Short Definitions in Constraint Languages

Jakub Bulín
Department of Theoretical Computer Science and Mathematical Logic,
Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic

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

Abstract

A first-order formula is called *primitive positive (pp)* if it only admits the use of existential quantifiers and conjunction. Pp-formulas are a central concept in (fixed-template) constraint satisfaction since CSP(Γ) can be viewed as the problem of deciding the primitive positive theory of Γ, and pp-definability captures gadget reductions between CSPs.

An important class of tractable constraint languages Γ is characterized by having *few subpowers*, that is, the number of n-ary relations pp-definable from Γ is bounded by $2^{p(n)}$ for some polynomial $p(n)$. In this paper we study a restriction of this property, stating that every pp-definable relation is definable by a pp-formula of polynomial length. We conjecture that the existence of such *short definitions* is actually equivalent to Γ having few subpowers, and verify this conjecture for a large subclass that, in particular, includes all constraint languages on three-element domains. We furthermore discuss how our conjecture imposes an upper complexity bound of co-NP on the subpower membership problem of algebras with few subpowers.

2012 ACM Subject Classification Theory of computation → Complexity theory and logic

Keywords and phrases constraint satisfaction, primitive positive definability, few subpowers, polynomially expressive, relational clone, subpower membership

Digital Object Identifier 10.4230/LIPIcs.MFCS.2023.28

Related Version *Previous Version*: https://arxiv.org/abs/2305.01984v1

Funding *Jakub Bulín*: Supported by the Charles University project UNCE/SCI/004 and the MŠMT ČR INTER-EXCELLENCE project LTAUSA19070.
Michael Kompatscher: Supported by the Charles University project UNCE/SCI/022 and the MŠMT ČR INTER-EXCELLENCE project LTAUSA19070.

Acknowledgements The authors would like to thank Dmitriy Zhuk for inspiring discussions about critical relations and the anonymous reviewers for their valuable suggestions.

1 Introduction

Constraint satisfaction is a unifying framework for expressing a wide range of computational tasks coming from a smorgasbord of real-life applications and theoretical contexts. In a CSP instance, the goal is to assign values to variables subject to a list of *constraints* to be satisfied. In the most general setting, a constraint consists of a tuple of variables (its *scope*) and a list of admissible evaluations of the scope (i.e., tuples of values, forming the *constraint relation*). Usually, the set of variables, the sets of admissible values for every variable (its *domain*), and the list of input constraints are all finite. This simple formulation strikes a "perfect balance between generality and structure" [3].

In this general formulation, the CSP is an NP-complete problem: for example, SAT or graph 3-colorability are easily expressible in this framework. However, many problems subsumed by it are tractable, e.g., 2-SAT, Horn-SAT, or checking the consistency of a system of linear equations over $\mathbb{Z}_p$. A natural way to explore the complex landscape of the CSP,

justified by applications as well as theory [13], is to fix a finite domain A and a finite set of relations Γ on A that are allowed to appear as constraints (a *constraint language*) [25]; such fragment of the CSP is usually denoted by $\mathrm{CSP}(\Gamma)$. (Sometimes, constraint languages on infinite domains, or with infinitely many relations are considered. But for simplicity, following, e.g., [3], we keep the two standard finiteness assumptions throughout the paper.)

The CSP dichotomy theorem [10, 27, 28] states that for every constraint language Γ, $\mathrm{CSP}(\Gamma)$ is in P or NP-complete. To tame the vast landscape of constraint languages, it was immensely helpful to realize that various ad hoc "gadget" complexity reductions share a common explanation using the notion of *primitive positive* (*pp*-) *definability* (i.e., the usual first-order logic definability restricted to $\{\exists, \wedge, =\}$-formulas) [19, 18] and the more general notions of *pp-interpretability* and *pp-constructibility* [4]. In fact, the CSP dichotomy theorem implies that $\mathrm{CSP}(\Gamma)$ is NP-complete if and only if Γ pp-constructs *every* finite constraint language. Moreover, pp-definability and its generalizations have an external characterization via so-called *polymorphisms* ("multivariate homomorphisms") [15, 7, 17]. For an introduction to the area, see [3].

Constraint languages Γ for which $\mathrm{CSP}(\Gamma)$ is solvable by a certain algorithmic approach involving computing with *compact representations* of solution sets (generalizing *bases* of vector spaces or *strong generating sets* from the Schreier-Sims algorithm for permutation groups [26]) were characterized in [6, 16] as those that have *few subpowers*, that is, the number of n-ary relations pp-definable from Γ is bounded by $2^{p(n)}$ for some polynomial $p(n)$. This property, also called *polynomial expressiveness* [12], is equivalent to having either of the following two properties where *small* means of size bounded by a polynomial in the arity n:

- *small generating sets*, i.e., every relation pp-definable from Γ has a small subset that is not contained in any proper pp-definable subset,
- *small independent sets*, i.e., sets of tuples such that every tuple can be separated from the remaining tuples by a pp-formula, are small.

The equivalence of those properties was established in [6, Proposition 1.4].

In this paper, we study another measure of "smallness" of a constraint language, that we call *short pp-definitions*: every n-ary relation pp-definable from Γ is definable by some primitive positive formula of polynomial length. Examples include constraint languages encoding 2-SAT, or the consistency of linear systems over $\mathbb{Z}_p$.

A simple cardinality argument shows that a constraint language with short pp-definitions must have few subpowers. We conjecture that the converse is also true and thus the two properties are equivalent.

▶ **Conjecture 1.** *A constraint language has short pp-definitions, if and only if it has few subpowers.*

We remark that exponential-length pp-definitions are needed for constraint languages without few subpowers (cf. [6, Theorem 3.12]).

In Section 2 we give a formal definition of short pp-definitions, examples, and an exposition of related properties. An equivalent condition (definability by pp-formulas with polynomially many existential quantifiers) was studied in [23] for Boolean constraint languages (i.e., constraint languages on a two-element domain) under the name *polynomial closedness*. It can be easily seen that Conjecture 1 is true in the Boolean case, as stated in [23, Corollary 1].

In Section 4 we prove Theorem 21, the main result of our paper, which confirms the conjecture for a substantial class of constraint languages, namely those whose *polymorphism algebra* generates a *residually finite variety*. This, in particular, implies that Conjecture 1 also holds for constraint languages on three-element domains (Corollary 23). The proof

proceeds by first reducing to the case of *critical* relations (see [29]), and then employing structural theorems from universal algebra, in a similar fashion as in [9]. We explain some necessary background from universal algebra in Section 3.

Apart from being a natural property in constraint satisfaction, in Section 5 we argue that short pp-definitions have further applications in the study of the *subpower membership problem* SMP($\mathbf{A}$) over an algebraic structure $\mathbf{A}$ (see [21, 24, 9]), i.e., the problem of deciding whether a given list of tuples over $\mathbf{A}$ generates another tuple over $\mathbf{A}$. For algebras $\mathbf{A}$ with few subpowers, *compact representations* provide a natural certificate for "Yes"-instances; this was used to show that SMP($\mathbf{A}$) $\in$ NP [20]. We argue that short pp-definitions can serve as a natural certificate for "No"-instances. In particular, we show how short pp-definitions impose an upper complexity bound of co-NP on the subpower membership problem. Thus, Conjecture 1 would imply that SMP($\mathbf{A}$) $\in$ NP $\cap$ co-NP, for all algebras $\mathbf{A}$ with few subpowers.

