GROWTH RATES OF ALGEBRAS, II:
WIEGOLD DICHOTOMY

KEITH A. KEARNES, EMIL W. KISS, AND ÁGNES SZENDREI

ABSTRACT. We investigate the function $d_{\mathbf{A}}(n)$, which gives the size of a least size generating set for $\mathbf{A}^n$, in the case where $\mathbf{A}$ has a cube term. We show that if $\mathbf{A}$ has a k-cube term and $\mathbf{A}^k$ is finitely generated, then $d_{\mathbf{A}}(n) \in O(\log(n))$ if $\mathbf{A}$ is perfect and $d_{\mathbf{A}}(n) \in O(n)$ if $\mathbf{A}$ is imperfect. When $\mathbf{A}$ is finite, then one may replace "Big O" with "Big Theta" in these estimates.

1. INTRODUCTION

For an algebraic structure $\mathbf{A}$, write $d_{\mathbf{A}}(n) = g$ if g is the least size of a generating set for the direct power $\mathbf{A}^n$. We call the function $d_{\mathbf{A}}(n)$ the *growth rate* of $\mathbf{A}$. The study of this function originated in group theory, and some of its history is surveyed in the preceding paper in this series, [3]. In the present paper we pursue a thread that may also be viewed as originating in group theory, but is directly motivated by some of the results in [3].

James Wiegold proved in [7] that the growth rate of a finite perfect group is logarithmic ($d_{\mathbf{A}}(n) \in \Theta(\log(n))$), and that the growth rate of a finite imperfect group is linear ($d_{\mathbf{A}}(n) \in \Theta(n)$). This result, herein called *Wiegold dichotomy*, was extended by Martyn Quick and Nik Ruškuc in [6] to several kinds of algebras that have underlying group structure. Namely, Quick and Ruškuc showed that a finite algebra $\mathbf{A}$ satisfies $d_{\mathbf{A}}(n) \in \Theta(\log(n))$ if $\mathbf{A}$ is a perfect ring, module, Lie algebra or k-algebra over a field k, and that $d_{\mathbf{A}}(n) \in \Theta(n)$ if $\mathbf{A}$ is an imperfect algebra of one of these types.

To put these results in a broader context, call a term t *basic* if it is a variable, a constant, or a function symbol applied to variables and constants. Call an identity $s \approx t$ basic if both s and t are. Say that a set Σ of identities is *realized* in an algebra $\mathbf{A}$ if it is possible to interpret each function symbol appearing in Σ as a term of $\mathbf{A}$ and each constant as an element of $\mathbf{A}$ so that all identities in Σ are satisfied by $\mathbf{A}$.

1991 *Mathematics Subject Classification.* 08A30 (08B05, 08B10).

Key words and phrases. Growth rate, Wiegold dichotomy, perfect algebra, basic identity, cube term, parallelogram term, maximal subalgebra.

This material is based upon work supported by the Hungarian National Foundation for Scientific Research (OTKA) grant no. K77409, K83219, and K104251.

For example, every algebra $\mathbf{A}$ that has underlying group structure realizes the (basic) identities

$$F(x,y,y) \approx x \qquad \text{and} \qquad F(y,y,x) \approx x, \tag{1.1}$$

because these identities hold in $\mathbf{A}$ for the group term $F(x_1,x_2,x_3) = x_1x_2^{-1}x_3$. A term for which the identities (1.1) hold in $\mathbf{A}$ is called a *Maltsev term* for $\mathbf{A}$.

Our paper [3] asks the question: Which sets Σ of basic identities impose a restriction on growth rates of algebras? Phrased differently: For which sets Σ is there an algebra $\mathbf{A}$ such that its growth rate $d_{\mathbf{A}}(n)$ does not occur as the growth rate of any algebra realizing Σ? The answer is: exactly those Σ which entail the existence of a pointed cube term. A *pointed cube term* is a term $F(x_1,\dots,x_m)$ with respect to which $\mathbf{A}$ satisfies an array of identities of the form

$$\begin{aligned} F(\mathbf{y}_1) &\approx x, \\ &\vdots \\ F(\mathbf{y}_k) &\approx x, \end{aligned} \tag{1.2}$$

where each of the elements of each tuple $\mathbf{y}_i$ is a variable or an element of $\mathbf{A}$, and a further condition is satisfied. The condition is that, when (1.2) is written as a matrix equation, $F(M) = \mathbf{x}$, with

$$M = \begin{pmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_k \end{pmatrix} \quad \text{and} \quad \mathbf{x} = \begin{pmatrix} x \\ \vdots \\ x \end{pmatrix},$$

then each *column* of M contains a symbol (a variable or constant) that is different from x. The term F is a *p-pointed, k-cube term* if the matrix M contains p distinct elements of $\mathbf{A}$ and k rows. Here are three basic examples: a binary term $F(x_1,x_2)$ for which some element $1 \in A$ is a left and right unit element is a 1-pointed, 2-cube term for $\mathbf{A}$, since $\mathbf{A}$ satisfies the row equations of

$$F\begin{pmatrix} 1 & x \\ x & 1 \end{pmatrix} \approx \begin{pmatrix} x \\ x \end{pmatrix}.$$

