \Sigma} \alpha_\sigma$ satisfies the conditions of the first part of the lemma and so $\mathcal{B}_\Sigma = \mathcal{B}/\alpha_\Sigma$ is an affine structure. If $\Sigma \subseteq \Sigma'$ then $\alpha_{\Sigma'}$ refines α_Σ . Hence, by the tree property, if $\mathcal{B}$ is finitely generated, there must be a finite Σ so that $\alpha_\Sigma = \bigcap_\sigma \alpha_\sigma$. The right-hand side is clearly a congruence and $\mathcal{B}/\alpha_\Sigma$ is affine. $\square$

Lemma 2.12. *Suppose $\mathcal{V}$ is amenable. If $\mathcal{A} \in \mathcal{V}$ is finitely generated and $\beta \in \text{Con}(\mathcal{A})$ so that $\mathcal{A}/\beta$ is affine then there is a kernel defined congruence $\bar{\beta} \subseteq \beta$ so that $\mathcal{A}/\bar{\beta}$ is affine.*

More precisely, $\bar{\beta}$ is the intersection of finitely many kernels of polynomials $h(x, y)$ so that for every $\mathcal{B} \in \mathcal{V}$, $a, b \in B$, $h(B, a) = h(B, b)$.

Proof. Suppose $g(x, y, z)$ is the Malcev term on $\mathcal{A}/\beta$. Choose n so that

$$g_x^n(g(B, c, a), c, a) = g_x^n(g(B, c, b), c, a)$$

for all $\mathcal{B} \in \mathcal{V}$ and $a, b, c \in B$. (Use Lemma 2.7.) Fix $a, c \in A$. Let $h(u, v) = g_x^n(g(u, c, v), c, a)$ and $\alpha = \ker_v h(u, v)$.

$\alpha \subseteq \beta$ since any iteration of g is one to one in all variables on $\mathcal{A}/\beta$. By Lemma 2.11, we can find $\bar{\beta} \subseteq \alpha$ so that $\mathcal{A}/\bar{\beta}$ is affine. $\square$

Definition 2.13. We say that a four variable relation $E(x, y; u, v)$ on an algebra $\mathcal{A}$ is a Malcev relation if it is a kernel defined equivalence relation with $\mathcal{A}$ satisfying $E(x, x; y, y)$. A kernel defined four variable relation is said to be a Malcev relation for the variety $\mathcal{V}$ if it defines a Malcev relation on every algebra in $\mathcal{V}$.

If E is a Malcev relation then the kernel of E is the relation $\theta_E(x, y) = E(x, x; x, y)$.

Remark 2.14. Note that if $\mathcal{V}$ is abelian then the kernel of a Malcev relation is kernel defined.

Further note that if $\mathcal{A}$ is an algebra in an amenable variety and α is a kernel defined equivalence on A with Malcev term g , then the relation $E(x, y; u, v)$ defined by $g(x, y, u)\alpha v$ is a Malcev relation. The original affine structure on A/α can be recovered from the relation E .

Lemma 2.15. *If $\mathcal{V}$ is abelian and E is Malcev then θ_E is an equivalence relation.*

Proof. $E(x, x; x, x)$ always holds so θ_E is reflexive. If $E(x, x; x, y)$ holds then since $\mathcal{V}$ is abelian and E is kernel defined $E(y, x; y, y)$ holds so by the symmetry of E , θ_E is symmetric. For essentially the same reason, θ_E is transitive. $\square$

Lemma 2.16. *Assume $\mathcal{V}$ is amenable.*

1. *If $E(x, y; u, v)$ is a Malcev relation and g is a unary polynomial then $E(g(x), g(y); g(u), g(v))$ is a Malcev relation.*
2. *The intersection of finitely many Malcev relations is Malcev.*
3. *For any Malcev relation E , there is a term $g(y, x, z)$ so that*

$$E(x, y; z, g(y, x, z))$$

holds in $\mathcal{V}$. Notice the order of the variables in g .

4. *If E is Malcev on $\mathcal{A}$ then $\mathcal{A}/\theta_E$ is an affine structure.*

Proof. The first two are easy. For the third, notice that

$$\mathcal{V} \models \exists w E(x, x; x, w) \wedge \exists w E(x, x; z, w) \wedge \exists w E(x, y; x, w).$$

By the normality of ppfs, $\mathcal{V} \models \exists w E(x, y; z, w)$. In the free algebra then we can find the required term g to witness this w .

For the fourth, the g from 3 provides the Malcev term. Note that

$$E(x, x; z, z) \quad \text{and} \quad E(x, x; z, g(x, x, z))$$

hold and so $g(x, x, z)\theta_E z$. Also, $E(z, x; z, g(x, z, z))$ and $E(z, x; z, x)$ hold. Since $\mathcal{V}$ is abelian then $E(x, x; x, g(x, z, z))$ also holds. This shows that $g(x, z, z)\theta_E x$. It is clear that g respects θ_E and so g is a Malcev term on $\mathcal{A}/\theta_E$. $\square$

Proposition 2.17. *If $\mathcal{V}$ satisfies the normality condition and $\theta = \bigcap \{\theta_E : E \text{ is Malcev in } \mathcal{V}\}$ then for all $\mathcal{A} \in \mathcal{V}$,*

1. *θ is a congruence and*
2. *θ is the maximal combinatorial congruence on $\mathcal{A}$.*

In addition, if $\mathcal{V}$ is amenable and $\mathcal{A}$ is finitely generated then θ is the minimal co-affine congruence on $\mathcal{A}$ and for some term $g(u, v)$, $\theta = \ker_v(g(u, v))$ on $\mathcal{A}$.

More precisely, if $\sigma(x, y, z)$ is a Malcev term for $\mathcal{A}/\theta$ and n is such that

$$\sigma_x^n(\sigma(B, d, b), d, a) = \sigma_x^n(\sigma(B, d, c), d, a)$$

for all $\mathcal{B} \in \mathcal{V}$ and $a, b, c, d \in B$ and if $g(u, v) = \sigma_x^n(\sigma(u, w, v), w, z)$ then $\theta = \ker_v(g(u, v))$.

Proof. By Lemma 2.16 and the fact the θ is the intersection of all θ_E , it is clear that θ is a congruence.

Now we need to show that θ is combinatorial. Suppose we have term $\tau(x, y_1, \dots, y_n)$. By the normality of pp formulas, we know that

$$\mathcal{V} \models \forall abc\bar{d}(\tau(a, \bar{c}) = \tau(b, \bar{d}) \rightarrow \forall e\exists f\tau(e, \bar{c}) = \tau(f, \bar{d})).$$

It follows that there is a term $f(a, b, \bar{c}, \bar{d}, e)$ so that

$$\mathcal{V} \models \forall abc\bar{d}e(\tau(a, \bar{c}) = \tau(b, \bar{d}) \rightarrow \tau(e, \bar{c}) = \tau(f(a, b, \bar{c}, \bar{d}, e), \bar{d})).$$

Consider the equivalence relations

$$E = \ker_{ab} \tau(f(a, b, \bar{c}, \bar{d}, e), \bar{w}) \quad \text{and} \quad E_i = \ker_{c_i d_i} \tau(f(a, b, \bar{c}, \bar{d}, e), \bar{w}).$$

By the abelian property and the definition of f , E and E_i are Malcev.

