The Subpower Membership Problem of 2-Nilpotent Algebras

Michael Kompatscher

Department of Algebra, Faculty of Mathematics and Physics,
Charles University, Prague, Czech Republic

Abstract

The subpower membership problem SMP(**A**) of a finite algebraic structure **A** asks whether a given partial function from A^n to A can be interpolated by a term operation of **A**, or not. While this problem can be EXPTIME-complete in general, Willard asked whether it is always solvable in polynomial time if **A** is a Mal'tsev algebra. In particular, this includes many important structures studied in abstract algebra, such as groups, quasigroups, rings, Boolean algebras. In this paper we give an affirmative answer to Willard's question for a big class of 2-nilpotent Mal'tsev algebras. We furthermore develop tools that might be essential in answering the question for general nilpotent Mal'tsev algebras in the future.

2012 ACM Subject Classification Theory of computation → Algebraic complexity theory

Keywords and phrases subpower membership problem, Mal'tsev algebra, compact representation, nilpotence, clonoids

Digital Object Identifier 10.4230/LIPIcs.STACS.2024.46

Funding This paper was supported by projects PRIMUS/24/SCI/008 and UNCE/SCI/022 of Charles University.

Acknowledgements I would like to thank Peter Mayr for introducing me to difference clonoids and giving several helpful comments on earlier versions of this paper. Furthermore I would like to thank the anonymous referees for their many helpful remarks.

1 Introduction

It is a recurring and well-studied problem in algebra to describe the closure of a given list of elements under some algebraic operations (let us only mention the affine and linear closure of a list of vectors, or the ideal generated by a list of polynomials). But also in a computational context, this problem has a rich history, appearing in many areas of computer science. In its formulation as *subalgebra membership problem*, the task is to decide whether a given finite list of elements of an algebraic structure generates another element or not.

Depending on the algebraic structures studied, a variety of different problems emerges. One of the most well-known examples is the *subgroup membership problem*, in which the task is to decide, if for a given set of permutations $\alpha_1, \dots, \alpha_n$ on a finite set X, another permutation β belongs to the subgroup generated by $\alpha_1, \dots, \alpha_n$ in S_X. This problem can be solved in polynomial-time by the famous Schreier-Sims algorithm [30], whose runtime was analysed in [15] and [19]. The existence of such efficient algorithms is however not always guaranteed: if the symmetric group S_X is for instance replaced by the full transformation semigroup on X, the corresponding membership problem is PSPACE-complete [22].

A common feature of many algorithms for the subalgebra membership problem is to generate canonical generating sets of some sorts (such as computing the basis of a vector space via Gaussian elimination, or computing a Gröbner basis via Buchberger's algorithm to solve the ideal membership problem [6]). But, in general, this is where the similarities

41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024).
Editors: Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov;
Article No. 46; pp. 46:1–46:17
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

end - depending on the algebraic structure, and the encoding of the input, the problem can range over a wide range of complexities, and have applications in vastly different areas such as cryptography [28, 29], computer algebra [6, 24], or proof complexity [22, 21].

In this paper, we study a version of the subalgebra membership problem that is called the *subpower membership problem.* For a fixed, finite algebraic structure $\mathbf{A}$ (henceforth also just called an *algebra*) its subpower membership problem SMP($\mathbf{A}$) is the problem of deciding if a given tuple $\mathbf{b} \in \mathbf{A}^k$ is in the subalgebra of $\mathbf{A}^k$ generated by some other input tuples $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbf{A}^k$ (here n and k are not fixed, but part of the input). This is equivalent to checking, whether the n-ary partial function that maps $\mathbf{a}_1, \dots, \mathbf{a}_n$ component-wise to $\mathbf{b}$ can be interpolated by a term function of $\mathbf{A}$. For example, if p is a prime, SMP($\mathbb{Z}_p$) is the problem of checking whether some vector $\mathbf{b} \in \mathbb{Z}_p^k$ is in the linear closure of $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{Z}_p^k$; this can easily be solved by Gaussian elimination. More general, for any finite group $\mathbf{G}$, SMP($\mathbf{G}$) can be solved in polynomial time by a version of the Schreier-Sims algorithm [32].

Besides being a natural problem in algebra, the subpower membership problem found some applications in some learning algorithms [7, 12, 17]. Moreover, an efficient algorithm for SMP($\mathbf{A}$) implies that it is also feasible to represent the relations invariant by some generating set of tuples. It was in particular remarked (see e.g. [9]), that a polynomial-time algorithm for SMP($\mathbf{A}$) would allow to define infinitary constraint satisfaction problems, in which the constraint relations are given by some generating tuples (with respect to $\mathbf{A}$). This infinitary version of CSPs has the benefit that most of the algebraic machinery to CSPs (see e.g. [3]) still applies.

Exhaustively generating the whole subalgebra generated by $\mathbf{a}_1, \dots, \mathbf{a}_n$ in $\mathbf{A}^k$ gives an exponential time algorithm for SMP($\mathbf{A}$). And, in general, we cannot expect to do better: In [23] Kozik constructed a finite algebra $\mathbf{A}$ for which SMP($\mathbf{A}$) is EXP-complete. Even semigroups can have PSPACE-complete subpower membership problem [8].

However, for so called *Mal'tsev algebras*, better upper bounds are known. Mal'tsev algebras are algebras defined by having a *Mal'tsev term* m, i.e. a term satisfying the identities $y = m(x, x, y) = m(y, x, x)$ for all x, y. Mal'tsev algebras lie at the intersection of many areas of mathematics: they include algebraic structures of ubiquitous importance (groups, fields, vector spaces), but also appear in logic (Boolean algebras, Heyting algebras), commutative algebra (rings, modules, K-algebras), and non-associative mathematics (quasigroups, loops). Mayr showed in [25] that the subpower membership problem of every Mal'tsev algebra is in NP. His proof is based on the fact that every subalgebra $\mathbf{R} \leq \mathbf{A}^n$ has a small generating set, which generates every element of $\mathbf{R}$ in a canonical way (a so-called *compact representation*). Thus, to solve the subpower membership problem, one can “guess” a compact representation of the subalgebra generated by $\mathbf{a}_1, \dots, \mathbf{a}_k$, and then check in polynomial time if it generates $\mathbf{b}$. If such a compact representation can be moreover found in *deterministic* polynomial time, then SMP($\mathbf{A}$) is in P; this is, in fact, the dominant strategy to prove tractability.

So far, the existence of such polynomial time algorithms was verified for groups and rings [32, 15], supernilpotent algebras [25], and algebras that generate residually finite varieties [9]. On the other hand, no examples of NP-hard or intermediate complexity are known. This leads to the question whether SMP($\mathbf{A}$) $\in$ P for *all* finite Mal'tsev algebras $\mathbf{A}$ [32]. On a broader scale, this question was also posed for algebras with *few subpowers* [17, Question 8].

An elementary class of Mal'tsev algebras, for which the question still remains open, are *nilpotent* algebras. In fact, they can also be seen as an important stepping stone in answering [17, Question 8], as nilpotent Mal'tsev algebras coincide with nilpotent algebras with few subpowers. Generalizing the concept of nilpotent groups, nilpotent algebras are defined by having a central series of congruences. While they have several nice structural properties, in general nilpotent algebras do not satisfy the two finiteness conditions mentioned above (supernilpotence, residual finiteness), thus they are a natural starting point when trying to generalize known tractability results. But even for 2-nilpotent algebras not much is known: all polynomial-time algorithms were only constructed by ad-hoc arguments for concrete examples (such as Vaughan-Lee's 12-element loop [26]).

The first contribution of this paper is to prove that all 2-nilpotent algebras of size $p \cdot q$ for two primes $p \neq q$ have a tractable subpower membership problem. In fact, we prove an even stronger result in Theorem 20: $\mathrm{SMP}(\mathbf{A})$ is in P, whenever $\mathbf{A}$ has a central series $0_{\mathbf{A}} < \rho < 1_{\mathbf{A}}$ such that $|\mathbf{A}/\rho| = p$ is a prime, and the blocks of ρ have size coprime to p.

While this is still a relatively restricted class of nilpotent algebras, our methods have the potential to generalize to all 2-nilpotent Mal'tsev algebras and beyond. Thus, our newly developed tools to analyze SMP can be regarded as the second main contribution. More specifically, in Theorem 11 we show that whenever $\mathbf{L} \otimes \mathbf{U}$ is a *wreath product* (see Section 3), such that $\mathbf{U}$ is supernilpotent, then $\mathrm{SMP}(\mathbf{L} \otimes \mathbf{U})$ reduces to $\mathrm{SMP}(\mathbf{L} \times \mathbf{U})$ (which is polynomial-time solvable by [25]) and a version of the subpower membership problem for a multi-sorted algebraic object called a *clonoid* from $\mathbf{U}$ to $\mathbf{L}$. This reduction in particular applies to all 2-nilpotent algebras; an analysis of clonoids between affine algebras then leads to Theorem 20. If, in future research, we could get rid of the condition of $\mathbf{U}$ being supernilpotent, this would provide a strong tool in studying general Mal'tsev algebras, as every Mal'tsev algebra with non-trivial center can be decomposed into a wreath product.

Our paper is structured as follows: Section 2 contains preliminaries and some background on universal algebra. In Section 3 we discuss how Mal'tsev algebras with non-trivial center can be represented by a wreath product and we introduce the concept of *difference clonoid* of such a representation. In Section 4 we discuss some situations, in which the subpower membership problem of a wreath product can be reduced to the membership problem of the corresponding difference clonoid. In particular, we prove Theorem 11. Section 5 contains an analysis of clonoids between $\mathbb{Z}_p$ and coprime Abelian groups, which then leads to the proof of our main result, Theorem 20. In Section 6 we discuss some possible directions for future research.

2 Preliminaries

In the following, we are going to discuss some necessary notions from universal algebra. For more general background we refer to the textbooks [4, 11]. For background on commutator theory we refer to [14] and [2]. For an introduction to Malt'sev algebras and compact representations we refer to [5, Chapters 1.7-1.9].

In this paper, we are going to denote tuples by lower case bold letters, e.g. $\mathbf{a} \in A^k$. In order to avoid double indexing in some situations, we are going to use the notation $\mathbf{a}(i)$ to denote the i-th entry of $\mathbf{a}$, i.e. $\mathbf{a} = (\mathbf{a}(1), \mathbf{a}(2), \dots, \mathbf{a}(k))$. However, otherwise we are going to follow standard notation as used e.g. in [4].

2.1 Basic notions for general algebras

An *algebra* $\mathbf{A} = (A; (f_i^{\mathbf{A}})_{i \in I})$ is a first-order structure in a purely functional language $(f_i)_{i \in I}$ (where each symbol f_i has an associated *arity*). We say $\mathbf{A}$ is finite if its domain A is finite. A *subalgebra* $\mathbf{B} = (B; (f_i^{\mathbf{B}})_{i \in I})$ of an algebra $\mathbf{A} = (A; (f_i^{\mathbf{A}})_{i \in I})$ (denoted $\mathbf{B} \leq \mathbf{A}$) is an algebra obtained by restricting all *basic operations* $f_i^{\mathbf{A}}$ to a subset $B \subseteq A$ that is invariant under all $f_i^{\mathbf{A}}$'s. The subalgebra generated by a list of elements $a_1, \dots, a_n$, denoted by $\mathrm{Sg}_{\mathbf{A}}(a_1, \dots, a_n)$

is the smallest subalgebra of $\mathbf{A}$ that contains $a_1, \dots, a_n$. The *product* $\prod_{i \in I} \mathbf{A}_i$ of a family of algebras $(\mathbf{A}_i)_{i \in I}$ in the same language is defined as the algebra with domain $\prod_{i \in I} A_i$, whose basic operations are defined coordinate-wise. The power $\mathbf{A}^n$ is the product of n-many copies of $\mathbf{A}$. Subalgebras of (finite) powers of $\mathbf{A}$ are sometimes also called *subpowers* of $\mathbf{A}$, which motivates the name "subpower membership problem". So, formally the subpower membership problem of $\mathbf{A}$ can be stated as follows:

SMP($\mathbf{A}$)
Input: $\mathbf{b}, \mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbf{A}^k$ for some $n, k \in \mathbb{N}$
Question: Is $\mathbf{b} \in \mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1, \dots, \mathbf{a}_n)$?