A Maltsev term $F(x_1,x_2,x_3)$ for $\mathbf{A}$ is a 0-pointed, 2-cube term for $\mathbf{A}$, since the identities in (1.1) can be rewritten as the row equations of

$$F\begin{pmatrix} x & y & y \\ y & y & x \end{pmatrix} \approx \begin{pmatrix} x \\ x \end{pmatrix}.$$

A *majority term* for $\mathbf{A}$ is a 0-pointed, 3-cube term $F(x_1,x_2,x_3)$ for which $\mathbf{A}$ satisfies the row equations of

$$F\begin{pmatrix} x & x & y \\ x & y & x \\ y & x & x \end{pmatrix} \approx \begin{pmatrix} x \\ x \\ x \end{pmatrix}.$$

A 0-pointed k-cube term is usually referred to as a *k-cube term*.

We prove in [3] that if Σ is a set of basic identities which entails no pointed cube term, then for any algebra $\mathbf{A}$ there is an algebra $\mathbf{B}$ that realizes Σ and has the same growth rate as $\mathbf{A}$. Thus the realization of Σ imposes no restriction on growth rates. On the other hand, if Σ entails a p-pointed, k-cube term and $\mathbf{A}$ is a (possibly infinite) algebra for which $\mathbf{A}^{p-1+k}$ (if $p > 0$) or $\mathbf{A}^k$ (if $p = 0$) is finitely generated, then $d_{\mathbf{A}}(n)$ is bounded above by a polynomial function of n. This is a restriction.

In the current paper we use different techniques to establish stronger results for algebras with (0-pointed) cube terms, namely we establish that Wiegold dichotomy holds for such algebras. We show that if $\mathbf{A}$ has a k-cube term and $\mathbf{A}^k$ is finitely generated, then $d_{\mathbf{A}}(n) = O(\log(n))$ if $\mathbf{A}$ is perfect and $d_{\mathbf{A}}(n) = O(n)$ if $\mathbf{A}$ is imperfect. ("Big O" can be strengthened to "Big Theta" when $\mathbf{A}$ is finite.) In this statement the word "perfect" is used with respect to the modular commutator (see [2]), namely an algebra is perfect if it has no nontrivial abelian homomorphic image.

Our approach will be through an analysis of maximal subalgebras of powers of $\mathbf{A}$. Cube terms were discovered and investigated first in [1], while an equivalent type of term was discovered independently and investigated in [5]. It is the results of the latter paper that are applicable to the analysis of maximal subalgebras of powers.

2. Preliminaries

$[n]$ denotes the set $\{1, \dots, n\}$. A tuple in A^n may be denoted $(a_1, \dots, a_n)$ or $\mathbf{a}$. A tuple $(a, a, \dots, a) \in A^n$ with all coordinates equal to a may be denoted $\hat{a}$. The size of a set A, the length of a tuple $\mathbf{a}$, and the length of a string σ are denoted $|A|$, $|\mathbf{a}|$ and $|\sigma|$. Structures are denoted in bold face font, e.g. $\mathbf{A}$, while the universe of a structure is denoted by the same character in italic font, e.g., A. The subuniverse of $\mathbf{A}$ generated by a subset $G \subseteq A$ is denoted $\langle G \rangle$.

We will use Big O notation. If f and g are real-valued functions defined on some subset of the real numbers, then $f \in O(g)$ and $f = O(g)$ both mean that there are positive constants M and N such that $|f(x)| \leq M|g(x)|$ for all $x > N$. We write $f \in \Omega(g)$ and $f = \Omega(g)$ to mean that there are positive constants M and N such that $|f(x)| \geq M|g(x)|$ for all $x > N$. Finally, $f \in \Theta(g)$ and $f = \Theta(g)$ mean that both $f \in O(g)$ and $f \in \Omega(g)$ hold.

Our focus in this paper is on obtaining good upper bounds for $d_{\mathbf{A}}(n)$ whether $\mathbf{A}$ is finite or infinite. When $\mathbf{A}$ is finite, the upper bounds we obtain are asymptotically equal to the easily-proved lower bounds mentioned here:

Theorem 2.1. *If $\mathbf{A}$ is a finite algebra of more than one element, then*

(1) $d_{\mathbf{A}}(n) \in \Omega(\log(n))$.
(2) $d_{\mathbf{A}}(n) \in \Omega(n)$ *if $\mathbf{A}$ is imperfect.*

Proof. Item (1) is proved in Theorem 2.2.2 of [3]. Item (2) follows from Corollary 2.2.5 (2) of [3]. □

We need one preliminary result for the case when $\mathbf{A}$ is infinite.