2 Preliminaries

Let A be a finite set. An n-ary relation R on A is any subset of n-tuples $R \subseteq A^n$. By a *constraint language* on A (its *domain*) we mean any finite set $\Gamma = \{R_1, \dots, R_m\}$ of relations on A of arbitrary, but finite arities.

A relation R is *primitive positive definable* (or *pp-definable* for short) from Γ, if it is definable in first-order logic by a formula using only the relations from Γ, the equality relation, conjunction, and existential quantification. Equivalently, in prenex normal form:

$$R(x_1, \dots, x_n) \leftrightarrow \exists y_1 \exists y_2 \dots \exists y_k \bigwedge_{i \in \{1,\dots,C\}} S_i(z_1^i, \dots, z_{r_i}^i)$$

where S_i is an r_i-ary relational symbol representing a relation from $\Gamma \cup \{=_A\}$ and $z_j^i \in \{x_1, \dots, x_n, y_1, \dots, y_k\}$. We remark that CSP($\Gamma$) can be defined as the problem of deciding the primitive positive fragment of the first-order theory of Γ.

The set of all relations pp-definable from Γ, denoted by $\langle\Gamma\rangle$, forms a *relational clone*, i.e., a set of relations on A containing the identity relation and closed under intersections, direct products, projections, and permutations of coordinates. Any constraint language Γ that generates a relational clone $\mathcal{R} = \langle\Gamma\rangle$ is called a *relational basis* of $\mathcal{R}$. Let us denote by $\langle\Gamma\rangle_n$ the set of all n-ary relations pp-definable from Γ.

The usefulness of pp-definability for the CSPs is summarized in the following theorem going back to [19]. For a modern exposition as well as generalizations see [3] and [4].

▶ **Theorem 2.** *If Γ and Δ are constraint languages such that $\Delta \subseteq \langle\Gamma\rangle$, then there is a logspace reduction from* CSP(Δ) *to* CSP(Γ).

In order to put Conjecture 1 on a firm footing, let us next formally define the notion of *few subpowers* [6, 16] and the central concept of the present paper, *short pp-definitions*.

▶ **Definition 3.** *A constraint language Γ has* few subpowers, *if there exists a polynomial $p(n)$ such that $|\langle\Gamma\rangle_n| \leq 2^{p(n)}$ for all $n > 0$.*

▶ **Definition 4.** *Let Γ be a constraint language. We say that Γ has:*

- pp-definitions of length [at most] $f(n)$, *if for every $n > 0$ and every $R \in \langle\Gamma\rangle_n$, R is definable from Γ by a primitive positive formula ϕ of length $|\phi| \leq f(n)$.*
- short pp-definitions *if Γ has pp-definitions of length $p(n)$ for some polynomial $p(n)$.*

Here we consider the length $|\phi|$ to be simply the number of symbols in some syntactical representation of the formula. In the definition of *short pp-definitions*, one could alternatively bound the number of atomic formulas in ϕ, or the number k of existentially quantified

variables by a polynomial $p(n)$. (The latter option was used in [23] in the notion of *polynomial closedness* of $\langle\Gamma\rangle$.) Note that, since Γ is fixed and finite, these three possible definitions coincide.

Clearly, having few subpowers is a property of the relational clone $\langle\Gamma\rangle$, independent of the choice of the relational basis Γ. We observe that the same is true for short pp-definitions. In fact, up to multiplication by a scalar, this is true for any bound $f(n)$ on the length of pp-definitions:

▶ **Lemma 5.** *Let Γ and Δ be constraint languages such that $\langle\Gamma\rangle = \langle\Delta\rangle$. If Γ has pp-definitions of length $f(n)$, then Δ has pp-definitions of length $O(f(n))$. In particular, Γ has short pp-definitions if and only if Δ does.*

Proof. Let $R \in \langle\Delta\rangle_n = \langle\Gamma\rangle_n$. By assumption, R has a pp-definition ϕ_R from Γ of length at most $f(n)$. Since $\Gamma \subseteq \langle\Gamma\rangle = \langle\Delta\rangle$, every relation $S \in \Gamma$ can be defined from Δ by some pp-formula ψ_S. Let $c = \max\{|\psi_S| \colon S \in \Gamma\}$. If we replace every atomic formula $S_i(z_1^i, \dots, z_{r_i}^i)$ in ϕ_R by a suitable variant of the formula ψ_{S_i}, we obtain a pp-definition of R from Δ of length at most $c \cdot f(n)$. ◀

Central to the algebraic approach to the CSP is the idea that constraint languages up to pp-definability (that is, relational clones) can be characterized by their *polymorphisms*. Following the terminology from [3], a k-ary operation $f : A^k \to A$ is *compatible* with an n-ary relation $R \subseteq A^n$, and R is *invariant* under f, if f applied coordinate-wise to any k n-tuples from R yields an n-tuple that is also in R. A *polymorphism* of a constraint language Γ is then any function on the domain that is compatible with all relations from Γ. As is usual, we write $\mathrm{Pol}(\Gamma)$ to denote the set of all polymorphisms of Γ and, similarly, $\mathrm{Inv}(\mathcal{F})$ for the set of all relations on the domain A invariant under a set of operations $\mathcal{F}$. The key connection between polymorphisms and pp-definability can be summarized in the following lemma.

▶ **Lemma 6** ([15, 7, 17])**.** *For any constraint language Γ, $\langle\Gamma\rangle = \mathrm{Inv}(\mathrm{Pol}(\Gamma))$.*

Few subpowers can be characterized by the existence of an *edge polymorphism*, that is, a polymorphism satisfying certain algebraic identities (under all evaluations of variables in the domain). Such characterizations are typical in the algebraic approach to the CSP.

▶ **Theorem 7** ([6, 16])**.** *A constraint language Γ has few subpowers, if and only if for some $k \geq 2$ there exists a k-*edge polymorphism $e \in \mathrm{Pol}(\Gamma)$*, that is, a $(k+1)$-ary operation $e \colon A^{k+1} \to A$ satisfying the following identities:*

$$
\begin{aligned}
e(y, y, x, x, x, \dots, x) &\approx x \\
e(y, x, y, x, x, \dots, x) &\approx x \\
e(x, x, x, y, x, \dots, x) &\approx x \\
e(x, x, x, x, y, \dots, x) &\approx x \\
&\vdots \\
e(x, x, x, x, x, \dots, y) &\approx x
\end{aligned}
$$

In this case $\mathrm{CSP}(\Gamma) \in \mathrm{P}$.