Note that the subpowers of $\mathbf{A}$ are exactly the relations on A that are invariant under $\mathbf{A}$. A *congruence* α of $\mathbf{A}$ is an equivalence relation on A that is invariant under $\mathbf{A}$. We write $\mathrm{Con}(\mathbf{A})$ for the lattice of all congruence of $\mathbf{A}$. We denote the minimal and maximal element of this lattice by $0_{\mathbf{A}} = \{(x, x) \mid x \in A\}$ and $1_{\mathbf{A}} = \{(x, y) \mid x, y \in A\}$. For every congruence $\alpha \in \mathrm{Con}(\mathbf{A})$, one can form a quotient algebra $\mathbf{A}/\alpha$ in the natural way.

The *term operations* of an algebra $\mathbf{A}$ are all finitary operations that can be defined by a composition of basic operations of $\mathbf{A}$. Two standard ways to represent them is by terms or circuits in the language of $\mathbf{A}$. For a term or circuit $t(x_1, \dots, x_n)$ in the language of $\mathbf{A}$, we write $t^{\mathbf{A}}(x_1, \dots, x_n)$ for the induced term operation on A. Occasionally, if it is clear from the context, we are not going to distinguish between a term/circuit and the corresponding term operation. The term operations of an algebra $\mathbf{A}$ are closed under composition and contain all projections, therefore they form an algebraic object called a *clone*. For short, we denote this *term clone* of an algebra $\mathbf{A}$ by $\mathsf{Clo}(\mathbf{A})$. Note that $\mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1, \dots, \mathbf{a}_n) = \{t(\mathbf{a}_1, \dots, \mathbf{a}_n) \mid t \in \mathsf{Clo}(\mathbf{A})\}$.

We call a ternary operation $m^{\mathbf{A}}(x, y, z) \in \mathsf{Clo}(\mathbf{A})$ a *Mal'tsev term* if it satisfies the identities $m^{\mathbf{A}}(y, x, x) = m^{\mathbf{A}}(x, x, y) = y$ for all $x, y \in A$, and call $\mathbf{A}$ a *Mal'tsev algebra* if it has a Mal'tsev term. For instance, every group is a Mal'tsev algebra with Mal'tsev term $m(x, y, z) = xy^{-1}z$. Mal'tsev terms are a classic topic of study in universal algebra (see e.g. [4, Chapter 7]), and are in particular known to characterize congruence permutable varieties.

2.2 Clonoids

We are also going to rely on a multi-sorted generalisation of clones, so-called *clonoids* that were first introduced in [1] (in a slightly less general way). For a set of operations between two sets $\mathcal{C} \subseteq \{f \colon A^n \to B \mid n \in \mathbb{N}\}$, and $k \in \mathbb{N}$ let us write $\mathcal{C}^{(k)} = \{f \colon A^k \to B \mid f \in \mathcal{C}\}$ for the subset of k-ary functions. Then, for two algebras $\mathbf{A} = (A, (f_i)_{i \in I})$, $\mathbf{B} = (B, (g_j)_{j \in J})$ (in possibly different domains and languages), a set $\mathcal{C} \subseteq \{f \colon A^n \to B \mid n \in \mathbb{N}\}$ is called a *clonoid from* $\mathbf{A}$ *to* $\mathbf{B}$, or $(\mathbf{A}, \mathbf{B})$-*clonoid*, if it is closed under composition with term operations of $\mathbf{A}$ from the inside, and $\mathbf{B}$ from the outside, i.e.: $\forall n, k \in \mathbb{N}$

(1) $f \in \mathcal{C}^{(n)}, t_1, \dots, t_n \in \mathsf{Clo}(\mathbf{A})^{(k)} \Rightarrow f \circ (t_1, \dots, t_n) \in \mathcal{C}^{(k)}$

(2) $s \in \mathsf{Clo}(\mathbf{B})^{(n)}, f_1, \dots, f_n \in \mathcal{C}^{(k)} \Rightarrow s \circ (f_1, \dots, f_n) \in \mathcal{C}^{(k)}$.

2.3 Commutator theory

Commutator theory is the subfield of universal algebra that tries to generalise notions such as central subgroups, nilpotence, or solvability from group theory to general algebras. The most commonly used framework is based on so-called term-conditions, which we outline in the following.

Let $\mathbf{A}$ be an algebra. For congruences $\alpha, \beta, \gamma \in \mathrm{Con}(\mathbf{A})$ we say that α *centralizes* β *modulo* γ (and write $C(\alpha, \beta; \gamma)$) if and only if for all $p(\mathbf{x}, \mathbf{y}) \in \mathsf{Clo}(\mathbf{A})$, and all tuples $\mathbf{a}, \mathbf{b} \in A^n$, $\mathbf{c}, \mathbf{d} \in A^m$, such that $a_i \sim_\alpha b_i$ for $i = 1, \dots, n$ and $c_j \sim_\beta d_j$ for $j = 1, \dots, m$, the implication

$$p(\mathbf{a}, \mathbf{c}) \sim_\gamma p(\mathbf{a}, \mathbf{d}) \Rightarrow p(\mathbf{b}, \mathbf{c}) \sim_\gamma p(\mathbf{b}, \mathbf{d})$$

holds. A congruence α is called *central* if $C(\alpha, 1_\mathbf{A}; 0_\mathbf{A})$ holds. The *center* is the biggest central congruence. An algebra $\mathbf{A}$ is called *n-nilpotent* if there is a *central series of length n*, i.e. a series of congruences $0_\mathbf{A} = \alpha_0 \leq \alpha_1 \leq \cdots \leq \alpha_n = 1_\mathbf{A}$, such that $C(\alpha_{i+1}, 1_\mathbf{A}; \alpha_i)$ for $i = 0, \dots, n-1$. An algebra $\mathbf{A}$ is called *Abelian*, if it is 1-nilpotent, i.e. $C(1_\mathbf{A}, 1_\mathbf{A}; 0_\mathbf{A})$ holds.

We are, however, not going to work directly with these definitions. There is a rich structural theory in the special case of Mal'tsev algebras (and, more general, in congruence modular varieties [14]) that gives us very useful characterizations of many commutator theoretical properties.

By a result of Herrmann [16], a Mal'tsev algebra $\mathbf{A}$ is Abelian if and only if it is *affine*, i.e. all of its term operations are affine combination $\sum_{i=1}^{n} \alpha_i x_i + c$ over some module; in particular the Mal'tsev term is then equal to $x - y + z$. More generally, we are going to use a result of Freese and McKenzie [14] that states that a Mal'tsev algebra $\mathbf{A}$ with a central congruence ρ can always be written as a *wreath product* $\mathbf{L} \otimes \mathbf{U}$, such that $\mathbf{L}$ is affine and $\mathbf{U} = \mathbf{A}/\rho$. We are going to discuss such wreath product representations in Section 3.

Lastly, we want to mention that the definition of the relation C naturally generalizes to higher arities $C(\alpha_1, \dots, \alpha_n, \beta; \gamma)$. This notion was first introduced by Bulatov; we refer to [14] and [2] to more background on *higher commutators*. In particular, an algebra is called *k-supernilpotent* if $C(1_\mathbf{A}, \dots, 1_\mathbf{A}; 0_\mathbf{A})$, where $1_\mathbf{A}$ appears $k+1$ times. There are several known characterizations of supernilpotent Mal'tsev algebras. We are mainly going to use the following:

▶ **Theorem 1** (Proposition 7.7. in [2]). *Let $\mathbf{A}$ be a k-supernilpotent Mal'tsev algebra, $0 \in A$ a constant and t, s two n-ary terms in the language of $\mathbf{A}$.Then $t^\mathbf{A} = s^\mathbf{A}$ if and only if they are equal on all tuples from the set $S = \{\mathbf{a} \in A^n \mid |\{i : \mathbf{a}(i) \neq 0\}| \leq k\}$. (In fact, $\mathbf{A}$ is k-supernilpotent iff that equivalence holds for all terms t, s).*

2.4 Compact representations and SMP

For any subset $R \subseteq A^n$, we define its *signature* $\mathrm{Sig}(R)$ to be the set of all triples $(i, a, b) \in \{1, \dots, n\} \times A^2$, such that there are $\mathbf{t}_a, \mathbf{t}_b \in R$ that agree on the first $i-1$ coordinates, and $t_a(i) = a$ and $t_b(i) = b$; we then also say that $\mathbf{t}_a, \mathbf{t}_b$ are *witnesses* for $(i, a, b) \in \mathrm{Sig}(R)$.

If $\mathbf{A}$ is a Mal'tsev algebra, and $\mathbf{R} \leq \mathbf{A}^n$, then it is known that $\mathbf{R}$ is already generated by every subset $S \subseteq \mathbf{R}$ with $\mathrm{Sig}(S) = \mathrm{Sig}(\mathbf{R})$ [5, Theorem 1.8.2.]. In fact, $\mathbf{R}$ is then equal to the closure of S under the Mal'tsev operation m alone, and a tuple $\mathbf{a}$ is in $\mathbf{R}$ if $\mathbf{a}$ can be written as $m(\dots m(\mathbf{a}_1, \mathbf{b}_2, \mathbf{a}_2), \dots, \mathbf{b}_n, \mathbf{a}_n)$, for some $\mathbf{a}_i, \mathbf{b}_i \in S$. For given $\mathbf{a} \in \mathbf{R}$ such elements $\mathbf{a}_i, \mathbf{b}_i \in S$ can be found polynomial time in $|S|$, by picking $\mathbf{a}_1$ such that $\mathbf{a}_1(1) = \mathbf{a}(1)$, and $\mathbf{a}_i, \mathbf{b}_i \in S$ are witnesses for $a(i)$ and $m(\dots m(\mathbf{a}_1, \mathbf{b}_2, \mathbf{a}_2), \dots, \mathbf{b}_{i-1}, \mathbf{a}_{i-1}))(i)$ at position i.

A *compact representation* of $\mathbf{R} \leq \mathbf{A}^n$ is a subset $S \subset \mathbf{R}$ with $\mathrm{Sig}(S) = \mathrm{Sig}(\mathbf{R})$ and $|S| \leq 2|\mathrm{Sig}(\mathbf{R})| \leq 2n|A|^2$. So, informally speaking, compact representations are small generating sets of $\mathbf{R}$ with the same signature. It is not hard to see that compact representations always exist. Generalizations of compact representations exist also for relations on different domains ($\mathbf{R} \leq \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$), and relations invariant under algebras with few subpowers, we refer to [5, Chapter 2] for more background.