Theorem 2.2. *If $\mathbf{A}^k$ is a finitely generated algebra with a k-cube term, then $d_{\mathbf{A}}(n) \in O(n^{k-1})$.*

Proof. This is the special case of Corollary 5.2.4 of [3] for (0-pointed) k-cube terms. □

In particular, if $\mathbf{A}$ has a k-cube term and $\mathbf{A}^k$ is finitely generated, then all finite powers of $\mathbf{A}$ are finitely generated.

Theorem 2.2 implies that if $\mathbf{A}$ has a Maltsev term (i.e., a 2-cube term) and $\mathbf{A}^2$ is finitely generated, then $d_{\mathbf{A}}(n) \in O(n)$. We will apply this fact when $\mathbf{A}$ is an *affine* algebra (i.e., an abelian algebra with a Maltsev term).

3. Maximal subuniverses of powers

In this section we relate arbitrary maximal subuniverses of $\mathbf{A}^n$ to critical maximal subuniverses. The results of this section require no assumptions on $\mathbf{A}$.

Definition 3.1. If R is a subuniverse of an algebra $\mathbf{B}$ and $\varphi\colon \mathbf{B} \to \mathbf{C}$ is a surjective homomorphism such that $R = \varphi^{-1}(\varphi(R))$, we will say that R is *induced by the homomorphism* φ.

Lemma 3.2. *If M is a maximal subuniverse of $\mathbf{A}^n$ and $\varphi\colon \mathbf{A}^n \to \mathbf{C}$ is a surjective homomorphism such that $\varphi(M) \neq C$, then $\varphi(M)$ is a maximal subuniverse of $\mathbf{C}$ and M is induced by φ.*

Proof. If S is a proper subuniverse of $\mathbf{C}$ containing $\varphi(M)$, then $M \subseteq \varphi^{-1}(\varphi(M)) \subseteq \varphi^{-1}(S) \subsetneq A^n$, where $\subsetneq$ holds, because $\varphi(\varphi^{-1}(S)) = S \subsetneq C = \varphi(A^n)$. Hence the maximality of M forces that $M = \varphi^{-1}(\varphi(M))$ and $M = \varphi^{-1}(S)$. The first equality proves that M is induced by φ, while the second equality implies that $\varphi(M) = \varphi(\varphi^{-1}(S)) = S$, so $\varphi(M)$ is a maximal subuniverse of $\mathbf{C}$. □

Definitions 3.3. [5] A *compatible n-ary relation* of $\mathbf{A}$ is a subuniverse of $\mathbf{A}^n$.

A compatible relation R is *critical* if it is completely $\cap$-irreducible in the subuniverse lattice of $\mathbf{A}^n$ and directly indecomposable as a relation. (The latter means that R is not of the form $S \times T$ for subsets $S \subseteq A^U$ and $T \subseteq A^V$, where $\{U, V\}$ is a partition of $[n]$ into two cells.)

Any maximal subuniverse M of $\mathbf{A}^n$ is completely $\cap$-irreducible in the subuniverse lattice of $\mathbf{A}^n$, so a critical maximal subuniverse of $\mathbf{A}^n$ is just a maximal subuniverse that is directly indecomposable as a relation.

Definition 3.4. If M is a subuniverse of $\mathbf{A}^n$, then a *support* of M is a subset $U \subseteq [n]$ such that $\pi_U(M) \neq A^U$, where $\pi_U\colon \mathbf{A}^n \to \mathbf{A}^U$ is the projection homomorphism.

Lemma 3.5. *If M is a maximal subuniverse of $\mathbf{A}^n$, then M has a unique minimal support. If U is the minimal support of M, then $M_U := \pi_U(M)$ is a critical maximal subuniverse of $\mathbf{A}^U$, $M = M_U \times A^{U'}$ for $U' = [n] \setminus U$, and M is induced by the projection π_U. In particular, M itself is critical if and only if its unique support is $[n]$.*

Proof. For any set $U \subseteq [n]$ let $M_U := \pi_U(M)$. If U is a support of M, then Lemma 3.2 applied to π_U yields that M_U is a maximal subuniverse of $\mathbf{A}^U$, M is induced by π_U, and hence $M = \pi_U^{-1}(M_U) = M_U \times A^{U'}$ for $U' = [n] \setminus U$.