The following two special cases were important intermediate steps towards Theorem 7 (as well as the CSP dichotomy theorem) and, in particular, cover all *Boolean* (i.e., where $A = \{0, 1\}$) constraint languages with few subpowers:

- A *Mal'tsev* operation is a ternary operation $m\colon A^3 \to A$ satisfying the identities $m(x,x,y) \approx m(y,x,x) \approx y$. If m is a Mal'tsev operation, then $e(x_1,x_2,x_3) = m(x_2,x_1,x_3)$ is a 2-edge term.
- A *near-unanimity* operation (of arity $k \geq 3$) is an operation t satisfying the following identities:

$$x \approx t(y,x,\dots,x) \approx t(x,y,x,\dots,x) \approx t(x,\dots,x,y,x) \approx t(x,\dots,x,y)$$

Then $e(x_1,x_2,\dots,x_{k+1}) = t(x_2,\dots,x_{k+1})$ is a k-edge term. A ternary near-unanimity is called a *majority*.

Let us now give two examples of constraint languages with short pp-definitions; Example 8 is invariant under a Mal'tsev operation, while Example 11 has a majority polymorphism.

▶ **Example 8.** The problem of checking consistency of a linear system over $\mathbb{Z}_2$ can be encoded as CSP(Γ) for $\Gamma = \{R_{\mathrm{Lin}}, C_0, C_1\}$ where $R_{\mathrm{Lin}} = \{(0,0,0),(0,1,1),(1,0,1),(1,1,0)\}$ encodes "$x_1 + x_2 = x_3$", $C_0 = \{0\}$, and $C_1 = \{1\}$. Indeed, any linear equation $x_1 + x_2 + \dots + x_n = b$ can be encoded using auxiliary variables $y_1,\dots,y_{n-1}$ and three-variable equations

$$x_1 + x_2 = y_1, \quad y_1 + x_3 = y_2, \quad \dots, \quad y_{n-2} + x_n = y_{n-1}, \quad y_{n-1} = b$$

thus providing a pp-definition:

$$\exists y_1 \dots \exists y_{n-1}\, (R_{\mathrm{Lin}}(x_1,x_2,y_1) \wedge \dots \wedge R_{\mathrm{Lin}}(y_{n-2},x_n,y_{n-1}) \wedge C_b(y_{n-1}))$$

The relational clone $\langle\Gamma\rangle$ then consists of all affine subspaces of $\mathbb{Z}_2^n$, for any $n > 0$. The relations from the relational basis Γ are all affine and it is easy to see that relations pp-definable from affine subspaces are also affine subspaces. On the other hand, an affine subspace R of $\mathbb{Z}_2^n$ can be described by at most n linear equations and the conjunction of the corresponding pp-formulas clearly defines R. The length of this conjunction is in $O(n^2)$ and therefore Γ has not only few subpowers but also short (quadratic) pp-definitions. Γ is arguably one of the easiest examples of a constraint language with a Malt'sev polymorphism; in fact $\langle\Gamma\rangle = \mathrm{Pol}(\{m\})$ for the Mal'tsev operation $m(x,y,z) = x + y + z \bmod 2$.

While Conjecture 1 is open even under the presence of a Mal'tsev polymorphism, the above example can be generalized to a *central* Mal'tsev polymorphism, that is, one which is compatible with its own function graph, i.e., the 4-ary relation $R = \{(x,y,z,m(x,y,z)) \mid x,y,z \in A\}$. (The reason is that the *polymorphism algebra* is then *affine*, i.e., polynomially equivalent to a module.)

For the second example, we need the following characterization of relations invariant under near-unanimity operations:

▶ **Theorem 9** ([2])**.** *If a relation $R \subseteq A^n$ is invariant under a $(k+1)$-ary near-unanimity operation t, then it is pp-definable from its projections to at most k-ary subsets of coordinates by the following formula:*

$$R(x_1,\dots,x_n) \leftrightarrow \bigwedge_{\substack{I=\{i_1,\dots,i_l\}\\ I\subseteq[n],|I|\leq k}} \mathrm{proj}_I\, R(x_{i_1},\dots,x_{i_l})$$

Here, for a relation $R \subseteq A^n$ and a subset of coordinates $I \subseteq \{1,2,\dots,n\}$, the *projection* of R to $I = \{i_1,i_2,\dots,i_k\}$ (where $i_1 < i_2 < \dots < i_k$) is the k-ary relation $\mathrm{proj}_I R = \{(a_{i_1},\dots,a_{i_k}) \mid (a_1,\dots,a_n) \in R\} \in \langle\Gamma\rangle$. Thus, if a constraint language Γ has a $(k+1)$-ary near-unanimity polymorphism t, then every relation $R \in \langle\Gamma\rangle_n$ can be written as the conjunction of $\binom{n}{k}$ relations of arity k (for $n \geq k$). As a direct corollary of Theorem 9 we obtain:

▶ **Corollary 10.** *Let Γ be a constraint language with a $(k+1)$-ary near-unanimity polymorphism. Then Γ has pp-definitions of length $O(n^k)$.*

▶ **Example 11.** Corollary 10 can be exemplified by 2-SAT. The standard way to encode 2-SAT as a CSP is by using the constraint language $\Gamma_{\text{2-SAT}} = \{R_{00}, R_{01}, R_{10}, R_{11}\}$ where $R_{ij} = \{0,1\}^2 \setminus \{(i,j)\}$ represent each clause type (see [3, Example 2.2]). It is well known that the relations pp-definable from $\Gamma_{\text{2-SAT}}$ are exactly those invariant under the majority operation, i.e., the unique 3-ary near-unanimity operation on $\{0,1\}$. By Corollary 10, $\Gamma_{\text{2-SAT}}$ has quadratic pp-definitions.

▶ **Theorem 12** (see [23, Corollary 1])**.** *Let Γ be a Boolean constraint language with few subpowers, then Γ has quadratic pp-definitions. Thus Conjecture 1 holds for constraint languages Γ over Boolean domains.*

Proof. By the classification of Post's lattice, every Boolean constraint language with few subpowers has either the Mal'tsev polymorphism $x + y + z \bmod 2$ or the (unique) majority polymorphism (this was observed, e.g., in [12]). In both cases, we obtain quadratic pp-definitions, as in Examples 8 and 11. ◀

We remark that, in general, the situation is much more complicated than in Theorem 12: Already in the 3-element case there are constraint languages that have few subpowers, but neither a Mal'tsev, nor a near-unanimity polymorphism (e.g. [8, Examples 2.1.1 and 2.1.2]).

3 Universal Algebra

In the following, we are going to introduce some basic notions from universal algebra that will allow us to state our main result (Theorem 21) in its full generality. In Section 3.1 we furthermore discuss how short pp-definitions behave with respect to basic algebraic constructions. For more background in universal algebra we refer to the textbooks [5, 11].