By the above, $\mathrm{SMP}(\mathbf{A})$ reduces in polynomial time to the problem of finding a compact representation of $\mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1,\dots,\mathbf{a}_n)$ for some input tuples $\mathbf{a}_1,\dots,\mathbf{a}_n \in A^k$. We are going to denote this problem by $\mathrm{CompRep}(\mathbf{A})$. Conversely, it was shown in [9] that finding a compact representations has a polynomial Turing reduction to $\mathrm{SMP}(\mathbf{A})$. Note further that, to solve $\mathrm{CompRep}(\mathbf{A})$ it is already enough to find a subset $S \subseteq R$ with $\mathrm{Sig}(S) = \mathrm{Sig}(R)$ of polynomial size, since such a set S can then be thinned out to a compact representation.

Let us call a set of pairs $\{(\mathbf{c}, p_{\mathbf{c}}) \mid \mathbf{c} \in S\}$ an *enumerated compact representation* of $\mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1,\dots,\mathbf{a}_n)$, if S is a compact representation of $\mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1,\dots,\mathbf{a}_n)$, and every $p_{\mathbf{c}}$ is a circuit in the language of $\mathbf{A}$ of polynomial size (in n and k), such that $p_{\mathbf{c}}(\mathbf{a}_1,\dots,\mathbf{a}_n) = \mathbf{c}$. Enumerated compact representations were already (implicitly) used in several proofs. In [9, Theorem 4.13.] it was shown that, for algebras with few subpowers, enumerated compact representations always exist; this was used to prove that $\mathrm{SMP}(\mathbf{A}) \in \mathsf{NP}$. Moreover, all of the known polynomial time algorithms for $\mathrm{CompRep}(\mathbf{A})$, in fact, compute enumerated compact representations. We are in particular going to need the following result that follows from [25]:

▶ **Theorem 2** ([25]). *Let $\mathbf{A}$ be a finite supernilpotent Mal'tsev algebra. Then, there is a polynomial time algorithm that computes an enumerated compact representations of $\mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1,\dots,\mathbf{a}_n)$, for given $\mathbf{a}_1,\dots,\mathbf{a}_n \in A^k$.*

Theorem 2 can be seen as a generalization of Gaussian elimination from affine to supernilpotent algebras. We remark that Theorem 2, although not explicitly stated as such in [25], follows directly from Algorithm 6 in [25], which computes so-called *group representations* $(T_1, T_2, \dots, T_k)$ of $\mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1,\dots,\mathbf{a}_n)$ and the fact that for such a group representation, there is a constant q such that $T = (T_1 + q\cdot T_2 + \cdots + q\cdot T_k)$ has the same signature as $\mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1,\dots,\mathbf{a}_n)$ (see Lemma 3.1. in [25]). Thus, T together with its defining circuits forms an enumerated compact representation of $\mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1,\dots,\mathbf{a}_n)$.

We are furthermore going to use that there is an algorithm that allows us to fix some values of a relation given by enumerated compact representation:

▶ **Lemma 3.** *Let $\mathbf{A}$ be a Mal'tsev algebra. Then, there is a polynomial-time algorithm `Fix-values`$(R, a_1, \dots, a_m)$ that, for a given compact representation R of $\mathbf{R} = \mathrm{Sg}_{\mathbf{A}^k}(X)$, and constants $a_1,\dots,a_m \in A$, returns a compact representation R' of $\{\mathbf{x} \in \mathbf{R} \mid \mathbf{x}(1) = a_1, \dots, \mathbf{x}(m) = a_m\}$ (or $\emptyset$ if the relation is empty). If R is moreover enumerated then `Fix-values` also computes polynomial size circuits defining the elements of R' from X.*

The existence of such a `Fix-values` algorithm for compact representation is a well-known result ([7], see also [5, Algorithm 5]); the additional statement about *enumerated* compact representation follows easily from bookkeeping the defining circuits. We prove Lemma 3 in Appendix A.

3 Wreath products and difference clonoids

In this section, we discuss how to represent Mal'tsev algebras with non-trivial center by a so-called *wreath product* $\mathbf{L} \otimes \mathbf{U}$, and associate to it its *difference clonoid*, which gives us a measure on how far it is from being the direct product $\mathbf{L} \times \mathbf{U}$.

▶ **Definition 4.** *Let $\mathbf{U} = (U, (f^{\mathbf{U}})_{f\in F})$ and $\mathbf{L} = (L, (f^{\mathbf{L}})_{f\in F})$ be two algebras in the same language F, such that $\mathbf{L}$ is affine. Furthermore, let $0 \in L$ and $T = (\hat{f})_{f\in F}$ be a family of operations $\hat{f}\colon U^n \to L$, for each $f \in F$ of arity n. Then we define the* wreath product *$\mathbf{L} \otimes^{T,0} \mathbf{U}$ as the algebra $(L \times U, (f^{\mathbf{L}\otimes^T \mathbf{U}})_{f\in F})$ with basic operations*

$$f^{\mathbf{L}\otimes^{T,0}\mathbf{U}}((l_1,u_1),\dots,(l_n,u_n)) = (f^{\mathbf{L}}(l_1,\dots,l_n) + \hat{f}(u_1,\dots,u_n), f^{\mathbf{U}}(u_1,\dots,u_n)),$$

(where $+$ is the addition on $\mathbf{L}$ with respect to neutral element 0). For simplicity, we are going to write $\mathbf{L} \otimes \mathbf{U}$, if T and 0 are clear from the context.

The name *wreath product* refers to the fact that this is a special case of VanderWerf's wreath products [31]. We remark that recently also alternative names for $\mathbf{L} \otimes \mathbf{U}$ were suggested, such as *central extension* (by Mayr) and *semidirect product* (by Zhuk). By a result of Freese and McKenzie we can represent Mal'tsev algebras with non-trivial centers as wreath products:

▶ **Theorem 5** (Proposition 7.1. in [14]). *Let $\mathbf{A}$ be a Mal'tsev algebra with a central congruence α, and let $\mathbf{U} = \mathbf{A}/\alpha$. Then there is an affine algebra $\mathbf{L}$, an element $0 \in L$ and a set of operations T, such that $\mathbf{A} \cong \mathbf{L} \otimes^{T,0} \mathbf{U}$.*

Note that, for a fixed quotient $\mathbf{U} = \mathbf{A}/\alpha$, there is still some freedom in how to choose the operations $f^{\mathbf{L}}$ of $\mathbf{L}$, and the operations $\hat{f} \colon U^n \to L$ in T (by adding/subtracting constants). To get rid of this problem, we are from now on always going to assume that $\mathbf{L}$ preserves 0, i.e. $f^{\mathbf{L}}(0, 0, \dots, 0) = 0$ for all $f \in F$. It is then easy to observe that wreath products $\mathbf{L} \otimes^{T,0} \mathbf{U}$ behaves nicely with respect to the direct product $\mathbf{L} \times \mathbf{U}$ in the same language:

▶ **Observation 6.** *Let $\mathbf{A}$ be a Mal'tsev algebra with wreath product representation $\mathbf{A} = \mathbf{L} \otimes^{T,0} \mathbf{U}$. Then $t^{\mathbf{A}} = s^{\mathbf{A}} \Rightarrow t^{\mathbf{L}\times\mathbf{U}} = s^{\mathbf{L}\times\mathbf{U}}$.*

Proof. Note that, for every term t in the language of $\mathbf{A}$:

$$t^{\mathbf{A}}((l_1, u_1), \dots, (l_n, u_n)) = (t^{\mathbf{L}}(l_1, \dots, l_n) + \hat{t}(u_1, \dots, u_n), t^{\mathbf{U}}(u_1, \dots, u_n)),$$

for some $\hat{t} \colon U^n \to L$ (this can be shown by induction over the height of the term tree). Clearly $t^{\mathbf{A}} = s^{\mathbf{A}}$ implies $t^{\mathbf{U}} = s^{\mathbf{U}}$, and $t^{\mathbf{L}} - s^{\mathbf{L}} = c$, $\hat{t} - \hat{s} = -c$ for some constant $c \in L$. Since, by our assumptions, the operations of $\mathbf{L}$ preserve 0, we get $t^{\mathbf{L}} = s^{\mathbf{L}}$ and $\hat{t} = \hat{s}$. Thus $t^{\mathbf{L}\times\mathbf{U}} = s^{\mathbf{L}\times\mathbf{U}}$. ◀

In other terminology, the map $t^{\mathbf{A}} \mapsto t^{\mathbf{L}\times\mathbf{U}}$ is a surjective *clone homomorphism* from $\mathsf{Clo}(\mathbf{A})$ to $\mathsf{Clo}(\mathbf{L} \times \mathbf{U})$, i.e. a map that preserves arities, projections and compositions. The converse of Observation 6 does however not hold, since this map is usually not injective. We define the *difference clonoid* $\mathrm{Diff}_0(\mathbf{A})$ as the kernel of the clone homomorphisms in the following sense:

▶ **Definition 7.** *Let $\mathbf{A} = \mathbf{L} \otimes^{T,0} \mathbf{U}$ be a Mal'tsev algebra given as a wreath product.*

(1) *We define the equivalence relation $\sim$ on $\mathsf{Clo}(\mathbf{A})$ by*

$$t^{\mathbf{A}} \sim s^{\mathbf{A}} :\Leftrightarrow t^{\mathbf{L}\times\mathbf{U}} = s^{\mathbf{L}\times\mathbf{U}}$$

(2) *the* difference clonoid $\mathrm{Diff}_0(\mathbf{A})$ *is defined as the set of all operation $\hat{r} \colon U^n \to L$, such that there are $t^{\mathbf{A}} \sim s^{\mathbf{A}} \in \mathsf{Clo}(\mathbf{A})$ with:*

$$t^{\mathbf{A}}((l_1, u_1), \dots, (l_n, u_n)) = (t^{\mathbf{L}}(\mathbf{l}) + \hat{t}(\mathbf{u}), t^{\mathbf{U}}(\mathbf{u})) \tag{1}$$

$$s^{\mathbf{A}}((l_1, u_1), \dots, (l_n, u_n)) = (t^{\mathbf{L}}(\mathbf{l}) + \hat{t}(\mathbf{u}) + \hat{r}(\mathbf{u}), t^{\mathbf{U}}(\mathbf{u})) \tag{2}$$

▶ Notation 8. In the following, we will stick to the following convention: Function symbols with a hat will always denote operations from some power of U to L. For operations $t, s \colon A^n \to A$, and $\hat{r} \colon U^n \to L$ such as in (1) and (2) we are slightly going to abuse notation, and write $s = t + \hat{r}$ and $\hat{r} = (s - t)$.

We next show that $\mathrm{Diff}_0(\mathbf{A})$ is indeed a clonoid from $\mathbf{U}$ to $\mathbf{L}$ (extended by the constant 0).

▶ **Lemma 9.** *Let $\mathbf{A} = \mathbf{L} \otimes^{0,T} \mathbf{U}$ be a Mal'tsev algebra given as wreath product. Then:*

(1) *For all $t \in \mathsf{Clo}(\mathbf{A})$, $\hat{r} \in \mathrm{Diff}_0(\mathbf{A})$ also $t + \hat{r} \in \mathsf{Clo}(\mathbf{A})$,*

(2) $\mathrm{Diff}_0(\mathbf{A})$ *is a $(\mathbf{U}, (\mathbf{L}, 0))$-clonoid.*

Here $(\mathbf{L}, 0)$ denotes the extension of $\mathbf{L}$ by the basic operation 0.