Now assume that $U, V \subseteq [n]$ are distinct minimal supports of the maximal subuniverse $M \leq \mathbf{A}^n$. U and V must be incomparable under inclusion. We shall view elements of $\mathbf{A}^n$ as functions from $[n]$ to A. In this language, M is a proper subset of the set of all functions, M_U is the set of restrictions to U of the functions in M, and M contains all functions whose restriction to U belongs to M_U. Similarly, M_V is the set of restrictions to V of the functions in M, and M contains all functions whose restriction to V belongs to M_V. Since $M_U \neq A^U$, there is a function $f\colon U \to A$ that is not in M_U. Since V is a minimal support and $U \cap V$ is properly contained in V, it follows that every function $U \cap V \to A$ is the restriction of some function in M. In particular, $f|_{U\cap V} = g|_{U\cap V}$ for some $g \in M$. Let $h \in A^n$ be any function that agrees with f on U and g on V. Then $h|_U = f \notin M_U$, so $h \notin M$. Yet $h|_V = g|_V \in M_V$, so $h \in M$, a contradiction. This shows that M has a unique minimal support.

Let U be the minimal support of M. The second statement of the lemma, except for the criticality of M_U, follows from the first paragraph of this proof. To show that M_U is a critical, assume that $M_U = S \times T$, where $S \leq \mathbf{A}^X$ and $T \leq \mathbf{A}^Y$ for some partition $\{X, Y\}$ of U. Since $M_U \neq A^U$, either $A^X \neq S = \pi_X^{\mathbf{A}^U}(M_U) = \pi_X^{\mathbf{A}^n}(M)$, or $A^Y \neq T = \pi_Y^{\mathbf{A}^U}(M_U) = \pi_Y^{\mathbf{A}^n}(M)$. Either way, one obtains that X or Y is a proper subset of U that is a support of M, contradicting the minimality of U.

For the final statement of the lemma, if the minimal support of M is $[n]$, then $\pi_{[n]}(M) = M$ is critical by the second statement of the lemma. Conversely, assume that M is critical and $U \subseteq [n]$ is its minimal support. Since $M = M_U \times A^{U'}$ and M is directly indecomposable as a relation, we get $U' = \emptyset$, equivalently $[n] = U$. □

4. The parallelogram property for critical relations

In the preceding section we showed that all maximal subuniverses of $\mathbf{A}^n$ are induced by critical maximal subuniverses on projections $\mathbf{A}^U$ of $\mathbf{A}^n$. In this section we show that the critical maximal subuniverses of $\mathbf{A}^U$ have a special structure when $\mathbf{A}$ has a k-cube term.

Definition 4.1. [5] Given a partition $\{S, T\}$ of $[n]$ into two cells, write $\mathbf{xy}$ for a tuple in A^n to mean that $\mathbf{x} \in A^S$ and $\mathbf{y} \in A^T$. A compatible n-ary relation R satisfies the *parallelogram property* if, for any partition $\{S, T\}$ of $[n]$, $\mathbf{au}, \mathbf{av}, \mathbf{bv} \in R$ implies $\mathbf{bu} \in R$.

Theorem 3.5 and Theorem 3.6 (3) of [5] together prove the following theorem.

Theorem 4.2. *A variety $\mathcal{V}$ has a k-cube term if and only if every member $\mathbf{A} \in \mathcal{V}$ has the property that any critical relation of $\mathbf{A}$ of arity at least k has the parallelogram property.*

It follows from this theorem and Lemma 3.5 that if $\mathbf{A}$ has a k-cube term, and M is a maximal subuniverse of $\mathbf{A}^n$ with minimal support $U \subseteq [n]$, then $M = M_U \times A^{U'}$ ($U' = [n] \setminus U$), and either M_U has arity less than k or M_U is a critical maximal subuniverse of $\mathbf{A}^U$ that has the parallelogram property. (In the latter case, M itself will also have the parallelogram property.) Our next step is to investigate the structure of maximal subuniverses with the parallelogram property.

The paper [5] analyzes arbitrary compatible relations with the parallelogram property in congruence modular varieties. It is shown in [1] that any algebra with a k-cube term generates a congruence modular variety, so the results of [5] apply here. The first step in the analysis is the "reduction" of a relation, which we describe next.

Suppose that $R \leq \mathbf{A}^n$ is a compatible relation with the parallelogram property; as a special case, suppose that $M \leq \mathbf{A}^n$ is a maximal critical subuniverse with the parallelogram property. For the first step in the reduction, realize R as a subdirect product $R \leq_{\text{sd}} \prod_{i=1}^n \mathbf{A}_i$, where $\mathbf{A}_i := \pi_i(R) \leq \mathbf{A}$. In the special case involving the maximal subuniverse M we will have $\mathbf{A}_i = \pi_i(M) = \mathbf{A}$ unless the projection of M onto one single coordinate is not surjective. This happens only if M has a support of size one, which, by criticality, implies that M is a unary relation. We henceforth consider only M of arity at least two, so that in our special case $\pi_i(M) = A$ for all i. Thus, in the first step in reduction, nothing happens if M is maximal and of arity greater than one.

