THE SUBPOWER MEMBERSHIP PROBLEM FOR FINITE ALGEBRAS WITH CUBE TERMS

ANDREI BULATOV, PETER MAYR, AND ÁGNES SZENDREI

Abstract. Our main result is that the subpower membership problem $\mathrm{SMP}(\mathcal{K})$ is in P if $\mathcal{K}$ is a finite set of finite algebras with a cube term in a residually small variety. We also prove that for any finite set of finite algebras $\mathcal{K}$ in a variety with a cube term, the following three problems are polynomial time equivalent, and lie in NP: $\mathrm{SMP}(\mathcal{K})$, $\mathrm{SMP}(\mathbb{HSK})$, and finding compact representations for subpowers in $\mathcal{K}$.

1. Introduction

Throughout the introduction we will only consider (classes of) algebras in a finite language.

The *subpower membership problem* for a finite algebra $\mathbf{A}$ is the following combinatorial decision problem: given finitely many elements $a_1, \dots, a_k$ and b in $\mathbf{A}^n$, determine whether b lies in the subalgebra $\mathbf{B}$ of $\mathbf{A}^n$ generated by $\{a_1, \dots, a_k\}$. A naive algorithm for solving this problem is to generate all elements of $\mathbf{B}$, and then check whether b is in $\mathbf{B}$. Since the size of the input $a_1, \dots, a_k, b$ is $(k+1)n$, while the best upper bound for the size of $\mathbf{B}$ is $|A|^n$, the time complexity of the naive algorithm is exponential. In fact, it turns out that without further restrictions on $\mathbf{A}$, one cannot do better than the naive algorithm; in fact, it follows from the main result of M. Kozik [11] that there exists a finite algebra $\mathbf{A}$ such that the subpower membership problem for $\mathbf{A}$ is EXPTIME-complete.

In contrast, for finite algebras $\mathbf{A}$ in many familiar classes, the subpower membership problem is in P, that is, there is a polynomial time algorithm for solving the problem. For example, if $\mathbf{A}$ is a group, then a variant of Sims' algorithm (cf. [7]) solves the problem in polynomial time. Other simple algorithms work if $\mathbf{A}$ is a finite lattice or a finite lattice with additional operations (use the Baker–Pixley Theorem [1]) or if $\mathbf{A}$ is a finite semilattice. A recent result of A. Bulatov, P. Mayr, and M. Steindl [3] extends this observation on semilattices to any finite commutative

This material is based upon work supported by the Austrian Science Fund (FWF) grant no. P24285, the National Science Foundation grant no. DMS 1500254, the Hungarian National Foundation for Scientific Research (OTKA) grant no. K104251 and K115518, and an NSERC Discovery Grant.

semigroup $\mathbf{A}$ by showing that if $\mathbf{A}$ embeds into a direct product of a Clifford semigroup and a nilpotent semigroup, then the subpower membership problem for $\mathbf{A}$ is in P, and for all other finite commutative semigroups $\mathbf{A}$ the problem is NP-complete.

Extending the result on groups mentioned earlier R. Willard [13] proved that the subpower membership problem is in P for every finite algebra $\mathbf{A}$ that is an expansion of a group by multilinear operations. In particular, this is the case for every finite ring, finite module, and finite K-algebra. It is not known whether this statement remains true if the word 'multilinear' is omitted.

A well-studied class of algebras that includes expanded groups and expanded lattices is the class of algebras which have cube terms.

Question 1. [9] Is the subpower membership problem for $\mathbf{A}$ in P if $\mathbf{A}$ is a finite algebra with a cube term?

The answer to this question is not only of theoretical interest, it features prominently in other problems in computer science; for example, constraint satisfaction problems and problems on learnability (see, e.g., [5, 9]).

Question 1 will remain unresolved in this paper, but we will prove (see Theorem 6.4) is that the answer to Question 1 is YES, provided $\mathbf{A}$ belongs to a residually small variety. The proof relies on a structure theorem proved in [10] for the subalgebras of finite powers of an algebra $\mathbf{A}$ with a cube term (or equivalently, parallelogram term). The application of this structure theorem leads us to considering subalgebras of finite products $\mathbf{S}_1 \times \cdots \times \mathbf{S}_n$ where the factors $\mathbf{S}_1, \ldots, \mathbf{S}_n$ come from the finite collection $\mathbb{HS}\mathbf{A}$ of homomorphic images of subalgebras of $\mathbf{A}$. As a consequence, it is natural for us to expand the scope of the *subpower membership problem*, and define it for any finite set $\mathcal{K}$ of finite algebras as follows:

SMP($\mathcal{K}$):

- INPUT: $a_1, \ldots, a_k, b \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ with $\mathbf{A}_1, \ldots, \mathbf{A}_n \in \mathcal{K}$.
- QUESTION: Is b in the subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by $\{a_1, \ldots, a_k\}$?

There is a more important issue, which is raised by our passage from SMP($\mathbf{A}$) to SMP($\mathbb{HS}\mathbf{A}$): Does the subpower membership problem get harder when $\mathbf{A}$ is replaced by $\mathbb{S}\mathbf{A}$ or $\mathbb{H}\mathbf{A}$? It is easy to see that for any finite algebra $\mathbf{A}$, SMP($\mathbf{A}$) and SMP($\mathbb{S}\mathbf{A}$) are essentially the same problem. However, this is not the case for homomorphic images. A surprising result of M. Steindl [12] shows that there exists a 10-element semigroup $\mathbf{S}$ with a 9-element quotient semigroup $\overline{\mathbf{S}}$ such that SMP($\mathbf{S}$) is in P while SMP($\overline{\mathbf{S}}$) is NP-complete. This shows that the problem SMP($\mathbb{HS}\mathbf{A}$) may be harder that SMP($\mathbf{A}$) (provided $\mathsf{P} \neq \mathsf{NP}$), and therefore poses the following question for us:

Question 2. Are the problems SMP($\mathcal{K}$) and SMP($\mathbb{HS}\mathcal{K}$) polynomial time equivalent if $\mathcal{K}$ is a finite set of finite algebras in a variety with a cube term?

We will prove (see Theorem 4.8) that the answer to Question 2 is YES. The proof uses the techniques of compact representations developed in [2]. We will also show that, given a finite set of finite algebras $\mathcal{K}$ in a variety with a cube term, the problem of finding compact representations for subalgebras of products of algebras in $\mathcal{K}$ is polynomial time equivalent to the subpower membership problem $\mathrm{SMP}(\mathcal{K})$ (Theorem 4.6), and both problems are in NP (Theorem 4.5 and Corollary 4.7).

2. Preliminaries

For every natural number m, we will use the notation $[m]$ for the set $\{1, 2, \dots, m\}$. The collection of all k-element subsets of a set S will be denoted by $\binom{S}{k}$.

Algebras will be denoted by boldface letters, their universes by the same letters in italics. For arbitrary algebras $\mathbf{A}$ and $\mathbf{B}$, we will write $\mathbf{B} \leq \mathbf{A}$ to indicate that $\mathbf{B}$ is a subalgebra of $\mathbf{A}$. For any algebra $\mathbf{A}$, $\mathrm{Con}(\mathbf{A})$ will denote the congruence lattice of $\mathbf{A}$, and the top and bottom elements of $\mathrm{Con}(\mathbf{A})$ are denoted 1 and 0, respectively. We will use the notation $\mathrm{Irr}(\mathbf{A})$ for the set of all meet irreducible congruences of $\mathbf{A}$, excluding 1. Thus, the subdirectly irreducible quotients of $\mathbf{A}$ are exactly the algebras $\mathbf{A}/\sigma$, $\sigma \in \mathrm{Irr}(\mathbf{A})$.

Let ϑ be a congruence of an algebra $\mathbf{A}$. The ϑ-class of an element $a \in A$ is usually denoted by a/ϑ, and we will often write $a \equiv_\vartheta b$ instead of $(a, b) \in \vartheta$. If $\mathbf{B}$ is a subalgebra of $\mathbf{A}$, we will say that $\mathbf{B}$ is *saturated with respect to* ϑ, or $\mathbf{B}$ is a ϑ*-saturated subalgebra* of $\mathbf{A}$, if $b \in B$ and $b \equiv_\vartheta a$ imply $a \in B$ for all $a \in A$. In other words, $\mathbf{B}$ is ϑ-saturated if and only if its universe is a union of ϑ-classes of $\mathbf{A}$. For arbitrary subalgebra $\mathbf{B}$ of $\mathbf{A}$ there exists a smallest ϑ-saturated subalgebra of $\mathbf{A}$ that contains $\mathbf{B}$, which we denote by $\mathbf{B}[\vartheta]$; the universe of $\mathbf{B}[\vartheta]$ is $B[\vartheta] := \bigcup_{b \in B} b/\vartheta$. Denoting the restrictions of ϑ to $\mathbf{B}$ and $\mathbf{B}[\vartheta]$ by $\vartheta_{\mathbf{B}}$ and $\vartheta_{\mathbf{B}[\vartheta]}$, respectively, we get from the second isomorphism theorem that the map $\mathbf{B}/\vartheta_{\mathbf{B}} \to \mathbf{B}[\vartheta]/\vartheta_{\mathbf{B}[\vartheta]}$, $b/\vartheta_{\mathbf{B}} \mapsto b/\vartheta_{\mathbf{B}[\vartheta]} (= b/\vartheta)$ is an isomorphism.

For a product $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ of algebras and for any set $I \subseteq [n]$, the projection homomorphism

$$\mathbf{A}_1 \times \cdots \times \mathbf{A}_n = \prod_{i \in [n]} \mathbf{A}_i \to \prod_{i \in I} \mathbf{A}_i, \qquad (a_i)_{i \in [n]} \mapsto (a_i)_{i \in I}$$

will be denoted by proj_I. For a subalgebra $\mathbf{B}$ or for an element b of $\prod_{i \in [n]} \mathbf{A}_i$, we will write $\mathbf{B}|_I$ or $b|_I$ for $\mathrm{proj}_I(\mathbf{B})$ or $\mathrm{proj}_I(b)$, respectively. If $I = \{j_1, \dots, j_k\}$, then the notation $|_{\{j_1,\dots,j_k\}}$ will be simplified to $|_{j_1,\dots,j_k}$.

2.1. Cube Terms and Parallelogram Terms.

Let $\mathcal{K}$ be a class of algebras in the same language. For any n-ary term t, if M is an $m \times n$ matrix of variables and $\vec{v}$ is an $m \times 1$ matrix of variables,

$$\mathcal{K} \models t(M) = \vec{v} \tag{2.1}$$

will denote that the m identities represented by the rows in (2.1) are true in $\mathcal{K}$. For example,

$$\mathcal{K} \models t \begin{pmatrix} x & x & y \\ y & x & x \end{pmatrix} = \begin{pmatrix} y \\ y \end{pmatrix} \tag{2.2}$$

expresses that t is a Maltsev term for $\mathcal{K}$.

Now let us fix an integer d (> 1). A d-*cube term* for $\mathcal{K}$ is a term t satisfying a set of identities of the form (2.1) in two variables x, y, where M is a matrix with d rows such that every column of M contains at least one x, and $\vec{v}$ consists of y's only. As (2.2) shows, a Mal'tsev term is a 2-cube term.

Cube terms were introduced in [2] to show that a finite algebra $\mathbf{A}$ has few subpowers — i.e., $\mathbf{A}$ the property that, for some polynomial p, the number of subalgebras of $\mathbf{A}^n$ is bounded above by $2^{p(n)}$ — if and only if $\mathbf{A}$ has a cube term. More manageable terms that are equivalent to cube terms (e.g., edge terms, star terms) were also found in [2]. In this paper we will use another family of equivalent terms, called parallelogram terms, which were introduced in [10].

Let m and n be positive integers and let $\mathsf{d} = m + n$. An (m,n)-*parallelogram term* for $\mathcal{K}$ is a $(\mathsf{d}+3)$-ary term $P_{m,n}$ such that

$$\mathcal{K} \models P_{m,n} \left(\begin{array}{ccc|cccccccc} x & x & y & z & y & \cdots & y & y & \cdots & y & y \\ x & x & y & y & z & & y & y & & y & y \\ & \vdots & & \vdots & & \ddots & & & & & \vdots \\ x & x & y & y & y & & z & y & & y & y \\ \hline y & x & x & y & y & & y & z & & y & y \\ & \vdots & & \vdots & & & & & \ddots & & \vdots \\ y & x & x & y & y & & y & y & & z & y \\ y & x & x & y & y & \cdots & y & y & \cdots & y & z \end{array} \right) = \begin{pmatrix} y \\ y \\ \vdots \\ y \\ y \\ \vdots \\ y \\ y \end{pmatrix}. \tag{2.3}$$

Here the rightmost block of variables is a $\mathsf{d} \times \mathsf{d}$ array, the upper left block is $m \times 3$ and the lower left block is $n \times 3$.

It is easy to see from these definitions that an (m,n)-parallelogram term that is independent of its last d variables is a Maltsev term, and an (m,n)-parallelogram term that is independent of its first 3 variables is a d-ary near unanimity term.

The theorem below summarizes the facts we will need later on about cube terms and parallelogram terms.

Theorem 2.1 (See [2],[10]). *Let $\mathcal{V}$ be a variety, and let $\mathsf{d}\,(> 1)$ be an integer.*

(1) *The following conditions are equivalent:*
   (a) *$\mathcal{V}$ has a d-cube term,*
   (b) *$\mathcal{V}$ has an (m,n)-parallelogram term for all $m, n \geq 1$ with $m + n = \mathsf{d}$,*
   (c) *$\mathcal{V}$ has an (m,n)-parallelogram term for some $m, n \geq 1$ with $m + n = \mathsf{d}$.*

(2) *If $\mathcal{V}$ has a cube term, then $\mathcal{V}$ is congruence modular.*

The equivalence of conditions (a)–(c) in statement (1) follows by combining results from [2, Theorem 4.4] and [10, Theorem 3.5]. Proofs for statement (2) can be found in [2, Theorem 2.7] and [4, Theorem 3.2].

In view of the the equivalence of conditions (a) and (c) in Theorem 2.1, when we consider classes of algebras with a d-cube term, we will work with a $(1,\mathsf{d}-1)$-parallelogram term $P=P_{1,\mathsf{d}-1}$, and we will also use the following terms derived from P:

$$\begin{aligned} s(x_1,\dots,x_{\mathsf{d}}) &:= P(x_1,x_2,x_2,x_1,\dots,x_{\mathsf{d}}),\\ p(x,u,y) &:= P(x,u,y,x,y,\dots,y). \end{aligned} \tag{2.4}$$

For any class $\mathcal{K}$ of algebras with a $(1,\mathsf{d}-1)$-parallelogram term P one can easily deduce from the $(1,\mathsf{d}-1)$-parallelogram identities that

$$\begin{aligned} \mathcal{K}\ \models\ y = p(x,x,y),&\\ p(x,y,y) = s(x,y,y,\dots,y),&\\ s(y,x,y,\dots,y) = y,&\\ \vdots\qquad&\\ s(y,y,y,\dots,x) = y.& \end{aligned} \tag{2.5}$$

To simplify notation, we also define

$$x^y := p(x,y,y)\ (= s(x,y,\dots,y)),$$

and

$$s^\ell(x_1,\dots,x_{\mathsf{d}}) := s(s^{\ell-1}(x_1,x_2\dots,x_{\mathsf{d}}),x_2,\dots,x_{\mathsf{d}})\quad \text{for all } \ell\ge 1,$$

where $s^0 := x_1$. So, $s^1 = s$ and s^ℓ is the ℓ-th iterate of s in the first variable.

2.2. **Congruence Modular Varieties: the Commutator and Residual Smallness.** Let $\mathcal{V}$ be an arbitrary congruence modular variety. For the definition and basic properties of the commutator operation $[\ ,\ ]$ on congruence lattices of algebras in $\mathcal{V}$ the reader is referred to [6]. A congruence $\alpha\in\mathrm{Con}(\mathbf{A})$ of an algebra $\mathbf{A}\in\mathcal{V}$ is called *abelian* if $[\alpha,\alpha]=0$, and the *centralizer* of a congruence $\alpha\in\mathrm{Con}(\mathbf{A})$, denoted $(0:\alpha)$, is the largest congruence $\gamma\in\mathrm{Con}(\mathbf{A})$ such that $[\alpha,\gamma]=0$.

Recall that $\mathcal{V}$ has a *difference term* (see [6, Theorem 5.5]), which we will denote by d. In the last two sections of this paper we will need some properties of abelian congruences, which can be summarized informally as follows: the difference term d induces abelian groups on the blocks of all abelian congruences α of all algebras $\mathbf{A}\in\mathcal{V}$; moreover, the term operations of $\mathbf{A}$ are 'linear between the blocks' of α with respect to these abelian groups. The theorem below gives a more precise formulation of these facts.

Theorem 2.2 (From [6, Section 9])**.** *Let $\mathcal{V}$ be a congruence modular variety with a difference term d, let $\mathbf{A}\in\mathcal{V}$, and let α be an abelian congruence of $\mathbf{A}$.*

(1) *For every $o \in A$, the α-class containing o is an abelian group $(o/\alpha; +_o, -_o, o)$ with zero element o for the operations $+_o$ and $-_o$ defined by*

$$x +_o y := d(x, o^{\mathbf{C}}, y) \quad \text{and} \quad -_o x := d(o^{\mathbf{C}}, x, o^{\mathbf{C}}) \quad \text{for all } x, y \in o/\alpha.$$

(2) *For every term $g(x_1, \ldots, x_k)$ in the language of $\mathcal{V}$, for arbitrary elements $o_1, \ldots, o_k, o \in A$ such that $g(o_1, \ldots, o_k) \equiv_\alpha o$, and for any tuple $(a_1, \ldots, a_k) \in (o_1/\alpha) \times \cdots \times (o_k/\alpha)$,*

$$g(a_1, a_2, \ldots, a_k) = g(a_1, o_2, \ldots, o_k) +_o g(o_1, a_2, o_3, \ldots, o_k) +_o \ldots \\ +_o g(o_1, \ldots, o_{k-1}, a_k) -_o (k-1)g(o_1, o_2, \ldots, o_k).$$

We will refer to the abelian groups described in statement (1) as the *induced abelian groups* on the α-classes of $\mathbf{A}$.

In any congruence modular variety $\mathcal{V}$, an equivalence relation, called *similarity*, is defined on the class of subdirectly irreducible algebras in $\mathcal{V}$. The definition may be found in [6, Definition 10.7], but for our purposes here it will be more convenient to use the following characterization given in [6, Theorem 10.8]: two subdirectly irreducible algebras $\mathbf{B}, \mathbf{C} \in \mathcal{V}$ are similar if and only if there exists an algebra $\mathbf{E} \in \mathcal{V}$ (which can be taken to be a subdirect subalgebra of $\mathbf{B} \times \mathbf{C}$) and there exist congruences $\beta, \gamma, \delta, \varepsilon \in \mathrm{Con}(\mathbf{E})$ such that $\mathbf{E}/\beta \cong \mathbf{B}$, $\mathbf{E}/\gamma \cong \mathbf{C}$ and there is a projectivity $\beta^*/\beta \searrow \varepsilon/\delta \nearrow \gamma^*/\gamma$ in $\mathrm{Con}(\mathbf{E})$, where β^* and γ^* are the unique upper covers of β and γ respectively.

For a cardinal c, a variety $\mathcal{V}$ is called *residually less than c* if every subdirectly irreducible algebra in $\mathcal{V}$ has cardinality $< c$; $\mathcal{V}$ is called *residually small* if it is residually less than some cardinal.

Theorem 2.3 (From [6]). *Let $\mathbf{A}$ be a finite algebra that generates a congruence modular variety $\mathcal{V}(\mathbf{A})$. Then the following conditions are equivalent:*

(a) *$\mathcal{V}(\mathbf{A})$ is residually small,*
(b) *$\mathcal{V}(\mathbf{A})$ is residually $< q$ for some natural number q,*
(c) *the congruence identity $[x \wedge y, y] = x \wedge [y, y]$ holds in the congruence lattice of every subalgebra of $\mathbf{A}$,*
(d) *the implication $x \leq [y, y] \to x = [x, y]$ holds in the congruence lattice of every subalgebra of $\mathbf{A}$,*
(e) *for every subdirectly irreducible algebra $\mathbf{S} \in \mathbb{HS}(\mathbf{A})$ with abelian monolith μ, the centralizer $(0 : \mu)$ of μ is an abelian congruence of $\mathbf{S}$.*

The equivalence of conditions (a), (b), and (d) is proved in [6, Theorem 10.15]. The equivalence of (c) and (d) is established in [6, Theorem 8.1]. To show that condition (e) is also equivalent to (c), one can apply the congruence identity in (c) (and basic properties of the commutator) directly to deduce (c) ⇒ (e). Finally, to prove that if (c) fails, then so does (e), one can use the specific failure of (c) produced in the first paragraph of the proof of [6, Theorem 10.14].

Now let $\mathcal{K} = \{\mathbf{A}_1, \dots, \mathbf{A}_n\}$ be a finite set of finite algebras, and let $\mathcal{V}(\mathcal{K})$ be the variety generated by $\mathcal{K}$; then $\mathcal{V}(\mathcal{K})$ is generated by the finite algebra $\mathbf{A}' := \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$, so Theorem 2.3 applies to $\mathcal{V}(\mathcal{K}) = \mathcal{V}(\mathbf{A}')$. Let $\mathcal{H}$ denote the subclass of $\mathcal{V}(\mathcal{K})$ consisting of all algebras $\mathbf{A} \in \mathcal{V}(\mathcal{K})$ satisfying condition (c). It follows from [6, Theorem 8.1] that $\mathcal{H}$ is closed under the formation of quotient algebras, subalgebras, and finite direct products. Therefore, condition (c) from Theorem 2.3 holds for the generator $\mathbf{A}' = \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ of $\mathcal{V}(\mathcal{K})$ if and only if it holds for each one of $\mathbf{A}_1, \dots, \mathbf{A}_n$. Thus, we get the following.