An *algebra* $\mathbf{A} = (A; (f_i)_{i\in I}^{\mathbf{A}})$ 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* $\mathbf{B} = (B; (f_i)_{i\in I}^{\mathbf{B}})$ of an algebra $\mathbf{A} = (A; (f_i)_{i\in I}^{\mathbf{A}})$ (denoted $\mathbf{B} \leq \mathbf{A}$) is an algebra obtained by restricting all *basic operations* $f_i^{\mathbf{A}}$ to an invariant subset $B \subseteq A$. 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. A *homomorphism* $h\colon \mathbf{A} \to \mathbf{B}$ between algebras is defined as a map that preserves all basic operations, i.e., $h(f_i^{\mathbf{A}}(a_1, \dots, a_n)) = f_i^{\mathbf{B}}(h(a_1), \dots, h(a_n))$ for all $i \in I$. The *kernel* of every homomorphism (i.e., the relation defined by $(x,y) \in \theta \leftrightarrow h(x) = h(y)$) is a *congruence*, that is, an equivalence relation invariant under $\mathbf{A}$. Conversely, for every congruence α of $\mathbf{A}$, it is easy to see, that one can construct a *quotient algebra* $\mathbf{A}/\alpha$, as the homomorphic image of the quotient mapping $x \mapsto x/\alpha$. Under the inclusion order, the set of all congruence of an algebra $\mathbf{A}$ forms the *congruence lattice* $\mathrm{Con}(\mathbf{A})$. The minimal element of this lattice is always the trivial congruence $0_{\mathbf{A}} = \{(x,x) \mid x \in A\}$. An algebra $\mathbf{A}$ is called *subdirectly irreducible* if $0_{\mathbf{A}}$ has a unique cover in $\mathrm{Con}(\mathbf{A})$, i.e., there is a unique minimal non-trivial congruence.

By H, S, and P we denote the closure of a set of algebras under homomorphic images, subalgebras, and products respectively. It is well-known that the closure of any set of algebras under HSP is a *variety*, i.e., a class of algebras defined by a set of identities (by Birkhoff's theorem, see, e.g., [5]). A variety is called *residually finite* if (up to isomorphism) it only contains finitely many subdirectly irreducible algebras, all of which are finite.

3.1 Algebras with short pp-definitions

If we assign a function symbol to every element of Pol(Γ), for a constraint language Γ, then we can also regard Pol(Γ) as an algebra (the *polymorphism algebra* of Γ). On the other hand, for every algebra $\mathbf{A}$, its invariant relations Inv($\mathbf{A}$) form a relational clone. Thus, it makes sense to say that a finite algebra $\mathbf{A}$ has *few subpowers* if Inv($\mathbf{A}$) has few subpowers.

Note that a relation R is invariant under $\mathbf{A}$ if and only if $R \leq \mathbf{A}^n$ for some n, i.e., R is a subalgebra of some power of $\mathbf{A}$ (such R is also called a *subpower of* $\mathbf{A}$, which motivates the notion of having "few subpowers").

By a (non-constructive) proof of Aichinger, Mayr and McKenzie [1], for every algebra $\mathbf{A}$ with few subpowers there exists a finite relational basis Γ of Inv($\mathbf{A}$), i.e., a constraint language, such that Inv($\mathbf{A}$) = ⟨Γ⟩ (for general algebras $\mathbf{A}$ this is not the case). Thus, it makes sense to define the following:

▶ **Definition 13.** *An algebra* $\mathbf{A}$ *has* pp-definitions of length $f(n)$*, if there exists a constraint language* Γ *such that* Inv($\mathbf{A}$) = ⟨Γ⟩*, and* Γ *has pp-definitions of length* $f(n)$*. An algebra* $\mathbf{A}$ *has* short pp-definitions*, if it has pp-definitions of length* $p(n)$*, for some polynomial* p*.*

Note that, by Lemma 5, having short pp-definitions is independent of the choice of the relational basis Γ. By the following lemma, having short pp-definitions is also preserved under forming finite powers of algebras. The proof is provided in Section A.1 in the Appendix.

▶ **Lemma 14.** *Let* $\mathbf{A}$ *be an algebra and* $\mathbf{B} = \mathbf{A}^k$ *for some* $k > 1$*. Then* $\mathbf{B}$ *has pp-definitions of length* $O(f(n))$ *if and only if* $\mathbf{A}$ *has pp-definitions of length* $O(f(\lceil \frac{n}{k} \rceil))$*.*

In the following, we will also work with multi-sorted relations, as this provides a natural framework for our proof in Section 4. More specifically, if $\mathcal{A}$ is a finite set of finite algebras of the same language, then it still makes to study the set of all invariant relations $R \leq \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ for $\mathbf{A}_1, \ldots, \mathbf{A}_n \in \mathcal{A}$. In particular, the set of all such relations will still form a relational clone (where variables have possibly different domains A_i). If, furthermore, all elements of $\mathcal{A}$ have few subpowers, the finite relational basis result of Aichinger, Mayr and McKenzie [1] still applies, and it makes sense to define the property of having short pp-definitions for $\mathcal{A}$ (we refrain from giving technical details here). We remark that studying constraint languages in which the variables can come from different domains is a fairly standard viewpoint in CSP; it was, for example, used in the proof of the CSP dichotomy theorem by Zhuk [28].

This multisorted approach allows us to consider relations over the closure HS($\mathbf{A}$) of $\mathbf{A}$ under homomorphic images and subalgebras instead of only $\mathbf{A}$ itself. This is justified by the following lemma; the proof is provided in Section A.2 of the Appendix.

▶ **Lemma 15.** *An algebra* $\mathbf{A}$ *has pp-definitions of length* $O(f(n))$*, if and only if the family of algebras* HS($\mathbf{A}$) *has (multisorted) pp-definitions of length* $O(f(n))$*.*

In particular, Lemma 15 implies that $\mathbf{A}$ has short pp-definitions, if and only if HS($\mathbf{A}$) has (multisorted) short pp-definitions. We remark that Lemma 15 does *not* imply that any *single* algebra $\mathbf{B} \in$ HS($\mathbf{A}$) has short pp-definitions if $\mathbf{A}$ does. In fact, we do not know if this is true (see Question 25 in the Discussion section).

4 Main result

In this section, we prove the main result of our paper. We first need to introduce some standard definitions that found prominent use in the universal algebraic approach to constraint satisfaction before (see, e.g., [8, Chapter 2]).

A relation $R \leq \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ is called *critical* if it is $\wedge$-*irreducible*, i.e., it cannot be written as the intersection of strictly bigger relations $Q \leq \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$, and it has *no dummy variables*, i.e., it depends on all of its inputs.

A binary relation $R \subseteq A \times B$ has the *parallelogram property* if $(a,c),(a,d),(b,c) \in R$ implies $(b,d) \in R$. An n-ary relation $R \subseteq A_1 \times A_2 \times \ldots \times A_n$ has the *parallelogram property*, if for all subsets $I \subset \{1,2,\ldots,n\}$ it has the parallelogram property when considered as a binary relation $R \subseteq (\prod_{i \in I} A_i) \times (\prod_{j \notin I} A_j)$. The *signature* of R is the following set of triples:

$$\mathrm{Sig}(R) = \{(i,a,b) \in [n] \times A_i^2 \mid \exists \bar{x}, \bar{y} \in R \text{ with } x_j = y_j \text{ for } j = 1,\ldots,i-1 \text{ and } x_i = a, y_i = b\}$$

In the proof of Lemma 17 below, we will need the following straightforward observation.