Suppose $\tau(a, \bar{c}) = \tau(b, \bar{d})$ with $a\theta b$ and $\bar{c}\theta\bar{d}$. At least then $aaEab$ and $d_i d_i E_i c_i d_i$. So

$$\begin{aligned}\tau(a, \bar{c}) &= \tau(f(a, b, \bar{c}, \bar{d}, a), \bar{d}) \\ &= \tau(f(a, a, \bar{c}, \bar{d}, a), \bar{d}) \\ &= \dots \\ &= \tau(f(a, a, \bar{d}, \bar{d}, a), \bar{d}) \\ &= \tau(a, \bar{d}).\end{aligned}$$

This is enough to show that θ is combinatorial.

Suppose that $\alpha \in \text{Con}(\mathcal{A})$ and α is combinatorial. We want to show that $\alpha \subseteq \theta$. Choose $a, b \in A$ so that $a\alpha b$. Now for any Malcev relation E , we have $aaEbb$. But since α is combinatorial, it follows that $aaEab$. That is, $a\theta_E b$. Since this is true for all such E , $\alpha \subseteq \theta$.

Now suppose in addition that $\mathcal{V}$ is amenable and $\mathcal{A}$ is finitely generated. Then by the tree property, θ is the intersection of finitely many (and hence one) θ_E 's for Malcev relations E . By Lemma 2.16, θ is a co-affine congruence.

Let us show that there is a minimal co-affine congruence on $\mathcal{A}$. By Lemma 2.12, the intersection of all co-affine congruences on $\mathcal{A}$ is equal to the intersection of all kernel defined co-affine congruences and moreover, kernel defined congruences with the property mentioned in Lemma 2.12. The intersection of such congruences is co-affine by Lemma 2.11 and the fact that $\mathcal{A}$ is finitely generated.

So now we have a minimal co-affine congruence on $\mathcal{A}$; call it α . We know $\alpha \subseteq \theta$ and we want to show $\alpha = \theta$. It is enough to show that α is the kernel of some Malcev relation. So suppose that $\sigma(x, y, z)$ is a Malcev term on $\mathcal{A}/\alpha$. Let n be chosen so that

$$\sigma_x^n(\sigma(A, c, a), c, a) = \sigma_x^n(\sigma(A, c, b), c, a)$$

for all $a, b, c \in A$. Let $g(u, v) = \sigma_x^n(\sigma(u, c, v), c, a)$. If $\beta = \ker_v(g(u, v))$ then by Lemma 2.11, $\mathcal{A}/\beta$ is an affine structure. Let $h(x, y, z)$ be a Malcev term for $\mathcal{A}/\beta$. Consider

$$E(x, y; x', y') := \forall uz(g(u, h(x, y, z)) = g(y, h(x', y', z))).$$

It is easy to prove that E is Malcev and $\theta_E \subseteq \beta \subseteq \alpha$ so in fact $\theta = \theta_E$.

In addition, we have also proved that $\theta = \beta$ which is exactly what the last line of the proposition states. $\square$

Now we are ready to prove Theorem 2.3.

Proof of Theorem 2.3. Proposition 2.17 is almost what we need at least for finitely generated algebras. The problem is that the g mentioned there may vary from algebra to algebra. Let us show this cannot happen.

Let $g_n(x)$ be a polynomial so that $\ker_x(g_n)$ is the minimal co-affine congruence on $\mathcal{F}_n$, the free algebra on n generators. Use the term provided by Lemma 2.17 to get such a polynomial.

Suppose $\alpha \in \text{Con}(\mathcal{F}_n)$ and $\theta = \ker_x(g_n)$. Clearly $\theta \vee \alpha$ is co-affine and by the minimality of θ , it is the minimal co-affine congruence above α . By Proposition 2.17 and our choice of g_n , it is clear that $\ker_x(g_n)$ is the minimal co-affine congruence in $\mathcal{F}_n/\alpha$.

Now consider $\mathcal{F}_2$. For any $n > 1$, we can consider $\mathcal{F}_2$ as a homomorphic image of $\mathcal{F}_n$. By the considerations above, we see that $\ker_x(g_2) = \ker_x(g_n)$ on $\mathcal{F}_2$.

Claim 2.18. $\ker_x(g_2) = \ker_x(g_n)$ in all of $\mathcal{V}$.

Proof. Since $\mathcal{V}$ is abelian, we need only consider 2-generated algebras. Suppose $\mathcal{B}$ is one. Again, by the considerations above, since $\mathcal{B}$ is a homomorphic image of $\mathcal{F}_n$ and $\mathcal{F}_2$, the minimal co-affine congruence on $\mathcal{B}$ is given by $\ker_x(g_2)$ and $\ker_x(g_n)$. Hence $\ker_x(g_2) = \ker_x(g_n)$ on $\mathcal{B}$. $\square$

Let $g = g_2$. It is important to remember where g comes from. By Proposition 2.17, g is an iterate of the Malcev term on $\mathcal{F}_2/\theta$. Suppose that $g(a)\theta g(b)$ for some $a, b \in F_2$. Then $g(a/\theta) = g(b/\theta)$ in $\mathcal{F}_2/\theta$. Since g is an iterate of the Malcev term, it is one-to-one on $\mathcal{F}_2/\theta$ so $a/\theta = b/\theta$. That is, $a\theta b$ in $\mathcal{F}_2$ i.e. $g(a) = g(b)$. Any iterate of the Malcev term is onto in $\mathcal{F}_2/\theta$ so we have proved that the range of g intersects each θ -class in $\mathcal{F}_2$ exactly once.

Let $\rho(x, y)$ be in the range of g and θ -related to x in $\mathcal{F}_2$. Hence we have $g(x) = g(\rho(x, y))$ and for some z we have $g(z) = \rho(x, y)$. Using these facts we see that $g(x) = g(\rho(\rho(x, y), y))$ and $\rho(\rho(x, y), y)$ is in the range of g . So in $\mathcal{F}_2$, $\rho(\rho(x, y), y) = \rho(x, y)$. Since $\mathcal{V}$ is abelian, $\rho(\rho(x, y), z) = \rho(x, z)$ holds in $\mathcal{V}$.

Claim 2.19. $\ker_x(g) = \ker_x(\rho(x, y))$ holds in $\mathcal{V}$.

Proof. Since $g(\rho(x, y)) = g(x)$ we already have one inclusion. For the other direction, suppose $\mathcal{B} \in \mathcal{V}$, $a, b \in B$ and $a\theta b$. Since $\mathcal{V}$ is abelian, we can assume $\mathcal{B}$ is generated by a and b . Hence $\mathcal{B}$ is a homomorphic image of $\mathcal{F}_2$. σ is the Malcev term on $\mathcal{F}_2/\theta$ so it is also the Malcev term on $\mathcal{B}/\theta$. g is an iterate of σ . Now for any $c \in B$, $\rho(a, c)\theta\rho(b, c)$ and they are in the range of g since this is true equationally in $\mathcal{F}_2$. By the argument used above then $\rho(a, c) = \rho(b, c)$. $\square$

Suppose σ is the Malcev term on $\mathcal{F}_2/\theta$. It is straightforward to write equations using σ and ρ so that in all of $\mathcal{V}$, $\ker_x(\rho)$ is a co-affine congruence with Malcev term σ . Moreover, since $\ker_x(\rho)$ is combinatorial in all finitely generated algebras, it is combinatorial in all algebras.

Now suppose that $\mathcal{A}$ is any algebra in $\mathcal{V}$, θ is $\ker_x(\rho)$ in $\mathcal{A}$ and α is co-affine in $\mathcal{A}$. We want to show $\theta \subseteq \alpha$. Suppose not. Choose a, b which are θ -related but not α -related. Let $\mathcal{A}' = \langle a, b \rangle$. If $\theta' = \theta \upharpoonright \mathcal{A}'$ and $\alpha' = \alpha \upharpoonright \mathcal{A}'$ then θ' is $\ker_x(\rho)$ on $\mathcal{A}'$ and α' is co-affine. So $\theta' \subseteq \alpha'$, a contradiction.