Corollary 2.4. *Let $\mathcal{K}$ be a finite set of finite algebras in a congruence modular variety. The variety generated by $\mathcal{K}$ is residually small if and only if*

(e)′ *for every subdirectly irreducible algebra $\mathbf{S} \in \mathbb{HS}(\mathcal{K})$ with abelian monolith μ, the centralizer $(0 : \mu)$ of μ is an abelian congruence of $\mathbf{S}$.*

3. Compact Representations

Throughout this section $\mathcal{V}$ will be a variety with a d-cube term ($\mathsf{d} > 1$). If $\mathcal{K} \subseteq \mathcal{V}$ is a finite set of finite algebras, $\mathsf{a}_{\mathcal{K}}$ will denote the maximum of the sizes of the algebras in $\mathcal{K}$.

Recall from Theorem 2.1 that if a variety has a d-cube term, it also has a $(1, \mathsf{d}-1)$-parallelogramm term P. Terms in the language of $\mathcal{V}$ which can be expressed using P only, will be referred to as *P-terms*. For example, the terms $s(x_1, \dots, x_{\mathsf{d}})$, $p(x, u, y)$, and x^y constructed earlier, which satisfy (2.5), are P-terms. By a *P-subalgebra* of an algebra $\mathbf{A} \in \mathcal{V}$ we mean a subalgebra of the reduct of $\mathbf{A}$ to the language $\{P\}$. We will say that an algebra $\mathbf{B} \in \mathcal{V}$ is *P-generated* by $R \subseteq B$ if R is a generating set for the reduct of $\mathbf{B}$ to the language $\{P\}$; or equivalently, if every element of $\mathbf{B}$ is of the form $t(r_1, \dots, r_m)$ for some $m \geq 0$, some elements $r_1, \dots .r_m \in R$, and some m-ary P-term t. The P-subalgebra of an algebra $\mathbf{B} \in \mathcal{V}$ generated by a set $S\,(\subseteq B)$ will be denoted by $\langle S \rangle_P$.

Now we will introduce a variant of the concept of 'compact representation' from [2]. One difference is that we will use a less restrictive notion of 'fork' than 'minority fork', because we want to avoid assuming finiteness of the algebras considered unless finiteness is necessary for the conclusions. Another difference is that we will consider subalgebras of products of algebras, rather than subalgebras of powers of a single algebra. Let $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{V}$, and let $B \subseteq A_1 \times \cdots \times A_n$. For $i \in [n]$ and $\gamma, \delta \in A_i^2$ we will say that (γ, δ) is a *fork in the i-th coordinate of* B if there exists $b, b' \in B$ such that

$$b|_{[i-1]} = b'|_{[i-1]} \quad \text{and} \quad b|_i = \gamma, \quad b'|_i = \delta. \tag{3.1}$$

The set of all forks in the i-th coordinate of B will be denoted by $\textsc{fork}_i(B)$. Tuples $b, b' \in B$ satisfying (3.1) will be referred to as *witnesses* for the fork $(\gamma, \delta) \in \textsc{fork}_i(B)$.

For each i and B as above and for every positive integer e, we define

$$\textsc{fork}_i^e(B) := \{(\gamma, \delta^{\gamma^e}) : (\gamma,\delta) \in \textsc{fork}_i(B)\}.$$

The elements of $\textsc{fork}_i^e(B)$ will be called *e-derived forks in the i-th coordinate of B*. In the case when $e = 1$ we will use the notation $\textsc{fork}'_i(B)$ instead of $\textsc{fork}_i^1(B)$, and will call the elements of $\textsc{fork}'_i(B)$ *derived forks in the i-th coordinate of B*. The next lemma shows that derived forks are indeed forks, and they are 'transferable', which does not hold for forks in general.

Lemma 3.1. *Let $\mathbf{A}_1, \dots, \mathbf{A}_n$ be algebras in a variety $\mathcal{V}$ with a d-cube term, and let $\mathbf{B}$ be a P-subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$. Then*

(1) *$\textsc{fork}_i(B) \supseteq \textsc{fork}'_i(B) \supseteq \cdots \supseteq \textsc{fork}_i^e(B) \supseteq \textsc{fork}_i^{e+1}(B) \supseteq \dots$ for all $i \in [n]$ and $e \geq 1$; moreover,*

(2) *for every $(\gamma,\delta) \in \textsc{fork}'_i(B)$ and for every $b \in B$ with $b|_i = \gamma$, there is an element $b' \in B$ such that (3.1) holds, that is, b and b' witness that $(\gamma,\delta) \in \textsc{fork}_i(B)$.*

Proof. Let $(\gamma,\delta) \in \textsc{fork}'_i(B)$. Then there exists $(\gamma,\beta) \in \textsc{fork}_i(B)$ such that $\delta = \beta^\gamma$. Let $c, c' \in B$ be witnesses for $(\gamma,\beta) \in \textsc{fork}_i(B)$; thus, $c|_{[i-1]} = c'|_{[i-1]}$ and $c|_i = \gamma$, $c'|_i = \beta$. It follows from the identities in (2.5) that for the element $b' := p(c', c, b) \in B$ we have

$$b'|_{[i-1]} = p(c'|_{[i-1]}, c|_{[i-1]}, b|_{[i-1]}) = b|_{[i-1]},$$

and

$$b'|_i = p(c'|_i, c|_i, b|_i) = p(\beta, \gamma, \gamma) = \beta^\gamma = \delta.$$

This proves (2), and also the inclusion $\textsc{fork}_i(B) \supseteq \textsc{fork}'_i(B)$ in (1). The inclusion $\textsc{fork}_i^e(B) \supseteq \textsc{fork}_i^{e+1}(B)$ for any $e \geq 1$ follows by the same argument, using $\delta = \beta^{\gamma^{e+1}} = (\beta^{\gamma^e})^\gamma$ and β^{γ^e} in place of $\delta = \beta^\gamma$ and β. □

Definition 3.2. For two sets $B, R \subseteq A_1 \times \cdots \times A_n$, we will say that R is a (d, e)-*representation for B* if the following three conditions are met:

(i) $R \subseteq B$;

(ii) $R|_I = B|_I$ for all $I \subseteq [n]$ with $|I| < \mathsf{d}$;

(iii) $\textsc{fork}_i(R) \supseteq \textsc{fork}_i^e(B)$ for all $i \in [n]$.

If $e = 1$ and the parameter d of the cube term of $\mathcal{V}$ is clear from the context, then reference to (d, e) will be omitted.

If the algebras $\mathbf{A}_1, \dots, \mathbf{A}_n$ all belong to a fixed finite set $\mathcal{K}$ of finite algebras in $\mathcal{V}$, then it is easy to see that every set $B \subseteq A_1 \times \cdots \times A_n$ has a (d, e)-representation R of size

$$|R| \leq \sum_{I \subseteq [n], |I| < k} \big|B|_I\big| + \sum_{i \in [n]} \textsc{fork}_i^e(B) \leq \binom{n}{\mathsf{d}-1} \mathsf{a}_{\mathcal{K}}^{\mathsf{d}-1} + 2n\mathsf{a}_{\mathcal{K}}^2.$$

A (d, e)-representation R for B of size $|R| \leq \binom{n}{\mathsf{d}-1}\mathsf{a}_{\mathcal{K}}^{\mathsf{d}-1} + 2n\mathsf{a}_{\mathcal{K}}^2$ is called a *compact* (d, e)*-representation* for B.

It was proved in [2] that for any subalgebra $\mathbf{B}$ of a finite power of a finite algebra $\mathbf{A} \in \mathcal{V}$, a compact representation (with minority forks) generates $\mathbf{B}$. The theorem below is essentially the same result, with one significant difference: we organize the proof so that we can get a useful upper bound on the length of computations needed to generate elements of $\mathbf{B}$ from elements of a compact representation for $\mathbf{B}$.

Theorem 3.3. *Let $\mathcal{V}$ be a variety with a d-cube term, let P be a $(1, \mathsf{d}-1)$-parallelogram term for $\mathcal{V}$, and let e be a positive integer. If $\mathbf{B}$ is a subalgebra of a product $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ of finitely many algebras $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{V}$, then $\mathbf{B}$ is P-generated by every (d, e)-representation $R \subseteq B$ for $\mathbf{B}$. In fact, every element of $\mathbf{B}$ can be obtained from elements of R by at most $3\binom{n+1}{\mathsf{d}+1}$ applications of P, p, or s^{e+1} (to elements of R or to elements obtained earlier in the process).*

First we will prove the following lemma.

Lemma 3.4. *Let $\mathcal{V}$ be a variety with a d-cube term, let P be a $(1, \mathsf{d}-1)$-parallelogram term for $\mathcal{V}$, and let e be a positive integer. For every positive integer $n \geq \mathsf{d}$ there exists a P-term $t_n = t_n\big(x, y, z, \overline{w_I}\big)$ with $\overline{w_I} := (w_I)_{I \in \binom{[n]}{\mathsf{d}-1}}$ such that t_n has the following properties.*

(i) *For every subset R of a product $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ with $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{V}$, and for every tuple $b = (b_1, \dots, b_{n-1}, \gamma)$ in $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$, if*

(a) *for each $I \in \binom{[n]}{\mathsf{d}-1}$ the set R contains a tuple b^I satisfying $b^I|_I = b|_I$, and*

(b) *for some element $b' = (b_1, \dots, b_{n-1}, \beta)$ of the P-subalgebra $\mathbf{R}^*$ of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by R, the set R contains tuples $u = (u_1, \dots, u_{n-1}, \gamma)$ and $\widehat{u} = (u_1, \dots, u_{n-1}, \beta^{\gamma^e})$ which are witnesses for the fork $(\gamma, \beta^{\gamma^e}) \in \mathrm{FORK}_n(R)$,*

then

$$b = t_n\big(b', \widehat{u}, u, \overline{b^I}\big) \quad \textit{where} \quad \overline{b^I} := (b^I)_{I \in \binom{[n]}{\mathsf{d}-1}}, \tag{3.2}$$

and therefore b is in $\mathbf{R}^*$.

(ii) *The right hand side of (3.2) can be computed from the input tuples $b' \in R^*$ and $\widehat{u}, u, b^I \in R$ $\big(I \in \binom{[n]}{\mathsf{d}-1}\big)$ in at most $3\binom{n}{\mathsf{d}}$ steps, where one step is a single application of P, p, or s^{e+1} to input tuples or tuples computed earlier.*

Proof. We start the proof by constructing a family of terms t_V $(V \subseteq [n])$ which 'approximate' t_n in the sense that for t_V in place of t_n the equality (3.2) holds in all coordinates in V (but may fail in other coordinates).

Claim 3.5. *For $\ell = n, n-1, \dots, \mathsf{d}-1$ and for every set $V = V' \cup ([n] \setminus [\ell])$ with $V' \in \binom{[\ell]}{\mathsf{d}-1}$ there exists a P-term $t_V = t_V(x, y, z, \overline{w_I}^V)$ with $\overline{w_I}^V := (w_I)_{I \in \binom{V}{\mathsf{d}-1}}$ such that*

(i) *t_V is independent of the choice of $\mathbf{A}_1, \dots, \mathbf{A}_n$, $\mathbf{R}^*$, and also of the choice of the elements b, b^I $\big(I \in \binom{[n]}{\mathsf{d}-1}\big)$, b', u, $\widehat{u}$;*

(ii) *for the element $b^V := t_V\big(b', \widehat{u}, u, \overline{b^I}^V\big)$, where $\overline{b^I}^V := (b^I)_{I \in \binom{V}{\mathsf{d}-1}}$, we have that $b^V|_V = b|_V$.*

Proof of Claim 3.5. We proceed by induction on $n-\ell$. For $n=\ell$ we have $V \in \binom{[n]}{\mathsf{d}-1}$, so we can choose $t_V = w_V$, and our claims (i)–(ii) are trivial.

Assume now that $\ell < n$ and that our claims are true for $\ell+1$. Let $V = V' \cup ([n] \setminus [\ell])$ with $V' = \{i_1, \dots, i_{\mathsf{d}-1}\} \in \binom{[\ell]}{\mathsf{d}-1}$, $i_1 < \dots < i_{\mathsf{d}-1}$. Let

$$V_j := V \setminus \{i_j\} = \{i_1, \dots, i_{j-1}, i_{j+1}, \dots, i_{\mathsf{d}-1}, \ell+1\} \cup ([n] \setminus [\ell+1]) \quad (j = 1, \dots, \mathsf{d}-1).$$

We will prove that the term

$$\begin{aligned} t_V = t_V\big(x, y, z, \overline{w_I}^V\big) \\ := P\Big(& s^{e+1}\big(x, (t_{V_j}(x, y, z, \overline{w_I}^{V_j}))_{j \in [\mathsf{d}-1]}\big), p\big(y, z, t_{V_1}(x, y, z, \overline{w_I}^{V_1})\big), \\ & t_{V_1}(x, y, z, \overline{w_I}^{V_1}), x, \big(t_{V_j}(x, y, z, x, \overline{w_I}^{V_j})\big)_{j \in [\mathsf{d}-1]}\Big) \end{aligned}$$

has the desired properties. It is clear that (i) holds for t_V. To establish (ii), notice that by the induction hypothesis, the elements $b^{V_j} := t_{V_j}\big(b', \widehat{u}, u, \overline{b^I}^{V_j}\big)$ satisfy the condition $b^{V_j}|_{V_j} = b|_{V_j}$ for all $j \in [\mathsf{d}-1]$; that is, $b^{V_j}|_V$ has the form

$$b^{V_j}|_V = (b_{i_1}, \dots, b_{i_{j-1}}, \zeta_j, b_{i_{j+1}}, \dots, b_{i_{\mathsf{d}-1}}, b_{\ell+1}, \dots, b_{n-1}, \gamma)$$

for some element $\zeta_j \in A_{i_j}$ in the j-th coordinate. To compute $b^V|_V = t_V(b', \widehat{u}, u, \overline{b^I}^V)|_V$, let's start with evaluating the first two arguments of P on the right hand side. Using the identities for s in (2.5) we get that

$$\begin{aligned} s\big(b', \big(t_{V_j}(b', \widehat{u}, u, \overline{b^I}^{V_j})\big)_{j \in [\mathsf{d}-1]}\big)\big|_V &= s\big(b'|_V, \big(t_{V_j}(b', \widehat{u}, u, \overline{b^I}^{V_j})|_V\big)_{j \in [\mathsf{d}-1]}\big) \\ &= s(b'|_V, b^{V_1}|_V, \dots, b^{V_{\mathsf{d}-1}}|_V) = s\begin{pmatrix} b_{i_1} & \zeta_1 & b_{i_1} & \dots & b_{i_1} \\ b_{i_2} & b_{i_2} & \zeta_2 & \dots & b_{i_2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ b_{i_{\mathsf{d}-1}} & b_{i_{\mathsf{d}-1}} & b_{i_{\mathsf{d}-1}} & \dots & \zeta_{\mathsf{d}-1} \\ b_{\ell+1} & b_{\ell+1} & b_{\ell+1} & \dots & b_{\ell+1} \\ \vdots & \vdots & \vdots & & \vdots \\ b_{n-1} & b_{n-1} & b_{n-1} & \dots & b_{n-1} \\ \beta & \gamma & \gamma & \dots & \gamma \end{pmatrix} = \begin{pmatrix} b_{i_1} \\ b_{i_2} \\ \vdots \\ b_{i_{\mathsf{d}-1}} \\ b_{\ell+1} \\ \vdots \\ b_{n-1} \\ \beta^\gamma \end{pmatrix}, \end{aligned}$$

so by repeating the same computation e more times so that every time the tuple just obtained is placed in the first argument of s in the next computation, we obtain that (3.3)

$$s^{e+1}\big(b', \big(t_{V_j}(b', \widehat{u}, u, \overline{b^{I}}^{V_j})\big)_{j\in[\mathsf{d}-1]}\big)\big|_V = s^{e+1}(b'|_V, b^{V_1}|_V, \ldots, b^{V_{\mathsf{d}-1}}|_V) = \begin{pmatrix} b_{i_1} \\ b_{i_2} \\ \vdots \\ b_{i_{\mathsf{d}-1}} \\ b_{\ell+1} \\ \vdots \\ b_{n-1} \\ \beta^{\gamma^{e+1}} \end{pmatrix}.$$

The identities for p in (2.5) yield that

$$\begin{aligned} (3.4)\quad p\big(\widehat{u}, u, t_{V_1}(b', \widehat{u}, u, \overline{b^{I}}^{V_1})\big)|_V &= p\big(\widehat{u}|_V, u|_V, t_{V_1}(b', \widehat{u}, u, \overline{b^{I}}^{V_1})|_V\big) \\ &= p\big(\widehat{u}|_V, u|_V, b^{V_1}|_V\big) = p\begin{pmatrix} u_{i_1} & u_{i_1} & \zeta_1 \\ u_{i_2} & u_{i_2} & b_{i_2} \\ \vdots & \vdots & \vdots \\ u_{i_{\mathsf{d}-1}} & u_{i_{\mathsf{d}-1}} & b_{i_{\mathsf{d}-1}} \\ u_{\ell+1} & u_{\ell+1} & b_{\ell+1} \\ \vdots & \vdots & \vdots \\ u_{n-1} & u_{n-1} & b_{n-1} \\ \beta^{\gamma^{e}} & \gamma & \gamma \end{pmatrix} = \begin{pmatrix} \zeta_1 \\ b_{i_2} \\ \vdots \\ b_{i_{\mathsf{d}-1}} \\ b_{\ell+1} \\ \vdots \\ b_{n-1} \\ \beta^{\gamma^{e+1}} \end{pmatrix}. \end{aligned}$$

Combining these results with the definition of t_V, and using the $(1, \mathsf{d}-1)$-parallelogram identities, we obtain that

$$\begin{aligned} (3.5)\quad b^V|_V &= t_V(b', \widehat{u}, u, \overline{b^{I}}^{V})|_V \\ &= P\begin{pmatrix} b_{i_1} & \zeta_1 & \zeta_1 & b_{i_1} & \zeta_1 & b_{i_1} & \ldots & b_{i_1} \\ b_{i_2} & b_{i_2} & b_{i_2} & b_{i_2} & b_{i_2} & \zeta_2 & \ldots & b_{i_2} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ b_{i_{\mathsf{d}-1}} & b_{i_{\mathsf{d}-1}} & b_{i_{\mathsf{d}-1}} & b_{i_{\mathsf{d}-1}} & b_{i_{\mathsf{d}-1}} & b_{i_{\mathsf{d}-1}} & \ldots & \zeta_{\mathsf{d}-1} \\ b_{\ell+1} & b_{\ell+1} & b_{\ell+1} & b_{\ell+1} & b_{\ell+1} & b_{\ell+1} & \ldots & b_{\ell+1} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & & \vdots \\ b_{n-1} & b_{n-1} & u_{n-1} & b_{n-1} & b_{n-1} & b_{n-1} & \ldots & b_{n-1} \\ \beta^{\gamma^{e+1}} & \beta^{\gamma^{e+1}} & \gamma & \beta & \gamma & \gamma & \ldots & \gamma \end{pmatrix} = \begin{pmatrix} b_{i_1} \\ b_{i_2} \\ \vdots \\ b_{i_{\mathsf{d}-1}} \\ b_{\ell+1} \\ \vdots \\ b_{n-1} \\ \gamma \end{pmatrix} = b|_V. \end{aligned}$$

This completes the proof of (ii), and thereby the proof of Claim 3.5. ◇

The term $t_n := t_{[n]}$ constructed in Claim 3.5 for

$$V = [n] = \{1, \dots, \mathsf{d} - 1\} \cup ([n] \setminus [\mathsf{d} - 1])$$

clearly has the property claimed in statement (i) of Lemma 3.4.

To prove statement (ii), we will follow the calculations done in the proof of Claim 3.5. To evaluate $t_n(b', \widehat{u}, u, \overline{b^I})$, we need to compute each tuple $b^V = t_V(b', \widehat{u}, u, \overline{b^I}^V)$ (with V as in Claim 3.5) only once. No computation is needed for $|V| = \mathsf{d} - 1$. If $|V| \geq \mathsf{d}$, then $V = V' \cup ([n] \setminus [\ell])$ for some set $V' \in \binom{[\ell]}{\mathsf{d}-1}$ and for some $\ell \in \{n-1, n-2, \dots, \mathsf{d}-1\}$. Hence, the number of V's — that is, the number of (new) tuples b^V computed — is $\sum_{\ell=\mathsf{d}-1}^{n-1} \binom{\ell}{\mathsf{d}-1} = \binom{n}{\mathsf{d}}$. To compute each b^V ($|V| \geq \mathsf{d}$) from b', $\widehat{u}$, u, and either from b^I's $\big(I \in \binom{[n]}{\mathsf{d}-1}\big)$ (if $|V| = \mathsf{d}$) or from $b^{\widehat{V}}$'s with $\widehat{V} \subsetneq V$ computed earlier (if $|V| > \mathsf{d}$), we need one application of each of s^{e+1}, p, and P, and two auxiliary tuples are computed in the process. Thus, altogether, at most $3\binom{n}{\mathsf{d}}$ tuples in $\mathbf{R}^*$ need to be computed to evaluate $t_n(b', \widehat{u}, u, \overline{b^I})$, each one requiring one application of P, p, or s^{e+1} to tuples given or computed earlier. □

Now let $n \geq \mathsf{d}$, and let $t_n = t_n(x, y, z, \overline{w_I})$ be the P-term constructed in the proof of Lemma 3.4, where $\overline{w_I}$ is a tuple of variables $(w_I)_{I \in \binom{[n]}{\mathsf{d}-1}}$ indexed by all $(\mathsf{d}-1)$-element subsets of $[n]$. The analogous P-terms for $\mathsf{d} \leq m \leq n$ are $t_m = t_m(x, y, z, \overline{w_I}^{[m]})$ with $\overline{w_I}^{[m]} = (w_I)_{I \in \binom{[m]}{\mathsf{d}-1}}$. We will use these terms to define new P-terms

$$T_m(z^{(\mathsf{d})}, \widehat{z}^{(\mathsf{d})}, \dots, z^{(m)}, \widehat{z}^{(m)}, \overline{w_I}^{[m]})$$

for each $m = \mathsf{d} - 1, \mathsf{d}, \dots, n$ by recursion as follows: $T_{\mathsf{d}-1}(w_{[\mathsf{d}-1]}) := w_{[\mathsf{d}-1]}$, and for all m with $\mathsf{d} \leq m \leq n$,

$$\begin{aligned} &T_m(z^{(\mathsf{d})}, \widehat{z}^{(\mathsf{d})}, \dots, z^{(m)}, \widehat{z}^{(m)}, \overline{w_I}^{[m]}) \\ &\qquad := t_m\big(T_{m-1}(z^{(\mathsf{d})}, \widehat{z}^{(\mathsf{d})}, \dots, z^{(m-1)}, \widehat{z}^{(m-1)}, \overline{w_I}^{[m-1]}), z^{(m)}, \widehat{z}^{(m)}, \overline{w_I}^{[m]}\big). \end{aligned}$$

In particular, for $m = n$, the term is $T_n(z^{(\mathsf{d})}, \widehat{z}^{(\mathsf{d})}, \dots, z^{(n)}, \widehat{z}^{(n)}, \overline{w_I})$, because $\overline{w_I}^{[n]} = \overline{w_I}$.