▶ **Observation 16.** *If R has the parallelogram property, $R \subseteq S$, and $\mathrm{Sig}(R) = \mathrm{Sig}(S)$, then $R = S$.*

Using these notions, we can reduce the problem of finding short pp-definitions to critical relations with the parallelogram property:

▶ **Lemma 17.** *Let $\mathbf{A}$ be an algebra with a k-edge term. If all the critical relations $R \leq \mathbf{A}^n$ with parallelogram property have pp-definitions of length at most $f(n)$, then $\mathbf{A}$ has pp-definitions of length $O(n^k + n \cdot f(n))$.*

Proof. Clearly, every relation $R \leq \mathbf{A}^n$ is the intersection of the $\wedge$-irreducible relations above it (in the inclusion order), thus it can be written as a conjunction of critical relations. Hence we only need to give an upper bound on the number of such critical relations.

By [20, Theorem 3.6], the presence of a k-edge term implies that every critical relation is either of arity $\leq k$, or has the parallelogram property. This further implies (see, e.g., [8, Corollary 2.3.5.]) that $R = R' \wedge \bigwedge_{I \subseteq [n], |I| \leq k} \mathrm{proj}_I(R)$, where $R' \leq \mathbf{A}^n$ is the minimal invariant relation containing R and having the parallelogram property. Clearly, $\bigwedge_{I \subseteq [n], |I| \leq k} \mathrm{proj}_I(R)$ can be written as the conjunction of at most $c \cdot n^k$ many critical relations, for some $c > 0$.

The relation R', if not already $\wedge$-irreducible itself, is given by the intersection of all $\wedge$-irreducible relations $S > R'$ that have the parallelogram property. Denote by $\mathcal{S}$ the set of all such relations. For every $(i,a,b) \notin \mathrm{Sig}(R')$ where $a \in \mathrm{proj}_i R'$ choose $S_{(i,a,b)} \in \mathcal{S}$ such that $(i,a,b) \notin \mathrm{Sig}(S_{(i,a,b)})$. Then by Observation 16, $R' = \bigcap_{(i,a,b) \notin \mathrm{Sig}(R')} S_{(i,a,b)}$ which is an intersection of at most $n \cdot |A|^2$ $\wedge$-irreducible relations. (To see that such $S_{(i,a,b)}$ exists, let $\bar{x} \in R'$ be such that $x_i = a$ and let $\bar{y}$ be such that $y_i = b$, $y_j = x_j$ for $j \neq i$. We can choose $S_{(i,a,b)}$ to be a maximal relation containing R but omitting $\bar{y}$.)

Consequently, R' can be defined as a conjunction of at most $n \cdot |A|^2$ many critical relations with the parallelogram property which concludes the proof. ◀

We remark that an analogous statement to Lemma 17 also holds for multisorted relations $R \leq \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$, such that the sorts $\mathbf{A}_i$ come from a finite set of algebras that have a common k-edge term. In particular, this is the case for $\mathbf{A}_i \in \mathrm{HS}(\mathbf{A})$, if $\mathbf{A}$ has a k-edge term (cf. Lemma 15).

When dealing with multisorted relations R over $\mathrm{HS}(\mathbf{A})$, we can furthermore always restrict the domain $\mathbf{A}_i$ of the i-th variable of a relation R to its projection $\mathrm{proj}_i(R) \leq \mathbf{A}_i$. So, without loss of generality, we can assume that $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ is *subdirect*, i.e., its projection to every coordinate i is the full domain $\mathbf{A}_i$.

For a subdirect relation $R \leq_{sd} \mathbf{B} \times \mathbf{C}$ with the parallelogram property, let us define the *linkedness congruence* θ_B on $\mathbf{B}$ by $(x,y) \in \theta_B \leftrightarrow (\exists c \in C)(R(x,c) \wedge R(y,c))$. For a general relation $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ with the parallelogram property, and any proper subset $I \subset [n]$

of coordinates, we define the *linkedness congruence* θ_I on $\mathrm{proj}_I(R)$ analogously, where we consider R as a binary relation between $\mathrm{proj}_I(R)$ and $\mathrm{proj}_{[n]\setminus I}(R)$ (we write θ_i instead of $\theta_{\{i\}}$). It follows from the parallelogram property that θ_I is indeed a congruence of $\mathrm{proj}_I(R)$. A subdirect relation $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ is called *reduced* if every tuple $(a_1, \ldots, a_n) \in R$ is already uniquely determined by $(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$, for any coordinate i; in other words, θ_i is trivial, for every $i = 1, \ldots, n$. By the following lemma, we can reduce the quest for short pp-definitions to *reduced*, subdirect, critical relations with the parallelogram property:

▶ **Lemma 18.** *Let $\mathbf{A}$ be an algebra with a k-edge term. Assume that all relations $R \leq_{sd} \mathbf{A}_1 \times \cdots \times \mathbf{A}_n$ with $\mathbf{A}_i \in \mathrm{HS}(\mathbf{A})$ that are reduced, critical, and have the parallelogram property, have (multisorted) pp-definitions of length at most $f(n)$. Then $\mathbf{A}$ has pp-definitions of length $O(n^k + n \cdot f(n))$.*

Proof. We first prove that we can pp-define all critical relations $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ with $\mathbf{A}_i \in \mathrm{HS}(\mathbf{A})$ that have the parallelogram property (but are not necessarily reduced), by pp-definitions of length at most $O(f(n))$. Given such a relation, let us consider the linkedness congruence $\theta_i \in \mathrm{Con}(\mathbf{A}_i)$ for every coordinate i. Then let us define the quotient $R' = R/(\theta_1, \ldots, \theta_n)$. It is not hard to see that $R' \leq_{sd} \mathbf{A}_1/\theta_1 \times \cdots \times \mathbf{A}_n/\theta_n$, is also critical, and has the parallelogram property. Furthermore, by definition of θ_i, R' is reduced.