The fact that θ is the maximal combinatorial congruence in $\mathcal{A}$ is proved similarly. $\square$

Remark 2.20. We have introduced the property of being amenable in order to present a single proof that handles both the superstable variety case and the case of a variety having few countable models. This works well enough, but at times certain things become obscured. For example, the proof that the affine algebras in a superstable variety are closed under products is almost immediate and we will include the proof. Those reading only for Vaught's conjecture can safely skip over the following theorem.

The key point is that Lemma 2.11 holds for all $\beta \in \mathcal{V}$ if $\mathcal{V}$ is superstable and not just for finitely generated algebras.

Theorem 2.21. *If $\mathcal{V}$ is superstable then the affine algebras in $\mathcal{V}$ are closed under products.*

Proof. Suppose that $\mathcal{A}_i$ is an affine algebra in $\mathcal{V}$ for all $i \in I$. Let $g_i(x, y, z)$ be a Malcev term for $\mathcal{A}_i$. By superstability and Lemma 2.7, there is an n so that

$$(g_i^n)_x(g_i(A, y, z_1), y, z) = (g_i^n)_x(g_i(A, y, z_2), y, z)$$

for any $\mathcal{A} \in \mathcal{V}$ and $y, z_1, z_2, z \in A$. Let

$$h_i(u, v) = (g_i^n)_x(g_i(u, y, v), y, z).$$

We suppress mention of y and z . By Lemma 2.11, there is a kernel congruence $\alpha_i \subseteq \ker_v h_i(u, v)$ so that $\mathcal{A}/\alpha_i$ is affine for all $\mathcal{A} \in \mathcal{V}$. Moreover, notice that since g_i is the Malcev term on $\mathcal{A}_i$, $\ker_v h_i(u, v) = 0$ in $\mathcal{A}_i$.

By Lemma 2.11 and superstability, there is a finite $I_0 \subseteq I$ so that

$$\bigcap_{i \in I} \alpha_i = \bigcap_{i \in I_0} \alpha_i = \alpha$$

and $\mathcal{A}/\alpha$ is affine for all $\mathcal{A} \in \mathcal{V}$. But now if we consider $\prod_{i \in I} \mathcal{A}_i$ component by component, we see that α is 0 on this product. Hence, $\prod_{i \in I} \mathcal{A}_i$ is affine. $\square$

3. THE DECOMPOSITION

It is instructive to consider where we are now in the proof of the main theorem, Theorem 1.19. If $\mathcal{V}$ is an amenable variety then by Corollary 2.4, the affine algebras of $\mathcal{V}$ form a subvariety and by Corollary 2.5, the combinatorial algebras in $\mathcal{V}$ form a subvariety. Moreover, since the only algebra which is both combinatorial and affine is the trivial algebra, it follows that for every $\mathcal{A} \in \mathcal{V}$, there are minimal $\alpha, \beta \in \text{Con}(\mathcal{A})$ so that $\mathcal{A}/\alpha$ is combinatorial, $\mathcal{A}/\beta$ is affine and $\alpha \vee \beta = 1_A$. What we do not know now is whether $\alpha \wedge \beta = 0_A$ (that is, whether $\mathcal{V}$ is the join of the combinatorial and affine subvarieties) and whether α and β commute. As it turns out we will achieve these two goals simultaneously. The underlying reason is that α and β will both be kernel defined and we have the following fact which is easily proved.

Fact 3.1. If α and β are two pp defined equivalence relations on $\mathcal{A}$ and ppfs are normal in $\mathcal{A}$ then α and β commute.

We will not use this fact explicitly.

An example to keep in mind throughout this section is the following.

Example 3.2. Suppose $\mathcal{L}$ has function symbols $\{\rho, +\}$ and a constant 0. ρ is unary and $+$ is binary. The axioms for the variety say that ρ is idempotent and $+$ is associative as well as $\rho(0) = 0$, $\rho(x + 0) = \rho(x)$, $x + x = 0$ and $\rho(x + y) = \rho(x) + \rho(y) = x + y$.

The kernel of ρ is a congruence on any algebra in this variety. Modulo this congruence the algebra is affine and polynomially equivalent to a vector space over the two element field. The kernel of ρ is easily seen to be combinatorial.

This variety is unstable since there is no control over the size of the pre-image of any element in the range of ρ . This lack of connection between the congruence classes of the minimal co-affine congruence will be the only obstacle in the proof of Theorem 1.19.

Notation 3.3. Throughout this section, $\mathcal{V}$ will be an amenable variety. Theorem 2.3 guarantees the existence of a term ρ with the properties listed there. These will be fixed as well. In any algebra, θ will be $\ker_x(\rho)$. Hence $\mathcal{A}/\theta$ is affine for any $\mathcal{A} \in \mathcal{V}$. $\mathcal{F}_n$ will be the free algebra on the generators $x_1, \dots, x_n$.

Definition 3.4. we say that two equivalence relations α and β on a set A are cross cutting or α cross cuts β if $\alpha \cap \beta = 0_A$ and $\alpha \circ \beta = \beta \circ \alpha = 1_A$.

Remark 3.5. Two cross cutting equivalence relations are sometimes called a pair of complementary factor relations in the literature.

The main proposition before the proof of Theorem 1.19 is

Proposition 3.6. *There is a kernel defined equivalence relation α on $\mathcal{F}_2$ which cross cuts θ .*

We need a number of technical definitions and lemmas before the proof of Proposition 3.6.

Definition 3.7. If $\beta \in \text{Con}(\mathcal{A})$ and $d(x)$ is some polynomial then d is said to β -depend on x if there are $a, b \in A$ so that $a\beta b$ and $d(a) \neq d(b)$.

Definition 3.8. Suppose that β is a congruence on $\mathcal{A}$. We say that the polynomial $d(x_1, \dots, x_n)$ of $\mathcal{A}$ is a β -decomposition if

1. $d(x, \dots, x) = x$ holds in $\mathcal{A}$ and,
2. d β -depends on each x_i .

Remark 3.9. If $d(x_1, \dots, x_n)$ is a β -decomposition polynomial for the algebra $\mathcal{A}$, and β is a combinatorial congruence, then the choice of parameters used to define d only depend on the β -classes of the parameters, i.e., if for some term t and elements a_i from A we have $d(\bar{x}) = t(\bar{x}; a_1, \dots, a_m)$, and if $a_i\beta b_i$ for all $i \leq m$, then $d(\bar{x}) = t(\bar{x}; b_1, \dots, b_m)$. This follows since

$$t(\bar{x}; \bar{b}) = t(t(\bar{x}; \bar{b}), \dots, t(\bar{x}; \bar{b}); \bar{a}),$$

$\mathcal{A}$ is abelian and β is combinatorial.

Definition 3.10. Suppose that β is a congruence on $\mathcal{A}$ and $d(x_1, \dots, x_n)$ and $d'(x_1, \dots, x_{n+1})$ are β -decompositions. We say that $d < d'$ if there is an $i \leq n$ so that $d(x_1, \dots, x_n) = d'(x_1, \dots, x_i, x_i, \dots, x_n)$ holds in $\mathcal{A}$.