Lemma 3.6. *Let $\mathcal{V}$ be a variety with a d-cube term, let P be a $(1, \mathsf{d}-1)$-parallelogram term for $\mathcal{V}$, and let e be a positive integer.*

(i) *For every element b and every subset R of a product $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ with $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{V}$ ($n \geq \mathsf{d}$), if*

(a) *for each $I \in \binom{[n]}{\mathsf{d}-1}$ the set R contains a tuple b^I satisfying $b^I|_I = b|_I$, and*

(b) *for every m with $\mathsf{d} \leq m \leq n$ the set R contains tuples $u^{(m)}, \widehat{u}^{(m)}$ which are witnesses for the fork $(\gamma, \beta^{\gamma^e}) \in \textsc{fork}_m(R)$ where*

$$\gamma = b|_m \quad \textit{and} \quad \beta = T_{m-1}(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m-1)}, \widehat{u}^{(m-1)}, \overline{b^I}^{[m-1]})|_m, \tag{3.6}$$

then the following equalities hold:

$$b|_{[m]} = T_m(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m)}, \widehat{u}^{(m)}, \overline{b^I}^{[m]})|_{[m]} \quad \text{for all } m = \mathsf{d}-1, \mathsf{d}, \dots, n; \tag{3.7}$$

in particular,

$$b = T_n(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(n)}, \widehat{u}^{(n)}, \overline{b^I}). \tag{3.8}$$

(ii) *The right hand side of* (3.8) *can be computed from the input tuples* $u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(n)}, \widehat{u}^{(n)}$ *and* b^I $\big(I \in \binom{[n]}{\mathsf{d}-1}\big)$ *in* R *in at most* $3\binom{n+1}{\mathsf{d}+1}$ *steps, where one step is a single application of* P, p, *or* s^{e+1} *to input tuples or tuples computed earlier.*

Proof. Let $\mathbf{R}^*$ denote the P-subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by R, and assume that conditions (a)–(b) hold for b and R. For $m = \mathsf{d}-1, \mathsf{d}, \dots, n$ let

$$b^{(m)} := T_m(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m)}, \widehat{u}^{(m)}, \overline{b^I}^{[m]}).$$

We will proceed by induction to prove that (3.7) holds, that is, $b|_{[m]} = b^{(m)}|_{[m]}$ for all $m = \mathsf{d}-1, \mathsf{d}, \dots, n$. Then, for the case when $m = n$, (3.7) yields the equality (3.8).

To start the induction, let $m = \mathsf{d}-1$. Since $T_{\mathsf{d}-1}(w_{[\mathsf{d}-1]}) := w_{[\mathsf{d}-1]}$, we have that $b^{(\mathsf{d}-1)} = b^{[\mathsf{d}-1]}$, so $b|_{[\mathsf{d}-1]} = b^{(\mathsf{d}-1)}|_{[\mathsf{d}-1]}$ is clearly true by assumption (a).

Now assume that $m \geq \mathsf{d}$ and that the equality $b|_{[m-1]} = b^{(m-1)}|_{[m-1]}$ holds. Then the tuples $b|_{[m]} \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_m$ and $b^{(m-1)}|_{[m]} \in \mathbf{R}^*|_{[m]}$ have the form $(b_1, \dots, b_{m-1}, \gamma)$ and $(b_1, \dots, b_{m-1}, \beta)$, respectively, where γ and β are defined by (3.6). Assumption (b) implies that R contains tuples $u^{(m)}, \widehat{u}^{(m)} \in R$ which are witnesses for the fork $(\gamma, \beta^{\gamma^e}) \in \textsc{fork}^e_m(R)$. Then $u^{(m)}|_{[m]}$ and $\widehat{u}^{(m)}|_{[m]}$ are in $R|_{[m]}$, and they have the form $u^{(m)}|_{[m]} = (u_1, \dots, u_{m-1}, \gamma)$ and $\widehat{u}^{(m)}|_{[m]} = (u_1, \dots, u_{m-1}, \beta^{\gamma^e})$ for some $u_i \in \mathbf{A}_i$ ($i \in [m-1]$). The fact that R satisfies assumption (a) also implies that the elements $b^I|_{[m]} \in R|_{[m]}$ have the property $(b^I|_{[m]})|_I = (b|_{[m]})|_I$ for all $I \in \binom{[m]}{\mathsf{d}-1}$.

This shows that the assumptions of statement (i) in Lemma 3.4 hold for

- ◇ the subset $R|_{[m]}$ and element $b|_{[m]} = (b_1, \dots, b_{m-1}, \gamma)$ of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_m$,
- ◇ the element $b^{(m)}|_{[m]} = (b_1, \dots, b_{m-1}, \beta)$ in $\mathbf{R}^*|_{[m]}$, and
- ◇ the elements $u^{(m)}|_{[m]} = (u_1, \dots, u_{m-1}, \gamma)$, $\widehat{u}^{(m)}|_{[m]} = (u_1, \dots, u_{m-1}, \beta^{\gamma^e})$, and $b^I|_{[m]}$ $\big(I \in \binom{[m]}{\mathsf{d}-1}\big)$ in $R|_{[m]}$.

Thus, Lemma 3.4 (i) — combined with the definitions of $b^{(m-1)}$, T_m, and $b^{(m)}$ — implies that

$$\begin{aligned}
b|_{[m]} &= t_m(b^{(m-1)}|_{[m]}, \widehat{u}^{(m)}|_{[m]}, u^{(m)}|_{[m]}, \overline{b^I|_{[m]}}^{[m]}) \\
&= t_m(b^{(m-1)}, \widehat{u}^{(m)}, u^{(m)}, \overline{b^I}^{[m]})|_{[m]} \\
&= t_m\big(T_{m-1}(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m-1)}, \widehat{u}^{(m-1)}, \overline{b^I}^{[m-1]}), \widehat{u}^{(m)}, u^{(m)}, \overline{b^I}^{[m]}\big)|_{[m]} \\
&= T_m(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m)}, \widehat{u}^{(m)}, \overline{b^I}^{[m]})|_{[m]} \\
&= b^{(m)}|_{[m]}.
\end{aligned}$$

This completes the proof of statement (i).

To prove statement (ii), we follow the steps of the induction in the preceding paragraphs. At the start, getting $b^{(\mathsf{d}-1)}$ requires no computation. By Lemma 3.4 (ii), for each m ($\mathsf{d} \le m \le m$) the tuple $t_m(b^{(m-1)}|_{[m]}, \widehat{u}^{(m)}|_{[m]}, u^{(m)}|_{[m]}, \overline{b^I|_{[m]}}^{[m]})$ can be computed from the input tuples $b^{(m-1)}|_{[m]}$, $\widehat{u}^{(m)}|_{[m]}$, $u^{(m)}|_{[m]}$, $b^I|_{[m]}$ $\big(I \in \binom{[m]}{\mathsf{d}-1}\big)$ in at most $3\binom{m}{\mathsf{d}}$ steps where one step is a single application of P, p, or s^{e+1} to input tuples or tuples computed earlier in the process. The same computation applied to the full tuples $b^{(m-1)}$, $\widehat{u}^{(m)}$, $u^{(m)}$, and b^I $\big(I \in \binom{[m]}{\mathsf{d}-1}\big)$ yields $b^{(m)}$. Therefore it takes at most $3\binom{m}{\mathsf{d}}$ steps to compute $b^{(m)}$ from $b^{(m-1)}$ and the elements $\widehat{u}^{(m)}$, $u^{(m)}$, and b^I $\big(I \in \binom{[m]}{\mathsf{d}-1}\big)$ of R. Hence the right hand side of (3.8) can be computed from the input tuples $u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(n)}, \widehat{u}^{(n)}$ and b^I $(I \in \binom{[n]}{\mathsf{d}-1})$ in R in at most

$$\sum_{m=\mathsf{d}}^{n} 3\binom{m}{\mathsf{d}} = 3\binom{n+1}{\mathsf{d}+1}$$

steps. This proves statement (ii). □

Proof of Theorem 3.3. Let R be a (d, e)-representation for $\mathbf{B}$, and let $\mathbf{R}^*$ denote the P-subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by R. Since $R \subseteq B$, we have that $\mathbf{R}^*$ is a P-subalgebra of $\mathbf{B}$. Let $b \in B$. The desired conclusions that b belongs to $\mathbf{R}^*$ and b can be obtained from elements of R in at most $3\binom{n+1}{\mathsf{d}+1}$ steps will follow from Lemma 3.6 if we show that conditions (a)–(b) hold for b and R. Condition (a) clearly follows from our assumptions that $b \in B$ and R is a (d, e)-representation for $\mathbf{B}$. To verify condition (b) we proceed by induction on m to show that for every m ($\mathsf{d} \le m \le n$)

(b)$_m$ R contains tuples $u^{(m)}, \widehat{u}^{(m)}$ witnessing the fork $(\gamma, \beta^{\gamma^e}) \in \textsc{fork}_m(R)$ where γ and β are defined by (3.6).

Let $\mathsf{d} \le m \le n$, and assume that condition (b)$_i$ holds for $i = \mathsf{d}, \dots, m-1$; note that this assumption is vacuously true for the base case $m = \mathsf{d}$. Our goal is to show that (b)$_m$ also holds. Let

$$b^{(m-1)} := T_{m-1}(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m-1)}, \widehat{u}^{(m-1)}, \overline{b^I}^{[m-1]}),$$

and let γ and β be defined by (3.6); that is, $\gamma = b|_m$ and $\beta = b^{(m-1)}|_m$. Since $b^{(m-1)}$ involves only the elements b^I $\bigl(I \in \binom{[m-1]}{\mathsf{d}-1}\bigr)$ and $u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m-1)}, \widehat{u}^{(m-1)}$ of R, and since our induction hypothesis ensures that these elements satisfy conditions (a)–(b) in Lemma 3.6, we get from Lemma 3.6 that $b|_{[m-1]} = b^{(m-1)}|_{[m-1]}$. Here $b \in B$ and $b^{(m-1)} \in R^* \subseteq B$, therefore b and $b^{(m-1)}$ are witnesses in $\mathbf{B}$ for the fork $(\gamma, \beta) \in \textsc{fork}_m(B)$. Hence, by Lemma 3.1, $(\gamma, \beta^{\gamma^e}) \in \textsc{fork}^e_m(B)$. So, the fact that R is a (d, e)-representation for $\mathbf{B}$ implies that $(\gamma, \beta^{\gamma^e}) \in \textsc{fork}^e_m(R)$. Thus, (b)$_m$ holds, as we wanted to show. This proves that R satisfies condition (b), which completes the proof of Theorem 3.3. □

Lemma 3.7. *Let $\mathcal{V}$ be a variety with a d-cube term, let P be a $(1, \mathsf{d}-1)$-parallelogram term for $\mathcal{V}$, and let p be the ternary P-term defined in* (2.4). *Furthermore, let $\mathbf{R}^*$ be a P-subalgebra of $\mathbf{A}_1 \times \dots \times \mathbf{A}_n$ with $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{V}$. If (γ, δ) and (β, δ) are forks in $\textsc{fork}_m(R^*)$ witnessed by the pairs (v, v') and (u, u') in $\mathbf{R}^*$, respectively, then the pair*

$$\Bigl(p(p(v, v', u'), p(v, v', v'), v), \quad p(u, v, v)\Bigr) \tag{3.9}$$

in $\mathbf{R}^*$ is a witness for the fork $(\gamma, \beta^\gamma) \in \textsc{fork}_m(R^*)$.

Proof. The choice of u, u', v, v' implies that $u|_{[m-1]} = u'|_{[m-1]}$, $v|_{[m-1]} = v'|_{[m-1]}$, and $u|_m = \beta$, $u'|_m = \delta = v'|_m$, $v|_m = \gamma$. Hence,

$$\begin{aligned} &p(p(v, v', u'), p(v, v', v'), v)|_{[m-1]} \\ &\qquad = p(p(v|_{[m-1]}, v'|_{[m-1]}, u'|_{[m-1]}), p(v|_{[m-1]}, v'|_{[m-1]}, v'|_{[m-1]}), v|_{[m-1]}) \\ &\qquad\qquad = p(u|_{[m-1]}, v|_{[m-1]}, v|_{[m-1]}) = p(u, v, v)|_{[m-1]} \end{aligned}$$

and

$$\begin{aligned} p(p(v, v', u'), p(v, v', v'), v)|_m &= p(p(\gamma, \delta, \delta), p(\gamma, \delta, \delta), \gamma)|_m = \gamma, \\ p(u, v, v)|_m &= p(\beta, \gamma, \gamma) = \beta^\gamma. \end{aligned}$$

Clearly, $u, u', v, v' \in \mathbf{R}^*$ implies that the pair (3.9) also lies in $\mathbf{R}^*$, so the proof of the lemma is complete. □

We close this section by discussing the following question: given a generating set for a subdirect subalgebra $\mathbf{B}$ of a product $\mathbf{B}_1 \times \dots \times \mathbf{B}_n$ (in a variety with a cube term) and a product congruence θ of $\mathbf{B}_1 \times \dots \times \mathbf{B}_n$, how can one construct a generating set for the θ-saturation $\mathbf{B}[\theta]$ of $\mathbf{B}$? (For the definition θ-saturation, see Section 2.)

Theorem 3.8. *Let $\mathcal{V}$ be a variety with a d-cube term, and let $\mathbf{B}$ be a subalgebra of $\mathbf{B}_1 \times \cdots \times \mathbf{B}_n$ with $\mathbf{B}_1, \dots, \mathbf{B}_n \in \mathcal{V}$ and $n \geq \mathsf{d}$. Let $\theta_i \in \operatorname{Con}(\mathbf{B}_i)$ for all $i \in [n]$, let $\theta := \theta_1 \times \cdots \times \theta_n \in \operatorname{Con}(\mathbf{B}_1 \times \cdots \times \mathbf{B}_n)$, and let $\theta_{\mathbf{B}}$ be the restriction of θ to $\mathbf{B}$.*

(1) *For any element $b \in \mathbf{B}_1 \times \cdots \times \mathbf{B}_n$ we have that*

$$b/\theta \textit{ is in } \mathbf{B}/\theta_{\mathbf{B}} \ (\leq \mathbf{B}_1/\theta_1 \times \cdots \times \mathbf{B}_n/\theta_n) \quad \iff \quad b \textit{ is in } \mathbf{B}[\theta].$$

(2) *If G is a generating set for $\mathbf{B}$, then for any sets $L, F \subseteq B[\theta]$ satisfying conditions* (L) *and* (F) *below, $G \cup L \cup F$ is a generating set for $\mathbf{B}[\theta]$.*

(L) *For every $I \in \binom{[n]}{\mathsf{d}-1}$ and for every $\overline{d} \in B[\theta]|_I$ there exist $r_2[I,\overline{d}] \in L \cap B[\theta]$ and $r_1[I,\overline{b}] \in B$ such that*

$$r_2[I,\overline{d}]|_I = \overline{d}, \quad r_1[I,\overline{b}]|_I = \overline{b}, \quad \textit{and} \quad r_2[I,\overline{d}] \equiv_\theta r_1[I,\overline{b}].$$

(F) *For every $i \in [n]$ and $(\beta,\gamma) \in \theta_i$ with $\beta \in B|_i$, F contains elements $f_1[i,\beta,\gamma] \in B$ and $f_2[i,\beta,\gamma] \in B[\theta]$ such that*

$$f_1[i,\beta,\gamma]|_i = \beta, \quad f_2[i,\beta,\gamma]|_i = \gamma, \quad \textit{and} \quad f_1[i,\beta,\gamma]|_{[n]\setminus\{i\}} = f_2[i,\beta,\gamma]|_{[n]\setminus\{i\}}.$$

Proof. Statement (1) is an immediate consequence of the fact that $\mathbf{B}[\theta]$ is the full inverse image of $\mathbf{B}/\theta_{\mathbf{B}}$ under the natural homomorphism

$$\mathbf{B}_1 \times \cdots \times \mathbf{B}_n \twoheadrightarrow (\mathbf{B}_1 \times \cdots \times \mathbf{B}_n)/\theta \cong \mathbf{B}_1/\theta_1 \times \cdots \times \mathbf{B}_n/\theta_n.$$

To prove statement (2) let $L, F \subseteq B[\theta]$ satisfy conditions (L) and (F). Let $\mathbf{C}$ denote the subalgebra of $\mathbf{B}_1 \times \cdots \times \mathbf{B}_n$ generated by $G \cup L \cup F$. Clearly, $\mathbf{B} \leq \mathbf{C} \leq \mathbf{B}[\theta]$. To show that $\mathbf{C} = \mathbf{B}[\theta]$, choose an arbitrary element $d = (d_1, \dots, d_n)$ in $\mathbf{B}[\theta]$. Then there exists $b = (b_1, \dots, b_n)$ in $\mathbf{B}$ such that $b \equiv_\theta d$. Using that L satisfies condition (L), choose and fix elements $b^I \in B$ and $d^I \in L \cap B[\theta]$ such that

$$b^I|_I = b|_I, \quad d^I|_I = d|_I, \quad \text{and} \quad b^I \equiv_\theta d^I \quad \text{for each} \quad I \in \tbinom{[n]}{\mathsf{d}-1}.$$

We will use the P-terms in Lemma 3.6 with $e = 2$ to show that

$$d = T_n(v^{(\mathsf{d})}, \widehat{v}^{(\mathsf{d})}, \dots, v^{(n)}, \widehat{v}^{(n)}, \overline{d^I}^{[n]}) \tag{3.10}$$

holds for appropriately chosen elements $v^{(\mathsf{d})}, \widehat{v}^{(\mathsf{d})}, \dots, v^{(n)}, \widehat{v}^{(n)}$ in $\mathbf{C}$. Since all d^I are in L ($\subseteq \mathbf{C}$), this will show that $d \in \mathbf{C}$, and hence will complete the proof of (2).

By the definition of the term $T_{\mathsf{d}-1} := w^{[\mathsf{d}-1]}$ and by the choice of $b^{[\mathsf{d}-1]}$ and $d^{[\mathsf{d}-1]}$ we have that

$$b|_{[\mathsf{d}-1]} = b^{[\mathsf{d}-1]}|_{[\mathsf{d}-1]}, \quad d|_{[\mathsf{d}-1]} = d^{[\mathsf{d}-1]}|_{[\mathsf{d}-1]},$$

and

$$T_{\mathsf{d}-1}(b^{[\mathsf{d}-1]}) = b^{[\mathsf{d}-1]} \equiv_\theta d^{[\mathsf{d}-1]} = T_{\mathsf{d}-1}(d^{[\mathsf{d}-1]}).$$

Now we proceed by induction to show that for every $m = \mathsf{d}, \dots, n$ there exist tuples $u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m)}, \widehat{u}^{(m)}$ in $\mathbf{B}$ and tuples $v^{(\mathsf{d})}, \widehat{v}^{(\mathsf{d})}, \dots, v^{(m)}, \widehat{v}^{(m)}$ in $\mathbf{C}$ such that

$$b|_{[m]} = T_m(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m)}, \widehat{u}^{(m)}, \overline{b^I}^{[m]})|_{[m]}, \tag{3.11}$$

$$d|_{[m]} = T_m(v^{(\mathsf{d})}, \widehat{v}^{(\mathsf{d})}, \dots, v^{(m)}, \widehat{v}^{(m)}, \overline{d^I}^{[m]})|_{[m]}, \tag{3.12}$$

and

$$u^{(j)} \equiv_\theta v^{(j)}, \quad \widehat{u}^{(j)} \equiv_\theta \widehat{v}^{(j)} \qquad \text{for all } j = \mathsf{d}, \dots, m, \tag{3.13}$$

hence also

$$\begin{aligned}(B \ni)\ & T_m(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m)}, \widehat{u}^{(m)}, \overline{b^I}^{[m]}) \\ & \qquad\qquad \equiv_\theta T_m(v^{(\mathsf{d})}, \widehat{v}^{(\mathsf{d})}, \dots, v^{(m)}, \widehat{v}^{(m)}, \overline{d^I}^{[m]})\ (\in C).\end{aligned} \tag{3.14}$$

Then, equality (3.12) for $m = n$ yields the desired equality (3.10).