Since R has the parallelogram property, it is equal to the full preimage of R' under the quotient map $(x_1, \ldots, x_n) \mapsto (x_1/\theta_1, \ldots, x_n/\theta_n)$. Thus, every pp-definition $\phi'(x_1, \ldots, x_n)$ of R' gives rise to the pp-definition $\exists y_1, \ldots, y_n \, (\bigwedge_{i=1}^{n}(x_i/\theta_i = y_i) \wedge \phi'(y_1, \ldots, y_n))$ of length $O(f(n))$ which defines R; this proves our claim. The statement of the lemma then follows directly from (the multi-sorted version of) Lemma 17, and Lemma 15. ◀

Any relation $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ as in Lemma 18 comes with several nice properties (in algebraic terms, it is a *graph of a joint similarity* between the algebras $\mathbf{A}_i$, cf. [8, Section 2.3.1]). We are mainly going to need the following property in Lemma 19, respectively its generalization in Lemma 20:

▶ **Lemma 19** ([20, Lemma 2.4])**.** *Let $\mathbf{A}_1, \ldots, \mathbf{A}_n$ be algebras with few subpowers, and let $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ be a reduced, critical relation with the parallelogram property. Then every $\mathbf{A}_i$ is subdirectly irreducible.*

▶ **Lemma 20.** *Let $\mathbf{A}_1, \ldots, \mathbf{A}_n$ be algebras with few subpowers, and let $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ be a critical relation with the parallelogram property. For $I \subset [n]$, let θ_I be the linkedness congruence on $\mathrm{proj}_I(R)$ with respect to R. Then θ_I is $\wedge$-irreducible.*

Proof. Since R is $\wedge$-irreducible, there is a unique cover $R^* > R$ in the lattice of all subalgebras of $\mathbf{A}_1 \times \ldots \times \mathbf{A}_n$. A tuple $\bar{a} = (a_1, \ldots, a_n) \in R^* \setminus R$ is called a *key tuple* of R (cf. [29]). It follows from the criticality of R that for every $j = 1, \ldots, n$, there is a tuple $(a_1, \ldots b_j, \ldots, a_n) \in R$ that only differs from $\bar{a}$ at position j.

For simplicity, let us assume that $I = \{1, 2, \ldots, i\}$ with $i < n$. Then, the linkedness-congruence θ_I has an equivalence class containing all elements of the form $(a_1, \ldots b_j, \ldots, a_i)$ for $j = 1, \ldots, i$. Note that $(a_1, a_2, \ldots, a_i) \in \mathrm{proj}_I(R)$ is not an element of this class.

To prove that θ_I is $\wedge$-irreducible, let θ' be a congruence strictly above θ_I. We claim that then θ' must also contains the pair $((a_1, a_2, \ldots, a_i), (b_1, a_2, \ldots, a_i))$. To prove the claim, let us define $R'(\bar{x}) = \exists \bar{y}_I \left(\theta'(\bar{x}_I, \bar{y}_I) \wedge R(\bar{y}_I, \bar{x}_{[n]\setminus I})\right)$. As R' properly contains R, it also must contain its cover R^*, and thus the key tuple $(a_1, a_2, \ldots, a_n)$. Moreover, the linkedness congruence of R' on coordinates I is equal to θ', thus θ' must contain the pair $((a_1, a_2, \ldots, a_i), (b_1, a_2, \ldots, a_i))$. So θ_I has a unique cover θ_I^*, which is the congruence generated by $\theta_I \cup \{((a_1, a_2, \ldots, a_i), (b_1, a_2, \ldots, a_i))\}$; this finishes the proof. ◀

We are now ready to prove our main result:

▶ **Theorem 21.** *Let* $\mathbf{A}$ *be an algebra with a* k*-edge term, and assume that* $\mathrm{HSP}(\mathbf{A})$ *is residually finite. Then* $\mathbf{A}$ *has pp-definitions of length* $O(n^k)$*.*

Proof. Let $\mathcal{V}_{SI}$ consist of all subdirectly irreducible elements of $\mathrm{HSP}(\mathbf{A})$. Since $\mathrm{HSP}(\mathbf{A})$ is residually finite, $\mathcal{V}_{SI}$ contains up to isomorphism only finitely many algebras, all of which are finite. In particular $\mathcal{V}_{SI}$ is a subset of $\mathrm{HS}(\mathbf{A}^l)$, for some finite power l.

By Lemma 14, it is enough to prove pp-definitions of length $O(n^k)$ for $\mathbf{A}^l$. By Lemma 18, it suffices to prove that every reduced, critical relation $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ with $\mathbf{A}_i \in \mathrm{HS}(\mathbf{A}^l)$ that has the parallelogram property, has a (multisorted) pp-definition of linear length.

Let Γ be the set of all at most ternary invariant relations over $\mathrm{HS}(\mathbf{A}^l)$. We construct a pp-definition of R from Γ of length linear in n, by induction on n. For $n \leq 3$, R itself is in Γ.

For general $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ with the parallelogram property, recall the definition of the linkedness congruence θ_I. We then define the algebra $\mathbf{A}_{1,2} = \mathrm{proj}_{\{1,2\}}(R)/\theta_{\{1,2\}}$, and the relations $Q = \{(x_1, x_2, y_{1,2}) \in A_1 \times A_2 \times A_{1,2} \mid y_{1,2} = (x_1, x_2)/\theta_{1,2}\}$ and $R' = \{(y_{1,2}, x_3, \ldots, x_n) \mid \exists x_1, x_2\,(Q(x_1, x_2, y_{1,2}) \wedge R(x_1, x_2, x_3, \ldots, x_n))\}$. Note that $Q \leq \mathbf{A}_1 \times \mathbf{A}_2 \times \mathbf{A}_{1,2}$, and $R' \leq \mathbf{A}_{1,2} \times \mathbf{A}_3 \times \cdots \times \mathbf{A}_n$. Since R has the parallelogram property, R can be defined by the pp-formula $(\exists y_{1,2} \in \mathbf{A}_{1,2})\,(Q(x_1, x_2, y_{1,2}) \wedge R'(y_{1,2}, x_3, \ldots, x_n))$.

By Lemma 20, $\theta_{\{1,2\}}$ is $\wedge$-irreducible. This implies that $\mathbf{A}_{1,2}$ is subdirectly irreducible and hence an element of $\mathcal{V}_{SI} \subseteq \mathrm{HS}(\mathbf{A}^l)$. In particular, this means that Q is a relation from our relational basis Γ. The relation R' is of arity $n-1$, and thus, by induction assumption, has a pp-definition of linear length. This finishes our proof. ◀

Note that, although in the proof of Theorem 21 we found a *ternary* constraint language Γ defining the reduced critical relations $R \leq_{sd} \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ with parallelogram property, the same may not be true for the original algebra $\mathbf{A}$. An explicit bound on the maximal required arity is given by $3l$, where l is such that all subdirectly irreducible elements of $\mathrm{HSP}(\mathbf{A})$ are contained in $\mathrm{HS}(\mathbf{A}^l)$. We are not aware of any better bound than the double exponential $l \leq |A|^{|A|^{|A|+1}+1}$ [14] (see also [8, Theorem A.5.27.]).