Second, define relations, called *coordinate kernels* in [5],

$$\theta_i = \{(a, b) \in A_i^2 \mid \exists \mathbf{c} \in \prod_{j \neq i} \mathbf{A}_j \ (a\mathbf{c} \in R \ \& \ b\mathbf{c} \in R)\}.$$

It is proved in Lemma 2.3 of [5] that (i) each θ_i is a congruence on $\mathbf{A}_i$, and (ii) R is induced by the homomorphism $\psi\colon \prod \mathbf{A}_i \to \prod \mathbf{A}_i/\theta_i$ that is the natural map in each coordinate. The relation $\overline{R} = \psi(R)$ is the *reduction* of R.

In our special case $M \leq \mathbf{A}^n$ is critical and maximal, therefore by Lemma 3.2 its reduction $\overline{M} = \psi(M)$ is a maximal subuniverse of $\prod \mathbf{A}/\theta_i$.

The next result is a specialization of (some parts of) Theorem 2.5 of [5] to the case where M is a critical maximal subuniverse of $\mathbf{A}^n$ and $n > 1$. We maintain the numbering of [5], but omit the unused parts of the theorem.

Theorem 4.3. *Let M be a critical maximal subuniverse of $\mathbf{A}^n$ that satisfies the parallelogram property, and let $\overline{M} \leq \prod \mathbf{A}/\theta_i$ be its reduction. If $n > 1$ and $\mathbf{A}$ lies in a congruence modular variety, then the following hold.*

(1) $\overline{M} \leq \prod \mathbf{A}/\theta_i$ *is a representation of* $\overline{M}$ *as a subdirect product of subdirectly irreducible algebras.*
(5)* *If* $n > 2$*, then the monolith of* $\mathbf{A}/\theta_i$ *is the total relation; i.e.* $\mathbf{A}/\theta_i$ *is simple.*
(7)* *If* $n > 2$*, then each simple algebra* $\mathbf{A}/\theta_i$ *is abelian.*

Here, items (5) and (7) are marked with asterisks, because we have altered the statement of (5) from [5] in order to take into account that $\overline{M}$ is a *maximal* subuniverse of $\prod \mathbf{A}/\theta_i$ and we have altered the statement of (7) in order to take into account the conclusion from (5)* that $\mathbf{A}/\theta_i$ is simple.

We explain what this theorem contributes to our current investigation. Suppose that $\mathbf{A}$ has a k-cube term. Suppose also that $M \leq \mathbf{A}^n$ is maximal, U is the minimal support of M, and $M = M_U \times A^{U'}$ is induced by $\pi_U \colon \mathbf{A}^n \to \mathbf{A}^U$. If $|U|$ is at least as large as $\max\{3, k\}$, then the theorem proves that M_U is induced by a homomorphism $\psi \colon \mathbf{A}^U \to \prod_U \mathbf{A}/\theta_i$ where each factor $\mathbf{A}/\theta_i$ is a simple abelian algebra. Thus, M itself is induced by the composition of the surjective homomorphisms

$$\mathbf{A}^n \xrightarrow{\pi_U} \mathbf{A}^U \longrightarrow (\mathbf{A}/[1,1])^U \longrightarrow \prod_U \mathbf{A}/\theta_i,$$

where the last two maps are a factorization of the map $\psi \colon \mathbf{A}^U \to \prod_U \mathbf{A}/\theta_i$ which induces M_U, and these two maps are defined coordinatewise by the natural maps $\mathbf{A} \to \mathbf{A}/[1,1] \to \mathbf{A}/\theta_i$. (We have $\theta_i \geq [1,1]$, since $\mathbf{A}/\theta_i$ is abelian.) Hence M is induced by the sub-composition $\mathbf{A}^n \xrightarrow{\pi_U} \mathbf{A}^U \longrightarrow (\mathbf{A}/[1,1])^U$, which may be factored another way as $\mathbf{A}^n \xrightarrow{\eta} (\mathbf{A}/[1,1])^n \xrightarrow{\pi_U} (\mathbf{A}/[1,1])^U$. Hence M is induced by the single map η, which maps $\mathbf{A}^n$ onto its abelianization. Altogether this proves the desired result:

Theorem 4.4. *Assume that* $\mathbf{A}$ *has a* k*-cube term. If* $M \leq \mathbf{A}^n$ *is a maximal subuniverse, then either*

(π) M *is induced by a projection* $\pi_U \colon \mathbf{A}^n \to \mathbf{A}^U$ *for some subset* $U \subseteq [n]$ *satisfying* $|U| < \max\{3, k\}$*, or*
(η) M *is induced by* $\eta \colon \mathbf{A}^n \to (\mathbf{A}/[1,1])^n$*.*

5. A solution to a combinatorial problem

To derive our result on growth rates from Theorem 4.4, we will use a solution to the following problem: If B is a finite set and $n \geq k > 1$ are integers, then how small can a set $G \subseteq B^n$ be if its projection onto any subset of k coordinates is surjective?