Our induction hypothesis is that the statement in the preceding paragraph is true for $m-1$, that is, there exist tuples $u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m-1)}, \widehat{u}^{(m-1)}$ in $\mathbf{B}$ and tuples $v^{(\mathsf{d})}, \widehat{v}^{(\mathsf{d})}, \dots, v^{(m-1)}, \widehat{v}^{(m-1)}$ in $\mathbf{C}$ such that (3.11)–(3.14) hold for $m-1$ in place of m. To simplify notation, let

$$\begin{aligned} b^{(m-1)} &:= T_{m-1}(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m-1)}, \widehat{u}^{(m-1)}, \overline{b^I}^{[m-1]}), \\ d^{(m-1)} &:= T_{m-1}(v^{(\mathsf{d})}, \widehat{v}^{(\mathsf{d})}, \dots, v^{(m-1)}, \widehat{v}^{(m-1)}, \overline{d^I}^{[m-1]}). \end{aligned}$$

Then we have that $b^{(m-1)}|_{[m-1]} = b|_{[m-1]}$ and $d^{(m-1)}|_{[m-1]} = d|_{[m-1]}$. Let $\gamma := b_m = b|_m$, $\beta := b^{(m-1)}|_m$, $\tau := d_m = d|_m$, and $\sigma := d^{(m-1)}|_m$. If we can show the existence of a pair of witnesses $u^{(m)}, \widehat{u}^{(m)}$ in $\mathbf{B}$ for the fork $(\gamma, \beta^{\gamma^2}) \in \textsc{fork}_m(B)$ and a pair of witnesses $v^{(m)}, \widehat{v}^{(m)}$ in $\mathbf{C}$ for the fork $(\tau, \sigma^{\tau^2}) \in \textsc{fork}_m(C)$ such that $u^{(m)} \equiv_\theta v^{(m)}$ and $\widehat{u}^{(m)} \equiv_\theta \widehat{v}^{(m)}$, then (3.13)–(3.14) will follow for m, and by Lemma 3.6, (3.11)–(3.12) will also hold for m.

To prove the existence of such $u^{(m)}, \widehat{u}^{(m)}$ and $v^{(m)}, \widehat{v}^{(m)}$, notice first that our induction hypothesis that (3.14) holds for $m-1$ in place of m has the effect that $(B \ni)\ b^{(m-1)} \equiv_\theta d^{(m-1)}\ (\in C)$, so in particular, $(B|_m \ni)\ \beta \equiv_{\theta_m} \sigma$. By our choice of b and d, we also have that $(B|_m \ni)\ \gamma = b_m \equiv_{\theta_m} d_m = \tau$. Thus, by assumption, F contains witnesses $b_\beta := f_1[m, \beta, \sigma] \in B$ and $c_\sigma := f_2[m, \beta, \sigma] \in C$ for the fork $(\beta, \sigma) \in \textsc{fork}_m(C)$ such that

$$b_\beta|_m = \beta, \quad c_\sigma|_m = \sigma, \quad \text{and} \quad b_\beta|_{[n]\setminus\{m\}} = c_\sigma|_{[n]\setminus\{m\}}; \tag{3.15}$$

F also contains witnesses $b_\gamma := f_1[m, \gamma, \tau] \in B$ and $c_\tau := f_2[m, \gamma, \tau] \in C$ for the fork $(\gamma, \tau) \in \textsc{fork}_m(C)$ such that

$$b_\gamma|_m = \gamma, \quad c_\tau|_m = \tau, \quad \text{and} \quad b_\gamma|_{[n]\setminus\{m\}} = c_\tau|_{[n]\setminus\{m\}}. \tag{3.16}$$

Let $b_{\beta^\gamma} := p(b_\beta, b_\gamma, b_\gamma)$, $c_{\beta^\tau} := p(b_\beta, c_\tau, c_\tau)$, and $c_{\sigma^\tau} := p(c_\sigma, c_\tau, c_\tau)$. Clearly, $b_{\beta^\gamma} \in B$ and $c_{\beta^\tau}, c_{\sigma^\tau} \in C$. Furthermore, the equalities in (3.15)–(3.16) imply that

$$\begin{aligned} b_{\beta^\gamma}|_m = \beta^\gamma, \quad c_{\beta^\tau}|_m = \beta^\tau, \quad c_{\sigma^\tau}|_m = \sigma^\tau, \\ \text{and} \quad b_{\beta^\gamma}|_{[n]\setminus\{m\}} = c_{\beta^\tau}|_{[n]\setminus\{m\}} = c_{\sigma^\tau}|_{[n]\setminus\{m\}}. \end{aligned} \tag{3.17}$$

It follows that c_{β^τ} and c_{σ^τ} are witnesses in $\mathbf{C}$ for the fork $(\beta^\tau, \sigma^\tau) \in \textsc{fork}_m(C)$. In addition, we get from (3.15)–(3.17) that

$$b_\beta \equiv_\theta c_\sigma, \quad b_\gamma \equiv_\theta c_\tau, \quad \text{and} \quad b_{\beta^\gamma} \equiv_\theta c_{\beta^\tau} \equiv_\theta c_{\sigma^\tau}. \tag{3.18}$$

Since b and $b^{(m-1)}$ are witnesses in B for the fork $(\gamma, \beta) \in \textsc{fork}_m(B) \subseteq \textsc{fork}_m(C)$, we can apply Lemma 3.7 first to (τ, γ) and (β, γ) to obtain witnesses

$$\begin{aligned} z &:= p(p(c_\tau, b_\gamma, b), p(c_\tau, b_\gamma, b_\gamma), c_\tau) \ (\in C) \\ z' &:= p(b^{(m-1)}, c_\tau, c_\tau) \ (\in C) \end{aligned}$$

for the fork $(\tau, \beta^\tau) \in \textsc{fork}_m(C)$. The analogous construction for the forks (γ, γ) and (β, γ) yields witnesses

$$\begin{aligned} w &:= p(p(b_\gamma, b_\gamma, b), p(b_\gamma, b_\gamma, b_\gamma), b_\gamma) = p(b, b_\gamma, b_\gamma) \ (\in B) \\ w' &:= p(b^{(m-1)}, b_\gamma, b_\gamma) \ (\in B) \end{aligned}$$

for the fork $(\gamma, \beta^\gamma) \in \textsc{fork}_m(B)$. Since w, w' are obtained from z, z' by replacing c_τ with b_γ, the relation $b_\gamma \equiv_\theta c_\tau$ in (3.18) implies that $z \equiv_\theta w$ and $z' \equiv_\theta w'$. Applying Lemma 3.7 again, now to the forks (τ, β^τ) and $(\sigma^\tau, \beta^\tau)$, we obtain witnesses

$$\begin{aligned} v^{(m)} &:= p(p(z, z', c_{\beta^\tau}), p(z, z', z'), z) \ (\in C) \\ \widehat{v}^{(m)} &:= p(c_{\sigma^\tau}, z, z) \ (\in C) \end{aligned}$$

for the fork $(\tau, \sigma^{\tau^2}) = (\tau, (\sigma^\tau)^\tau) \in \textsc{fork}_m(C)$. Similarly, for the corresponding forks (γ, β^γ) and $(\beta^\gamma, \beta^\gamma)$, we get witnesses

$$\begin{aligned} u^{(m)} &:= p(p(w, w', b_{\beta^\gamma}), p(w, w', w'), w) \ (\in B) \\ \widehat{u}^{(m)} &:= p(b_{\beta^\gamma}, w, w) \ (\in B) \end{aligned}$$

for the fork $(\gamma, \beta^{\gamma^2}) = (\gamma, (\beta^\gamma)^\gamma) \in \textsc{fork}_m(B)$. Since $u^{(m)}$ and $\widehat{u}^{(m)}$ are obtained from $v^{(m)}$ and $\widehat{v}^{(m)}$ by replacing z with w, z' with w', and $c_{\sigma^\tau}, c_{\beta^\tau}$ with b_{β^γ}, the relations $z \equiv_\theta w$, $z' \equiv_\theta w'$ proved earlier, and the relations $b_{\beta^\tau} \equiv_\theta c_{\beta^\tau} \equiv_\theta c_{\sigma^\tau}$ in (3.18) imply that $v^{(m)} \equiv_\theta u^{(m)}$ and $\widehat{v}^{(m)} \equiv_\theta \widehat{u}^{(m)}$. This completes the proof of Theorem 3.8. □

4. Algorithms Involving Compact Representations

Let $\mathcal{V}$ be a variety in a finite language with a d-cube term, and let $\mathcal{K}$ be a finite set of finite algebras in $\mathcal{V}$.

Our aim in this section is to use the results of Section 3 to show that the problems $\mathrm{SMP}(\mathcal{K})$ and $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$ are polynomial time equivalent (Theorem 4.8), and $\mathrm{SMP}(\mathcal{K})$ is also polynomial time equivalent to the problem of finding compact representations for algebras in $\mathbb{SP}_{\mathrm{fin}}\mathcal{K}$ given by their generators (Theorem 4.6). The latter problem can be described in more detail as follows:

CompactRep($\mathcal{K}$):

- INPUT: $a_1, \dots, a_k \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ with $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$.
- OUTPUT: A compact representation for the subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by $\{a_1, \dots, a_k\}$.

In addition, we will present a nondeterministic, polynomial time algorithm for solving CompactRep($\mathcal{K}$) (Theorem 4.5), which will show that both CompactRep($\mathcal{K}$) and $\mathrm{SMP}(\mathcal{K})$ are in NP.

To set up some terminology and notation, let $\mathbf{B}$ be a subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ with $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$, and let R be a representation for B (with $e = 1$, see Definition 3.2). Condition (ii) in the definition makes sure that for every $I \subseteq [n]$ with $|I| = \min\{n, \mathsf{d}-1\}$ and for every tuple $\bar{b} \in B|_I$ there exists an element $r[I, \bar{b}]$ in R such that $r[I, \bar{b}]|_I = \bar{b}$. Similarly, condition (iii) makes sure that for every derived fork $(\gamma, \delta) \in \textsc{fork}'_i(B)$ there exists a pair $(u[i, \gamma, \delta], v[i, \gamma, \delta])$ of elements in R which witness that $(\gamma, \delta) \in \textsc{fork}_i(R)$. For some of the algorithms we are going to discuss, it will be convenient to fix these choices for witnesses in R. Therefore we introduce the following definition.

Definition 4.1. Under the same assumptions as in Definition 3.2 with $e = 1$, we will call a representation R for B a *standardized representation* if the following two conditions are satisfied.

(ii)′ For every $I \subseteq [n]$ with $|I| = \min\{n, \mathsf{d}-1\}$ and for every tuple $\bar{b} \in B|_I$, an element $r[I, \bar{b}]$ of R is fixed so that $r[I, \bar{b}]|_I = \bar{b}$; this element will be referred to as *the designated witness in R for $\bar{b} \in R|_I$.*

(iii)′ For every $i \in [n]$ there is a set F_i with $\textsc{fork}'_i(B) \subseteq F_i \subseteq \textsc{fork}_i(R)$ such that for every $(\gamma, \delta) \in F_i$ a pair $(u[i, \gamma, \delta], v[i, \gamma, \delta])$ in R^2 is fixed which witnesses that $(\gamma, \delta) \in \textsc{fork}_i(R)$; this element will be referred to as *the designated witness in R for $(\gamma, \delta) \in \textsc{fork}_i(R)$.*

(iv) Every element of R has at least one designation.

Remark 4.2. It is clear that if the algebras $\mathbf{A}_1, \dots, \mathbf{A}_n$ all belong to our fixed finite set $\mathcal{K}$ of finite algebras in $\mathcal{V}$, then a standardized representation for any $B \subseteq A_1 \times \cdots \times A_n$ is compact. Moreover, it follows from the proof of Theorem 3.3 that if

R is a standardized representation for $\mathbf{B}$, then the equality (3.8) in Lemma 3.6 (with $e = 1$) holds provided

- $\overline{b^I}$ is an enumeration of the elements b^I of R designated to witness $b^I|_I = b|_I$ for all $I \in \binom{[n]}{\mathsf{d}-1}$; and
- for every $\mathsf{d} \leq m \leq n$, the pair $(u^{(m)}, \widehat{u}^{(m)})$ is the designated witness in R for the fork $(\gamma, \beta^\gamma) \in \textsc{fork}_m(R)$ where γ and β are determined by (3.6).

When we construct standardized representations, the following concepts will be useful.

Definitions 4.3. Let $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$ with $n \geq \mathsf{d}$, and let $R \subseteq \mathbf{A}_1 \times \dots \times \mathbf{A}_n$ and $b \in \mathbf{A}_1 \times \dots \times \mathbf{A}_n$.

1. We will call R a *partial standardized representation* if every element of R is designated to witness either $\bar{b} \in R|_I$ for some $I \in \binom{[n]}{\mathsf{d}-1}$ or a fork in $\textsc{fork}_m(R)$ for some m.

2. For a partial standardized representation R, we will say that b *is representable by* R if the following conditions are met:

(i) for each $I \in \binom{[n]}{\mathsf{d}-1}$, R contains elements b^I designated to witness $b^I|_I = b|_I$, and

(ii) for every $\mathsf{d} \leq m \leq n$, R contains designated witnesses $u^{(m)}, \widehat{u}^{(m)}$ for the fork (γ, β^γ) where β, γ are as defined in (3.6) (with $e = 1$).

It is clear from Lemma 3.6 that if b is representable by R, then the equality (3.8) (with $e = 1$) holds for b.

The proof of Lemma 3.6 can easily be turned into a polynomial time algorithm for solving the following problem:

IsRepresentable($\mathcal{K}$):

- INPUT: $b \in \mathbf{A}_1 \times \dots \times \mathbf{A}_n$ and a partial standardized representation $R \subseteq \mathbf{A}_1 \times \dots \times \mathbf{A}_n$ ($\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$) such that R contains elements b^I designated to witness $b^I|_I = b|_I$ for each $I \in \binom{[n]}{\mathsf{d}-1}$.
- OUTPUT: (YES, $\emptyset$, S) or (NO, S', S) where S', S ($\subseteq \langle R \cup \{b\} \rangle_P$) are lists of designated witnesses (missing from R) for the derived forks and for the forks that are not derived forks, respectively, so that b becomes representable by the partial standardized representation $R \cup S'$.

Note that the designated witnesses (for forks that are note derived forks) collected in the set S do not play a role in determining whether or not b is representable, but they will be useful in other algorithms that call IsRepresentable($\mathcal{K}$).

Lemma 4.4. *Let $\mathcal{V}$ be a variety in a finite language with a d-cube term. For any finite set $\mathcal{K}$ of finite algebras in $\mathcal{V}$, Algorithm 1 solves* IsRepresentable($\mathcal{K}$) *in polynomial time.*

Algorithm 1: For IsRepresentable($\mathcal{K}$)

Input: $b \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ and a partial standardized representation $R \subseteq \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ ($\mathbf{A}_1, \ldots, \mathbf{A}_n \in \mathcal{K}$) such that R contains elements b^I designated to witness $b^I|_I = b|_I$ for each $I \in \binom{[n]}{\mathsf{d}-1}$.
Output: (YES, $\emptyset$, S) or (NO, S', S) where S', S ($\subseteq \langle R \cup \{b\} \rangle_P$) are lists of designated witnesses (missing from R) for the derived forks and for the forks that are not derived forks, respectively, so that b becomes representable by the partial standardized representation $R \cup S'$.

```
1. $S = \emptyset$, $S' = \emptyset$,
2. $b^{(\mathsf{d}-1)} = b^{[\mathsf{d}-1]}$
3. for $m = \mathsf{d}, \ldots, n$ do
   3.1. $\beta = b^{(m-1)}|_m$, $\gamma = b|_m$, $c = p(b^{(m-1)}, b, b)$
   3.2. if $R$ has no designated witnesses for $(\gamma, \beta^\gamma) \in \textsc{fork}_m(R)$ then
            add $b, c$ to $S'$ as designated witnesses for $(\gamma, \beta^\gamma) \in \textsc{fork}_m(R)$
        end if
   3.3. if $R \cup S'$ has no designated witness for $(\gamma, \beta) \in \textsc{fork}_m(R)$ then
            add $b, b^{(m)}$ to $S$ as designated witnesses for $(\gamma, \beta) \in \textsc{fork}_m(R)$
        end if
   3.4. let $u, \widehat{u} \in R \cup S'$ be the designated witnesses for $(\gamma, \beta^\gamma) \in \textsc{fork}_m(R)$
   3.5. $b^{(m)} = t_m(b^{(m-1)}, \widehat{u}, u, \overline{b^I}^{[m]})$
   end for
4. if $S' = \emptyset$ then
      return (YES, $S'$, $S$) else
      return (NO, $S'$, $S$)
```

Table 1

Proof. First we prove that Algorithm 1 is correct. In Step 2, $b^{[\mathsf{d}-1]} \in R$ is the designated witness for $b^{[\mathsf{d}-1]}|_{[\mathsf{d}-1]} = b|_{[\mathsf{d}-1]}$. Step 3 follows the induction step in the proof of Lemma 3.6 (with $e = 1$). For each $m = \mathsf{d}, \ldots, n$, if $b^{(m-1)}$ has been constructed such that $b^{(m-1)} \in \langle R \rangle_P$ and $b^{(m-1)}|_{[m-1]} = b|_{[m-1]}$, then to construct $b^{(m)} = t_m(b^{(m-1)}, \widehat{u}, u, \overline{b^I}^{[m]})$ we need designated witnesses u, $\widehat{u}$ for the fork $(\gamma, \beta^\gamma) \in \textsc{fork}_m(R)$ where $\gamma = b|_m$ and $\beta = b^{(m-1)}|_m$. If such designated witnesses exist in R, then $b^{(m)} \in \langle R \rangle_P$ is computed, and the loop starts over, except when $m = n$. Therefore, if R contains designated witnesses for the appropriate forks for every m, then b is representable, and the algorithm stops in step 4 with the correct output: (YES, $\emptyset$, S).

If, for some m, R fails to contain designated witnesses for the fork (γ, β^γ) in the m-th coordinate, then it is clear that b is not representable by R. Furthermore, in this case $b, b^{(m)}$ witness the fork $(\gamma, \beta) \in \textsc{fork}_m(\langle R \cup \{b\}\rangle_B)$. It follows that b and $c = p(b^{(m-1)}, b, b)$ witness derived fork $(\gamma, \beta^\gamma) \in \textsc{fork}'_m(\langle R \cup \{b\}\rangle_P)$, because $c|_{[m-1]} = p(b^{(m-1)}|_{[m-1]}, b|_{[m-1]}, b|_{[m-1]} = p(b|_{[m-1]}, b|_{[m-1]}, b|_{[m-1]}) = b|_{[m-1]}$ and $b|_m = \gamma$, $c|_m = p(b^{(m-1)}|_m, b|_m, b|_m) = p(\beta, \gamma, \gamma) = \beta^\gamma$. Thus, after performing the designations in Steps 3.2 and 3.3 the algorithm can continue as before until it finishes, in step 4, with the correct output, (NO, S', S), where S' is a list of designated derived forks such that b is representable by the partial standardized representation $R \cup S'$. This proves the correctness of Algorithm 1.

To estimate the time complexity of Algorithm 1, notice that step 1 requires constant time, while steps 2 and 4 can be done in time $O(n|R|)$ and $O(n|S| + n|S'|)$, respectively, where $|R| + |S| + |S'| \leq O(n^{\mathsf{d}-1})$. Finally, by Lemma 3.6 (with $e = 1$), if b is representable by R, the computations in the loop in step 3 require at most $4\binom{n+1}{\mathsf{d}+1}$ applications of P so step 3 (including the search for witnesses for forks in R) can be done in $O(n^{\mathsf{d}+2})$ time. If b is not representable by R, essentially the same computation is performed, so the bound $O(n^{\mathsf{d}+2})$ applies in this case as well. This proves that Algorithm 1 runs in time $O(n^{\mathsf{d}+2})$. □

Theorem 4.5. *Let $\mathcal{V}$ be a variety in a finite language with a d-cube term. For any finite set $\mathcal{K}$ of finite algebras in $\mathcal{V}$, there is a nondeterministic algorithm which solves* CompactRep($\mathcal{K}$) *in polynomial time.*

Proof. We claim that nondeterministic Algorithm 2 runs in polynomial time and solves CompactRep($\mathcal{K}$).

To prove the correctness of Algorithm 2, let $a_1, \dots, a_k \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ ($\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$) be an arbitrary input for CompactRep($\mathcal{K}$), and let $\mathbf{B}$ denote the subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by $a_1, \dots, a_k$. We have to show that the set R returned by Algorithm 2 is a standardized representation for $\mathbf{B}$.

The set $R = R_0$ produced in steps 1–3 clearly satisfies $R_0 \subseteq B$, and contains one designated witness $r[I, \bar{b}]$ for each $\bar{b} \in \mathbf{B}|_I$ and every $I \in \binom{[n]}{\mathsf{d}-1}$.

In step 4, by applying IsRepresentable($\mathcal{K}$) k times, finitely many new pairs of elements from $\langle R \cup \{a_1, \dots, a_k\}\rangle_P$ ($\subseteq B$) are added to R to be designated witnesses for forks in $\langle R \cup \{a_1, \dots, a_k\}\rangle_P$ ($\subseteq B$) so that $a_1, \dots, a_k$ become representable by R. At this point it is clear that R is a partial standardized representation for B; however, R may not contain designated witnesses for all derived forks in B. Notice also that R will remain a partial standardized represenration for B as long as the elements added to R later on in the algorithm all come from B, and only designated elements are added to R.