Proof. To prove (1), let $t \in \mathsf{Clo}(\mathbf{A})$ and $\hat{r} \in \mathrm{Diff}_0(\mathbf{A})$. By definition of the difference clonoid, $\hat{r} = s_1 - s_2$ for two terms $s_1, s_2 \in \mathsf{Clo}(\mathbf{A})$, with $s_1 \sim s_2$. In particular, $s_1^{\mathbf{U}} = s_2^{\mathbf{U}}$. For any Mal'tsev term $m \in \mathsf{Clo}(\mathbf{A})$, necessarily $\hat{m}(u, u, v) = \hat{m}(v, u, u) = 0$ holds. This implies that

$$t + \hat{r} = m(t, s_2, s_1) \in \mathsf{Clo}(\mathbf{A}).$$

We next prove (2). So we only need to verify that $\mathrm{Diff}_0(\mathbf{A})$ is closed under composition with $\mathsf{Clo}(\mathbf{U})$ (from the inside), respectively $\mathsf{Clo}((\mathbf{L}, 0))$ (from the outside).

To see that $\mathrm{Diff}_0(\mathbf{A})$ is closed under $(\mathbf{L}, 0)$, note that $0 \in \mathrm{Diff}_0(\mathbf{A})$, as $t - t = 0$, for every term $t \in \mathsf{Clo}(\mathbf{A})$. Further $\mathrm{Diff}_0(\mathbf{A})$ is closed under $+$; for this, let $\hat{r}_1, \hat{r}_2 \in \mathrm{Diff}_0(\mathbf{A})$. By (1), we know that $t + \hat{r}_1 \in \mathsf{Clo}(\mathbf{A})$, for some term $t \in \mathsf{Clo}(\mathbf{A})$. Again, by (1) also $(t + \hat{r}_1) + \hat{r}_2)) \in \mathsf{Clo}(\mathbf{A})$, which shows that $\hat{r}_1 + \hat{r}_2 \in \mathrm{Diff}_0(\mathbf{A})$. For all unary $e^{\mathbf{L}} \in \mathsf{Clo}(\mathbf{L})$, and $t \sim s$ with $\hat{r} = t - s$, note that $e^{\mathbf{A}} t - e^{\mathbf{A}} s = e^{\mathbf{L}} \circ \hat{r} \in \mathrm{Diff}_0(\mathbf{A})$. Since $\mathbf{L}$ is affine, $\mathsf{Clo}(\mathbf{L}, 0)$ is generated by $+$ and its unary terms, thus $\mathrm{Diff}_0(\mathbf{A})$ is closed under $(\mathbf{L}, 0)$.

To see that $\mathrm{Diff}_0(\mathbf{A})$ is closed under $\mathbf{U}$ from the inside, simply notice that $t(x_1, \dots, x_n) \sim s(x_1, \dots, x_n)$ implies $t(f_1(\mathbf{x}), \dots, f_n(\mathbf{x})) \sim s(f_1(\mathbf{x}), \dots, f_n(\mathbf{x}))$, for all terms $f_1, \dots, f_n$. If $\hat{r} = t^{\mathbf{A}} - s^{\mathbf{A}}$, then $\hat{r} \circ (f_1^{\mathbf{U}}, \dots, f_n^{\mathbf{U}}) = t \circ (f_1^{\mathbf{U}}, \dots, f_n^{\mathbf{U}}) - s \circ (f_1^{\mathbf{U}}, \dots, f_n^{\mathbf{U}}) \in \mathrm{Diff}_0(\mathbf{A})$. ◀

We remark that the choice of the constant $0 \in L$ is not relevant in this construction: since for every $c \in L$ the map $\hat{r} \mapsto \hat{r} + c$ is an isomorphism between the $(\mathbf{U}, (\mathbf{L}, 0))$-clonoid $\mathrm{Diff}_0(\mathbf{A})$ and the $(\mathbf{U}, (\mathbf{L}', c))$-clonoid $\mathrm{Diff}_c(\mathbf{A})$ (where $f^{\mathbf{L}'}(\mathbf{l}) = f^{\mathbf{L}}(\mathbf{l} - (c, c \dots, c)) + c$).

Our goal in the next section is to reduce the subpower membership problem to a version of the subpower membership problem for the difference clonoid in which we ask for membership of a tuple $\mathbf{l} \in L^k$ in the subalgebra of $\mathbf{L}$ given by the image of $\mathbf{u}_1, \dots, \mathbf{u}_n \in U^k$ under the clonoid. In fact, it will be more convenient for us to ask for a compact representation, that's why we define the following problem, for a clonoid $\mathcal{C}$ from $\mathbf{U}$ to $\mathbf{L}$.

CompRep($\mathcal{C}$):
Input: A list of tuples $\mathbf{u}_1, \dots, \mathbf{u}_n \in U^k$.
Output: A compact representation of $\mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n) = \{f(\mathbf{u}_1, \dots, \mathbf{u}_n) \mid f \in \mathcal{C}\} \leq \mathbf{L}^k$

In the case of the difference clonoid $\mathcal{C} = \mathrm{Diff}_0(\mathbf{A})$ the image algebra $\mathbf{L}$ is affine and contains a constant 0. Thus then this problem is then equivalent to finding generating set of $\mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n)$ as a subgroup of $(L, +, 0, -)^k$ of polynomial size. By then running Gaussian elimination (generalized to finite Abelian groups), or simply applying Theorem 2 one can then compute a compact representation of $\mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n)$.

4 The subpower membership problem of wreath products

In this section we discuss our main methodological results. We show that, in some cases the subpower membership problem SMP($\mathbf{L} \otimes \mathbf{U}$) of a wreath product can be reduced to CompRep($\mathbf{L} \times \mathbf{U}$) and CompRep($\mathcal{C}$). We first show how such a reduction can be achieved relatively easily in the case where $\mathsf{Clo}(\mathbf{L} \otimes \mathbf{U})$ contains $\mathsf{Clo}(\mathbf{L} \times \mathbf{U})$ (i.e. the identity map is a retraction of the clone homomorphism from Observation 6):

▶ **Theorem 10.** *Let* $\mathbf{A} = \mathbf{L} \otimes^{T,0} \mathbf{U}$ *be a finite Mal'tsev algebra, and let* $\mathcal{C} = \mathrm{Diff}_0(\mathbf{A})$*. Further assume that* $\mathsf{Clo}(\mathbf{L} \times \mathbf{U}) \subseteq \mathsf{Clo}(\mathbf{A})$*. Then* $\mathrm{CompRep}(\mathbf{A})$ *(and hence also* $\mathrm{SMP}(\mathbf{A})$*) reduces in polynomial time to* $\mathrm{CompRep}(\mathbf{L} \times \mathbf{U})$ *and* $\mathrm{CompRep}(\mathcal{C})$*.*

Proof. Let $\mathbf{a}_1, \dots, \mathbf{a}_n \in A^k$ an instance of $\mathrm{CompRep}(\mathbf{A})$; our goal is to find a compact representation of $\mathbf{B} = \mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1, \dots, \mathbf{a}_n)$. Let us write $\mathbf{l}_i$ and $\mathbf{u}_i$ for the projection of $\mathbf{a}_i$ to L^k and U^k respectively. Let us further define $\mathbf{B}^+ = \mathrm{Sg}_{(\mathbf{L}\times\mathbf{U})^k}(\mathbf{a}_1, \dots, \mathbf{a}_n)$. Then

$$\mathbf{B} = \{(t^{\mathbf{L}}(\mathbf{l}_1, \dots, \mathbf{l}_n) + \hat{t}(\mathbf{u}_1, \dots, \mathbf{u}_n), t^{\mathbf{U}}(\mathbf{u}_1, \dots, \mathbf{u}_n) \mid t \text{ is } F\text{-term}\}, \text{ and}$$
$$\mathbf{B}^+ = \{(t^{\mathbf{L}}(\mathbf{l}_1, \dots, \mathbf{l}_n), t^{\mathbf{U}}(\mathbf{u}_1, \dots, \mathbf{u}_n) \mid t \text{ is } F\text{-term}\}.$$

Since $\mathsf{Clo}(\mathbf{L} \times \mathbf{U}) \subseteq \mathsf{Clo}(\mathbf{A})$, we can pick a Mal'tsev term of $\mathbf{A}$ that is of the form $m^{\mathbf{A}}((l_1, u_1), (l_2, u_2), (l_3, u_3)) = (l_1 - l_2 + l_3, m^{\mathbf{U}}(u_1, u_2, u_3))$. Moreover, by Lemma 9, every term $t^{\mathbf{A}} \in \mathsf{Clo}(\mathbf{A})$ can be uniquely written as the sum of $t^{\mathbf{L}\times\mathbf{U}}$ (which by assumption is also in $\mathsf{Clo}(\mathbf{A})$) and some $\hat{t} \in \mathcal{C}$. Thus, every element of $\mathbf{B}$ is equal to the sum of an element of $\mathbf{B}^+$ and an expression $\hat{t}(\mathbf{u}_1, \dots, \mathbf{u}_n)$.

Let C^+ be a compact representation of $\mathbf{B}^+$, and $\hat{C}$ a compact representation of $\mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n)$. Then, it follows that every tuple in $\mathbf{B}$ can be written as

$$m(\dots, m(\mathbf{c}_1, \mathbf{d}_2, \mathbf{c}_2), \dots \mathbf{d}_n, \mathbf{c}_n) + \hat{\mathbf{r}}_1 - \hat{\mathbf{s}}_2 + \hat{\mathbf{r}}_2 - \dots - \hat{\mathbf{s}}_n + \hat{\mathbf{r}}_n, \tag{3}$$

for $\mathbf{c}_i, \mathbf{d}_i \in C^+$ and $\hat{\mathbf{r}}_i, \hat{\mathbf{s}}_i \in \hat{C}$. (We are aware that tuples in C^+ an $\hat{C}$ have different domains; here we follow the same convention as in Notation 8). Moreover, in formula (3), any pair $\mathbf{c}_i, \mathbf{d}_i$ (respectively $\hat{\mathbf{r}}_i, \hat{\mathbf{s}}_i$) witnesses a fork in the i-th coordinate. By our choice of m it is easy to see that formula (3) can be rewritten to

$$m(\dots, m(\mathbf{c}_1 + \hat{\mathbf{r}}_1, \mathbf{d}_2 + \hat{\mathbf{s}}_2, \mathbf{c}_2 + \hat{\mathbf{r}}_2), \dots \mathbf{d}_n + \hat{\mathbf{s}}_n, \mathbf{c}_n + \hat{\mathbf{r}}_n),$$

Thus the elements $\mathbf{c}_i + \hat{\mathbf{r}}_i, \mathbf{d}_i + \hat{\mathbf{s}}_i$ witness forks of $\mathbf{B}$ in the i-th coordinate. If we define $D = \{\mathbf{c} + \hat{\mathbf{r}} \mid \mathbf{c} \in C, \hat{\mathbf{r}} \in \hat{C}\}$, then it follows that $\mathrm{Sig}(D) = \mathrm{Sig}(\mathbf{B})$. Moreover $D \subset \mathbf{B}$, and it is of polynomial size in n and k, as $|D| \leq |C| \cdot |\hat{C}|$. Thus D can be thinned out in polynomial time to a compact representation of $\mathbf{B}$, which finishes our proof. ◀

We remark that, by following the proof of Theorem 10, also finding *enumerated* compact representations in $\mathbf{A}$ can be reduced to finding *enumerated* compact representations in $\mathbf{L} \times \mathbf{U}$ and $\mathcal{C}$ (if $\mathcal{C}$ is given by some finite set of operations that generate it as a clonoid).