If B is finite, $G \subseteq B^n$ and $|G| = g$, then G can be linearly ordered and taken to be the sequence of rows of a $g \times n$ matrix of elements of B, say $[b_{i,j}]$. If

$$\sigma \colon \quad 1 \leq j(1) < \cdots < j(k) \leq n$$

is a selection of k numbers between 1 and n, then the projection of G onto the coordinates in σ is the set of row vectors $(b_{1,j(1)}, \ldots, b_{1,j(k)}), \ldots, (b_{g,j(1)}, \ldots, b_{g,j(k)})$

which occur as the set of rows of the $g \times k$ minor of $[b_{i,j}]$ whose column indices are the indices in σ. G projects surjectively onto each k coordinates of B^n if and only if, for each choice σ of k column indices, the set of row vectors of the corresponding $g \times k$ minor of $[b_{i,j}]$ exhausts B^k. Therefore, call a $g \times k$ matrix of elements of B a *bad minor* (or *bad matrix*) if its rows fail to exhaust B^k. The desired property of G is that its associated matrix has no bad minors.

Theorem 5.1. *Let B be a finite set of size $|B| = b > 1$. Let $n \geq k > 1$ be natural numbers, and set $u = b^k/(b^k - 1)$. If $g \geq k \log_u(n) + \log_u\left(b^k/k!\right)$, then there is a matrix in $B^{g\times n}$ with no bad minors.*

Proof. This is a probabilistic proof. Our sample space is the set $B^{g\times n}$ of all $g \times n$ matrices of elements of B. Our probability distribution is the uniform one, so each individual matrix $M \in B^{g\times n}$ has probability $P(M) = |B^{g\times n}|^{-1} = b^{-gn}$. For each matrix $M \in B^{g\times n}$ and each sequence of k column indices,

$$\sigma\colon \quad 1 \leq j(1) < \cdots < j(k) \leq n,$$

let M_σ denote the $g \times k$ minor of M whose column indices are those enumerated by σ (called the *σ-minor* of M). Let X_σ be the random variable whose value at the element $M \in B^{g\times n}$ is 1 if M_σ is a bad minor and 0 otherwise, i.e., X_σ is the indicator variable for bad σ-minors.

Claim 5.2. *For any σ, the expected value of X_σ satisfies*

$$E(X_\sigma) \leq b^k \left(b^k - 1\right)^g b^{-gk}. \tag{5.1}$$

The expectation is computed

$$\begin{array}{rcl} E(X_\sigma) & = & \sum_{M\in B^{g\times n}} (X_\sigma(M) \cdot P(M)) \\ & = & \sum_{M\in B^{g\times n}} (X_\sigma(M) \cdot b^{-gn}) \\ & = & \left(\sum_{M\in B^{g\times n}} X_\sigma(M)\right) b^{-gn}, \end{array}$$

where the sum $\sum_{M\in B^{g\times n}} X_\sigma(M)$ on the last line represents the number of matrices in $B^{g\times n}$ whose σ-minor is bad. By definition, a $g \times k$ matrix is bad if some tuple $\mathbf{b} \in B^k$ does not appear among its rows. So, for each $\mathbf{b} \in B^k$, let $\mathcal{U}_\mathbf{b}$ denote the set of all $g \times k$ matrices where $\mathbf{b}$ does not appear among the rows. $|\mathcal{U}_\mathbf{b}|$ can be computed by noting that the g rows of a matrix in $\mathcal{U}_\mathbf{b}$ may be freely chosen from the set $B^k - \{\mathbf{b}\}$, which has size $b^k - 1$, so $|\mathcal{U}_\mathbf{b}| = (b^k - 1)^g$. The bad $g \times k$ matrices are those from $\bigcup_{\mathbf{b}\in B^k} \mathcal{U}_\mathbf{b}$. Since the cardinality of the union is no more than the sum of the individual cardinalities, and these summands have the same size, we get that the number of bad $g \times k$ matrices is no more than $|B^k| \cdot |\mathcal{U}_\mathbf{b}| = b^k(b^k - 1)^g$. Each bad $g \times k$ matrix N can be extended in $b^{g(n-k)}$ ways to a matrix $M \in B^{g\times n}$ whose σ-minor satisfies $M_\sigma = N$, so the number of matrices in $B^{g\times n}$ with a bad σ-minor is

no more than $b^k(b^k-1)^g b^{g(n-k)}$. Hence

$$E(X_\sigma) = \left(\sum_{M \in B^{g\times n}} X_\sigma(M)\right) b^{-gn} \leq b^k(b^k-1)^g b^{g(n-k)} b^{-gn} = b^k(b^k-1)^g b^{-gk},$$

as claimed.