Algorithm 2: For CompactRep($\mathcal{K}$) (nondeterministic)

Input: $a_1, \dots, a_k \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ with $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$ $(n \geq \mathsf{d})$.
Output: Standardized representation R for the subalgebra $\mathbf{B}$ of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by $a_1, \dots, a_k$.

```
1. $R_0 = \emptyset$
2. for $I \in \binom{[n]}{\mathsf{d}-1}$ do
  2.1. generate $\mathbf{B}|_I$ by $a_1|_I, \dots, a_k|_I$, and simultaneously,
  2.2. for each new $\overline{b} \in \mathbf{B}|_I$, add to $R_0$ an element $r[I, \overline{b}] \in \mathbf{B}$ that is designated to
       witness $r[I, \overline{b}]|_I = \overline{b}$.
  end for
3. $R = R_0$
4. for $b \in \{a_1, \dots, a_k\}$ do
  4.1. run Algorithm 1 for IsRepresentable($\mathcal{K}$) with input $b, R$ to get output
       (YES, $S'$, $S$) or (NO, $S'$, $S$)
  4.2. $R = R \cup S \cup S'$
  end for
5. closed := false
6. while not closed do
  6.1. closed := true
  6.2. for (nondeterministically chosen) basic operation symbol $f$ (say, $t$-ary) do
    6.2.1. for (nondeterministically chosen) tuples $b_i^I \in R_0$ $\big(i \in [t],\ I \in \binom{[n]}{\mathsf{d}-1}\big)$
           and $u_i^{(\mathsf{d})}, \widehat{u}_i^{(\mathsf{d})}, \dots, u_i^{(n)}, \widehat{u}_i^{(n)} \in R \setminus R_0$ such that each pair $u_i^{(m)}, \widehat{u}_i^{(m)}$ is a
           designated witness for a fork in $\textsc{fork}'_m(R)$ do
      6.2.1.1. $b_i = T_n(u_i^{(\mathsf{d})}, \widehat{u}_i^{(\mathsf{d})}, \dots, u_i^{(n)}, \widehat{u}_i^{(n)}, \overline{b_i^I})$ $(i \in [t])$
      6.2.1.2. $b = f(b_1, \dots, b_t)$
      6.2.1.3. run Algorithm 1 for IsRepresentable($\mathcal{K}$) with input $b, R$, to get
               output (YES, $S'$, $S$) or (NO, $S'$, $S$)
      6.2.1.4. if $S \cup S' \neq \emptyset$ then
               closed := false;  $R = R \cup S \cup S'$
               end if
           end for
    end for
  end while
7. for $m = \mathsf{d}, \dots, n$ do
  7.1. for all $(\gamma, \delta), (\beta, \delta) \in \textsc{fork}_m(R)$ which have designated witnesses in $R$ do
    7.1.1. if $R$ has no designated pair of witnesses for $(\gamma, \beta^\gamma)$ then
      7.1.1.1. find the pairs of designated witnesses $(v, v')$, $(u, u')$ for the forks
               $(\gamma, \delta), (\beta, \delta) \in \textsc{fork}_m(R)$
      7.1.1.2. add the pair $\big(p(p(v, v', u'), p(v, v', v'), v),\ p(u, v, v)\big)$ to $R$, and designate it to witness the fork $(\gamma, \beta^\gamma) \in \textsc{fork}_m(R)$.
           end if
    end for
  end for
8. return $R$
```

Table 2

Therefore, we will be done if we show that the rest of the algorithm, namely steps 6–7, add to R a pair of designated witnesses from B for every derived fork in B which has not been witnessed yet. First, the loop in step 6 is run as many times as necessary to make sure — by adding new elements to R, using IsRepresentable($\mathcal{K}$) — that

($*$) the set $\widetilde{R}$ of all elements representable by R is a subuniverse of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$.

All pairs of designated witnesses for forks that are added to R in step 6 belong to $\langle R \cup \{b\}\rangle_P \subseteq B$ (with the current $b \in B$ and R), therefore $R \subseteq B$ throughout step 6. So, ($*$) implies that $\widetilde{R}$ is a subuniverse of $\mathbf{B}$. On the other hand, by step 4, the generators $a_1, \dots, a_k$ of $\mathbf{B}$ are all in $\widetilde{R}$, therefore B is contained in $\widetilde{R}$. Thus, $B = \widetilde{R}$, and every element of B is representable by R.

By Lemma 3.7, the designated witnesses for forks added to R in step 7 all belong to the P-subalgebra generated by the preceding version of R, and hence all belong to B. Moreover, they are witnesses for derived forks in B.

It remains to prove that by the end of step 7, R contains designated witnesses for all derived forks of B in coordinates $\geq \mathsf{d}$. Let $m \geq \mathsf{d}$, and let $(\gamma, \sigma) \in \textsc{fork}'_m(B)$. Then there exists $(\gamma, \beta) \in \textsc{fork}_m(B)$ such that $\sigma = \beta^\gamma$. Hence, there exist $b, b' \in B$ such that $b|_{[m-1]} = b'|_{[m-1]}$ and $b|_m = \beta$, $b'|_m = \gamma$. By step 6, both b and b' are representable by R. Since $b|_{[m-1]} = b'|_{[m-1]}$, the standardized representations of b and b' agree up to step $m-1$; in particular,

$$b^{(m-1)} = T_{m-1}(u^{(\mathsf{d})}, \widehat{u}^{(\mathsf{d})}, \dots, u^{(m-1)}, \widehat{u}^{(m-1)}, \overline{b^I}^{[m-1]}) = (b')^{(m-1)}.$$

Let $\delta := b^{(m-1)}|_m$. Since b and b' are representable by R, R contains designated witnesses (u, u') and (v, v') for the forks $(\beta, \delta), (\gamma, \delta) \in \textsc{fork}_m(R)$. Step 7 makes sure that in this situation, a designated pair of witnesses for the fork $(\gamma, \beta^\gamma) = (\gamma, \sigma)$ gets into R (if it was not there before). This proves that by the end of step 7, the set R returned by Algorithm 2 is a standardized representation for B. Hence R is a compact representation for B.

To bound the time complexity of Algorithm 2 we have to look at the run times of steps 2, 4, 6, and 7. In step 2, the subalgebras $\mathbf{B}|_I$ can be generated in a constant number of steps that depeneds on $\mathcal{K}$ only (and is independent of the size of the input), therefore the time complexity is determined by the number $\binom{n}{\mathsf{d}-1}$ of iterations of the loop in step 2 and the time needed for computing the designated witnesses that are added to R in each loop, which is bounded above by $O(n)$. Thus, step 2 runs in $O(n^{\mathsf{d}})$ time.

The time required by Step 4 is dominated by the k calls of Algorithm 1 for IsRepresentable($\mathcal{K}$), which is $O(kn^{\mathsf{d}+2})$.

The nondeterministic loop in step 6 is repeated as long as a nonempty list of pairs of witnesses for forks needs to be added to R. Since there are at most $O(n)$ forks to be witnessed, the loop will be repeated at most $O(n)$ times. In each loop, computing every b_i requires at most $3\binom{n+1}{\mathsf{d}+1}$ applications of P (by Lemma 3.6, $e = 1$), b can

Algorithm 3: Reduction of CompactRep($\mathcal{K}$) to SMP($\mathcal{K}$)

```
Input: $a_1, \dots, a_k \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ with $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$.
Output: Standardized representation $R$ for the subalgebra $\mathbf{B}$ of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$
generated by $a_1, \dots, a_k$.
    1. $R_0 = \emptyset$
    2. for $I \in \binom{[n]}{\mathsf{d}-1}$ do
       2.1. generate $\mathbf{B}|_I$ by $a_1|_I, \dots, a_k|_I$, and simultaneously,
       2.2. for each new $\bar{b} \in \mathbf{B}|_I$, add to $R_0$ an element $r[I, \bar{b}] \in \mathbf{B}$ that is designated
            to witness $r[I, \bar{b}]|_I = \bar{b}$.
       end for
    3. $R = R_0$
    4. for $i = \mathsf{d}, \dots, n$ and $\gamma \in \mathbf{B}|_i$ do
       4.1. find $b \in R_0$ with $b|_i = \gamma$
       4.2. for $\beta \in \mathbf{B}|_i$ do
            4.2.1. let $c \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_i$ be such that $c|_{[i-1]} = b|_{[i-1]}$ and $c|_i = \beta^\gamma$
            4.2.2. run SMP($\mathcal{K}$) with input $a_1|_{[i]}, \dots, a_k|_{[i]}, c \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_i$
            4.2.3. if answer is YES then
                   for $j = i + 1, \dots, n$ do
               4.2.3.1. find $c_j \in A_j$ such that SMP($\mathcal{K}$) with input
                        $a_1|_{[j]}, \dots, a_k|_{[j]}, (c, c_j) \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_j$ answers YES
               4.2.3.2. $c = (c, c_j)$
                   end for
                   end if
            4.2.4. add the pair $(b, c)$ to $R$ and designate it to witness the fork $(\gamma, \beta^\gamma) \in$
                   $\text{FORK}_i(R)$.
            end for
       end for
    5. return $R$
```

Table 3

be obtained in $O(n)$ time, and then Algorithm 1 for IsRepresentable($\mathcal{K}$) runs in $O(n^{\mathsf{d}+2})$ time. Thus, the time required by step 6 is $O(n^{\mathsf{d}+3})$.

Finally, in step 7 the loop is repeated $O(n)$ times, and each loop requires no more than $O(n)$ time (including the search for designated witnesses for forks).

Thus, the overall run time for Algorithm 2 is $O(kn^{\mathsf{d}+3})$. □

Algorithm 4: Reduction of $\mathrm{SMP}(\mathcal{K})$ to $\textsc{CompactRep}(\mathcal{K})$

```
Input: $a_1, \dots, a_k, b \in \mathbf{A}_1 \times \dots \times \mathbf{A}_n$ with $\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$.
Question: Is $b$ in the subalgebra $\mathbf{B}$ of $\mathbf{A}_1 \times \dots \times \mathbf{A}_n$ generated by $a_1, \dots, a_k$?
    1. Run $\textsc{CompactRep}(\mathcal{K})$ with input $a_1, \dots, a_k$, and let $R$ be its output (a
       standardized representation for $\mathbf{B}$)
    2. for $I \in \binom{[n]}{\mathsf{d}-1}$ do
       2.1. if $R$ contains no designated witness for $b|_I$ then
              return NO
            end if
       end for
    3. Run Algorithm 1 for $\textsc{IsRepresentable}(\mathcal{K})$ with input $b, R$, to get output
       $(\mathsf{A}, S', S)$ with $\mathsf{A}$ = YES or $\mathsf{A}$ = NO
    4. return $\mathsf{A}$
```

Table 4

Theorem 4.6. *Let $\mathcal{V}$ be a variety in a finite language with a d-cube term. For any finite set $\mathcal{K}$ of finite algebras in $\mathcal{V}$, the decision problem $\mathrm{SMP}(\mathcal{K})$ and the computational problem $\textsc{CompactRep}(\mathcal{K})$ are polynomial time reducible to one another.*

Proof. First we will prove that Algorithm 3 solves $\textsc{CompactRep}(\mathcal{K})$ by repeated calls of $\mathrm{SMP}(\mathcal{K})$ in polynomial time.

For the correctness notice that steps 1–3 of Algorithm 3 are the same as those of Algorithm 2, so by the end of step 3, $R = R_0$ contains a designated witness $r[I, \bar{b}] \in \mathbf{B}$ for $r[I, \bar{b}]|_I = \bar{b}$ for each $\bar{b} \in \mathbf{B}|_I$ and for all $I \in \binom{[n]}{\mathsf{d}-1}$. Elements without designations are not added to R during this process or later in step 4, therefore we will be done if we show that step 4 adds to R a pair of designated witnesses (from $\mathbf{B}$) for all derived forks of $\mathbf{B}$ in coordinates $\geq \mathsf{d}$.

Lines 4, 4.1–4.2, and 4.2.1–4.2.2 show that step 4 examines each pair $(\gamma, \beta) \in \mathbf{B}|_i \times \mathbf{B}|_i$ for every $i = \mathsf{d}, \dots, n$, finds $b = (b_1, \dots .b_n) \in R_0$ $(\subseteq R \subseteq B)$ such that $b_i = \gamma$, and checks — using $\mathrm{SMP}(\mathcal{K})$ — whether or not the tuple $c = (b_1, \dots, b_{i-1}, \beta^\gamma)$ is in the subalgebra $\mathbf{B}|_{[i]}$ generated by the elements $a_1|_{[i]}, \dots, a_k|_{[i]}$. If the answer is YES, then B contains a tuple of the form $(c, c_{i+1}, \dots, c_n) = (b_1, \dots, b_{i-1}, \beta^\gamma, c_{i+1}, \dots, c_n)$ for some $c_j \in A_j$ $(j = i+1, \dots, n)$, which will be found, coordinate-by-coordinate, by repeated applications of $\mathrm{SMP}(\mathcal{K})$ in 4.2.3. Moreover, in this case it is clear that (b, c) witnesses that (γ, β^γ) is a derived fork in B, so this witness is correctly added to R. If the instance of $\mathrm{SMP}(\mathcal{K})$ run in 4.2.2 gives the answer NO, then we have that the tuple $c = (b_1, \dots, b_{i-1}, \beta^\gamma)$ is not in $\mathbf{B}|_{[i]}$. An application of Lemma 3.1(2) with

$\delta := \beta^\gamma$ shows that in this case (γ, β^γ) is not a derived fork in the i-th coordinate of $\mathbf{B}$. Hence step 4 correctly adds no witnesses for (γ, β^γ) to R in this case. This completes the proof of correctness of Algorithm 3.

As we saw in the proof of Theorem 4.5 the run time of steps 1–3 is $O(n^{\mathsf{d}})$. Step 4 requires running $\mathrm{SMP}(\mathcal{K})$ $O(n^2)$ times on inputs not larger than the input for Algorithm 3, and adding one pair of witnesses to R no more than $O(n)$ times. This show that Algorithm 3 reduces $\textsc{CompactRep}(\mathcal{K})$ to $\mathrm{SMP}(\mathcal{K})$ in polynomial time.

For the reverse direction we will argue that Algorithm 4 reduces $\mathrm{SMP}(\mathcal{K})$ to $\textsc{CompactRep}(\mathcal{K})$ in polynomial time. The algorithm starts with computing a standardized representation R for the algebra $\mathbf{B}$ generated by the input tuples $a_1, \ldots, a_k$ — using $\textsc{CompactRep}(\mathcal{K})$. A necessary condition for the input tuple b to be in $\mathbf{B}$ is that $b|_I \in \mathbf{B}|_I$ for all $I \in \binom{[n]}{\mathsf{d}-1}$. Since R contains designated witnesses $r[I, \bar{b}]$ for all $\bar{b} \in \mathbf{B}|_I$ and $I \in \binom{[n]}{\mathsf{d}-1}$, b will satisfy this necessary condition if and only if R contains designated witnesses for all projections $b|_I$ $\bigl(I \in \binom{[n]}{\mathsf{d}-1}\bigr)$ of b. This is being checked in step 2 of Algorithm 4; if the condition fails for some $I \in \binom{[n]}{\mathsf{d}-1}$, the algorithm returns the correct answer NO, meaning, $b \notin \mathbf{B}$.

If the algorithm passes step 2 without returning NO, then b, R is a correct input for $\textsc{IsRepresentable}\,\mathcal{K}$, which checks in step 3 whether b is representable by R. Since every tuple representable by R must be in $\mathbf{B}$, and conversely, by Remark 4.2, every element of $\mathbf{B}$ is representable by R, we get that the YES/NO answer provided by $\textsc{IsRepresentable}(\mathcal{K})$ is the correct answer to $\mathrm{SMP}(\mathcal{K})$ for the given input. This shows the correctness of Algorithm 4.

Step 2 of Algorithm 4 runs in $O(n^{\mathsf{d}-1})$ time, while $\textsc{IsRepresentable}(\mathcal{K})$ in step 3 requires $O(n^{\mathsf{d}+2})$ time. Thus, Algorithm 4 reduces $\mathrm{SMP}(\mathcal{K})$ to $\textsc{CompactRep}(\mathcal{K})$ in $O(n^{\mathsf{d}+2})$ time. □

The following statement is now an immediate consequence of Theorems 4.5–4.6.

Corollary 4.7. *If $\mathcal{V}$ is a variety in a finite language with a d-cube term, and $\mathcal{K}$ is a finite set of finite algebras in $\mathcal{V}$, then* $\mathrm{SMP}(\mathcal{K}) \in \mathsf{NP}$.

Now we prove that for a finite set $\mathcal{K}$ of finite algebras in a variety (in a finite language) with a cube term, the problems $\mathrm{SMP}(\mathcal{K})$ and $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$ are polynomial time equivalent.

Theorem 4.8. *Let $\mathcal{V}$ be a variety in a finite language with a d-cube term. For any finite set $\mathcal{K}$ of finite algebras in $\mathcal{V}$, the decision problems* $\mathrm{SMP}(\mathcal{K})$ *and* $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$ *are polynomial time equivalent.*

Proof. $\mathrm{SMP}(\mathcal{K})$ is a subproblem of $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$, so $\mathrm{SMP}(\mathcal{K})$ is clearly polynomial time reducible to $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$. For the converse we will show that Algorithm 5 reduces $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$ to $\mathrm{SMP}(\mathcal{K})$ in polynomial time.

Algorithm 5: Reduction of $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$ to $\mathrm{SMP}(\mathcal{K})$

```
Input: $c_1, \dots, c_k, c_{k+1} \in \mathbf{C}_1 \times \cdots \times \mathbf{C}_n$ with $\mathbf{C}_1, \dots, \mathbf{C}_n \in \mathbb{HS}\mathcal{K}$.
Question: Is $c_{k+1}$ in the subalgebra $\mathbf{D}$ of $\mathbf{C}_1 \times \cdots \times \mathbf{C}_n$ generated by $c_1, \dots, c_k$?
    1. for $j = 1, \dots, n$ do
       find $\mathbf{A}_j \in \mathcal{K}$, $\mathbf{B}_j \leq \mathbf{A}_j$ and $\theta_j \in \mathrm{Con}(\mathbf{B}_j)$ such that $\mathbf{C}_j = \mathbf{B}_j/\theta_j$
       end for
    2. $G = \emptyset$, $F = \emptyset$, $L = \emptyset$
    3. for $i = 1, \dots, k+1$ do
       3.1. $a_i = ()$
       3.2. for $j = 1, \dots, n$ do
                find $a_{ij} \in B_j$ with $a_{ij}/\theta_j = c_i|_j$
                $a_i = (a_i, a_{ij})$
                end for
       3.3. if $i \leq k$ then $G = G \cup \{a_i\}$ end if
       end for
    4. for $j = 1, \dots, n$ do
       4.1. generate $\mathbf{B}|_j$ by $a_{1j}, \dots, a_{kj}$, and simultaneously,
       4.2. for each new $\beta$ in $\mathbf{B}|_j$ do
            4.2.1. find an element $f_1[j, \beta, *]$ generated by $a_1, \dots, a_k$ satisfying
                   $f_1[j, \beta, *]|_j = \beta$
            4.2.2. for all $\gamma \equiv_{\theta_j} \beta$ $(\gamma \in B_j)$ do
                   add to $F$ the tuple $f_2[j, \beta, \gamma]$ satisfying $f_2[j, \beta, \gamma]|_j = \gamma$ and
                   $f_2[j, \beta, \gamma]|_{[n]\setminus\{j\}} = f_1[j, \beta, *]|_{[n]\setminus\{j\}}$
                   end for
            end for
       end for
    5. for $I \in \binom{[n]}{\mathsf{d}-1}$ do
       5.1. generate $\mathbf{B}|_I$ by $(G \cup F)|_I$, and simultaneously,
       5.2. for each new $\overline{b} \in \mathbf{B}|_I$ do
            5.2.1. find an element $r_1[I, \overline{b}] \in \mathbf{B}$ such that $r_1[I, \overline{b}]|_I = \overline{b}$
            5.2.2. for all $\overline{d}$ such that $\overline{d}|_j \equiv_{\theta_j} \overline{b}|_j$ for all $j \in I$ do
                   add to $L$ the tuple $r_2[I, \overline{d}]$ satisfying $r_2[I, \overline{d}]|_I = \overline{d}$ and
                   $r_2[I, \overline{d}]|_{[n]\setminus I} = r_1[I, \overline{b}]|_{[n]\setminus I}$
                   end for
            end for
       end for
    6. run $\mathrm{SMP}(\mathcal{K})$ with the input $G \cup L \cup F \subseteq \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ and $a_{k+1} \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$
       $(\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K})$, to get an answer $\mathsf{A} = $ YES or $\mathsf{A} = $ NO
    7. return $\mathsf{A}$
```

TABLE 5

In step 1 Algorithm 5 finds the algebras $\mathbf{A}_i \in \mathcal{K}$, their subalgebras $\mathbf{B}_i$ and their congruences θ_i such that the algebras $\mathbf{C}_i$ in the input are $\mathbf{C}_i = \mathbf{B}_i/\theta_i$ ($i \in [n]$). Step 2 initializes the computation of three sets G, F, and L, which will be completed in steps 3, 4, and 5, respectively.

In step 3 tuples $a_1, \dots, a_{k+1} \in \mathbf{B}_1 \times \cdots \times \mathbf{B}_n$ ($\leq \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$) are found such that for the product congruence $\theta = \theta_1 \times \cdots \times \theta_n$ we have $c_j = a_j/\theta$ for all $j \in [k+1]$. Thus, the set $G = \{a_1, \dots, a_k\}$ obtained at the end of step 3 is a generating set for a subalgebra $\mathbf{B}$ of $\mathbf{B}_1 \times \dots \mathbf{B}_n$ such that $\mathbf{B}/\theta|_{\mathbf{B}} = \mathbf{D}$.