Fact 3.11. Suppose $\beta \in \text{Con}(\mathcal{A})$, $d(x_1, \dots, x_n)$ is a β -decomposition and β is combinatorial. If $a_j^i \in A$ for $1 \leq i, j \leq n$ and $a_j^i\beta a_k^i$ for all $1 \leq i, j, k \leq n$ then

$$d(d(a_1^1, \dots, a_n^1), \dots, d(a_1^n, \dots, a_n^n)) = d(a_1^1, \dots, a_n^n).$$

Proof. Similar to the proof of Lemma 1.5 from [11]. $\square$

Lemma 3.12. *Suppose $\mathcal{V}$ is amenable and $\mathcal{A}$ is finitely generated. If β is a congruence defined by a product formula and β is combinatorial in $\mathcal{A}$ then there is a maximal β -decomposition. That is, there is a β -decomposition d so that there is no β -decomposition d' with $d < d'$.*

Proof. Otherwise we contradict the tree condition; not unlike the proof of Theorem 1.7 from [11]. $\square$

Remark 3.13. 1. Suppose σ is the Malcev term on $\mathcal{F}_4/\theta$. Let γ be the congruence generated by identifying $\rho(x_4, x_1)$ and $\rho(\sigma(x_2, x_1, x_3), x_1)$. Let $\mathcal{G} = \mathcal{F}_4/\gamma$. Fix $\mathcal{G}$, γ and σ for the next lemma. The subalgebra generated by x_1/γ , x_2/γ and x_3/γ is isomorphic to $\mathcal{F}_3$. To see this suppose $\mathcal{B} \in \mathcal{V}$ and $a_1, a_2, a_3 \in \mathcal{B}$. Let $a_4 = \sigma(a_2, a_1, a_3)$ and let h be the homomorphism from $\mathcal{F}_4$ to $\mathcal{B}$ extending the map sending x_i to a_i . Since $h(x_4) = h(\sigma(x_2, x_1, x_3))$, certainly γ is contained in the kernel of h . Hence, h factors through $\mathcal{G}$. This is enough to show that the subalgebra generated by x_1/γ , x_2/γ and x_3/γ is isomorphic to $\mathcal{F}_3$. We will identify it with $\mathcal{F}_3$.

2. We will also use $\mathcal{G}$ in the following way. Suppose $\mathcal{L}' = \mathcal{L} \cup \{u_1, u_2, u_3, u_4\}$ where the u_i 's are new constants. Consider the axioms of $\mathcal{V}$ together with the new axiom $\rho(u_4, u_1) = \rho(\sigma(u_2, u_1, u_3), u_1)$. If $\mathcal{F}^*$ is the free algebra generated by the constants of this theory in $\mathcal{L}'$ then $\mathcal{G}$ is just its reduct to $\mathcal{L}$.

We will prove

Lemma 3.14. *On $\mathcal{G}$ there is a kernel defined equivalence relation α which cross cuts θ .*

Proof. Let $\tau(x)$ be the polynomial $\rho(x, x_1)$ (recall that we have identified the element x_1/γ in $\mathcal{G}$ with the free generator x_1 in $\mathcal{F}_3$). Then from Theorem 2.3 we know for all $a, b \in G$ that $a\theta\tau(a)$ and if $a\theta b$ then $\tau(a) = \tau(b)$. Let $\mu(x, y, z) = \tau(\sigma(x, y, z))$. It follows that μ is a Malcev term on the range of τ .

If the congruence $\theta = \theta_G$, then $\mathcal{G}$ is affine and we choose α to be 1_G . Otherwise $\theta \neq 0_G$ and so there will be some maximal θ -decomposition polynomial $d(x_1, \dots, x_n)$ in $\mathcal{G}$. Let us show that we can choose d so that the only parameters needed to define it are the elements x_1 and x_2 . Suppose that $d(\bar{x}) = t(\bar{x}; x_1, x_2, x_3, u)$ for some term t . Then we have the following equality in $\mathcal{G}$,

$$t(x_1, \dots, x_1; x_1, x_2, x_3, u) = x_1.$$

If we substitute x_1 for x_3 and x_2 for u in the above equality, then we get that

$$t(x_1, \dots, x_1; x_1, x_2, x_1, x_2) = x_1$$

holds in $\mathcal{G}$. This follows from the remarks contained in 3.9 and the fact that $\rho(x_2, x_1) = \rho(\sigma(x_2, x_1, x_1), x_1)$. So,

$$t(x_1, \dots, x_1; x_1, x_2, x_3, u) = t(x_1, \dots, x_1; x_1, x_2, x_1, x_2)$$

in $\mathcal{G}$ and thus by the abelian property we can conclude that

$$d(\bar{x}) = t(\bar{x}; x_1, x_2, x_1, x_2)$$

for all $\bar{x}$ in G .

For each $i \leq n$, let

$$\bar{d}_i(x) = d(\tau(x), \dots, \underset{\substack{\uparrow \\ i}}{x}, \dots, \tau(x))$$

and let $\bar{D}_i$ be the range of $\bar{d}_i$. Notice that by Fact 3.11, $\bar{d}_i(\bar{d}_i(x)) = \bar{d}_i(x)$.

Claim 3.15. There are polynomials $\varepsilon_i(x)$ for $i < n$ so that

1. if $\langle u, v \rangle \in \ker_x \varepsilon_i(x) \cap \theta \cap \overline{D}_i^2$ then $u = v$ and
2. for every $a \in \overline{D}_i$ and $b \in G$ there is $c \in \overline{D}_i$ so that $c\theta b$ and $\langle a, c \rangle \in \ker_x \varepsilon_i(x)$.

Assume the claim and let us prove the lemma. Let $\alpha = \bigcap_{i=1}^n \ker_x \varepsilon_i(\overline{d}_i(x))$. Suppose that $\langle a, b \rangle \in \alpha \cap \theta$. $\overline{d}_i(a)$ and $\overline{d}_i(b)$ are in $\overline{D}_i$ and, since a and b are θ -related, they are θ -related. $\langle a, b \rangle \in \alpha$ so $\langle \overline{d}_i(a), \overline{d}_i(b) \rangle \in \ker_x \varepsilon_i(x)$. By part 1 of the claim, it follows that $\overline{d}_i(a) = \overline{d}_i(b)$ for all i . Of course then $d(\overline{d}_1(a), \dots, \overline{d}_n(a)) = d(\overline{d}_1(b), \dots, \overline{d}_n(b))$. But $a\theta\tau(a)$, $b\theta\tau(b)$, θ is combinatorial and d is a θ -decomposition so by Fact 3.11 the first term is a and the second is b . This shows that $\theta \cap \alpha = 0_G$.

Now suppose that we have any pair a and b . We want to find c so that $a\alpha c$ and $b\theta c$ and this will tell us that θ and α are cross cutting. By the second part of the claim, for every i , there is c_i so that $c_i\theta b$ and $\langle \overline{d}_i(a), \overline{d}_i(c_i) \rangle \in \ker_x \varepsilon_i(x)$. Let $c = d(\overline{d}_1(c_1), \dots, \overline{d}_n(c_n))$. Since θ is a congruence, $c\theta b$ and since θ is combinatorial $c\alpha a$. Now let us prove the claim.

Proof of Claim 3.15. To prove the claim we must construct the ε_i 's. The cases are all identical so to spare notation, we will concentrate on the first variable of d . Let $\overline{d}(x) = d(x, \tau(x), \dots, \tau(x))$ and write a^* for $\overline{d}(a)$. Let $u = x_4/\gamma$.