If $X := \sum_\sigma X_\sigma$ is the sum of all X_σ as σ ranges over all $\binom{n}{k}$ choices of k column indices and $M \in B^{g\times n}$, then $X(M)$ equals the number of bad $g \times k$ minors of M. Since expectation is linear, and since $\binom{n}{k} < n^k/k!$ when $n \geq k > 1$, we get from (5.1) that

$$E(X) = \sum_\sigma E(X_\sigma) \leq \binom{n}{k} b^k(b^k-1)^g b^{-gk} < n^k(b^k/k!)(b^k-1)^g b^{-gk}.$$

If it is the case that

$$n^k(b^k/k!)(b^k-1)^g b^{-gk} \leq 1, \tag{5.2}$$

then we will have $E(X) < 1$, meaning that the expected number of bad minors in an element of $B^{g\times n}$ is strictly less than 1. This can happen only if matrices without bad minors exist. Rewriting (5.2) as

$$n^k \leq \left(\frac{b^k}{(b^k-1)}\right)^g (b^k/k!)^{-1} = u^g (b^k/k!)^{-1},$$

using the definition $u = b^k/(b^k-1)$, we can solve for g to get

$$g \geq k\log_u(n) + \log_u(b^k/k!). \tag{5.3}$$

When this inequality holds we get that (5.2) holds, so a matrix with no bad minors exists. This is exactly the statement of the theorem. □

Corollary 5.3. *Let B be a finite set of size $|B| = b > 1$. Let $n \geq k > 1$ be natural numbers, and set $u = b^k/(b^k-1)$. If $g = \lceil k\log_u(n) + \log_u\left(b^k/k!\right)\rceil$, then there exists a subset $G \subseteq B^n$ of size g whose projection onto any k coordinates of B^n is surjective.*

Proof. By the theorem, there is a matrix in $B^{g\times n}$ with no bad minors. The set $G \subseteq B^n$ consisting of the rows of this matrix has size g and projects surjectively onto any k coordinates of B^n. □

Corollary 5.4. *Let $\mathbf{A}$ be an algebra of more than one element, and suppose that for some $k > 1$ the algebra $\mathbf{A}^k$ is generated by a finite set $H \subseteq A^k$. Let $B \subseteq A$ be the set of elements of A that appear in the coordinates of tuples in H. Set $b = |B|$ and $u = b^k/(b^k-1)$. For any $n \geq k$, if $g = \lceil k\log_u(n) + \log_u\left(b^k/k!\right)\rceil$, then there exists a subset $G \subseteq B^n \subseteq A^n$ of size g such that the subalgebra $\mathbf{S} = \langle G\rangle \leq \mathbf{A}^n$ has the property that the projection of $\mathbf{S}$ onto any k coordinates of $\mathbf{A}^n$ is surjective.*

Proof. Since $\mathbf{A}$ has more than one element and $k > 1$, the algebra $\mathbf{A}^k$ is not generated by diagonal tuples. Therefore the generating set H has a nondiagonal tuple. It follows that the set B of elements of A that appear in the coordinates of tuples in H is finite and has more than one element.

We choose $G \subseteq B^n$ as in Corollary 5.3 so that the projection of G onto any k coordinates of B^n is surjective. Any projection $\pi_U \colon \mathbf{S} = \langle G \rangle \to \mathbf{A}^U$ of the subalgebra $\mathbf{S} \leq \mathbf{A}^n$ onto a k-element set $U \subset [n]$ contains the projection of the subset $B^n \subseteq S$ onto those k coordinates, and $\pi_U(B^n) = B^U$ is a generating set for $\mathbf{A}^U$, because B^U contains a copy of H. Hence $\mathbf{S} \leq \mathbf{A}^n$ has the property that the projection onto any k coordinates of $\mathbf{A}^n$ is surjective. □

6. Growth rates of algebras with a cube term

In this section we combine the preceding results to obtain the following.

Theorem 6.1. *Suppose that* $\mathbf{A}$ *has a* k*-cube term and that* $\mathbf{A}^k$ *is finitely generated. If* $\mathbf{A}$ *is perfect, then* $d_{\mathbf{A}}(n) \in O(\log(n))$*. If* $\mathbf{A}$ *is imperfect, then* $d_{\mathbf{A}}(n) \in O(n)$*.*

Proof. According to Theorem 2.2, the fact that $\mathbf{A}^k$ is finitely generated implies that $\mathbf{A}^n$ is finitely generated for all finite n. Hence any proper subuniverse of $\mathbf{A}^n$ is contained in a maximal subuniverse of $\mathbf{A}^n$.