By Theorem 3.8(1),

$$c_{k+1} \in \mathbf{D} \quad \Longleftrightarrow \quad a_{k+1} \in \mathbf{B}[\theta]. \tag{4.1}$$

Therefore, Algorithm 5 gives the correct answer in steps 6–7, provided the set $G \cup L \cup F$ produced earlier in the process is a generating set for $\mathbf{B}[\theta]$. By Theorem 3.8(2) it suffices to check that the set L constructed in step 5 satisfies condition (L), while the set F constructed in step 4 satisfies condition (F) in Theorem 3.8. For L this is straightforward to check. For F note that, given $j \in [n]$ and $\beta \in B|_j$ as in step 4.2, the tuple $f_1[j, \beta, *]$ obtained in step 4.2.1 belongs to $\mathbf{B}$. Hence, when the 'for' loop in step 4.2.2 is performed for $\gamma = \beta$, we get the tuple $f_2[j, \beta, \beta] = f_1[j, \beta, *] \in B$, which is added to F. This tuple can serve as the tuple denoted $f_1[j, \beta, \gamma]$ in condition (F) for every $f_2[j, \beta, \gamma]$ added to F in step 4.2.2. This shows that the set F constructed in step 4 satisfies condition (F) in Theorem 3.8, and hence finishes the proof of the correctness of Algorithm 5.

Steps 1–3 run in $O(kn)$ time, step 4 in $O(n^2)$ time, while steps 5 in $O(n^{\mathsf{d}})$ time. So, the reduction of $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$ to $\mathrm{SMP}(\mathcal{K})$ takes $O(kn^{\mathsf{d}})$ time. For an input of size $O(kn)$ of $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$ we get an input of size $O(kn^{\mathsf{d}})$ for $\mathrm{SMP}(\mathcal{K})$. □

5. Structure Theory and the Subpower Membership Problem

The main result of [10] is a structure theorem for the critical subalgebras of finite powers of algebras with cube (or parallelogram) terms. In this section we adapt the structure theorem from [10] to find a new representation (different from compact representations) for subalgebras of products of algebras in a variety $\mathcal{V}$ with a cube term. In the next section, this representation will be used to prove the main result of the paper.

To restate the result from [10] that we need here, we introduce some terminology and notation. Let $\mathcal{V}$ be a variety with a d-cube term (or equivalently, a d-parallelogram term, see Theorem 2.1), and let $\mathbf{R}$ be a subalgebra of a product $\mathbf{A}^{(1)} \times \cdots \times \mathbf{A}^{(n)}$ of some algebras $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(n)} \in \mathcal{V}$. Let $\mathbf{A}_i := \mathbf{R}|_i$ for each $i \in [n]$, and let $\mathbf{C} := \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$. So, $\mathbf{R}$ is a subdirect product of the subalgebras $\mathbf{A}_i$ of $\mathbf{A}^{(i)}$ ($i \in [n]$), and $\mathbf{R}$ is a subalgebra of $\mathbf{C}$.

We say that $\mathbf{R}$ is a *critical subalgebra of* $\mathbf{A}^{(1)} \times \cdots \times \mathbf{A}^{(n)}$ if

- $\mathbf{R}$ is completely $\cap$-irreducible in the lattice of subalgebras of $\mathbf{A}^{(1)} \times \cdots \times \mathbf{A}^{(n)}$, and
- $\mathbf{R}$ is *directly indecomposable* in the following sense: $[n]$ cannot be partitioned into two nonempty sets I and J such that $\mathbf{R}$ and $\mathbf{R}|_I \times \mathbf{R}|_J$ differ only by a permutation of coordinates.

Now let us assume that $\mathbf{R}$ is a critical subalgebra of $\mathbf{A}^{(1)} \times \cdots \times \mathbf{A}^{(n)}$. Choose and fix $\delta_i \in \mathrm{Con}(\mathbf{A}_i)$ $(i \in [n])$ such that $\delta := \delta_1 \times \cdots \times \delta_n$ is the largest product congruence of $\mathbf{C}$ with the property that $\mathbf{R}$ is a δ-saturated subalgebra of $\mathbf{C}$. (Such a congruence exists, because the join of product congruences is a product congruence, and if $\mathbf{R}$ is saturated with respect to a family of congruences of $\mathbf{C}$, then it is saturated with respect to their join.) With this notation, let $\overline{\mathbf{R}} := \mathbf{R}/(\delta{\restriction_{\mathbf{R}}})$, and let $\overline{\mathbf{A}}_i := \mathbf{A}_i/\delta_i$ $(i \in [n])$; we call $\overline{\mathbf{R}}$ *the reduced representation of* $\mathbf{R}$.

Theorems 2.5 and 4.1 of [10] yield a structure theorem for the critical subalgebras of finite powers $\mathbf{A}^n$ of an arbitrary algebra $\mathbf{A} \in \mathcal{V}$. The relevant proofs in [10], namely the proofs of Theorem 2.5 (and its preparatory Lemmas 2.1, 2.3, 2.4) and Theorem 3.6 (part (3), implication $\Rightarrow$), carry over without any essential changes to the more general situation when instead of subalgebras of powers $\mathbf{A}^n$ with $\mathbf{A} \in \mathcal{V}$ we consider subalgebras of products $\mathbf{A}^{(1)} \times \cdots \times \mathbf{A}^{(n)}$ with $\mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(n)} \in \mathcal{V}$. Thus, we get the theorem below, where we state only those parts of the structure theorem that we need later on, retaining the numbering from [10, Theorem 2.5], but replacing 'd-parallelogram term' by 'd-cube term'. The superscript $^\flat$ in $(6)^\flat$ indicates that instead of the original condition (6) we state a weaker condition which is sufficient for our purposes.

Theorem 5.1 (Cf. [10])**.** *Let $\mathcal{V}$ be a variety with a d-cube term, let $\mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(n)} \in \mathcal{V}$, and let $\overline{\mathbf{R}}$ be the reduced representation of a critical subalgebra $\mathbf{R}$ of $\mathbf{A}^{(1)} \times \cdots \times \mathbf{A}^{(n)}$. If $n \geq \mathsf{d}$, then the following hold.*

(1) *$\overline{\mathbf{R}} \leq \prod_{i=1}^n \overline{\mathbf{A}}_i$ is a representation of $\overline{\mathbf{R}}$ as a subdirect product of subdirectly irreducible algebras $\overline{\mathbf{A}}_i$.*

$(6)^\flat$ *$\overline{\mathbf{A}}_i$ and $\overline{\mathbf{A}}_j$ are similar for any $i, j \in [n]$.*

(7) *If $n > 2$, then each $\overline{\mathbf{A}}_i$ has abelian monolith μ_i $(i \in [n])$.*

(8) *For the centralizers $\rho_\ell := (0 : \mu_\ell)$ of the monoliths μ_ℓ $(\ell \in [n])$, the image of the composite map*

$$\overline{\mathbf{R}} \overset{\mathrm{proj}_{ij}}{\to} \overline{\mathbf{A}}_i \times \overline{\mathbf{A}}_j \twoheadrightarrow \overline{\mathbf{A}}_i/\rho_i \times \overline{\mathbf{A}}_j/\rho_j.$$

is the graph of an isomorphism $\overline{\mathbf{A}}_i/\rho_i \to \overline{\mathbf{A}}_j/\rho_j$ for any $i, j \in [n]$.

Note that the homomorphism in part (8) is the same as in [10, Theorem 2.5 (8)]; the slightly different description presented here will be more convenient later on.

Now we are ready to discuss our representation theorem. Let $\mathbf{B}_1, \ldots, \mathbf{B}_n$ be nontrivial algebras in $\mathcal{V}$, and let $\mathbf{B}$ be a subdirect subalgebra of $\mathbf{B}_1 \times \cdots \times \mathbf{B}_n$.

First, we replace each $\mathbf{B}_j$ ($j \in [n]$) by its image under the embedding

$$\mathbf{B}_j \hookrightarrow \prod_{\sigma \in \mathrm{Irr}(\mathbf{B}_j)} \mathbf{B}_j/\sigma, \qquad x_j \mapsto (x_j/\sigma)_{\sigma \in \mathrm{Irr}(\mathbf{B}_j)}; \tag{5.1}$$

thus, each $\mathbf{B}_j$ is replaced by a subdirect product of all of its subdirectly irreducible quotients $\mathbf{B}_j/\sigma$. To set up a more convenient notation, let

$$W := \{(j,\sigma) : j \in [n], \sigma \in \mathrm{Irr}(\mathbf{B}_j)\},$$

and for each $j \in [n]$, let $W_j := \{j\} \times \mathrm{Irr}(\mathbf{B}_j)$; thus, $W = W_1 \cup \dots \cup W_n$. Furthermore, for each $w = (j,\sigma) \in W$ let $\widehat{\mathbf{B}}_w := \mathbf{B}_j/\sigma$, and for every element $x_j \in \mathbf{B}_j$ let $\widehat{x}_w := x_j/\sigma$. Then the product of the embeddings (5.1) for all $j \in [n]$ yields an embedding

$$\widehat{\ }: \prod_{j \in [n]} \mathbf{B}_j \hookrightarrow \prod_{w \in W} \widehat{\mathbf{B}}_w \qquad x = (x_j)_{j \in [n]} \mapsto \widehat{x} := (\widehat{x}_w)_{w \in W}.$$

We will denote the image of $\mathbf{B}$ under this embedding $\widehat{\ }$ by $\widehat{\mathbf{B}}$. By construction, $\widehat{\mathbf{B}}$ is a subdirect product of the subdirectly irreducible algebras $\widehat{\mathbf{B}}_w$ ($w \in W$); equivalently, $\widehat{\mathbf{B}}_w = \widehat{\mathbf{B}}|_w$ for all $w \in W$. For each $w \in W$ let μ_w denote the monolith of $\widehat{\mathbf{B}}_w$ and ρ_w its centralizer $(0 : \mu_w)$.

Next we define a relation $\sim$ on W as follows: we require $\sim$ to be reflexive, and for distinct $v, w \in W$ we define $v \sim w$ to hold if and only if

- the subdirectly irreducible algebras $\widehat{\mathbf{B}}_v$ and $\widehat{\mathbf{B}}_w$ are similar with abelian monoliths μ_v and μ_w, and
- the image of $\widehat{\mathbf{B}}|_{vw}$ under the natural map $\widehat{\mathbf{B}}_v \times \widehat{\mathbf{B}}_w \twoheadrightarrow (\widehat{\mathbf{B}}_v/\rho_v) \times (\widehat{\mathbf{B}}_w/\rho_w)$ is the graph of an isomorphism $\widehat{\mathbf{B}}_v/\rho_v \to \widehat{\mathbf{B}}_w/\rho_w$.

It is easy to see that $\sim$ is an equivalence relation on W.

Our representation theorem describes the algebra $\mathbf{B}$ in terms of its image $\widehat{\mathbf{B}}$, namely, it shows that $\widehat{\mathbf{B}}$ is determined by its projections onto small sets of coordinates (i.e., small subsets of W) and by its projections onto the blocks of $\sim$. A block of $\sim$ may be large, but the image of $\widehat{\mathbf{B}}$ under a projection onto a block of $\sim$ has a special structure.

Theorem 5.2. *Let $\mathcal{V}$ be a variety with a d-cube term, let $\mathbf{B}_1, \dots, \mathbf{B}_n$ be nontrivial algebras in $\mathcal{V}$, and let $\mathbf{B}$ be a subdirect subalgebra of $\mathbf{B}_1 \times \dots \times \mathbf{B}_n$. Furthermore, let W, $\widehat{\ }$, $\widehat{\mathbf{B}}_w$ ($w \in W$), $\widehat{\mathbf{B}}$, and $\sim$ be as defined above. Then, for any tuple $c \in \mathbf{B}_1 \times \dots \times \mathbf{B}_n$, the following conditions are equivalent:*

(a) $c \in \mathbf{B}$.

(b) *c satisfies $c|_I \in \mathbf{B}|_I$ for all $I \subseteq [n]$ such that $|I| \leq \max\{\mathsf{d}, 3\}$, and the image $\widehat{c} \in \prod_{w \in W} \widehat{\mathbf{B}}_w$ of c under the map $\widehat{\ }$ satisfies $\widehat{c}|_U \in \widehat{\mathbf{B}}|_U$ for all blocks U ($\subseteq W$) of $\sim$ of size $|U| \geq \max\{\mathsf{d}, 3\}$.*

(c) *the image* $\widehat{c} \in \prod_{w\in W}\widehat{\mathbf{B}}_w$ *of* c *under the map* $\widehat{\ }$ *satisfies* $\widehat{c}|_U \in \widehat{\mathbf{B}}|_U$ *for all* $U \subseteq W$ *such that*
- $|U| < \max\{\mathsf{md}, 3\}$ *where* $\mathsf{m} = \max\{|\operatorname{Irr}(\mathbf{B}_j)| : j \in [n]\}$*, or*
- U *is a block of* $\sim$ *of size* $|U| \geq \max\{\mathsf{md}, 3\}$.

Remark 5.3. The equivalence of conditions (a) and (c) in Theorem 5.2 can be restated as follows: $\widehat{\mathbf{B}}$ is the intersection of the subalgebras

$$\operatorname{proj}_U^{-1}\bigl(\operatorname{proj}_U(\widehat{\mathbf{B}})\bigr) = \operatorname{proj}_U^{-1}\bigl(\widehat{\mathbf{B}}|_U\bigr)$$

of $\prod_{w\in W}\widehat{\mathbf{B}}_w$ as U runs over the subsets of W listed in (c).

Proof of Theorem 5.2. Since $\widehat{\ }$ and $|_U$ ($U \subseteq W$) are homomorphisms, it is clear that $c \in \mathbf{B}$ implies $\widehat{c}|_U \in \widehat{\mathbf{B}}|_U$ for all $U \subseteq W$. This proves (a) ⇒ (c).

For the implication (c) ⇒ (b), assume that (c) holds. Then $\widehat{c}|_U \in \widehat{\mathbf{B}}|_U$ for all blocks U of $\sim$, so the second statement in (b) holds. To establish the first statement, choose $I \subseteq [n]$ such that $|I| < \max\{\mathsf{d}, 3\}$, and let $W_I := \bigcup_{j\in I} W_j$. Since the product of the isomorphisms $\mathbf{B}_j \to \widehat{\mathbf{B}}|_{W_j}$, $x_j \mapsto (\widehat{x}_w)_{w\in W_j}$ (induced by the embeddings in (5.1)) yields an isomorphism

$$\prod_{j\in I}\mathbf{B}_j \to \prod_{j\in I}\widehat{\mathbf{B}}|_{W_j}\Bigl(\leq \prod_{w\in W_I}\widehat{\mathbf{B}}_w\Bigr), \quad (x_j)_{j\in I} \mapsto (x_w)_{w\in W_I} \qquad (x_j \in \mathbf{B}_j),$$

which maps $\mathbf{B}|_I$ onto $\widehat{\mathbf{B}}|_{W_I}$, we get that $c|_I \in \mathbf{B}|_I$ holds if and only if $\widehat{c}|_{W_I} \in \widehat{\mathbf{B}}|_{W_I}$. The latter follows from assumption (c), because $|W_I| \leq \sum_{j\in I}|W_j| = \sum_{j\in I}|\operatorname{Irr}(\mathbf{B}_j)| \leq \mathsf{m}|I|$. This completes the proof of (c) ⇒ (b).

The remaining implication (b) ⇒ (a) is the heart of Theorem 5.2, which we will prove now. Assume that $c \notin \mathbf{B}$, but $c|_I \in \mathbf{B}|_I$ for all $I \subseteq [n]$ with $|I| < \max\{\mathsf{d}, 3\}$. We have to show that $\widehat{c}|_U \notin \widehat{\mathbf{B}}|_U$ for some block $U\,(\subseteq W)$ of $\sim$ of size $|U| \geq \max\{\mathsf{d}, 3\}$.

Using the assumption $c \notin \mathbf{B}$, we first choose and fix a subalgebra $\mathbf{M}$ of $\mathbf{B}_1\times\cdots\times\mathbf{B}_n$ containing $\mathbf{B}$, which is maximal for the property that it fails to contain c. Then $\mathbf{M}$ is completely $\cap$-irreducible in the lattice of subalgebras of $\mathbf{B}_1\times\cdots\times\mathbf{B}_n$. Let $\{T_1, \dots, T_\ell\}$ be a partition of $[n]$ such that $\mathbf{M}|_{T_i}$ is directly indecomposable for every $i \in [\ell]$, and $\mathbf{M}$ differs from $\mathbf{M}|_{T_1}\times\cdots\times\mathbf{M}|_{T_\ell}$ by a permutation of variables only; we will denote this fact by $\mathbf{M} \approx \mathbf{M}|_{T_1}\times\cdots\times\mathbf{M}|_{T_\ell}$. We must have $c|_T \notin \mathbf{M}|_T$ for at least one block $T := T_i$, because otherwise $\mathbf{M} \approx \mathbf{M}|_{T_1}\times\cdots\times\mathbf{M}|_{T_\ell}$ would imply that $c \in \mathbf{M}$, contradicting the choice of $\mathbf{M}$. Let us fix such a T for the rest of the proof. Note that $|T| > 1$, because $\mathbf{M}$ is a subdirect product of $\mathbf{B}_1, \dots, \mathbf{B}_n$ (as $\mathbf{B} \leq \mathbf{M}$), so we have $c|_j = c_j \in \mathbf{B}_j = \mathbf{M}|_j$ for every one-element set $\{j\} \subseteq [n]$.

It follows from $\mathbf{M} \approx \mathbf{M}|_{T_1} \times \cdots \times \mathbf{M}|_{T_\ell}$ that $\mathbf{M}$ is the intersection of two subalgebras of $\mathbf{B}_1 \times \cdots \times \mathbf{B}_n$, as shown below:

$$\mathbf{M} \approx \Big(\mathbf{M}|_T \times \prod_{j\in[n]\setminus T} \mathbf{B}_j\Big) \cap \Big(\prod_{j\in T} \mathbf{B}_j \times \mathbf{M}|_{[n]\setminus T}\Big).$$

Since $\mathbf{M}$ is a completely $\cap$-irreducible subalgebra of $\mathbf{B}_1 \times \cdots \times \mathbf{B}_n$, we get that $\mathbf{M} \approx \mathbf{M}|_T \times \prod_{j\in[n]\setminus T} \mathbf{B}_j$ or $\mathbf{M} \approx \prod_{j\in T} \mathbf{B}_j \times \mathbf{M}|_{[n]\setminus T}$. The latter is impossible, because $\mathbf{M}|_T \lneq \prod_{j\in T} \mathbf{B}_j$ (as witnessed by $c|_T$). Hence, for the subalgebra $\mathbf{R} := \mathbf{M}|_T$ of $\prod_{j\in T} \mathbf{B}_j$ we get that $\mathbf{M} \approx \mathbf{R} \times \prod_{j\in[n]\setminus T} \mathbf{B}_j$. Furthermore, $\mathbf{R}$ is both $\cap$-irreducible (because $\mathbf{M}$ is) and directly indecomposable (by construction), therefore $\mathbf{R}$ is a critical subalgebra of $\prod_{j\in T} \mathbf{B}_j$. Our construction also implies that $\mathbf{B}|_T \leq \mathbf{R}$ and $c|_T \notin \mathbf{R}$. Hence, in particular, $c|_T \notin \mathbf{B}|_T$. Thus, our assumption that $c|_I \in \mathbf{B}|_I$ holds for all $I \subseteq [n]$ with $|I| < \max\{\mathsf{d}, 3\}$ forces that $|T| \geq \max\{\mathsf{d}, 3\}$.

Now we can apply Theorem 5.1 to the variety $\mathcal{V}$, the algebras $\mathbf{B}_t$ $(t \in T)$ in $\mathcal{V}$, and the critical subalgebra $\mathbf{R}$ of $\prod_{t\in T} \mathbf{B}_t$ where the number of factors in the product is $|T| \geq \max\{\mathsf{d}, 3\}$. Since $\mathbf{R}$ is a subdirect product of the algebras $\mathbf{B}_t$ $(t \in T)$, the reduced representation $\overline{\mathbf{R}}$ of $\mathbf{R}$ is the quotient algebra $\mathbf{R}/(\delta{\restriction}_{\mathbf{R}})$ where $\delta = \prod_{t\in T} \delta_t$ $(\delta_t \in \mathrm{Con}(\mathbf{B}_t))$ is the largest product congruence of $\prod_{t\in T} \mathbf{B}_t$ for which $\mathbf{R}$ is δ-saturated. Let $\overline{\mathbf{B}}_t := \mathbf{B}_t/\delta_t$ for every $t \in T$. So, the conclusions of Theorem 5.1 can be restated as follows:

(1) $\overline{\mathbf{R}} \leq \prod_{t\in T} \overline{\mathbf{B}}_t$ is a representation of $\overline{\mathbf{R}}$ as a subdirect product of subdirectly irreducible algebras $\overline{\mathbf{B}}_t$.

$(6)^\flat$ $\overline{\mathbf{B}}_s$ and $\overline{\mathbf{B}}_t$ are similar for any $s, t \in T$.

(7) Each $\overline{\mathbf{B}}_t$ has abelian monolith μ_t $(t \in T)$.

(8) For the centralizers $\rho_\ell := (0 : \mu_\ell)$ of the monoliths μ_ℓ $(\ell \in T)$, the image of the composite map

$$\overline{\mathbf{R}} \overset{\mathrm{proj}_{st}}{\to} \overline{\mathbf{B}}_s \times \overline{\mathbf{B}}_t \twoheadrightarrow \overline{\mathbf{B}}_s/\rho_s \times \overline{\mathbf{B}}_t/\rho_t \tag{5.2}$$

is the graph of an isomorphism $\overline{\mathbf{B}}_s/\rho_s \to \overline{\mathbf{B}}_t/\rho_t$ for any $s, t \in T$.