As an immediate consequence of Theorem 21 we get that every 3-element algebra with few subpowers has short pp-definitions, confirming Conjecture 1 for the 3-element case. It is well known that few subpowers imply *congruence modularity* [6, Theorem 4.2]; thus we can use the following fact:

▶ **Theorem 22** ([8, Corollary A.5.31.])**.** *Let* $\mathbf{A}$ *be an algebra on a 3-element set, such that* $\mathrm{HSP}(\mathbf{A})$ *is congruence modular. Then* $\mathrm{HSP}(\mathbf{A})$ *is residually finite.*

▶ **Corollary 23.** *Let* Γ *be a constraint language on a 3-element domain. Then* Γ *has short pp-definitions if and only if* Γ *has few subpowers. More precisely,* Γ *has pp-definitions of length* $O(n^k)$*, where* k *is the minimal number such that* Γ *has a* k*-edge polymorphism.*

Proof. Let us assume that Γ is a constraint language with a k-edge polymorphism, and let $\mathbf{A} = \mathrm{Pol}(\Gamma)$ be its polymorphism algebra. Since the existence of an edge operation implies that $\mathrm{HSP}(\mathbf{A})$ is congruence modular [6, Theorem 4.2], by Theorem 22, $\mathrm{HSP}(\mathbf{A})$ is residually finite. By Theorem 21, $\mathbf{A}$, and thus also Γ, has pp-definitions of length $O(n^k)$.

If Γ has few subpowers, then by Theorem 7, it has a k-edge polymorphism for some k. Thus Γ has short pp-definitions if it has few subpowers. ◀

5 The Subpower Membership Problem

The *Subpower Membership Problem* SMP($\mathbf{A}$) of a finite algebra $\mathbf{A}$ is the computational problem in which the input consists of a list of tuples $\bar{b}, \bar{a}_1, \dots, \bar{a}_k \in A^n$, for arbitrary $n \geq 1$, and one needs to decide whether $\bar{b}$ lies in the subalgebra $\mathrm{Sg}_{\mathbf{A}^n}(\bar{a}_1, \dots, \bar{a}_k)$ generated by $\bar{a}_1, \dots, \bar{a}_k$, i.e., in the smallest $R \leq \mathbf{A}^n$ that contains $\bar{a}_1, \dots, \bar{a}_k$.

The existence of an efficient algorithm for SMP($\mathbf{A}$) implies that it is feasible to represent the relations in Inv($\mathbf{A}$) by some generating set of tuples. In particular, in the context of constraint satisfaction, it was remarked by several authors (see, e.g., [9]) that a polynomial-time algorithm for SMP($\mathbf{A}$) would allow us to define constraint satisfaction problems over *infinite* constraint languages $\Gamma \subseteq \mathrm{Inv}(\mathbf{A})$, where the constraint relations in Γ are encoded by generating tuples. In [16], any algebra $\mathbf{A}$ with SMP($\mathbf{A}$) in P was referred to as *polynomially evaluable*.

While there are algebras for which SMP($\mathbf{A}$) is EXPTIME-complete [22], it was asked in [16, Question 3] whether all algebras with few subpowers are polynomially evaluable. An affirmative answer was given for several special cases [24, 9], but the question still remains open in general. The best general upper bound on the complexity of SMP($\mathbf{A}$) for algebras with few subpowers is NP [9]. This bound is based on the fact that membership of an element in a relation $R \leq \mathbf{A}^n$ can always be witnessed by a *compact representation* of R, i.e., a small, canonical generating set. The difficulty in finding *deterministic* polynomial algorithms lies in efficiently computing such compact representations of R from an arbitrary generating set.

Note that for an algebra $\mathbf{A}$ with $\mathrm{Inv}(\mathbf{A}) = \langle\Gamma\rangle$, the *non*-membership of a tuple $\bar{b}$ in a relation $\mathrm{Sg}_{\mathbf{A}^n}(\bar{a}_1, \dots, \bar{a}_k)$ can be witnessed by a pp-formula $\phi(\bar{x})$ over Γ, such that ϕ holds for all tuples $\bar{a}_1, \dots, \bar{a}_k$, but not for $\bar{b}$.

If Γ has short pp-definitions, we can guess such a certificate ϕ for "No"-instances of SMP($\mathbf{A}$), and verify it in polynomial time. As a direct consequence of this (and the fact that short pp-definitions imply few subpowers), we obtain the following:

▶ **Theorem 24.** *Let* $\mathbf{A}$ *be an algebra with short pp-definitions. Then* $\mathrm{SMP}(\mathbf{A}) \in \mathrm{NP} \cap \text{co-NP}$.

In particular, Conjecture 1 would imply that $\mathrm{SMP}(\mathbf{A}) \in \mathrm{NP} \cap \text{co-NP}$ for every algebra $\mathbf{A}$ with few subpowers. Note, however, that in the setting of our main result (Theorem 21), this does not provide any progress on the subpower membership problem, since it was shown in [9] that SMP($\mathbf{A}$) is even in P for every algebra $\mathbf{A}$ with few subpowers that generates a residually finite variety.

6 Discussion

By Theorem 21, constraint languages Γ with few subpowers, whose polymorphism algebra generates a residually finite variety, have short pp-definitions. While this confirms Conjecture 1 for a large subclass of constraint languages, much work remains to prove the conjecture in full generality.

The condition of residual finiteness does not bear much importance in constraint satisfaction, it is mainly used in purely algebraic contexts. Important steps to connect short pp-definitions closer to the theory of constraint satisfaction would be to extend our results to specific tractability classes (such as constraint languages with Mal'tsev polymorphisms), and to show invariance under *pp-interpretations*:

▶ **Question 25.** *Let* Γ *and* Δ *be two constraint languages, such that* Δ *is pp-interpretable in* Γ *and* Γ *has short pp-definitions. Then, does* Δ *also have short pp-definitions?*

Pp-interpretations are a generalization of pp-definitions, that describe certain gadget reductions between constraint languages on different domains. All standard tractability classes (including few subpowers, Mal'tsev, near-unanimity) are closed under pp-interpretations, which motivates Question 25 (note that Conjecture 1 implies a positive answer). We remark that we do not know the answer to this question even for Δ being *pp-definable* in Γ (as Lemma 5 assumes *pp-interdefinability*). In the special case of *pp bi-interpretable* structures, Question 25 has a positive answer by a straigthforward generalization of the proof of Lemma 5, we thank the anonymous reviewer for this observation. In algebraic terms, Question 25 asks whether for $\mathbf{A}$ with short pp-definitions, it is the case that also every *extension* of every algebra $\mathbf{B} \in \mathrm{HSP}^{\mathrm{fin}}(\mathbf{A})$ has short pp-definitions; for finite powers $\mathrm{P}^{\mathrm{fin}}$, we verified the statement in Lemma 14.

As discussed in Section 5, it is also essential for progress on the Subpower Membership Problem to extend our results to algebras $\mathbf{A}$ that do not generate residually finite varieties. While we did not present any results on this in this paper, we are aware of singular examples of such algebras with short pp-definitions (such as the 4-element algebra, described by Brady in [8, Example 2.3.2.]; short pp-definitions follow directly from his analysis).