According to Theorem 4.4, if $M \leq \mathbf{A}^n$ is a maximal subuniverse, then either

- (π) M is induced by a projection $\pi_U \colon \mathbf{A}^n \to \mathbf{A}^U$ for some subset $U \subseteq [n]$ satisfying $|U| < \max\{3, k\}$, or
- (η) M is induced by $\eta \colon \mathbf{A}^n \to (\mathbf{A}/[1,1])^n$.

For each n, choose a subset $G_\pi \subseteq A^n$ of size $O(\log(n))$ such that the subalgebra $\langle G_\pi \rangle$ of $\mathbf{A}^n$ has the property that its projection onto any $\max\{3, k\}$ coordinates of $\mathbf{A}^n$ is surjective. The existence of such a set is guaranteed by Corollary 5.4. Clearly G_π is contained in no maximal subuniverse of $\mathbf{A}^n$ that is induced by a projection onto any subset of less than $\max\{3, k\}$ coordinates.

The algebra $\mathbf{A}/[1,1]$ is abelian and has a cube term, so $\mathbf{A}/[1,1]$ is affine by [2, Corollary 5.9]. According to the remarks following Theorem 2.2, $(\mathbf{A}/[1,1])^n$ contains a set of generators of size $O(n)$. For each n, choose a set $G_\eta \subseteq A^n$ of size $O(n)$ such that $\eta(G_\eta)$ generates $(\mathbf{A}/[1,1])^n$. Then G_η is contained in no maximal subuniverse of $\mathbf{A}^n$ induced by η.

We now have that $G_\pi \cup G_\eta$ is a set of size $O(n)$ that is contained in no maximal subuniverse of $\mathbf{A}^n$, hence $G_\pi \cup G_\eta$ is a generating set for $\mathbf{A}^n$ of size $O(n)$.

When $\mathbf{A}$ is perfect, then $\mathbf{A}^n$ has no maximal subuniverses induced by η, so G_π is a generating set for $\mathbf{A}^n$ of size $O(\log(n))$. □

Corollary 6.2. *Suppose that* $\mathbf{A}$ *is finite, has more than one element, and has a* k*-cube term. If* $\mathbf{A}$ *is perfect, then* $d_{\mathbf{A}}(n) \in \Theta(\log(n))$*. If* $\mathbf{A}$ *is imperfect, then* $d_{\mathbf{A}}(n) \in \Theta(n)$*.*

Proof. Combine the upper bounds of Theorem 6.1 with the lower bounds of Theorem 2.1. □

References

[1] Berman, Joel, Idziak, Paweł, Marković, Petar, McKenzie, Ralph, Valeriote, Matthew, Willard, Ross, *Varieties with few subalgebras of powers.* Trans. Amer. Math. Soc. **362** (2010), no. 3, 1445–1473.

[2] Freese, Ralph, McKenzie, Ralph, *Commutator Theory for Congruence Modular Varieties*, London Mathematical Society Lecture Note Series **125**, Cambridge University Press, Cambridge, 1987.

[3] Kearnes, Keith, Kiss, Emil, Szendrei, Ágnes, *Growth rates of finite algebras, I: pointed cube terms.* J. Austral. Math. Soc., to appear.

[4] Kearnes, Keith, Kiss, Emil, Szendrei, Ágnes, *Growth rates of finite algebras, III: solvable algebras.* Algebra Universalis, to appear.

[5] Kearnes, Keith, Szendrei, Ágnes, *Clones of algebras with parallelogram terms.* Internat. J. Algebra Comput. **22** (2012), no. 1, 1250005, 30 pp.

[6] Quick, Martyn, Ruškuc, Nik, *Growth of generating sets for direct powers of classical algebraic structures.* J. Austral. Math. Soc. **89** (2010), 105–126.

[7] Wiegold, James, *Growth sequences of finite groups.* Collection of articles dedicated to the memory of Hanna Neumann, VI. J. Austral. Math. Soc. **17** (1974), 133–141.

(Keith Kearnes) Department of Mathematics, University of Colorado, Boulder, CO 80309-0395, USA

E-mail address: Keith.Kearnes@Colorado.EDU

(Emil W. Kiss) Loránd Eötvös University, Department of Algebra and Number Theory, H–1117 Budapest, Pázmány Péter stny 1/c., Hungary

E-mail address: ewkiss@cs.elte.hu

(Ágnes Szendrei) Department of Mathematics, University of Colorado, Boulder, CO 80309-0395, USA

E-mail address: Agnes.Szendrei@Colorado.EDU