By conclusion (1), we have for each $t \in T$ that $\overline{\mathbf{B}}_t = \mathbf{B}_t/\delta_t$ is subdirectly irreducible, so $\delta_t \in \mathrm{Irr}(\mathbf{B}_t)$ and $(t, \delta_t) \in W$. Hence, the algebra $\overline{\mathbf{B}}_t = \mathbf{B}_t/\delta_t$ is one of the subdirect factors of $\widehat{\mathbf{B}}$, namely, $\overline{\mathbf{B}}_t = \mathbf{B}_t/\delta_t = \widehat{\mathbf{B}}_w$ for $w = (t, \delta_t)$. This implies also that μ_t, ρ_t are the congruences of $\overline{\mathbf{B}}_t = \widehat{\mathbf{B}}_w$ that we denoted earlier by μ_w, ρ_w.

Let $\widehat{T} := \{(t, \delta_t) : t \in T\}$. Next we want to show that any two elements of $\widehat{T}$ are related by $\sim$. Let $s, t \in T$, and let $v := (s, \delta_s)$, $w := (t, \delta_t)$. As we noticed in the preceding paragraph, we have that $\overline{\mathbf{B}}_s = \widehat{\mathbf{B}}_v$, $\mu_s = \mu_v$, $\rho_s = \rho_v$, and $\overline{\mathbf{B}}_t = \widehat{\mathbf{B}}_w$, $\mu_t = \mu_w$, $\rho_t = \rho_w$. If $s = t$, then $v = w$, and hence $v \sim w$ holds because $\sim$ is an equivalence relation. So, assume from now on that $s \neq t$. Hence $v \neq w$. In this case, checking whether $v \sim w$ holds involves two conditions. One is that the subdirectly irreducible algebras $\widehat{\mathbf{B}}_v$ and $\widehat{\mathbf{B}}_w$ are similar with abelian monoliths μ_v and μ_w, which follows from conclusions (6)$^\flat$–(7). The other is that the image of $\widehat{\mathbf{B}}|_{vw}$ under the natural map $\psi\colon \widehat{\mathbf{B}}_v \times \widehat{\mathbf{B}}_w \twoheadrightarrow (\widehat{\mathbf{B}}_v/\rho_v) \times (\widehat{\mathbf{B}}_w/\rho_w)$, is the graph of an isomorphism $\widehat{\mathbf{B}}_v/\rho_v \to \widehat{\mathbf{B}}_w/\rho_w$. We will establish this property by proving that $\widetilde{\mathbf{B}}_{vw} := \psi(\widehat{\mathbf{B}}|_{vw})$ is equal to the image of $\overline{\mathbf{R}}$ under the homomorphism in (5.2).

Let φ denote the natural homomorphism $\overline{\mathbf{B}}_s \times \overline{\mathbf{B}}_t \twoheadrightarrow \overline{\mathbf{B}}_s/\rho_s \times \overline{\mathbf{B}}_t/\rho_t$, and let $\widetilde{\mathbf{R}}_{st} := \varphi(\overline{\mathbf{R}}|_{st})$; thus, $\widetilde{\mathbf{R}}_{st}$ is the image of $\overline{\mathbf{R}}$ under the composite map in (5.2). Since $\widehat{\mathbf{B}}_v = \overline{\mathbf{B}}_s$, $\widehat{\mathbf{B}}_w = \overline{\mathbf{B}}_t$, $\rho_v = \rho_s$, and $\rho_w = \rho_r$, we have that $\varphi = \psi$. Therefore $\widetilde{\mathbf{B}}_{vw} = \varphi(\widehat{\mathbf{B}}|_{vw})$. Moreover, since $\widehat{\mathbf{B}}|_{vw}$ is a subdirect product of $\widehat{\mathbf{B}}_v$ and $\widehat{\mathbf{B}}_w$, we get that $\widetilde{\mathbf{B}}_{vw}$ is a subdirect product of $\widehat{\mathbf{B}}_v/\rho_v$ and $\widehat{\mathbf{B}}_w/\rho_w$. By the construction of $\mathbf{R}$, we have that $\mathbf{R} \geq \mathbf{B}|_T$, and $\geq$ is preserved under the natural homomorphism

$$\prod_{r\in T}\mathbf{B}_r \twoheadrightarrow \prod_{r\in T}\mathbf{B}_r/\delta_r = \prod_{u\in\widehat{T}}\widehat{\mathbf{B}}_u. \tag{5.3}$$

The images of $\mathbf{R}$ and $\mathbf{B}|_T$ under this homomorphism are $\overline{\mathbf{R}}$ and $\widehat{\mathbf{B}}|_{\widehat{T}}$, respectively, hence we conclude that $\overline{\mathbf{R}} \geq \widehat{\mathbf{B}}|_{\widehat{T}}$. Projecting further onto the coordinates s,t in T, and the corresponding coordinates $v = (s,\delta_s)$, $w = (t,\delta_t)$ in $\widehat{T}$, we get that $\overline{\mathbf{R}}|_{st} \geq \widehat{\mathbf{B}}|_{vw}$. Hence, it follows that $\widetilde{\mathbf{R}}_{st} = \varphi(\overline{\mathbf{R}}|_{st}) \geq \varphi(\widehat{\mathbf{B}}|_{vw}) = \widetilde{\mathbf{B}}_{vw}$. By conclusion (8) above, $\widetilde{\mathbf{R}}_{st}$ is the graph of an isomorphism $\overline{\mathbf{B}}_s/\rho_s \to \overline{\mathbf{B}}_t/\rho_t$, or equivalently, the graph of an isomorphism $\widehat{\mathbf{B}}_v/\rho_v \to \widehat{\mathbf{B}}_w\rho_w$. Combining this fact with the earlier observation that $\widetilde{\mathbf{B}}_{vw}$ is a subdirect product of $\widehat{\mathbf{B}}_v/\rho_v$ and $\widehat{\mathbf{B}}_w\rho_w$, we obtain that $\widetilde{\mathbf{R}}_{st}$ and $\widetilde{\mathbf{B}}_{vw}$ must be equal. This proves that $\widetilde{\mathbf{B}}_{vw}$ is the graph of an isomorphism $\widehat{\mathbf{B}}_v/\rho_v \to \widehat{\mathbf{B}}_w/\rho_w$, and hence finishes the proof of $v \sim w$.

Our arguments in the last two paragraphs show that $\widehat{T}$ is contained in one of the blocks U of $\sim$. We have $|U| \geq |\widehat{T}| = |T| \geq \max\{\mathsf{d},3\}$. It remains to verify that $\widehat{c}|_U \notin \widehat{\mathbf{B}}|_U$.

Assume, for a contradiction, that $\widehat{c}|_U \in \widehat{\mathbf{B}}|_U$. Then projecting further to $\widehat{T} \subseteq U$ yields that $\widehat{c}|_{\widehat{T}} \in \widehat{\mathbf{B}}|_{\widehat{T}}$. As we saw earlier, $\widehat{\mathbf{B}}|_{\widehat{T}} \leq \overline{\mathbf{R}}$, therefore we get that the tuple $\widehat{c}|_{\widehat{T}} = (c_r/\delta_r)_{r\in T}$ lies in $\overline{\mathbf{R}}$. Hence the tuple $c|_T = (c_r)_{r\in T}$ lies in the full inverse image of $\overline{\mathbf{R}}$ under the natural homomorphism (5.3). This inverse image is $\mathbf{R}$, because $\overline{\mathbf{R}} = \mathbf{R}/(\delta{\restriction}_{\mathbf{R}})$ with $\delta = \prod_{r\in T}\delta_r$, and $\mathbf{R}$ is δ-saturated in $\prod_{r\in T}\mathbf{B}_r$. Thus, we obtain that $c|_T \in \mathbf{R}$, which is impossible, since $\mathbf{R}$ was chosen so that $c|_T \notin \mathbf{R}$. This contradiction proves that $\widehat{c}|_U \notin \widehat{\mathbf{B}}|_U$, and completes the proof of Theorem 5.2. □

6. Algorithms Based on Theorem 5.2

In this section we will assume that $\mathcal{V}$ is a fixed variety (in a finite language) with a d-cube term, and $\mathcal{K}$ is a finite set of finite algebras in $\mathcal{V}$.

Definition 6.1. Let $a_1, \ldots, a_k, b \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ ($\mathbf{A}_1, \ldots, \mathbf{A}_n \in \mathcal{K}$) be an input for SMP($\mathcal{K}$) where $a_r = (a_{r1}, \ldots, a_{rn})$ ($r \in [k]$) and $b = (b_1, \ldots, b_n)$. We call this input d-*coherent* if the following conditions are satisfied:

(i) $n \geq \max\{\mathsf{d}, 3\}$;

(ii) $\mathbf{A}_1, \ldots, \mathbf{A}_n$ are similar subdirectly irreducible algebras, and each $\mathbf{A}_\ell$ has abelian monolith μ_ℓ;

(iii) $b|_I$ is in the subalgebra of $\prod_{i \in I} \mathbf{A}_i$ generated by $\{a_1|_I, \ldots, a_k|_I\}$ for all $I \subseteq [n]$, $|I| < \max\{\mathsf{d}, 3\}$; and

(iv) for the centralizers $\rho_\ell := (0 : \mu_\ell)$ of the monoliths μ_ℓ, the subalgebra of $\mathbf{A}_i/\rho_i \times \mathbf{A}_j/\rho_j$ generated by $\{(a_{1i}/\rho_i, a_{1j}/\rho_j), \ldots, (a_{ki}/\rho_i, a_{kj}/\rho_j)\}$ is the graph of an isomorphism $\mathbf{A}_i/\rho_i \to \mathbf{A}_j/\rho_j$ for any $i, j \in [n]$.

Definition 6.2. We define $\mathrm{SMP}_{\mathsf{d}\text{-coh}}(\mathcal{K})$ to be the restriction of SMP($\mathcal{K}$) to d-coherent inputs.

It is clear from Definition 6.1 that d-coherence for inputs of SMP($\mathcal{K}$) can be checked in polynomial time.

Theorem 6.3. *If $\mathcal{V}$ is a variety in a finite language with a d-cube term, then the decision problems* SMP($\mathcal{K}$) *and* $\mathrm{SMP}_{\mathsf{d}\text{-}coh}(\mathbb{HS}\mathcal{K})$ *are polynomial time equivalent for every finite family $\mathcal{K}$ of finite algebras in $\mathcal{V}$.*

Proof. By Theorem 4.8, SMP($\mathcal{K}$) is polynomial time equivalent to SMP($\mathbb{HS}\mathcal{K}$). Clearly, $\mathrm{SMP}_{\mathsf{d}\text{-coh}}(\mathbb{HS}\mathcal{K})$ is polynomial time reducible to SMP($\mathbb{HS}\mathcal{K}$), because it is a subproblem of SMP($\mathbb{HS}\mathcal{K}$). Therefore we will be done if we show that Algorithm 6 reduces SMP($\mathbb{HS}\mathcal{K}$) to $\mathrm{SMP}_{\mathsf{d}\text{-coh}}(\mathbb{HS}\mathcal{K})$ in polynomial time.

The correctness of Algorithm 6 is based on Theorem 5.2. Let $\mathbf{B}$ denote the subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by $b_1, \ldots, b_k$. In steps 1–2, the projections $\mathbf{B}|_I$ of $\mathbf{B}$ are computed for all $I \in \binom{[n]}{\mathsf{d}^-}$ if $n > \mathsf{d}^- := \max(\mathsf{d} - 1, 2)$, and for $I = [n]$ if $n \leq \mathsf{d}^-$. Furthermore, it is checked whether $b_{k+1}|_I \in \mathbf{B}|_I$ holds for all such I. If $b_{k+1}|_I \notin \mathbf{B}|_I$ for one of these I's, then clearly $b_{k+1} \notin \mathbf{B}$, so in this case the algorithm correctly returns the answer 'NO' in step 2.2. In steps 3–4 the algebras $\mathbf{B}_i := \mathbf{B}|_i$ are computed for every $i \in [n]$, and every coordinate i with $|B_i| = 1$ is omitted from the input tuples (but the earlier notation is kept for simplicity). Note that since b_{k+1} passed the tests in step 2.2, we had $b_{k+1}|_i \in \mathbf{B}_i$ for every deleted coordinate i. Therefore, deletion of the trivial coordinates does no affect whether or not $b_{k+1} \in \mathbf{B}$. It follows that by the end of step 4 we have that

(1) $\mathbf{B}$ is a subdirect subalgebra of $\mathbf{B}_1 \times \cdots \times \mathbf{B}_i$ where $\mathbf{B}_1, \ldots, \mathbf{B}_n$ are nontrivial, and

(2) $b_{k+1}|_I \in \mathbf{B}|_I$ for all $I \in \binom{[n]}{\mathsf{d}^-}$ if $n > \mathsf{d}^- := \max(\mathsf{d} - 1, 2)$, and for $I = [n]$ if $n \leq \mathsf{d}^-$.

In the latter case $\mathbf{B}|_I = \mathbf{B}|_{[n]} = \mathbf{B}$, so in step 5 the algorithm correctly returns the answer 'YES'. Thus, we may assume from now on that $n > \mathsf{d}^-$.

Algorithm 6: Reduction of $\mathrm{SMP}(\mathcal{K})$ to $\mathrm{SMP}_{\mathsf{d}\text{-coh}}(\mathbb{H}\mathbb{S}\mathcal{K})$

```
Input: $b_1, \dots, b_k, b_{k+1} \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ $(\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K})$
Question: Is $b_{k+1}$ in the subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by $b_1, \dots, b_k$?
    1. $\mathsf{d}^- = \max(\mathsf{d}-1, 2)$, $\mathsf{d}^* = \min(\mathsf{d}^-, n)$
    2. for $I \in \binom{[n]}{\mathsf{d}^*}$ do
        2.1. generate $\mathbf{B}|_I$ by $b_1|_I, \dots, b_k|_I$
        2.2. if $b_{k+1}|_I \notin \mathbf{B}|_I$ then
                return NO
             end if
       end for
    3. for $i \in [n]$ do
          $\mathbf{B}_i = \langle b_1|_i, \dots, b_k|_i \rangle$
       end for
    4. omit all coordinates $i$ from the input for which $|\mathbf{B}_i| = 1$ (but keep earlier notation)
    5. if $n \le \mathsf{d}^-$ then
          return YES
       end if
    6. $W := \{(j, \sigma) : j \in [n], \sigma \in \mathrm{Irr}(\mathbf{B}_j)\}$
    7. for $i = 1, \dots, k+1$ do
          $\widehat{b}_i = (\widehat{b}_{iw})_{w \in W}$ where for each $w = (j, \sigma)$, $\widehat{b}_{iw} = b_i|_j/\sigma$
       end for
    8. for $w = (j, \sigma) \in W$ do
          $\widehat{\mathbf{B}}_w = \mathbf{B}_j/\sigma$,  $\mu_w =$ monolith of $\widehat{\mathbf{B}}_w$,  $\rho_w = (0 : \mu_w)$
       end for
    9. for distinct $v, w \in W$ do
          $\widetilde{\mathbf{B}}_{vw} = \langle (\widehat{b}_{iv}/\rho_v, \widehat{b}_{iw}/\rho_w) : i \in [k] \rangle$
       end for
   10. compute the equivalence relation $\sim$ on $W$ determined by the following condition:
       for distinct $v, w \in W$,
          $v \sim w \quad \Leftrightarrow \quad \begin{cases} \mu_v \le \rho_v,\ \mu_w \le \rho_w,\ \widehat{B}_v \text{ and } \widehat{B}_w \text{ are similar, and} \\ \widetilde{\mathbf{B}}_{vw} \text{ is the graph of an isomorphism } \widehat{\mathbf{B}}_v/\rho_v \to \widehat{\mathbf{B}}_w/\rho_w \end{cases}$
   11. find the equivalence classes $E_1, \dots, E_\kappa$ of $\sim$ of size $> \mathsf{d}^-$
   12. if $\kappa > 0$ then
      12.1. for $\lambda = 1, \dots, \kappa$ do
          12.1.1 run $\mathrm{SMP}_{\mathsf{d}\text{-coh}}(\mathbb{H}\mathbb{S}\mathcal{K})$ with input $\widehat{b}_1|_{E\lambda}, \dots, \widehat{b}_k|_{E\lambda}, \widehat{b}_{k+1}|_{E\lambda} \in \prod_{w \in E_\lambda} \widehat{\mathbf{B}}_w$
                 to get answer $\mathsf{A} \in \{\mathrm{YES}, \mathrm{NO}\}$
         12.1.2. if $\mathsf{A} = \mathrm{NO}$ then
                 return NO
                 end if
             end for
       end if
   13. return YES
```

Table 6

(1) implies that the algebra $\mathbf{B}$ satisfies the assumptions of Theorem 5.2, and (2) shows that the tuple $c := b_{k+1} \in \mathbf{B}_1 \times \cdots \times \mathbf{B}_n$ satisfies the first part of condition (b) in Theorem 5.2.

In steps 6–11, Algorithm 6 computes the data needed to check whether or not $c = b_{k+1}$ satisfies the second part of condition (b) as well. In more detail, the algorithm computes the set W, the images $\widehat{b_i}$ of the input tuples b_i ($i \in [k+1]$) under the homomorphism $\widehat{\ }$, the subdirectly irreducible algebras $\widehat{B}_w$ ($w \in W$), their congruences μ_w (the monolith) and $\rho_w = (0 : \mu_w)$, and finally, the equivalence relation $\sim$ on W, and its equivalence classes $E_1, \ldots, E_\kappa$ of size $> \mathsf{d}^-$. All computations follow the definitions exactly, except step 10. We explain now why the conditions used in step 10 to compute $\sim$ are equivalent to the conditions in the definition of $\sim$ (stated right before Theorem 5.2):

- The inclusion $\mu_w \leq \rho_w$ ($w \in W$) is equivalent to the condition that μ_w is abelian, because the inclusion is true if μ_w is abelian, and $\rho_w = 0$ (and hence the inclusion fails) if μ_w is nonabelian.
- As discussed in the proof of Theorem 5.2, $\widetilde{B}_{vw}$ is the image of $\widehat{B}|_{vw}$ under the natural map $\widehat{\mathbf{B}}_v \times \widehat{\mathbf{B}}_w \twoheadrightarrow (\widehat{\mathbf{B}}_v/\rho_v) \times (\widehat{\mathbf{B}}_w/\rho_w)$.

Since $c = b_{k+1}$ satisfies the first part of condition (b) in Theorem 5.2, it follows from the equivalence of conditions (a) and (b) in Theorem 5.2 that $b_{k+1} \in \mathbf{B}$ if and only if $c = b_{k+1}$ satisfies the second part of condition (b) as well, that is,

(3) $\widehat{b}_{k+1}|_{E_\lambda} \in \widehat{\mathbf{B}}|_{E_\lambda}$ for all $\lambda \in [\kappa]$.

This is clearly equivalent to the condition that

(3)′ $\widehat{b}_{k+1}|_{E_\lambda}$ belongs to the subalgebra of $\prod_{w \in E_\lambda} \widehat{B}_w$ generated by the tuples $\widehat{b}_1|_{E_\lambda}, \ldots, \widehat{b}_k|_{E_\lambda}$ for all $\lambda \in [\kappa]$.

It follows from the construction that for each $\lambda \in [\kappa]$, $\widehat{b}_1|_{E_\lambda}, \ldots, \widehat{b}_k|_{E_\lambda}, \widehat{b}_{k+1}|_{E_\lambda}$ is a d-coherent input for $\mathrm{SMP}(\mathbb{HS}\mathcal{K})$, so condition (3)′ can be checked using $\mathrm{SMP}_{\mathsf{d\text{-}coh}}(\mathbb{HS}\mathcal{K})$. This is exactly what Algorithm 6 does in steps 12–13, and returns the correct answer: 'YES' if (3)′ holds (including the case when $\kappa = 0$), and 'NO' otherwise. This completes the proof of the correctness of Algorithm 6.

Now we show that Algorithm 6 reduces $\mathrm{SMP}(\mathcal{K})$ to $\mathrm{SMP}_{\mathsf{d\text{-}coh}}(\mathbb{HS}\mathcal{K})$ in polynomial time. Clearly, steps 1, 5, and 13 require constant time. Since each $\mathbf{B}_i$ ($i \in [n]$) is a subalgebra of some member of $\mathcal{K}$, we have $|\mathbf{B}_i| \leq \mathsf{a}_{\mathcal{K}}$ where the constant $\mathsf{a}_{\mathcal{K}}$ is independent of the input. The parameter d is also independent of the input, and so is $\mathsf{s} := \max\{|\operatorname{Irr}(\mathbf{A})| : \mathbf{A} \in \mathbb{S}\mathcal{K}\}$. It follows that $|W| \leq n\mathsf{s}$ and that each iteration of the 'for' cycles in steps 2, 3, 8, and 9 require constant time. Hence, steps 2, 3, 4, 6, 8, and 9 run in $O(n^{\mathsf{d}-1})$, $O(n)$, $O(n)$, $O(n)$, $O(n)$, $O(n^2)$ time, respectively. In step 7 the 'for' cycle is iterated $k+1$ times, and each iteration requires $O(n)$ time, so the total run time of step 7 is $O(nk)$.