Unfortunately, the conditions of Theorem 10 are not met for general wreath-products, not even if both $\mathbf{U}$ and $\mathbf{L}$ are affine (the dihedral group D_4 can be shown to be a counterexample). But, if $\mathbf{U}$ is supernilpotent, then we are able to prove the following reduction, independent of the conditions of Theorem 10:

▶ **Theorem 11.** *Let* $\mathbf{A} = \mathbf{L} \otimes \mathbf{U}$ *be a finite Mal'tsev algebra, and let* $\mathcal{C} = \mathrm{Diff}_0(\mathbf{A})$ *for some* $0 \in A$*. Further, assume that* $\mathbf{U}$ *is supernilpotent. Then* $\mathrm{SMP}(\mathbf{A})$ *reduces in polynomial time to* $\mathrm{CompRep}(\mathcal{C})$*.*

Proof. Let $\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{b} \in A^k$ an instance of $\mathrm{SMP}(\mathbf{A})$; our goal is to check whether $\mathbf{b} \in \mathbf{B} = \mathrm{Sg}_{\mathbf{A}^k}(\mathbf{a}_1, \dots, \mathbf{a}_n)$. Let us write $\mathbf{l}_i$ and $\mathbf{u}_i$ for the projection of $\mathbf{a}_i$ to L^k and U^k respectively, and $\mathbf{l}_b$ and $\mathbf{u}_b$ for the projections of $\mathbf{b}$ to L^k and U^k. Let F be the signature of $\mathbf{A}$ and $\mathbf{L} \times \mathbf{U}$, and let $\mathbf{B}^+ = \mathrm{Sg}_{(\mathbf{L}\times\mathbf{U})^k}(\mathbf{a}_1, \dots, \mathbf{a}_n)$. Then

$$\mathbf{B} = \{(t^{\mathbf{L}}(\mathbf{l}_1, \dots, \mathbf{l}_n) + \hat{t}(\mathbf{u}_1, \dots, \mathbf{u}_n), t^{\mathbf{U}}(\mathbf{u}_1, \dots, \mathbf{u}_n) \mid t \text{ is } F\text{-term}\}, \text{ and}$$
$$\mathbf{B}^+ = \{(t^{\mathbf{L}}(\mathbf{l}_1, \dots, \mathbf{l}_n), t^{\mathbf{U}}(\mathbf{u}_1, \dots, \mathbf{u}_n) \mid t \text{ is } F\text{-term}\}.$$

Recall the definition of $t^{\mathbf{A}} \sim s^{\mathbf{A}}$ from Definition 7. If T is a $\sim$-transversal set of $\{t^{\mathbf{A}} \in \mathsf{Clo}(\mathbf{A}) \mid t^{\mathbf{U}}(\mathbf{u}_1, \dots, \mathbf{u}_n) = \mathbf{u}_b\}$, then clearly $\mathbf{b} \in B$ iff $\exists t \in T$ and $\mathbf{d} \in \mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n)$, with $\mathbf{b} = t(\mathbf{a}_1, \dots, \mathbf{a}_n) + \mathbf{d}$. So, intuitively speaking, the goal of this proof is to first compute such a transversal set, by computing an enumerated compact representation of $\{(\mathbf{l}, \mathbf{u}) \in \mathbf{B}^+ \mid \mathbf{u} = \mathbf{u}_b\}$ and then use it together with a compact representation of $\mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n)$ to check membership of $\mathbf{b}$ in $\mathbf{B}$.

In practice we need however to consider a relation of higher arity than $\mathbf{B}^+$, since term operations of $\mathbf{L} \times \mathbf{U}$ are not uniquely determined by their value on $\mathbf{a}_1, \dots, \mathbf{a}_n$. So let S be the degree of supernilpotence of $\mathbf{U}$ (and hence also $\mathbf{L} \times \mathbf{U}$). If we think about $\mathbf{a}_1, \dots, \mathbf{a}_n$ as the columns of a matrix of dimension $k \times n$, then let $\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n \in A^l$ be its extension by rows that enumerate $H = \{(a_1, \dots, a_n) \in A^n \mid |\{i \colon a_i \neq 0\}| \leq S\}$ (hence $l \leq k + |A|^S \binom{n}{S}$).

It follows from Theorem 2 that we can compute an enumerated compact representation $\tilde{C}$ of $\mathrm{Sg}_{(\mathbf{L}\times\mathbf{U})^l}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n)$ in polynomial time in n and l. So, every element in $\tilde{B} = \mathrm{Sg}_{(\mathbf{L}\times\mathbf{U})^l}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n)$ can be written as $m(\dots m(\tilde{\mathbf{c}}_1, \tilde{\mathbf{d}}_2, \tilde{\mathbf{c}}_2) \dots \tilde{\mathbf{d}}_l, \tilde{\mathbf{c}}_l)$, for $(\tilde{\mathbf{c}}_i, p_{\tilde{\mathbf{c}}_i}), (\tilde{\mathbf{d}}_i, p_{\tilde{\mathbf{d}}_i}) \in \tilde{C}$, where $\tilde{C}$ is of size at most $2l|A|^2$, and every element of $\tilde{\mathbf{c}} \in \tilde{C}$ is equal to $p_{\tilde{\mathbf{c}}}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n) = \tilde{\mathbf{c}}$ for the given circuit $p_{\tilde{\mathbf{c}}}$ of polynomial size.

By Theorem 1, in an S-supernilpotent algebra, every term operation is already completely determined by its values on the subset H. It follows, that every n-ary term operation of $\mathbf{L} \times \mathbf{U}$ can be uniquely described by a circuit $m(\dots m(p_{\tilde{\mathbf{c}}_1}, p_{\tilde{\mathbf{d}}_2}, p_{\tilde{\mathbf{c}}_2}), \dots p_{\tilde{\mathbf{d}}_l}, p_{\tilde{\mathbf{c}}_l})$ for $\tilde{\mathbf{c}}_i, \tilde{\mathbf{d}}_i \in \tilde{C}$. By definition of $\sim$, it follows that also every n-ary term operation of $\mathbf{A}$ is $\sim$-equivalent to the operation given by the circuit described by a circuit $m(\dots m(p_{\tilde{\mathbf{c}}_1}, p_{\tilde{\mathbf{d}}_2}, p_{\tilde{\mathbf{c}}_2}), \dots p_{\tilde{\mathbf{d}}_l}, p_{\tilde{\mathbf{c}}_l})$ for $\tilde{\mathbf{c}}_i, \tilde{\mathbf{d}}_i \in \tilde{C}$.

We are however only interested in terms t such that $t^{\mathbf{U}}$ maps $\mathbf{u}_1, \dots, \mathbf{u}_n$ to the value $\mathbf{u}_b$. By Lemma 3, we can also compute an enumerated compact representation $\tilde{C}'$ of $\{(\tilde{\mathbf{l}}, \tilde{\mathbf{u}}) \in \mathrm{Sg}_{\mathbf{L}\times\mathbf{U}}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n) \mid \tilde{\mathbf{u}}(i) = \mathbf{u}_b(i) \text{ for all } i = 1, \dots, k\}$ in polynomial time. Note that this is possible, as $\{(\tilde{\mathbf{l}}, \tilde{\mathbf{u}}) \in \mathrm{Sg}_{\mathbf{L}\times\mathbf{U}}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n) \mid \tilde{\mathbf{u}}(i) = \mathbf{u}_b(i) \text{ for all } i = 1, \dots, k\}$ is closed under the Mal'tsev operation $m^{\mathbf{L}\times\mathbf{U}}$. Also, although we only prove Lemma 3 for fixing variables to constants, we remark that it can straightforwardly be generalized to fixing the value of the variables to domains $L \times \{u\}$ (alternatively, this can also be achieved by regarding $\mathrm{Sg}_{(\mathbf{L}\times\mathbf{U})^l}(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_n)$ as a subalgebra of $\mathbf{U}^l \times \mathbf{L}^l$, which however would require us to work with relations on different domains).

If $\tilde{C}' = \emptyset$, then we output "False", as then $\mathbf{u}_b \notin \mathrm{Sg}_{\mathbf{U}^k}(\mathbf{u}_1, \dots, \mathbf{u}_n)$. Otherwise, let $C = \{p^{\mathbf{A}}_{\tilde{\mathbf{c}}}(\mathbf{a}_1, \dots, \mathbf{a}_n) \mid \tilde{\mathbf{c}} \in \tilde{C}'\}$. Also, let $\hat{C}$ be a compact representation of $\mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n)$. By our proof, every element of $\{(\mathbf{l}, \mathbf{u}) \in \mathbf{B} \mid \mathbf{u} = \mathbf{u}_b\}$ is equal to the sum of an element $m^{\mathbf{A}}(\dots, m^{\mathbf{A}}(\mathbf{c}_1, \mathbf{d}_2, \mathbf{c}_2), \dots \mathbf{d}_n, \mathbf{c}_n)$ with $\mathbf{c}_i, \mathbf{d}_i \in C$ and an element of $\mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n)$. Since m is an affine Malt'sev operation when restricted to $\{(\mathbf{l}, \mathbf{u}) \in \mathbf{B} \mid \mathbf{u} = \mathbf{u}_b\}$ this means that $\mathbf{b} \in \mathbf{B}$ iff $\mathbf{l}_b$ is in the affine closure of all elements $\mathbf{c} + \hat{\mathbf{r}}$ with $\mathbf{c} \in C$ and $\hat{\mathbf{r}} \in \hat{C}$. But this can be checked in polynomial time (by generalized Gaussian elimination, or Theorem 2), which finishes the proof. ◀

5 Clonoids between affine algebras

We continue our paper with an analysis of clonoids between affine algebras to prove our main result, Theorem 20.

For a prime p, let us write $\mathbb{Z}_p$ for the cyclic group of order p, i.e. $\mathbb{Z}_p = (\{0, 1, \dots, p-1\}, +, 0, -)$. Let us further define the idempotent reduct $\mathbb{Z}_p^{id} = (\{0, 1, \dots, p-1\}, x - y + z)$. Using the unary terms $ax = x + \dots + x$ (a-times), for $a \in \mathbb{Z}_p$, we can regard $\mathbb{Z}_p$ as a vector space over the p-element field. More general, using this notation, we will also consider finite Abelian groups $(L, +, 0, -)$ as modules over $\mathbb{Z}_{|L|}$.

For short, we are going to denote constant 1-tuples by $\mathbf{1} = (1, 1, \dots, 1) \in \mathbb{Z}_p^n$. For two vectors $\mathbf{a}, \mathbf{x} \in \mathbb{Z}_p^n$, we further denote by $\mathbf{a} \cdot \mathbf{x} = \sum_{i=1}^n \mathbf{a}(i) \cdot \mathbf{x}(i)$ the standard inner product. Then $\mathsf{Clo}(\mathbb{Z}_p) = \{\mathbf{x} \mapsto \mathbf{a} \cdot \mathbf{x} \mid \mathbf{a} \in \mathbb{Z}_p^n\}$ and $\mathsf{Clo}(\mathbb{Z}_p^{id}) = \{\mathbf{x} \mapsto \mathbf{a} \cdot \mathbf{x} \mid \mathbf{a} \in \mathbb{Z}_p^n, \mathbf{a} \cdot \mathbf{1} = 1\}$.