Let $D_w = \{a^* : a\theta w\}$ and let D be the range of $\overline{d}$. Let β be the congruence obtained by collapsing D_{x_1} , D_{x_2} and D_{x_3} . Since D_{x_i} is contained in the θ -class of x_i and θ is a congruence, $\beta \subseteq \theta$. Hence in the algebra $\mathcal{G}/\beta$, θ/β is the minimal co-affine congruence. So if $v/\beta\theta x_i/\beta$ in $\mathcal{G}/\beta$ then it follows that $v\theta x_i$ in $\mathcal{G}$. Hence $v^* \in D_{x_i}$. So $\mathcal{G}/\beta$ satisfies the formulas saying that there is a unique element in D_w when w is any one of x_1/β , x_2/β or x_3/β . That is, if

$$\varphi(x, y, z) = \exists! w "w \in \text{rng}(\overline{d}) \text{ and } w\theta\mu(x, y, z)"$$

then if we write $\hat{x}_i$ for x_i/β ,

$$\mathcal{G}/\beta \models \varphi(\hat{x}_1, \hat{x}_1, \hat{x}_1) \wedge \varphi(\hat{x}_1, \hat{x}_1, \hat{x}_3) \wedge \varphi(\hat{x}_2, \hat{x}_1, \hat{x}_1).$$

φ is a pp!-formula so by normality of such formulas we have

$$\mathcal{G}/\beta \models \varphi(\hat{x}_2, \hat{x}_1, \hat{x}_3).$$

That is, u^* is β -related to $\mu(x_2, x_1, x_3)$.

This means that there is a sequence $v_0, \dots, v_k$ so that $v_0 = u^*$ and $v_k = \mu(x_2, x_1, x_3)$, polynomials f_i for $i = 1, \dots, k$ and pairs $\langle \alpha_i, \beta_i \rangle \in D_{x_1}^2 \cup D_{x_2}^2 \cup D_{x_3}^2$ so that $f_i(\alpha_i) = v_{i-1}$ and $f_i(\beta_i) = v_i$.

By replacing f_i with $\overline{d}(f_i)$, we can assume that all of the v_i are in D_u . Without loss, we can assume there is $\varepsilon^* \in D_{x_1}$ and polynomial h^* so that either

1. $h^*(\varepsilon^*) = u^*$ and $h^*(\tau(x_1)) \neq u^*$ or
2. $h^*(\tau(x_1)) = u^*$ and $h^*(\varepsilon^*) \neq u^*$.

Consider the first case. Suppose that h^* is $h^*(w, x_1, x_2, x_3, u)$ and that ε^* is $\varepsilon^*(x_1, x_2, x_3, u)$. We have

$$h^*(\varepsilon^*(x_1, x_2, x_3, u), x_1, x_2, x_3, u) = u^*$$

holds in $\mathcal{G}$. By 3.13, if we let $x_3 = x_1$ and $u = x_2$ then the equation

$$h^*(\varepsilon^*(x_1, x_2, x_1, x_2), x_1, x_2, x_1, x_2) = x_2^*$$

holds in $\mathcal{V}$. Let $\varepsilon(x_1, x_2) = \varepsilon^*(x_1, x_2, x_1, x_2)$. Clearly we have $\varepsilon(x_1, x_2)\theta x_1$. Let $h(w, x_1, x_2) = h^*(w, x_1, x_2, x_1, x_2)$. We have $h(\varepsilon(x_1, x_2), x_1, x_2) = x_2^*$. If $h(\varepsilon(x_1, x_2), x_1, x_2)$ θ -depends on both x_2 variables in $\mathcal{F}_2$ then it does in $\mathcal{G}$ which would contradict the maximality of d . Suppose it does not θ -depend on the first.

Claim 3.16. $h(a, x_1, -)$ is one-to-one from D to D .

Proof. Suppose $h(a, x_1, v) = h(a, x_1, v')$ where $v, v' \in D$. Then

$$v = h(\varepsilon(a, v), a, v) = h(\varepsilon(a, v), a, v').$$

But $\varepsilon(a, v)\theta\varepsilon(a, v')$ so $h(\varepsilon(a, v), a, v)\theta v'$. Therefore $v\theta v'$. But then

$$v = h(\varepsilon(a, v), a, v) = h(\varepsilon(a, v), a, v') = h(\varepsilon(a, v'), a, v') = v'$$

since there is no θ -dependence on the first x_2 variable. $\square$

Claim 3.17. For any $b, c \in G$, $h(b, x_1, D) = h(c, x_1, D)$.

Proof. Suppose not. By the normality of pp formulas, this means

$$h(b, x_1, D) \cap h(c, x_1, D) = \emptyset.$$

We will define a family of pp-defined sets U_η for $\eta \in \{b, c\}^{<\omega}$ so that for each $n \in \omega$, the sets U_η with $\text{len}(\eta) = n$ are defined as instances of the same formula. Moreover, if $\eta \in \{b, c\}^{<\omega}$ then $U_{\eta b}, U_{\eta c} \subseteq U_\eta$ and $U_{\eta b} \cap U_{\eta c} = \emptyset$. This will contradict the tree condition.

In fact, U_η will be the image of D under a polynomial. Define $j(v, w) = h(v, x_1, w)$. We define polynomials $j_n(v_0, \dots, v_{n-1}, w)$ inductively. Let j_0 be the identity and $j_1(v_0, w) = j(v_0, w)$. If j_n has been defined, let

$$j_{n+1}(v_0, \dots, v_n, w) = j_n(v_0, \dots, v_{n-1}, j(v_n, w)).$$

For $\eta \in \{b, c\}^{<\omega}$ of length n , let $U_\eta = j_n(\eta, D)$. Using the fact that j is one to one in its second variable, it is easy to verify that $\{U_\eta : \eta \in \{b, c\}^{<\omega}\}$ has the desired properties. $\square$

Using the claim now, we see that by Lemma 2.11, there is a co-affine congruence contained in $\ker_w h(w, x_1, x_3)$. Since θ is the minimal co-affine congruence on $\mathcal{G}$, it follows that $\theta \subseteq \ker_w h(w, x_1, x_3)$. But

$$\ker_w(h) = \ker_w(h^*(w, x_1, x_2, x_3, u)) \quad \text{and} \quad \varepsilon^*\theta\tau(x_1)$$

which contradicts $h^*(\varepsilon^*) \neq h^*(\tau(x_1))$ in $\mathcal{G}$.

Therefore there must be no θ -dependence in the second x_2 variable in $h(\varepsilon(x_1, x_2), x_1, x_2)$.

Claim 3.18. $\varepsilon(x_1, -)$ is one-to-one from D_{x_2} to D_{x_1} .

Proof. Suppose $v, v' \in D_{x_2}$ and $\varepsilon(x_1, v) = \varepsilon(x_1, v')$. Then

$$v = h(\varepsilon(x_1, v), x_1, v) = h(\varepsilon(x_1, v'), x_1, v) = h(\varepsilon(x_1, v'), x_1, v') = v'$$

by the lack of θ -dependence in the second x_2 variable. $\square$

ε is the function we are looking for. We have just proved that

$$\theta \cap \ker_w \varepsilon(x_1, w) \cap D^2 = 0.$$

Now we need to see that for every $b, c \in D$ there is $d \in D$ so that $\varepsilon(x_1, d) = \varepsilon(x_1, b)$ and $d\theta c$. This will follow if we can show that $\varepsilon(x_1, D_b) = \varepsilon(x_1, D_c)$ for any b and c .

Claim 3.19. $\varepsilon(x_1, D_b) = \varepsilon(x_1, D_c)$ for any $b, c \in G$.