In step 10, to determine whether $v \sim w$ holds for a particular pair of elements $v, w \in W$ requires constant time, because the condition only involves data on algebras in $\mathbb{HS}\mathcal{K}$ and on products of two such algebras. (In particular, recall from Section 2 that similarity of $\widehat{B}_v$ and $\widehat{B}_w$ can be checked by looking at congruences of subalgebras of $\widehat{B}_v \times \widehat{B}_w$.) Thus, step 10 runs in $O(n^2)$ time. Step 11 requires no more that $O(n^2)$ time. Since $E_1, \dots, E_\kappa$ are disjoint subsets of W and $|W| \le n\mathsf{s}$, we get that $\kappa \le n\mathsf{s}$ and each E_λ has size $|E_\lambda| \le n\mathsf{s}$. Thus, in step 12, $\mathrm{SMP}_{\mathsf{d\text{-}coh}}(\mathbb{HS}\mathcal{K})$ has to be run at most $O(n)$ times, and the input size of each run is $O(nk)$, approximately the same as the size of the original input.

This proves that Algorithm 6 reduces $\mathrm{SMP}(\mathcal{K})$ to $\mathrm{SMP}_{\mathsf{d\text{-}coh}}(\mathbb{HS}\mathcal{K})$ in polynomial time. □

Theorem 6.4. *If $\mathcal{V}$ is a residually small variety in a finite language with a* d*-cube term, then for every finite set $\mathcal{K}$ of finite algebras in $\mathcal{V}$,*

$$\mathrm{SMP}(\mathcal{K}) \in \mathsf{P}.$$

Proof. We will show that, under the assumption of the theorem, Algorithm 7 solves $\mathrm{SMP}(\mathcal{K})$ in polynomial time.

First we discuss the correctess of Algorithm 7. Let $a_1, \dots, a_k, b \in \mathbf{A}_1 \times \dots \mathbf{A}_n$ ($\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$) be a correct input for $\mathrm{SMP}_{\mathsf{d\text{-}coh}}(\mathcal{K})$ (i.e., a d-coherent input for $\mathrm{SMP}(\mathcal{K})$). Then conditions (i)–(iv) in Definition 6.1 hold. By condition (ii), the algebras $\mathbf{A}_j$ ($j \in [n]$) are subdirectly irreducible with abelian monoliths, so in step 1 of Algorithm 7 the monoliths μ_j and their centralizers ρ_j will be found. Moreover, since $\mathcal{K}$ is assumed to be in a residually small variety, we get from ... (CM) and ... (char.of.RS) that

(†) ρ_j is an abelian congruence of $\mathbf{A}_j$ for every $j \in [n]$.

Let $\mathbf{B}$ denote the subalgebra of $\mathbf{A}_1 \times \dots \times \mathbf{A}_n$ generated by the input tuples $a_1, \dots, a_k$, and let ρ denote the product congruence $\rho_1 \times \dots \times \rho_n$ on $\mathbf{A}_1 \times \dots \times \mathbf{A}_n$. The restriction of ρ to $\mathbf{B}$ will be denoted by $\rho_{\mathbf{B}}$. Condition (iv) in Definition 6.1 implies that

(‡) the map $\mathbf{B}/\rho_{\mathbf{B}} \to \mathbf{A}_j/\rho_j$, $(x_1, \dots, x_n)/\rho_{\mathbf{B}} \mapsto x_j/\rho_j$ is a bijection for every $j \in [n]$.

Finally, conditions (i) and (iii) together imply that for the input tuple b we have $b|_I \in \mathbf{B}|_I$ for all sets $I \in \binom{[n]}{2}$. Hence, we have $b|_{i,j}/(\rho_i \times \rho_j) \in \mathbf{B}|_{i,j}/(\rho_i \times \rho_j)$ for all $i, j \in [n]$. Since, by condition (iv), $\mathbf{B}|_{i,j}/(\rho_i \times \rho_j)$ is the graph of an isomorphism $\mathbf{A}_i/\rho_i \to \mathbf{A}_j/\rho_j$ for every pair $i, j \in [n]$, it follows that the tuple $b/\rho = (b|_1/\rho_1, \dots, b|_n/\rho_n)$ belongs to $\mathbf{B}/\rho_{\mathbf{B}}$. Hence,

($*$) b is an element of the algebra $\mathbf{B}[\rho]$, the ρ-saturation of $\mathbf{B}$.

Let's return to the analysis of Algorithm 7. After appropriately reindexing $a_1, \dots, a_k$, Steps 2–3 produce a subset $\mathcal{O}$ of $\mathbf{B}$ such that the first coordinates of the tuples in $\mathcal{O}$

Algorithm 7: For $\mathrm{SMP}_{\mathsf{d\text{-}coh}}(\mathbb{HS}\mathcal{K})$ if $\mathcal{K}$ is in a RS variety

Input: d-coherent $a_1, \dots, a_k, b \in \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ ($\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$) where $a_i = (a_{i1}, \dots, a_{in})$ for all $i \in [k]$

Question: Is b in the subalgebra of $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by $a_1, \dots, a_k$?

```
1. for $j \in [n]$ do
   $\mu_j$ = monolith of $\mathbf{A}_j$,  $\rho_j = (0 : \mu_j)$
   end for
2. reindex $a_1, \dots, a_k$ such that $a_{11}/\rho_1, \dots, a_{r1}/\rho_1$ are pairwise distinct and
   $\{a_{11}/\rho_1, \dots, a_{r1}/\rho_1\} = \{a_{11}/\rho_1, \dots, a_{k1}/\rho_1\}$
   let $\mathcal{O} = \{a_1, \dots, a_r\}$
3. generate $\mathbf{A}_1/\rho_1$ by $a_{11}/\rho_1, \dots, a_{r1}/\rho_1$, and simultaneously,
   3.1. for each new $a/\rho_1 = t(a_{11}/\rho_1, \dots, a_{r1}/\rho_1) \in \mathbf{A}_1/\rho_1$ do
          3.1.1 $\mathcal{O} = \mathcal{O} \cup \{t(a_{i_1}, \dots, a_{i_r})\}$
        end for
4. find the equivalence relation $\equiv$ on $[n]$ defined by
        $s \equiv t \quad \Leftrightarrow \quad \mathbf{A}_s = \mathbf{A}_t$ and $o|_s = o|_t$ for all $o \in \mathcal{O} \qquad (s, t \in [n])$
   let $T$ be a transversal for the blocks of $\equiv$, and let $\mathbf{A}_T = \prod_{j \in T} \mathbf{A}_j$
5. compute the subalgebra $P$ of $\mathbf{A}_T^{A_T}$ generated by the identity function $A_T \to A_T$
   and by the constant functions with value $o|_T$ ($o \in \mathcal{O}$) (so $P$ is a set of functions
   $A_T \to A_T$)
6. find $o \in \mathcal{O}$ such that $b \in o/\rho$; let $o = (o_1, \dots, o_n)$
7. $H = \emptyset$
8. for $p \in P$ do
   8.1 for $t \in T$ do
         8.1.1 compute the function $p_t \colon \mathbf{A}_t \to \mathbf{A}_t$, $x \mapsto p(\breve{x})|_t$ where $\breve{x}|_t = x$ and
               $\breve{x}|_s = o|_s$ for all $s \in T \setminus \{t\}$
       end for
   8.2 for $c \in \{a_1, \dots, a_k\}$ do
         8.2.1 $d = ()$
         8.2.2 for $j \in [n]$ do
                 8.2.2.1 find $t \in T$ with $t \equiv j$
                         let $d_j = p_t(c|_j)$
                 8.2.2.2 $d = (d, d_j)$
               end for
         8.2.3 if $d \in o/\rho$ then $H = H \cup \{d\}$
               end if
       end for
   end for
9. run Sims' algorithm for $\mathrm{SMP}(\mathcal{G})$ with the input $H \cup \{b\} \subseteq \mathbf{G}_1 \times \cdots \times \mathbf{G}_n$ where
   $\mathbf{G}_j$ is the group $(o_j/\rho_j; +_{o_j}, -_{o_j}, o_j)$ for each $j \in [n]$ and $\mathcal{G}$ is the family of all
   induced abelian groups on blocks of abelian congruences of algebras in $\mathcal{K}$; get
   answer $\mathsf{A} \in \{\mathrm{YES}, \mathrm{NO}\}$
10. return $\mathsf{A}$
```

Table 7

form a transversal for the ρ_1-classes of $\mathbf{A}_1$. Thus, it follows from (‡) that the tuples in $\mathcal{O}$ form a transversal for the $\rho_{\mathbf{B}}$-classes of $\mathbf{B}$ (and hence also for the ρ-classes of $\mathbf{B}[\rho]$). Let $|\mathcal{O}| = \ell$. Since $|\mathcal{O}| = |\mathcal{O}|_1$, we have $\ell \leq |\mathbf{A}_1/\rho_1| < |A_1|$.

Now let $\equiv$, T, and $\mathbf{A}_T$ be as defined (and computed) in step 4. It is easy to see that the set P of functions $A_T \to A_T$ computed in step 5 is

$$P = \{\mathsf{t}^{\mathbf{A}_T}(x, \mathcal{O}|_T) : \mathsf{t} \text{ is a } (1+\ell)\text{-ary term}\} \tag{6.1}$$

where we assume that an ordering of $\mathcal{O}$ has been fixed to ensure that its elements are always substituted into terms in that fixed order. Our observation $(*)$ shows that in step 6, Algorithm 7 will find an element $o \in \mathcal{O}$ such that $b \in o/\rho$.

To simplify notation on our discussion of steps 7–8 let $\mathbf{A} := \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$.

Claim 6.5. *The set H obtained by Algorithm 7 after completing steps 7–8 is*

$$H = \{\mathsf{t}^{\mathbf{A}}(a_i, \mathcal{O}) \in o/\rho : i \in [k],\ \mathsf{t} \textit{ is a } (1+\ell)\textit{-ary term}\}. \tag{6.2}$$

Proof of Claim 6.5. We will use the notation in step 8. Let $p \in P$ and $c \in \{a_1, \ldots, a_k\}$, say $c = (c_1, \ldots, c_n)$. Then, by (6.1), $p \in P$ if and only if p is a unary polynomial operation of $\mathbf{A}_T$ of the form $\mathsf{t}^{\mathbf{A}_T}(x, \mathcal{O}|_T)$ for some $(1+\ell)$-ary term t. Our goal is to show that for such p and c the tuple d computed in steps 8.1.1-8.1.2 is

$$d = \mathsf{t}^{\mathbf{A}}(c, \mathcal{O}). \tag{6.3}$$

This will prove that the tuples d computed in steps 8.1.1-8.1.2 are exactly the elements of $\mathbf{A}$ of the form $\mathsf{t}^{\mathbf{A}}(c, \mathcal{O})$ where $c \in \{a_1, \ldots, a_k\}$ and t is a $(1+\ell)$-ary term. Since such a tuple d is added to H in step 8.1.3 if and only if d also satisfies $d \in o/\rho$, the equality (6.2) will follow.

To verify (6.3) let $j \in [n]$ and let $t \in T$ be the unique transversal element such that $t \equiv j$. By the definition of $\equiv$ we have that and $o'|_t = o'|_j$ for all $o' \in \mathcal{O}$. The latter condition may be written as $\mathcal{O}|_t = \mathcal{O}|_j$ (with the fixed ordering of $\mathcal{O}$ in mind, these are tuples of elements in $\mathbf{A}_t = \mathbf{A}_j$). The function $p_t \colon \mathbf{A}_t \to \mathbf{A}_t$ computed in step 8.1 assigns to every $x \in \mathbf{A}_t$ the element

$$p_t(x) = p(\breve{x})|_t = \mathsf{t}^{\mathbf{A}_T}(\breve{x}, \mathcal{O}|_T)|_t = \mathsf{t}^{\mathbf{A}_t}(\breve{x}|_t, \mathcal{O}|_t) = \mathsf{t}^{\mathbf{A}_t}(x, \mathcal{O}|_t).$$

So, p_t is the polynomial function $\mathsf{t}^{\mathbf{A}_t}(x, \mathcal{O}|_t)$ of $\mathbf{A}_t$. Thus, using the equalities $\mathbf{A}_t = \mathbf{A}_j$ and $\mathcal{O}|_t = \mathcal{O}|_j$ we get that

$$d_j = p_t(c_j) = \mathsf{t}^{\mathbf{A}_t}(c_j, \mathcal{O}|_t) = \mathsf{t}^{\mathbf{A}_j}(c_j, \mathcal{O}|_j) = \mathsf{t}^{\mathbf{A}}(c, \mathcal{O})|_j.$$

This holds for every $j \in [n]$, so the proof of (6.3), and hence the proof of Claim 6.5, is complete. ◇

To establish the correctness of the last two steps of Algorithm 7 recall from (†) that ρ_j is an abelian congruence of $\mathbf{A}_j$ for every $j \in [n]$. Therefore, $\rho = \rho_1 \times \cdots \times \rho_n$ is an abelian congruence of $\mathbf{A} = \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$. Hence, by ...(see prelim), there is an induced abelian group $\mathbf{G} = (o/\rho; +_0, -_o, o)$ on the ρ-class o/ρ. Moreover, since ρ

is the product congruence $\rho_1 \times \cdots \times \rho_n$ of $\mathbf{A}$, we get that $\mathbf{G} = \mathbf{G} \times \cdots \times \mathbf{G}_n$ where $\mathbf{G}_j$ is the group $(o_j/\rho_j; +_{o_j}, -_{o_j}, o_j)$ for every $j \in [n]$. Recall also that $b \in o/\rho$, that is, $b \in \mathbf{G} = \mathbf{G}_1 \times \cdots \times \mathbf{G}_n$.

Claim 6.6. *The following conditions on b are equivalent:*

(a) *b is in the subalgebra $\mathbf{B}$ of $\mathbf{A} = \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ generated by $\{a_1, \dots, a_k\}$;*
(b) *b is in the subgroup of $\mathbf{G} = \mathbf{G}_1 \times \cdots \times \mathbf{G}_n$ generated by the set H.*

Proof of Claim 6.6. To prove the implication (a) ⇒ (b) assume that $b \in \mathbf{B}$, that is, $b = \mathsf{g}^{\mathbf{B}}(a_1, \dots, a_k)$ for some k-ary term g. For each $i \in [k]$, let $o^{(i)}$ denote the unique element of $\mathcal{O}$ in the $\rho_{\mathbf{B}}$-class of a_i. Since $\mathsf{g}^{\mathbf{B}}(a_1, \dots, a_k) = b \in o/\rho$, it follows from ... (term ops are linear) that

$$\mathsf{g}^{\mathbf{B}}(a_1, \dots, a_k) = \mathsf{g}^{\mathbf{B}}(a_1, o^{(2)} \dots, o^{(k)}) +_o \mathsf{g}^{\mathbf{B}}(o^{(1)}, a_2, o^{(3)} \dots, o^{(k)}) +_o \cdots \\ +_o \mathsf{g}^{\mathbf{B}}(o^{(1)}, \dots, o^{(k-1)}, a_k) -_o (k-1)\mathsf{g}^{\mathbf{B}}(o^{(1)}, \dots, o^{(k-1)}, o^{(k)}).$$

All $+_o$-summands on the right hand side belong to H, therefore b is in the subgroup of $\mathbf{G} = \mathbf{G}_1 \times \cdots \times \mathbf{G}_n$ generated by H.

For the reverse implication (b) ⇒ (a) notice first that $H \subseteq \mathbf{B}$, because the elements of H are obtained from $a_1, \dots, a_k \in \mathbf{B}$ by unary polynomials that are obtained from term operations using parameters from $\mathcal{O}$ only, and $\mathcal{O} \subseteq \mathbf{B}$. Since the group operations $+_o$, $-_o$, o of $\mathbf{G}$ are also polynomial operations of $\mathbf{B}$ obtained from term operations using parameters from $\mathcal{O}$ only, we get that the subgroup of $\mathbf{G}$ generated by H is contained in $\mathbf{B}$. ◇

As in step 8 of Algorithm 7, let $\mathcal{G}$ denote the set of all induced abelian groups on blocks of abelian congruences of algebras in $\mathcal{K}$. Then, clearly, $\mathbf{G}_1, \dots, \mathbf{G}_n \in \mathcal{G}$. Therefore, Claim 6.6 shows that SMP($\mathcal{K}$) run with the input $a_1, \dots, a_k, b$ in $\mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ ($\mathbf{A}_1, \dots, \mathbf{A}_n \in \mathcal{K}$) has the same answer as SMP($\mathcal{G}$) run with the input H and b in $\mathbf{G}_1 \times \cdots \times \mathbf{G}_n$ ($\mathbf{G}_1, \dots, \mathbf{G}_n \in \mathcal{G}$). Hence, Algorithm 7 finds the correct answer in steps 8–9, so the proof of the correctness of Algorithm 7 is complete.

To prove that Algorithm 7 runs in polynomial time, we will estimate the time complexity of each steps separately. Recall that $\mathsf{a}_{\mathcal{K}}$ denotes the maximum size of an algebra in $\mathcal{K}$.

Each one of Steps 1–3 runs in $O(n)$ time. In step 4, $\equiv$ and a transversal T for the blocks of $\equiv$ can be found in $O(n^2)$ time, since $\ell = |\mathcal{O}|$ is bounded above by a constant ($\leq |A_1| - 1 \leq \mathsf{a}_{\mathcal{K}}$) which is independent of the size of the input. Since $\equiv$ is the kernel of the map $[n] \to \mathcal{K} \times \bigcup_{j \in [n]} A_j^\ell$, $j \mapsto (\mathbf{A}_j, \mathcal{O}|_j)$, the number $|T|$ of the $\equiv$-blocks is at most $|\mathcal{K}|\mathsf{a}_{\mathcal{K}}^\ell$, which is a constant, independent of the input. Therefore, $\mathbf{A}_T$ can be computed in constant time, so step 4 altogether requires $O(n^2)$ time. For the same reason, $|\mathbf{A}_T^{A_T}|$ is also bounded above by a constant, independent of the input, therefore step 5 runs in constant time. Step 6 also runs in constant time, because, in view of (‡), $b \in o/\rho$ is equivalent to $b|_1 \in o|_1/\rho_1$. Clearly, step 7 also runs in constant time.

Using the previous estimates on $|P|$ ($\leq |\mathbf{A}_T^{A_T}|$) and $|T|$ we see that the number of iterations of the outer 'for' cycle (line 8) and the 'for' cycle on line 8.1 is bounded above by a constant. Step 8.1.1 also needs constant time only, therefore step 8.1 runs in constant time. In step 8.2 the outer 'for' cycle is iterated k times; in each iteration steps 8.2.1 and 8.2.3 require constant time, while in step 8.2.2 the 'for' cycle is repeated n times, and each time the computation requires constant time. Thus, step 8 runs in $O(nk)$ time.

In step 8, at most one element is added to H for each choice of $p \in P$ and $c \in \{a_1, \dots, a_k\}$. Since $|P|$ is bounded above by a constant independent of the input, we get that $|H|$ has size $O(k)$. Thus, in step 9, the size of the input $H \cup \{b\}$ for $\mathrm{SMP}(\mathcal{G})$ is $O(nk)$. Moreover, the size of each group in $\mathcal{G}$ is $\leq \mathsf{a}_{\mathcal{K}}$. Since Sims' algorithm for $\mathrm{SMP}(\mathcal{G})$ runs in $O(n^3k)$ time [check this!!!] on an input $H \cup \{b\} \subseteq \mathbf{G}_1 \times \cdots \times \mathbf{G}_n$ with $|H| = O(k)$, we get that step 9 of Algorithm 7 requires $O(n^3k)$ time. Clearly, step 10 runs in constant time.

Combining the time complexities of steps 1–10 we get that Algorithm 7 runs in $O(n^3k)$ time. This competes the proof of Theorem 6.4. □

References

[1] Baker, Kirby A.; Pixley, Alden F. *Polynomial interpolation and the Chinese remainder theorem for algebraic systems.* Math. Z. **143** (1975), no. 2, 165–174.

[2] 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.

[3] Bulatov, Andrei; Mayr, Peter; Steindl, Markus *The subpower membership problem for semigroups.* submitted.

[4] Dent, Topaz; Kearnes, Keith A.; Szendrei, Ágnes *An easy test for congruence modularity.* Algebra Universalis **67** (2012), no. 4, 375–392.

[5] Dalmau, Victor; Jeavons, Peter *Learnability of quantified formulas.* Theoret. Comput. Sci. **306** (2003), 485–511.

[6] Freese, Ralph; McKenzie, Ralph *Commutator theory for congruence modular varieties.* London Mathematical Society Lecture Note Series, **125**. Cambridge University Press, Cambridge, 1987.

[7] Furst, Merrick; Hopcroft, John; Luks, Eugene *Polynomial-time algorithms for permutation groups.* 21st Annual Symposium on Foundations of Computer Science (Syracuse, N.Y., 1980), pp. 36–41, IEEE, New York, 1980.

[8] Hobby, David; McKenzie, Ralph *The structure of finite algebras.* Contemporary Mathematics, **76**. American Mathematical Society, Providence, RI, 1988.

[9] Idziak, Paweł; Marković, Petar; McKenzie, Ralph; Valeriote, Matthew; Willard, Ross *Tractability and learnability arising from algebras with few subpowers.* SIAM J. Comput. **39** (2010), no. 7, 3023–3037.

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

[11] Kozik, Marcin *A finite set of functions with an EXPTIME-complete composition problem.* Theoretical Computer Science **407** (2008), 330–341.

[12] Steindl, Markus *Computational complexity of the subpower membership problem for semigroups.* Ph.D. Dissertation, Johannes Kepler Universität Linz, Austria, 2015.

[13] Willard, Ross *Four unsolved problems in congruence permutable varieties.* Talk at the Conference on Order, Algebra, and Logics, Nashville, 2007.

(Andrei Bulatov) School of Computing Science, Simon Fraser University, Burnaby BC, Canada V5A 1S6

E-mail address: abulatov@sfu.ca

(Peter Mayr) Institut für Algebra, Johannes Kepler Universität Linz, 4040 Linz, Austria, and Department of Mathematics, University of Colorado, Boulder CO, USA, 80309-0395

E-mail address: peter.mayr@jku.at, Peter.Mayr@Colorado.EDU

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

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