In this section, we are going to study clonoids between affine algebras $\mathbf{U}$ and $\mathbf{L}$, such that $|U| = p$ for some prime p, and $p \nmid |L|$. Since every such affine algebra $\mathbf{U}$ has $x - y + z$ as a term operation, it makes sense to study the special case $\mathbf{U} = \mathbb{Z}_p^{id}$. As we are in particular interested in difference clonoids, we furthermore can assume that $\mathbf{L}$ contains a constant operation 0 (see Lemma 9), and hence the operations of the Abelian group $(L, +, 0, -)$. We remark that our analysis is structurally similar to (but not covered by) Fioravanti's classification of $(\mathbb{Z}_p, \mathbb{Z}_q)$-clonoids [13].

5.1 $(\mathbb{Z}_p^{id}, \mathrm{L})$-clonoids satisfying $p \nmid |L|$ and $f(x, x, \dots, x) = 0$

Throughout this subsection, let p be a prime, and $\mathbf{L} = (L, +, 0, -)$ an Abelian group with $p \nmid |L|$, and $\mathcal{C}$ be a $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoid satisfying $f(x, x, \dots, x) = 0$ for all $f \in \mathcal{C}$ and $x \in \mathbb{Z}_p$. In other words, for every $n \in \mathbb{N}$, $\mathcal{C}$ maps all tuples from the diagonal $\Delta^n = \{(x, x \dots, x) \in \mathbb{Z}_p^n\}$ to 0. We are going to prove that $\mathcal{C}$ is generated by its binary elements, and therefore by any set of generators B of $\mathcal{C}^{(2)} \leq \mathbf{L}^{\mathbb{Z}_p^2}$. Moreover, from B, we are going to construct a canonical generating set of the n-ary functions $\mathcal{C}^{(n)} \leq \mathbf{L}^{\mathbb{Z}_p^n}$. We are, in particular going to use the following set of coefficient vectors for every $n > 2$:

$$C_n = \{\mathbf{a} \in \mathbb{Z}_p^n \mid \exists i > 1 \colon \mathbf{a}(1) = \mathbf{a}(2) = \ldots = \mathbf{a}(i-1) = 0, \mathbf{a}(i) = 1\}.$$

▶ **Observation 12.** *Every 2-dimensional subspace $V \leq \mathbb{Z}_p^n$ containing the diagonal Δ^n has a unique parameterization by the map*

$$e_{\mathbf{c}}(x, y) = x(\mathbf{1} - \mathbf{c}) + y\mathbf{c} = (x, \mathbf{c}(2)x + (1 - \mathbf{c}(2))y, \dots, \mathbf{c}(n)x + (1 - \mathbf{c}(n))y),$$

for some $\mathbf{c} \in C_n$, i.e. it is equal to the range of a unique such map.

Proof. To see this, note that V contains $\mathbf{1}$, and can be therefore parameterized by $e_{\mathbf{d}}(x, y)$, for some $\mathbf{d} \notin \Delta^n$. So there is an index i with $\mathbf{d}(1) = \ldots = \mathbf{d}(i-1) \neq \mathbf{d}(i)$. If $\mathbf{d} \notin C_n$, then we define $\mathbf{c} = (\mathbf{d}(i) - \mathbf{d}(1))^{-1}(\mathbf{d} - \mathbf{d}(1)\mathbf{1})$; clearly $\mathbf{c} \in C_n$, and $\mathbf{c}$ and $\mathbf{1}$ still generate V. It is further not hard to see that different elements of C_n generate different planes together with $\mathbf{1}$, thus we obtain a unique parameterization of V by $e_{\mathbf{c}}(x, y)$. ◀

▶ **Lemma 13.** *Let $f \in \mathcal{C}^{(2)}$. Then, there is a function $f_n \in \mathcal{C}^{(n)}$, such that*

$$f_n(x_1, x_2, \dots, x_n) = \begin{cases} f(x_1, x_2) \text{ if } x_2 = x_3 = \ldots = x_n \\ 0 \text{ else.} \end{cases}$$

Proof. We prove the lemma by induction on n. For $n = 2$, we simply set $f_2 = f$. For an induction step $n \to n+1$, we first define $t_{n+1}(x_1, x_2, \dots, x_n, x_{n+1})$ as the sum

$$\sum_{\mathbf{a} \in \mathbb{Z}_p^{n-1}} f_n(x_1, x_2 + \mathbf{a}(1)(x_{n+1} - x_n), \dots, x_n + \mathbf{a}(n-1)(x_{n+1} - x_n))$$
$$- \sum_{\mathbf{a} \in \mathbb{Z}_p^{n-1}} f_n(x_1, x_1 + \mathbf{a}(1)(x_{n+1} - x_n), \dots, x_1 + \mathbf{a}(n-1)(x_{n+1} - x_n)).$$

Note that, if $x_{n+1} \neq x_n$, then t_{n+1} evaluates to $\sum_{\mathbf{a} \in \mathbb{Z}_p^{n-1}} f(x_1, \mathbf{a}) - \sum_{\mathbf{a} \in \mathbb{Z}_p^{n-1}} f(x_1, \mathbf{a}) = 0$. On the other hand, if $x_n = x_{n+1}$, then the second sum is equal to 0, while the first one is equal to $p^{n-1} f_n(x_1, x_2, \dots, x_n)$. By the induction hypothesis, the function $f_{n+1} = p^{-(n-1)} t_{n+1}$ satisfies the statement of the lemma (note that $p^{-(n-1)}$ exist modulo $|L|$, since $p \nmid |L|$). ◀

We can prove an analogue statement for all 2-dimensional subspaces of $\mathbb{Z}_p^n$ containing Δ^n:

▶ **Lemma 14.** *Let $f \in \mathcal{C}^{(2)}$, and $\mathbf{c} \in C_n$. Then there is a function $f^{\mathbf{c}} \in \mathcal{C}^{(n)}$, such that*

$$f^{\mathbf{c}}(x_1, x_2, \ldots, x_n) = \begin{cases} f(x,y) \text{ if } (x_1, x_2, \ldots, x_n) = e_{\mathbf{c}}(x,y) \\ 0 \text{ else.} \end{cases}$$

Proof. Let $\mathbf{c} \in C^{(n)}$. It is easy to see that there is an invertible matrix $\mathbf{T} \in \mathbb{Z}_p^{n \times n}$, such that $\mathbf{T} \cdot \mathbf{1} = \mathbf{1}$ and $\mathbf{T} \cdot (\mathbf{1} - \mathbf{c}) = \mathbf{e}_1$. Let $T \colon \mathbb{Z}_p^n \to \mathbb{Z}_p^n$ be the corresponding linear map $T(\mathbf{x}) = \mathbf{T} \cdot \mathbf{x}$. Let f_n as in Lemma 13 and $f^{\mathbf{c}}(\mathbf{x}) := f_n \circ T$. Note that by the first condition, all rows of $\mathbf{T}$ sum up to 1, hence T can be expressed by terms of $\mathbb{Z}_p^{id}$. Then $f^{\mathbf{c}}(e_{\mathbf{c}}(x,y)) = f_n(T(x(\mathbf{1}-\mathbf{c}) + y\mathbf{c})) = f_n(x\mathbf{e}_1 + y(\mathbf{1} - \mathbf{e}_1)) = f(x,y)$, and $f^{\mathbf{c}}(\mathbf{x}) = 0$ for $\mathbf{x} \notin e_{\mathbf{c}}(\mathbb{Z}_p^2)$. ◀

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

▶ **Lemma 15.** *Let $\mathcal{C}$ be a $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoid satisfying $\forall f \in \mathcal{C}, x \in \mathbb{Z}_p \colon f(x, \ldots, x) = 0$, and let B be a generating set of $\mathcal{C}^{(2)} \leq \mathbf{L}^{\mathbb{Z}_p^2}$. Then*

(1) *$\mathcal{C}$ is the $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoid generated by B, and*

(2) *$B_n := \{f^{\mathbf{c}} \mid f \in B, \mathbf{c} \in C_n\}$ is a generating set of $\mathcal{C}^{(n)}$ in $\mathbf{L}^{\mathbb{Z}_p^n}$,*

Proof. For any $g \in \mathcal{C}^{(n)}$ and $\mathbf{c} \in C^{(n)}$, let us define the binary operation $g_{\mathbf{c}} = f(e_{\mathbf{c}}(x,y)) \in \mathcal{C}^{(2)}$. By Lemma 14, $g_{\mathbf{c}}$ generates a function $g_{\mathbf{c}}^{\mathbf{c}} \in \mathcal{C}^{(n)}$, that agrees with $f(x,y)$ on all tuples of the form $e_{\mathbf{c}}(x,y)$, and that is 0 else. Since every point of $\mathbb{Z}_p^n \setminus \Delta^n$ is in the image of a unique map $e_{\mathbf{c}}$, we get $g = \sum_{\mathbf{c} \in C_n} g_{\mathbf{c}}^{\mathbf{c}}$. Every element of the form $g_{\mathbf{c}}^{\mathbf{c}}$ can be clearly written as a linear combination of elements $f^{\mathbf{c}}$, where $f \in B$. It follows that B_n generates $\mathcal{C}^{(n)}$ in $\mathbf{L}^{\mathbb{Z}_p^n}$, and that the clonoid generated by B is $\mathcal{C}$. ◀

We remark that if $\mathbf{L} = \mathbb{Z}_q$ for a prime $q \neq p$, and B is a basis of the vector space $\mathcal{C}^{(2)} \leq \mathbf{L}^{\mathbb{Z}_p^2}$, then B_n is a basis of $\mathcal{C}^{(n)}$. The generating set B_n can be used to decide efficiently the following version of the subpower membership problem for $\mathcal{C}$:

▶ **Lemma 16.** *Let $\mathcal{C}$ be a $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoid satisfying $\forall f \in \mathcal{C}, x \in \mathbb{Z}_p \colon f(x, \ldots, x) = 0$. Then we can solve* CompRep($\mathcal{C}$) *in polynomial time.*

Proof. By Lemma 15, $\mathcal{C}^{(n)}$ is the linear closure of B_n. Thus $\mathcal{C}(\mathbf{u}_1, \ldots, \mathbf{u}_n)$ is equal to the linear closure of $B_n(\mathbf{u}_1, \ldots, \mathbf{u}_n) := \{f^{\mathbf{c}}(\mathbf{u}_1, \ldots, \mathbf{u}_n) \mid f \in B, \mathbf{c} \in C_n\}$.