Proof. It will suffice to show $\varepsilon(x_1, D_{x_1}) = \varepsilon(x_1, D_b)$ for any $b \in G$. Suppose not. Then by normality of pp formulas $\varepsilon(x_1, D_{x_1}) \cap \varepsilon(x_1, D_b) = \emptyset$. We will now contradict the tree condition. Note that by Claim 3.18 and abelianness, $\varepsilon(b, -)$ is one-to-one from D_{x_1} to D_b .

Write $\varepsilon_v(w)$ for $\varepsilon(v, w)$ and define polynomials ε_η and definable sets U_η for $\eta \in 2^{<\omega}$ inductively:

1. $\varepsilon_{\langle \rangle}$ is the identity.
2. $\varepsilon_{\eta 0} = \varepsilon_\eta \circ \varepsilon_{x_1}$, $\varepsilon_{\eta 1} = \varepsilon_\eta \circ \varepsilon_b$.
3. $U_{\langle \rangle} = G$.
4. $U_{\eta 0} = \varepsilon_\eta(D_{x_1})$ and $U_{\eta 1} = \varepsilon_\eta(D_b)$.

Using the injectivity of ε , it is easy to see that this uniformly pp-definable family of sets contradicts the tree condition. $\square$

We are now left with the second case. As in the first case, we can write h^* and display all its variables. We get $h^*(\tau(x_1, x_1, x_2, x_3, u) = u^*)$. Again, by letting $x_1 = x_3$ and $u = x_2$ we have $h^*(\tau(x_1), x_1, x_2, x_1, x_2) = x_2^*$ holds in $\mathcal{V}$.

This time though, there can only be θ -dependence in the last variable of h^* since this is the only place where u appears. Claims 3.16 and 3.17 now go through virtually unchanged and we arrive at the same contradiction. This finishes the proof of the lemma. $\square$

Proof of Proposition 3.6. There are two ways of viewing the relationship between $\mathcal{F}_2$ and $\mathcal{G}$. first, the subalgebra generated by x_1/γ and x_2/γ in $\mathcal{G}$ is isomorphic to $\mathcal{F}_2$ so we can think of $\mathcal{F}_2$ as a subalgebra of $\mathcal{G}$. Second, the homomorphism from $\mathcal{F}_4$ to $\mathcal{F}_2$ sending x_1 and x_3 to x_1 and x_2 and x_4 to x_2 factors through $\mathcal{G}$ so $\mathcal{F}_2$ is a homomorphic image of $\mathcal{G}$.

Now fix the α which cross cuts θ on $\mathcal{G}$. Using the second fact, it is clear that in $\mathcal{F}_2$, there is a δ so that $x_1\theta\delta$ and $x_2\alpha\delta$. Using the first fact, we see that this δ is unique. From this we conclude that α cross cuts θ on $\mathcal{F}_2$. $\square$

Theorem 1.19. *If $\mathcal{V}$ is amenable then $\mathcal{V}$ is structured.*

Proof. On $\mathcal{F}_2$, there is the pair of cross cutting equivalence relations α and θ as guaranteed by Proposition 3.6. Let $d(x_1, x_2)$ be the element of $\mathcal{F}_2$ which is θ -related to x_1 and α -related to x_2 .

It is immediate that $d(x, y)$ is a diagonal term for $\mathcal{V}$, that $\theta = \ker_x(d(x, y))$ and $\alpha = \ker_y(d(x, y))$. If a and b are in the same θ -class we say that they have the same first component and if they are in the same α -class we say they have the same second component.

We will now prove that α is a congruence. We need the following fact which is a restatement of Lemma 7.3 from [16].

Fact 3.20. Suppose $\mathcal{V}$ is abelian, $\mathcal{B} \in \mathcal{V}$ and $\tau(x, \bar{y})$ is a term. If $\beta \in \text{Con}(\mathcal{B})$ is combinatorial and $a, b, \bar{c}, \bar{d} \in B$ then if $\tau(a, \bar{c}) = \tau(b, \bar{d})$ and $\tau(a, \bar{c})\beta\tau(b, \bar{c})$ then $\tau(a, \bar{c}) = \tau(b, \bar{c})$.

Note that the conclusion of this fact is stronger than saying β is combinatorial since there is no restriction on $a, b, \bar{c}$ and $\bar{d}$.

Now suppose $a\alpha b$ in $\mathcal{F}_2$. That is, a and b have the same second component. Let h be a unary polynomial on $\mathcal{F}_2$ and let σ be the Malcev term on $\mathcal{F}_2/\theta$. We have

$$\begin{aligned} d(a, h(a)) &= d(a, h(d(\sigma(a, a, a), a))) \\ &= d(a, h(d(\sigma(a, b, b), a))) \end{aligned}$$

and

$$d(a, h(d(\sigma(a, a, a), a)))\theta d(a, h(d(\sigma(a, a, b), a)))$$

since they have the same first component. By Fact 3.20, we have

$$d(a, h(d(\sigma(a, a, a), a))) = d(a, h(d(\sigma(a, a, b), a))).$$

That is, $h(a)$ and $h(d(\sigma(a, a, b), a))$ have the same second component. But

$$d(\sigma(a, a, b), a) = d(b, a) = d(b, b) = b$$

so $h(a)\alpha h(b)$.

Now any algebra with minimal co-affine congruence 1 is combinatorial. Hence the combinatorial algebras form a subvariety satisfying $d(x, z) = d(y, z)$. For any algebra $\mathcal{A}$, $\mathcal{A}/\theta$ is affine and $\mathcal{A}/\alpha$ is combinatorial so $\mathcal{V}$ is the join of these subvarieties. Moreover, d is a diagonal term so we have shown that $\mathcal{V}$ is structured. $\square$

4. APPLICATIONS OF THE STRUCTURE THEOREM

4.1. The spectrum function. The calculation of the uncountable spectrum for varieties of countable type was achieved by Palyutin and announced in [20]. Indeed, Palyutin and Starchenko [21] calculated the spectrum for all Horn classes. Using Theorem 0.11, it is possible to give an alternative calculation which has a more algebraic component.

For this subsection, $\mathcal{L}$ is countable. Suppose $\mathcal{V}$ is a variety in $\mathcal{L}$. If $\mathcal{V}$ is unsuperstable then by [23], $I(\mathcal{V}, \lambda) = 2^\lambda$ for all $\lambda > \aleph_0$ (see Discussion 0.7 for the definition of $I(\mathcal{V}, \lambda)$). If $\mathcal{V}$ is superstable then by Theorem 0.11, there is an affine subvariety $\mathcal{A}$ and a combinatorial subvariety $\mathcal{C}$ so that $\mathcal{V} = \mathcal{A} \otimes \mathcal{C}$. By Fact 0.6 then for all $\lambda \geq \aleph_0$,

$$I(\mathcal{V}, \lambda) = \sum_{\kappa < \lambda} (I(\mathcal{A}, \kappa) \cdot I(\mathcal{C}, \lambda) + I(\mathcal{A}, \lambda) \cdot I(\mathcal{C}, \kappa)) + I(\mathcal{A}, \lambda) \cdot I(\mathcal{C}, \lambda).$$

It follows then that we need only calculate the spectrum functions for affine and combinatorial varieties.

The case of an affine variety has been handled by Baldwin and McKenzie in [2].

Theorem 4.1. *If $\mathcal{A}$ is a superstable affine variety then $I(\mathcal{A}, \aleph_\alpha)$ is exactly one of the following functions of α for $\alpha \geq 0$,*

1. *some fixed finite number,*
2. $2^{\aleph_0}$,
3. $|\alpha + \omega|$,
4. $|\alpha + 2^{\aleph_0}|$,
5. $|\alpha + \omega|^{\aleph_0}$.