Note that the i-th entry $f^{\mathbf{c}}(\mathbf{u}_1, \ldots, \mathbf{u}_n)(i)$ of such a generating element can only be different from 0 if $(\mathbf{u}_1, \ldots, \mathbf{u}_n)(i)$ lies in the 2-dimensional subspace generated by the diagonal Δ^n and $\mathbf{c}$. Thus, there are at most k many vectors $\mathbf{c} \in C_n$ such that $f^{\mathbf{c}}(\mathbf{u}_1, \ldots, \mathbf{u}_n) \neq \mathbf{0}$, let $\mathbf{c}_1, \ldots, \mathbf{c}_l$ be an enumeration of them. Clearly $D = \{f^{\mathbf{c}}(\mathbf{u}_1, \ldots, \mathbf{u}_n) \mid f \in B, \mathbf{c} \in \{\mathbf{c}_1, \ldots, \mathbf{c}_l\}\}$ generates $\mathcal{C}(\mathbf{u}_1, \ldots, \mathbf{u}_n)$; note that we can compute it in linear time $O(kn)$. From the generating set D we can compute a compact representation of $\mathcal{C}(\mathbf{u}_1, \ldots, \mathbf{u}_n)$ in polynomial time (by generalized Gaussian elimination, or Theorem 2). ◀

5.2 General $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoids satisfying $p \nmid |L|$

For an arbitrary $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoid $\mathcal{C}$, let us define the subclonoid $\mathcal{C}_\Delta = \{f \in \mathcal{C} \colon f(x, \ldots, x) = 0\}$. We then show, that every $f \in \mathcal{C}$ can be written in a unique way as the sum of an element of $\mathcal{C}_\Delta$, and a function that is generated by $\mathcal{C}^{(1)}$. For this, we need the following lemma:

▶ **Lemma 17.** *For any $f \in \mathcal{C}^{(n)}$, let us define $f'(\mathbf{x}) = f(x_1, x_1, \ldots, x_1)$. Then $f - f' \in \mathcal{C}_\Delta$, and f' is generated by $\mathcal{C}^{(1)}$.*

Proof. trivial. ◀

It follows in particular from Lemma 17 and Lemma 15 that every $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoid $\mathcal{C}$ is generated by any set $A \cup B$, such that A generates $\mathcal{C}^{(1)}$ in $\mathbf{L}^{\mathbb{Z}_p}$ and B generates $\mathcal{C}_\Delta^{(2)}$ in $\mathbf{L}^{\mathbb{Z}_p^2}$. Note that the clonoid generated by A does not need to be disjoint from $\mathcal{C}_\Delta$. We can, however, still prove results analogous to the previous section.

▶ **Lemma 18.** *Let $\mathcal{C}$ be a $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoid, let A be a generating set of $\mathcal{C}^{(1)} \leq \mathbf{L}^{\mathbb{Z}_p}$ and B a generating set of $\mathcal{C}_\Delta^{(2)} \leq \mathbf{L}^{\mathbb{Z}_p^2}$. For every n, let us define $A_n = \{\sum_{\mathbf{a} \in \mathbb{Z}_p^n, \mathbf{a} \cdot \mathbf{1} = 1} f(\mathbf{a} \cdot \mathbf{x}) \mid f \in A\}$ and let B_n be defined as in Lemma 15. Then $A_n \cup B_n$ is a generating set of $\mathcal{C}^{(n)}$ in $\mathbf{L}^{\mathbb{Z}_p^n}$.*

Proof. We already know from Lemma 15 that B_n, generates $\mathcal{C}_\Delta^{(n)} \leq \mathbf{L}^{\mathbb{Z}_p^n}$.

By Lemma 17, every element $f \in \mathcal{C}^{(n)}$ can be uniquely written as the sum f' and $f - f'$. Furthermore f', by definition, is generated by A_n, and $f - f'$ is in $\mathcal{C}_\Delta^{(n)}$, which finishes our proof. ◀

Lemma 18 allows us to straightforwardly generalize Lemma 16 to arbitrary $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoids:

▶ **Lemma 19.** *Let $\mathcal{C}$ be a $(\mathbb{Z}_p^{id}, \mathbf{L})$-clonoid. Then* $\mathrm{CompRep}(\mathcal{C}) \in \mathsf{P}$.

Proof. Let A_n and B_n be defined as in Lemma 18. Our goal is to compute a compact representation of $\mathcal{C}(\mathbf{u}_1, \dots, \mathbf{u}_n)$ for some given $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{Z}_p^k$. By Lemma 18, every $g \in \mathcal{C}$ decomposes into the sum of g' and $g - g'$, where g' is generated by A_n and $g - g'$ is generated by B_n. Thus any image $g(\mathbf{u}_1, \dots, \mathbf{u}_n)$ is in the linear closure of all tuples $f(\mathbf{u}_1, \dots, \mathbf{u}_n)$, for $f \in A_n$ and $B_n(\mathbf{u}_1, \dots, \mathbf{u}_n) = \{f(\mathbf{u}_1, \dots, \mathbf{u}_n) \mid f \in B, \in C_n\}$ in $\mathbf{L}^k$. There are at most $|A|$-many tuples of the first form. Furthermore, as in the proof of Lemma 16 we can compute a generating set of $B_n(\mathbf{u}_1, \dots, \mathbf{u}_n)$ in polynomial time. By generalized Gaussian elimination (or Theorem 2), we can obtain a compact representation from these generators in polynomial time. ◀

Lemma 19 allows us to finish the proof of our main result:

▶ **Theorem 20.** *Let $\mathbf{A}$ be a finite Mal'tsev algebra, with a central series $0_\mathbf{A} < \rho < 1_\mathbf{A}$ such that $|\mathbf{A}/\rho| = p$ is a prime, and the blocks of ρ are of size coprime to p. Then* $\mathrm{SMP}(\mathbf{A}) \in \mathsf{P}$.

Proof. By Theorem 5, $\mathbf{A}$ is isomorphic to a wreath product $\mathbf{L} \otimes \mathbf{U}$, such that $\mathbf{U}$, $\mathbf{L}$ are affine with $|U| = p$ and $|L|$ coprime to p (as $|L|$ is the size of every block of ρ). By Theorem 11, $\mathrm{SMP}(\mathbf{A})$ reduces to $\mathrm{CompRep}(\mathrm{Diff}_0(\mathbf{A}))$ in polynomial time. The difference clonoid is a clonoid from $\mathbf{U}$ to $(\mathbf{L}, 0)$. Since both $\mathbf{L}$ and $\mathbf{U}$ are affine, and therefore have term operations describing $x - y + z$, $\mathrm{Diff}_0(\mathbf{A})$ is also a clonoid from $\mathbb{Z}_p^{id}$ to $(L, +, 0, -)$. By Lemma 19, $\mathrm{CompRep}(\mathrm{Diff}_0(\mathbf{A}))$ is solvable in polynomial time, which finishes the proof. ◀

▶ **Corollary 21.** *For every nilpotent Mal'tsev algebra $\mathbf{A}$ with $|A| = pq$ for distinct primes $p \neq q$, we have* $\mathrm{SMP}(\mathbf{A}) \in \mathsf{P}$.

Proof. If $\mathbf{A}$ is affine, then the result holds by (generalized) Gaussian elimination. So assume that $\mathbf{A}$ is 2-nilpotent, but not affine. So $\mathbf{A}$ is isomorphic to $\mathbf{L} \otimes \mathbf{U}$, and wlog. $|L| = q$ and $|U| = p$. Then the result follows directly from Theorem 20. ◀

6 Discussion

In Theorem 20 we proved that every Mal'tsev algebra, which can be written as a wreath product $\mathbf{L} \otimes \mathbf{U}$ with $|U| = p$ and $p \nmid |L|$ has a tractable subpower membership problem. But, since the reduction discussed in Theorem 11 extends beyond this case, it is natural to ask, whether we can also extend the tractability also extends to all those cases:

▶ **Question 22.** *Is* $\mathrm{SMP}(\mathbf{L} \otimes \mathbf{U}) \in \mathsf{P}$ *for every supernilpotent Mal'tsev algebra* $\mathbf{U}$*?*

In particular, if $\mathbf{U}$ is affine, Question 22 asks, whether the subpower membership problem of all finite 2-nilpotent Mal'tsev algebras can be solved in polynomial time. By Theorem 11, this reduces to computing compact representations with respect the clonoids between affine algebras. Thus answering the question requires a better understanding of such clonoids.

The recent preprint [27] studies such clonoids in the case where $\mathbf{U}$ has a distributive congruence lattice, and $\mathbf{L}$ is coprime to $\mathbf{U}$. Such clonoids are always generated by functions of bounded arity (as in Lemma 14), thus we expect a similar argument as in Lemma 19 to work in solving $\mathrm{CompRep}(\mathcal{C})$. We remark that, in the case of the clonoid of *all* operations from $\mathbf{U}$ and $\mathbf{L}$ this was already implicitly shown in [18] to obtain a polynomial time algorithm for checking whether two circuits over a 2-nilpotent algebra are equivalent. However [27] does not cover all clonoids between affine algebras; e.g. for the case $\mathbf{U} = \mathbb{Z}_p \times \mathbb{Z}_p$ and coprime $\mathbf{L}$ nothing is known so far.

A reason why much emphasis is placed on coprime $\mathbf{U}$ and $\mathbf{L}$ is, that their wreath products $\mathbf{L} \otimes^{T,0} \mathbf{U}$ are not supernilpotent (for non-trivial operations T), and therefore not covered by Theorem 2. In fact, finite Mal'tsev algebras in finite language are supernilpotent if and only if they decompose into the direct product of nilpotent algebras of prime power size (see e.g. [2, Lemma 7.6.]). It is further still consistent with our current knowledge that the conditions of Theorem 10 are always met, for coprime $\mathbf{L}$ and $\mathbf{U}$. This naturally leads to the question:

▶ **Question 23.** *Is* $\mathsf{Clo}(\mathbf{L} \times \mathbf{U}) \subseteq \mathsf{Clo}(\mathbf{L} \otimes \mathbf{U})$*, for all finite nilpotent Mal'tsev algebras* $\mathbf{L} \otimes \mathbf{U}$ *where* $\mathbf{L}$ *and* $\mathbf{U}$ *have coprime size?*

In fact, in an unpublished proof [20], a positive answer to Question 23 is given in the case that $\mathsf{Clo}(\mathbf{L} \otimes \mathbf{U})$ contains a constant operation. A more general version of Question 23 would ask, whether every finite nilpotent Mal'tsev algebra $\mathbf{A}$ has a Mal'tsev term m, such that (A, m) is supernilpotent.

Lastly we would like to mention that recently the property of *short pp-defitions* was suggested as a witnesses for $\mathrm{SMP}(\mathbf{A}) \in \mathsf{coNP}$. While Mal'tsev algebras that generate residually finite varieties have short pp-definitions [10], it is not know whether this is true in the nilpotent case. Thus we ask:

▶ **Question 24.** *Does every finite nilpotent Mal'tsev algebras* $\mathbf{A}$ *have short pp-definitions (and hence* $\mathrm{SMP}(\mathbf{A}) \in \mathsf{NP} \cap \mathsf{coNP}$*)?*

Studying Question 24 might especially be a useful approach to discuss the complexity for algebras of high nilpotent degree, if studying the corresponding difference clonoids turns out to be too difficult or technical.

References

1 Erhard Aichinger and Peter Mayr. Finitely generated equational classes. *Journal of Pure and Applied Algebra*, 220(8):2816–2827, 2016. doi:10.1016/j.jpaa.2016.01.001.

2 Erhard Aichinger and Nebojša Mudrinski. Some applications of higher commutators in Mal'cev algebras. *Algebra universalis*, 63(4):367–403, 2010. doi:10.1007/s00012-010-0084-1.

3 Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them. In *Dagstuhl Follow-Ups*, volume 7. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/DFU.Vol7.15301.1.

4 Clifford Bergman. *Universal algebra: Fundamentals and selected topics*. CRC Press, 2011.

5 Zarathustra Brady. Notes on CSPs and Polymorphisms. arXiv preprint arXiv:2210.07383v1, 2022.

6 Bruno Buchberger. Bruno Buchberger's PhD thesis 1965: An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal. *Journal of symbolic computation*, 41(3-4):475–511, 2006. doi:10.1016/j.jsc.2005.09.007.