The case of a combinatorial variety was discussed in [11]. Although a method for determining the spectrum was discussed at the end of §2 of that paper, the list of functions given in §3 implicitly relied on [20]. To demonstrate this approach's independence from [20] and for completeness, we will include an outline of the calculation here.

Assume $\mathcal{C}$ is a superstable combinatorial variety. By comments made in the introduction or in [10] and in §2 of [11], for the purposes of calculating the spectrum function, we can assume $\mathcal{C}$ is a linear multi-sorted unary variety which is well founded.

To make the following description easier, we make the following conservative change to $\mathcal{C}$. Introduce a new constant for every term which is provably constant in $\mathcal{C}$ and interpret it as the value of this term in each algebra in $\mathcal{C}$.

Every algebra $\mathcal{B} \in \mathcal{C}$ is naturally associated to a tree; the tree of 1-generated subuniverses. Call this $\text{Sub}(\mathcal{B})$. The fact that $\mathcal{C}$ is linear guarantees that what would otherwise just be a partial order is a tree. Notice that if $a, b \in \text{Sub}(\mathcal{B})$ then $a \cap b$ is either empty or in $\text{Sub}(\mathcal{B})$. This follows from the well-foundedness of $\mathcal{C}$. Although this tree is a canonical choice for each algebra, it may have infinite descending chains which make it unsuitable for use in calculating the spectrum. We define another, less canonical, tree.

If $\mathcal{B} \in \mathcal{C}$ and $a, b, c \in \text{Sub}(\mathcal{B})$ then we say that b and c are in the same component above a if $b \cap c \not\subseteq a$. By the linearity of $\mathcal{C}$, this is an equivalence relation on those elements of $\text{Sub}(\mathcal{B})$ which are not contained in a .

Now suppose $\mathcal{B} \in \mathcal{C}$. We will build this new tree by levels and it will have elements of $\text{Sub}(\mathcal{B})$ as nodes. Let the constant subuniverse be the root or, if there are no constants, artificially let the empty set be the root.

On the first level, choose one representative from each component above the root. These will be the nodes on the first level.

Suppose a is on level $n - 1$ and b is a successor on level n . Pick one representative from component above b among elements which are in the same component as b above a . These will be the successors of b on level $n + 1$.

This defines a tree $\mathcal{P}$ from $\mathcal{B}$. Now suppose $a \in \mathcal{B}$. Is there a $b \in \mathcal{P}$ with $a \in b$? We can view $\bigcup \mathcal{P}$ as a subalgebra of $\mathcal{B}$. Consider $\langle a \rangle \cap \bigcup \mathcal{P}$. This equals d for some $d \in \text{Sub}(\mathcal{B})$. There is a $b \in \mathcal{P}$ least so that $d \subseteq b$. If $a \notin \bigcup \mathcal{P}$ then when we considered b , we did not choose anything from the component above b containing a which goes against the construction of $\mathcal{P}$. So $\mathcal{P}$ exhausts $\mathcal{B}$.

Now consider the rank of $\mathcal{P}$. If the rank is undefined this is equivalent to the condition called the ascending chain condition (Definition 2.19 from [11]). It was proved there (Theorem 2.20) that if $\mathcal{C}$ satisfies the ascending chain condition then $I(\mathcal{C}, \lambda) = 2^\lambda$ for all $\lambda \geq \aleph_0$. So assume that $\mathcal{C}$ satisfies the ascending chain condition i.e. the rank of $\mathcal{P}$ is always defined.

Let $\delta_{\mathcal{B}}$ be the supremum of all the ranks of all the trees which could be formed in this way from $\mathcal{B}$. By standard arguments in $\mathcal{L}_{\omega_1, \omega}$ (see for example [14]) if $\delta_{\mathcal{B}} \geq \omega_1$ then $\mathcal{C}$ satisfies the ascending chain condition so $\delta_{\mathcal{B}} < \omega_1$. In fact, if δ , the depth of $\mathcal{C}$, is the supremum of all $\delta_{\mathcal{B}}$ over all $\mathcal{B} \in \mathcal{C}$ then by the same argument $\delta < \omega_1$. Now let us argue that δ is not a limit ordinal.

How could δ be a limit ordinal? For this to happen, there would have to be $\mathcal{B}_i \in \mathcal{C}$ with associated trees $\mathcal{P}_i$ for $i < \omega$ so that the rank of $\mathcal{P}_i$ was δ_i and the limit of the δ_i 's and δ . Now in a unary (or multi-sorted unary)

variety, there is a natural way to “glue” algebras together. For suppose that $\mathcal{A}_0 \subseteq \mathcal{A}_1, \mathcal{A}_2$. Then we can form the disjoint union of $\mathcal{A}_1$ and $\mathcal{A}_2$ over $\mathcal{A}_0$ and obtain an algebra which is in any variety which contains $\mathcal{A}_1$ and $\mathcal{A}_2$. What we would like to do here is glue the B_i ’s together over the constant subuniverse. The only problem is that the constant subuniverses may not be isomorphic. However, if we consider the congruence which collapses all the constants in a single sort to one element in every sort which has a constant then this will only change the root of $\mathcal{P}_i$ and make inessential changes to the nodes. So without loss, we can assume that the constant subuniverse in the $\mathcal{B}_i$ ’s are the same and glue them together to form an algebra of depth at least δ . So the depth of the variety is at least $\delta + 1$.

Now if $\delta > \omega$ then the techniques used in §5 of [8] can easily be adapted to show that $I(\mathcal{C}, \aleph_\alpha)$ is the minimum of $2^{\aleph_\alpha}$ and $\beth_\delta(|\alpha + \omega|)$.

If $\delta < \omega$ then for each $\mathcal{B} \in \mathcal{C}$ we can return to $\text{Sub}(\mathcal{B})$. This now is a canonical choice for the tree associated with $\mathcal{B}$ and δ is just one more than the length of the longest chain of one-generated subuniverses in any algebra in $\mathcal{C}$. It is not hard to show that in this case $I(\mathcal{C}, \aleph_\alpha)$ is the minimum of $2^{\aleph_\alpha}$ and one of the following functions of α

1. some fixed finite number,
2. $2^{\aleph_0}$,
3. $\beth_{\delta-2}(|\alpha + \lambda|^\kappa)$ where $\lambda \in \{\omega, 2^{\aleph_0}\}$ and $\kappa \in \{1, \omega, 2^{\aleph_0}\}$.

Putting this all together then we get

Theorem 4.2. *If $\mathcal{V}$ is a countable variety then $I(\mathcal{V}, \aleph_\alpha)$ is the minimum of $2^{\aleph_\alpha}$ and exactly one of the following functions of α ,*

1. some fixed finite number,
2. $2^{\aleph_0}$,
3. some finite number for $\alpha < \omega$ and $|\alpha|$ for $\alpha \geq \omega$,
4. $\beth_\delta(|\alpha + \lambda|^\kappa)$ for some $\delta < \omega$, $\lambda \in \{\omega, 2^{\aleph_0}\}$ and $\kappa \in \{1, \omega, 2^{\aleph_0}\}$,
5. $\beth_\delta(|\alpha + \omega|)$, $\omega < \delta < \omega_1$ and δ not a limit ordinal,
6. $2^{\aleph_\alpha}$.