7 Andrei Bulatov and Víctor Dalmau. A simple algorithm for Mal'tsev constraints. *SIAM Journal on Computing*, 36(1):16–27, 2006. doi:10.1137/050628957.

8 Andrei Bulatov, Marcin Kozik, Peter Mayr, and Markus Steindl. The subpower membership problem for semigroups. *International Journal of Algebra and Computation*, 26(07):1435–1451, 2016. doi:10.1142/S0218196716500612.

9 Andrei Bulatov, Peter Mayr, and Ágnes Szendrei. The subpower membership problem for finite algebras with cube terms. *Logical Methods in Computer Science*, 15(1), 2019. doi:10.23638/LMCS-15(1:11)2019.

10 Jakub Bulín and Michael Kompatscher. Short definitions in constraint languages. In *Proceedings of the 48th International Symposium on Mathematical Foundations of Computer Science, (MFCS 2023)*, volume 272 of *LIPIcs*, pages 28:1–28:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.MFCS.2023.28.

11 Stanley Burris and Hanamantagouda Sankappanavar. *A course in universal algebra*, volume 78. Springer, 1981.

12 Victor Dalmau and Peter Jeavons. Learnability of quantified formulas. *Theoretical Computer Science*, 306(1-3):485–511, 2003. doi:10.1016/S0304-3975(03)00342-6.

13 Stefano Fioravanti. Closed sets of finitary functions between finite fields of coprime order. *Algebra universalis*, 81(4):Art. No 52, 2020. doi:10.1007/s00012-020-00683-5.

14 Ralph Freese and Ralph McKenzie. *Commutator theory for congruence modular varieties*, volume 125. CUP Archive, 1987.

15 Merrick Furst, John Hopcroft, and Eugene Luks. Polynomial-time algorithms for permutation groups. In *21st Annual Symposium on Foundations of Computer Science (sfcs 1980)*, pages 36–41. IEEE, 1980.

16 Christian Herrmann. Affine algebras in congruence modular varieties. *Acta Universitatis Szegediensis*, 41:119–11, 1979.

17 Paweł Idziak, Petar Marković, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Tractability and learnability arising from algebras with few subpowers. *SIAM Journal on Computing*, 39(7):3023–3037, 2010.

18 Piotr Kawałek, Michael Kompatscher, and Jacek Krzaczkowski. Circuit equivalence in 2-nilpotent algebras. arXiv preprint arXiv:1909.12256, accepted for publication in STACS 2024.

19 Donald E Knuth. Efficient representation of perm groups. *Combinatorica*, 11(1):33–43, 1991. doi:10.1007/BF01375471.

20 Michael Kompatscher, Peter Mayr, and Patrick Wynne. Personal communication. presented at AAA102, Szeged, 2022. URL: https://www.math.u-szeged.hu/aaa102/Slides/Mayr,%20Peter2.pdf.

21 Dexter Kozen. Complexity of finitely presented algebras. In *Proceedings of the 9th annual ACM Symposium on Theory of Computing (STOC)*, pages 164–177, 1977. doi:10.1145/800105.803406.

22 Dexter Kozen. Lower bounds for natural proof systems. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 254–266, 1977. doi:10.1109/SFCS.1977.16.

23 Marcin Kozik. A finite set of functions with an EXPTIME-complete composition problem. *Theoretical Computer Science*, 407(1-3):330–341, 2008.

24 Ernst W Mayr and Albert R Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in mathematics*, 46(3):305–329, 1982. doi:10.1016/0001-8708(82)90048-2.

25 Peter Mayr. The subpower membership problem for Mal'cev algebras. *International Journal of Algebra and Computation*, 22(07):1250075, 2012. doi:10.1142/S0218196712500750.

26 Peter Mayr. Vaughan–Lee's nilpotent loop of size 12 is finitely based. *Algebra universalis*, 85(1):1–12, 2024. doi:10.1007/s00012-023-00832-6.

27 Peter Mayr and Patrick Wynne. Clonoids between modules. arXiv preprint arXiv:2307.00076, 2023.

28 Vladimir Shpilrain and Alexander Ushakov. Thompson's group and public key cryptography. In *Applied Cryptography and Network Security: Third International Conference, ACNS 2005, New York, NY, USA, June 7-10, 2005. Proceedings 3*, pages 151–163. Springer, 2005. doi:10.1007/11496137_11.

29 Vladimir Shpilrain and Gabriel Zapata. Using the subgroup membership search problem in public key cryptography. *Contemporary Mathematics*, 418:169, 2006. doi:10.1090/conm/418/07955.

30 Charles C Sims. Computational methods in the study of permutation groups. In *Computational problems in abstract algebra*, pages 169–183. Elsevier, 1970. doi:10.1016/B978-0-08-012975-4.50020-5.

31 Joel VanderWerf. Wreath products of algebras: generalizing the Krohn-Rhodes theorem to arbitrary algebras. *Semigroup Forum*, 52(1):93–100, 1996. doi:10.1007/BF02574084.

32 Ross Willard. Four unsolved problems in congruence permutable varieties, 2007. Talk at the Conference on Order, Algebra, and Logics, Nashville. URL: https://www.math.uwaterloo.ca/~rdwillar/documents/Slides/willard_nashville_2007_slides.pdf.

A Proof of Lemma 3

In this appendix we prove the second statement of Lemma 3, i.e. we show that for a given *enumerated* compact representation R of a subpower $\mathbf{R} = \mathrm{Sg}_{\mathbf{A}^k}(X)$ of some Mal'tsev algebra, we can obtain an *enumerated* compact representation R' of $\mathbf{R} \cap \{\mathbf{x} \in A^k \mid \mathbf{x}(1) = a_1, \dots, \mathbf{x}(k) = a_k\}$ for a given list of constants $a_1, \dots, a_k$. In Algorithm 1 we describe the algorithm $\texttt{Fix-value}(R, a)$ that fixes the first coordinate of $\mathbf{R}$ to a; iterating this algorithm m times results in the statement of the Lemma.

We remark that $\texttt{Fix-value}(R, a)$ is based on the `Fix-values` algorithm in [5, Algorithm 5]); although, for simplicity, we only fix the value of one coordinate. Line 7 and 8 corresponds to the call of the subroutine `Nonempty` in [5, Algorithm 5]), with the difference that we compute all the elements of the set $T_j = \{(x,y) \in \mathrm{pr}_{1,j}\mathbf{R} \mid x = a\}$, instead of computing a witness for $(a,y) \in T_j$ once at a time.

We claim that the running time of $\texttt{Fix-Value}(R, a)$ is $O(|A|^2 \cdot n))$. For this note that the exhaustive search in line 7 and 8 will simply recursively apply m to elements from $\mathrm{pr}_{1,j}(R)$ until the set is closed under m, and then select all values with $x_1 = a$. Since $|\mathrm{pr}_{1,j}\mathbf{R}| \leq A^2$ this takes at most $|A|^2$ steps. For this reason also the size of the defining circuits C_j (when e.g. stored all together as a circuit with multiple out-gates) is bounded by $|R| + |A|^2$. Since the for-loop of line 6 has at most n iterations, it follows that both the running time of the algorithm and the size of the defining circuits in R' are bounded by $O(|A|^2 \cdot n)$.

If we then repeatedly call `Fix-value` to fix the value of the first m-many values of $\mathbf{R}$, this results in an algorithm that runs in time $O(|A|^2 \cdot nm))$.

Thus, the only thing that remains to prove is that the algorithm `Fix-Value` is correct. i.e. it indeed outputs an enumerated R' with $\mathrm{Sig}(R') = \mathrm{Sig}(\mathbf{R} \cap \{\mathbf{x} \in A^k \mid \mathbf{x}(1) = a\})$ (if the output is not empty). So assume that $(i, b, c) \in \mathrm{Sig}(\mathbf{R} \cap \{\mathbf{x} \in A^k \mid \mathbf{x}(1) = a\})$. If $i = 1$, then clearly $(i, b, c) = (1, a, a)$, which is in $\mathrm{Sig}(R')$. So let us assume wlog. that $i > 1$. Since R is a compact representation of $\mathbf{R}$, there exist tuples $\mathbf{r}_b, \mathbf{r}_c \in R$ (and defining circuits $p_{\mathbf{r}_b}$ and $p_{\mathbf{r}_c}$), witnessing that $(i, b, c) \in \mathrm{Sig}(R)$. Then R' contains the tuples $\mathbf{t}$ and $\mathbf{s} = m(\mathbf{t}, \mathbf{r}_b, \mathbf{r}_c)$, as constructed in line 12 and 13 of Algorithm 1. Since $\mathbf{r}_b$ and $\mathbf{r}_c$ agree on the first $i - 1$ coordinates also $\mathbf{t}$ and $\mathbf{s}$ do. Moreover $\mathbf{t}(1) = a$, $\mathbf{t}(i) = b$, and $\mathbf{s}(i) = m(b, b, c) = c$, thus $\mathbf{t}$ and $\mathbf{s}$ witness $(i, b, c) \in \mathrm{Sig}(\mathbf{R} \cap \mathbf{x}(1) = a)$. It follows that $\mathrm{Sig}(R') = \mathrm{Sig}(\mathbf{R} \cap \{\mathbf{x} \in A^k \mid \mathbf{x}(1) = a\})$, which is what we wanted to prove.

Algorithm 1 An algorithm that, for a given enumerated compact representations R of $\mathbf{R} = \mathrm{Sg}_{\mathbf{A}^k}(X)$ outputs an enumerated compact representation R' of the relation that fixes $x_1 = a$, (where the defining circuits of R' are evaluated on X).

```
1:  procedure Fix-Value($a \in A$, $R$ (enum. c.r. of $\mathbf{R} = \mathrm{Sg}_{\mathbf{A}^k}(X)$), Mal'tsev term $m$)
2:      if $(1, a, a) \notin \mathrm{Sig}(R)$ then return $\emptyset$
3:      else
4:          Let $(\mathbf{t}, p_{\mathbf{t}}) \in R$ be such that $(\mathbf{t}, \mathbf{t})$ is a witness of $(1, a, a) \in \mathrm{Sig}(R)$.
5:          $R' = \{(\mathbf{t}, p_{\mathbf{t}})\}$
6:          for $j > 1$ do
7:              Recursively apply $m$ to $\mathrm{pr}_{1,j}(R)$ to compute $T_j = \{(x, y) \in \mathrm{pr}_{1,j}(\mathbf{R}) \mid x = a\}$,
8:              and circuits $C_j = \{p_{(x,y)} \mid (x, y) \in T_j\}$ such that $\mathrm{pr}_{1,j}(p_{(x,y)}(X)) = (x, y)$.
9:              for $(j, b, c) \in \mathrm{Sig}(R)$ do
10:                 Let $(\mathbf{r}_b, p_{\mathbf{r}_b}), (\mathbf{r}_c, p_{\mathbf{r}_c}) \in R$ be witnesses of $(j, b, c) \in \mathrm{Sig}(R)$
11:                 if $(a, b) \in T_j$ then
12:                     Let $\mathbf{t} = p_{(a,b)}(X)$
13:                     $\mathbf{s} = m(\mathbf{t}, \mathbf{r}_b, \mathbf{r}_c)$ and $p_{\mathbf{s}} = m(p_{(a,b)}, p_{\mathbf{r}_b}, p_{\mathbf{r}_c})$
14:                     $R' = R' \cup \{(\mathbf{t}, p_{(a,b)}), (\mathbf{s}, p_{\mathbf{s}})\}$
15:     return $R'$
```