Remark 4.3. It is interesting to note that it follows from what we have said that if a variety is superstable then it has NDOP. That is, if $\mathcal{A} \in \mathcal{V}$ and $\mathcal{V}$ is superstable then $\text{Th}(\mathcal{A})$ has NDOP. Although this is a consequence of our theorem, we have no direct explanation for it. To be specific, if $\mathcal{A}$ is an algebra so that $\text{Th}(\mathcal{A})$ is superstable with DOP then we know that $\mathcal{V}(\mathcal{A})$ is unsuperstable but we do not know how to construct, from $\mathcal{A}$, an unsuperstable algebra.

4.2. If $\mathcal{A}$ is any structure and $T = \text{Th}(\mathcal{A})$ then the h -companion T^h of T is the theory of the reduced product $\mathcal{A}^\omega/\mathcal{F}$ where $\mathcal{F}$ is the Fréchet filter. T^h is a Horn theory. In [20], Palyutin states the following elimination of quantifiers result.

Theorem 4.4. *If T^h is superstable and has NDOP then T has elimination of quantifiers up to Boolean combinations of ppfs and formulas in one free variable.*

Now suppose $\mathcal{A}$ is an algebra and $\mathcal{V}(\mathcal{A})$ is superstable. Then of course $\mathcal{A}^\omega/\mathcal{F} \in \mathcal{V}(\mathcal{A})$ so by Palyutin’s theorem, $\text{Th}(\mathcal{A})$ has an elimination of quantifiers as stated. However, in the case of varieties, this is easy to see. Here is a sketch.

Any formula is equivalent in a superstable variety to a pair of formulas; one about the combinatorial factor and one about the affine factor. It is well known that any formula in an affine variety is equivalent to a Boolean combination of ppfs and sentences about the index of one ppf in another. In a superstable combinatorial variety (in fact in any variety equivalent to a linear multi-sorted unary variety) it is easy to show that any formula is equivalent to a Boolean combination of atomic formulas and formulas in one free variable. The latter corresponds to partial information about the isomorphism type of the tree above an element as described in the previous subsection.

4.3. Definition 4.5. An algebra is said to be quasi-affine if it is the subreduct of an affine algebra.

McKenzie had conjectured that if a countable variety had few models in some uncountable power then all of its algebras would be quasi-affine. The quasi-affine algebras in any language form a quasi-variety and Quackenbush in [22] has given a practical set of quasi-identities for these quasi-varieties. If $\mathcal{V}$ is a superstable variety then any algebra in $\mathcal{V}$ is the product of an affine algebra and a combinatorial algebra. Since quasi-affine algebras are closed under products, it suffices to show that any algebra in a superstable combinatorial variety is quasi-affine. In fact, the following is true

Fact 4.6. If $\mathcal{V}$ is equivalent to a multi-sorted unary variety then any $\mathcal{A} \in \mathcal{V}$ is a subreduct of the matrix power of a unary algebra.

It is fairly easy to show that a unary algebra is quasi-affine and that the matrix power of a quasi-affine algebra is quasi-affine so it follows that

Theorem 4.7. *If $\mathcal{V}$ is superstable and $\mathcal{A} \in \mathcal{V}$ then $\mathcal{A}$ is quasi-affine.*

This answers McKenzie's conjecture.

4.4. Open questions. 1. Is any stable variety structured?

2. Any locally finite variety satisfies the tree condition. By considering the conclusion of Theorem 0.15 and the description of combinatorial locally finite decidable varieties in [16], it is clear that a locally finite decidable abelian variety satisfies the normality condition. This is a somewhat tortuous proof.

Since any locally finite variety which satisfies the normality condition is amenable then by Theorem 1.19 it is structured. In [16] it is proved that any decidable locally finite solvable variety is abelian. Is it possible to give a direct proof that a decidable locally finite solvable variety satisfies the normality condition?

REFERENCES

1. J. Baldwin and A. Lachlan, *On universal Horn classes categorical in some infinite power*, Algebra Universalis **3** (1973), 98–111.
2. J. Baldwin and R. McKenzie, *Counting models in universal Horn classes*, Algebra Universalis **15** (1982), 359–384.
3. W. Baur, $\aleph_0$ -categorical modules, J. Symbolic Logic **40** (1975), 213–220.
4. Steve Buechler and Saharon Shelah, *On the existence of regular types*, Ann. Pure Appl. Logic **45** (1989), 277–308.
5. S. Garavaglia, *Decomposition of totally transcendental modules*, J. Symbolic Logic **45** (1980), 155–164.

6. S. Givant, *Universal Horn classes categorical or free in power*, Ann. Math. Logic **15** (1978), 1–53.
7. —, *A representation theorem for universal Horn classes categorical in power*, Ann. Math. Logic **17** (1979), 91–116.
8. L. Harrington and M. Makkai, *The main gap: Counting uncountable models of ω -stable and superstable theories*, Notre Dame J. Formal Logic **26** (1985), 139–177.
9. C. Herrmann, *Affine algebras in congruence modular varieties*, Acta. Sci. Math. **41** (1979), 119–125.
10. B. Hart and S. Starchenko, *Addendum to “A structure theorem for strongly abelian varieties”*, J. Symbolic Logic (to appear).
11. B. Hart and M. Valeriote, *A structure theorem for strongly abelian varieties with few models*, J. Symbolic Logic **56** (1991), 832–852.
12. Bradd Hart, *An exposition of OTOP*, Classification Theory: Chicago, 1985 (J. Baldwin, ed.), Lecture Notes in Math., vol. 1292, Springer-Verlag, Berlin and New York, 1987.
13. D. Hobby and R. McKenzie, *The structure of finite algebras*, Contemp. Math., vol. 76, Amer. Math. Soc., Providence, R.I., 1989.
14. D. Lascar, *Quelques précisions sur la DOP et la profondeur d’une théorie*, J. Symbolic Logic **50** (1985), 316–330.
15. R. McKenzie, *Categorical quasivarieties revisited*, Algebra Universalis **19** (1984), 273–303.
16. R. McKenzie and M. Valeriote, *The structure of locally finite varieties*, vol. 78, Progress in Math., Birkhäuser, Boston, Mass., 1989.
17. T. Mustafin, *The stability theory of polygons*, Akad. Nauk SSSR Sibirsk. Otdel. Trudy Inst. Mat. **8** (1988), 92–108.
18. E. Palyutin, *Categorical Horn classes 1*, Algebra i Logika **12** (1980).
19. —, *The description of categorical quasivarieties*, Algebra and Logic **14** (1976), 86–111.
20. —, *Spectra of varieties*, Soviet Math. Dokl. **39** (1989), 553–554.
21. E. Palyutin and S. Starchenko, *Spectra of Horn classes*, Soviet Math. Dokl. **34** (1987), 394–396.
22. R. Quackenbush, *Quasi-affine algebras*, Algebra Universalis **20** (1985), 318–327.
23. S. Shelah, *Classification theory and the number of non-isomorphic models*, 2nd ed., North-Holland, 1990.

DEPARTMENT OF MATHEMATICS AND STATISTICS, MCMASTER UNIVERSITY, HAMILTON, ONTARIO, CANADA, L8S 4K1

E-mail address: bradd@billy.math.mcmaster.ca

DEPARTMENT OF MATHEMATICS AND STATISTICS, MCMASTER UNIVERSITY, HAMILTON, ONTARIO, CANADA, L8S 4K1 AND INSTITUTE OF MATHEMATICS, NOVOSIBIRSK, RUSSIA

E-mail address: starchen@billy.math.mcmaster.ca

DEPARTMENT OF MATHEMATICS AND STATISTICS, MCMASTER UNIVERSITY, HAMILTON, ONTARIO, CANADA, L8S 4K1

E-mail address: matt@johnny.math.mcmaster.ca