In order to improve the complexity result of Theorem 24 and put $\mathrm{SMP}(\mathbf{A})$ in the class P, we would need an explicit method of efficiently computing a short pp-definition for a relation $R = \mathrm{Sg}_{\mathbf{A}^n}(\bar{a}_1, \ldots, \bar{a}_k)$ given by its generators $\bar{a}_1, \ldots, \bar{a}_k$. This motivates the following question:

▶ **Question 26.** *Let Γ be a constraint language with short pp-definitions. Is there a polynomial-time algorithm that computes a (short) pp-definition of a relation $R \leq \langle\Gamma\rangle_n$, given by a set of generators $\bar{a}_1, \ldots, \bar{a}_k$?*

We remark that over Boolean domains, Question 26 has a positive answer (see Examples 8 and 11).

Finally, recall that the bound from Theorem 21 is a polynomial of degree k if Γ has a k-edge polymorphism. It is therefore tempting to conjecture that the same degree could be enough in general for Conjecture 1. Note that the number of n-ary pp-definable relations, for Γ with a k-edge polymorphism, is known to be in $2^{O(n^k)}$ [16, Theorem 3.4].

References

1 Erhard Aichinger, Peter Mayr, and Ralph McKenzie. On the number of finite algebraic structures. *Journal of the European Mathematical Society*, 16(8):1673–1686, September 2014. doi:10.4171/jems/472.

2 Kirby A. Baker and Alden F. Pixley. Polynomial interpolation and the Chinese Remainder Theorem for algebraic systems. *Mathematische Zeitschrift*, 143(2):165–174, June 1975. doi:10.1007/BF01187059.

3 Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and How to Use Them. In Andrei Krokhin and Stanislav Zivny, editors, *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 1–44. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017. doi:10.4230/DFU.Vol7.15301.1.

4 Libor Barto, Jakub Opršal, and Michael Pinsker. The wonderland of reflections. *Israel Journal of Mathematics*, 223(1):363–398, 2018. doi:10.1007/s11856-017-1621-9.

5 Clifford Bergman. *Universal Algebra: Fundamentals and Selected Topics*. Chapman and Hall/CRC, New York, November 2011. doi:10.1201/9781439851302.

6 Joel Berman, Paweł Idziak, Petar Marković, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Varieties with few subalgebras of powers. *Transactions of the American Mathematical Society*, 362(3):1445–1473, March 2010. doi:10.1090/S0002-9947-09-04874-0.

7 V. G. Bodnarčuk, L. A. Kalužnin, V. N. Kotov, and B. A. Romov. Galois theory for Post algebras, part I and II. *Cybernetics*, 5:243–539, 1969.

8 Zarathustra Brady. Notes on CSPs and Polymorphisms, October 2022. arXiv:2210.07383 [cs, math]. doi:10.48550/arXiv.2210.07383.

9 Andrei Bulatov, Peter Mayr, and Ágnes Szendrei. The Subpower Membership Problem for Finite Algebras with Cube Terms. *Logical Methods in Computer Science*, Volume 15, Issue 1, February 2019. doi:10.23638/LMCS-15(1:11)2019.

10 Andrei A. Bulatov. A Dichotomy Theorem for Nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330, October 2017. doi:10.1109/FOCS.2017.37.

11 S. Burris and H. P. Sankappanavar. *A Course in Universal Algebra*. Springer New York, November 1981.

12 Hubie Chen. The expressive rate of constraints. *Annals of Mathematics and Artificial Intelligence*, 44(4):341–352, August 2005. doi:10.1007/s10472-005-7031-4.

13 Tomás Feder and Moshe Y. Vardi. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM Journal on Computing*, 28(1):57–104, January 1998. doi:10.1137/S0097539794266766.

14 Ralph Freese and Ralph McKenzie. *Commutator Theory for Congruence Modular Varieties*. CUP Archive, August 1987.

15 David Geiger. Closed systems of functions and predicates. *Pacific Journal of Mathematics*, 27:95–100, 1968.

16 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, January 2010. doi:10.1137/090775646.

17 Peter Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200(1):185–204, June 1998. doi:10.1016/S0304-3975(97)00230-2.

18 Peter Jeavons, David Cohen, and Martin C. Cooper. Constraints, consistency and closure. *Artificial Intelligence*, 101(1-2):251–265, May 1998.

19 Peter Jeavons, David Cohen, and Marc Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, July 1997. doi:10.1145/263867.263489.

20 Keith A. Kearnes and Ágnes Szendrei. Clones of algebras with parallelogram terms. *International Journal of Algebra and Computation*, 22(01):1250005, February 2012. doi:10.1142/S0218196711006716.

21 Dexter Kozen. Complexity of finitely presented algebras. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, STOC '77, pages 164–177, New York, NY, USA, May 1977. Association for Computing Machinery. doi:10.1145/800105.803406.

22 Marcin Kozik. A finite set of functions with an EXPTIME-complete composition problem. *Theoretical Computer Science*, 407(1):330–341, November 2008. doi:10.1016/j.tcs.2008.06.057.

23 Victor Lagerkvist and Magnus Wahlström. Polynomially Closed Co-clones. In *2014 IEEE 44th International Symposium on Multiple-Valued Logic*, pages 85–90, May 2014. doi:10.1109/ISMVL.2014.23.

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

25 Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, STOC '78, pages 216–226, New York, NY, USA, May 1978. Association for Computing Machinery. doi:10.1145/800133.804350.

26 Charles C. Sims. Computational methods in the study of permutation groups. In John Leech, editor, *Computational Problems in Abstract Algebra*, pages 169–183. Pergamon, January 1970. doi:10.1016/B978-0-08-012975-4.50020-5.

27 Dmitriy Zhuk. A Proof of CSP Dichotomy Conjecture. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 331–342, October 2017. doi:10.1109/FOCS.2017.38.

28 Dmitriy Zhuk. A Proof of the CSP Dichotomy Conjecture. *Journal of the ACM*, 67(5):30:1–30:78, August 2020. doi:10.1145/3402029.

29 Dmitriy N. Zhuk. Key (critical) relations preserved by a weak near-unanimity function. *Algebra universalis*, 77(2):191–235, April 2017. doi:10.1007/s00012-017-0426-3.

A Omitted proofs of technical lemmata

A.1 Proof of Lemma 14

Proof. It is straightforward to see that any relation $R \subseteq B^n$ is invariant under $\mathbf{B}$ if and only if it is invariant under $\mathbf{A}$, when interpreted as an kn-ary relation on A.

We are first going to prove the "only if" direction. Let Δ be a relational basis of $\mathrm{Inv}(\mathbf{B})$. By interpreting every m-ary relation $Q \in \Delta$ as a km-ary relation $Q' \leq \mathbf{A}^{km}$, we obtain a relational basis $\Delta' = \{Q' \mid Q \in \Delta\}$ of $\mathrm{Inv}(\mathbf{A})$.

Let $R \leq \mathbf{A}^n$. By adding dummy variables, we can assume without loss of generality that $n = k\ell$ for $\ell = \lceil \frac{n}{k} \rceil$. Then by assumption, R considered as an ℓ-ary relation over B has a pp-definition $\phi(x_1, \dots, x_\ell)$ from Δ of length in $O(f(\ell))$. If we substitute each (B-valued) variable in ϕ by a k-tuple of (A-valued) variables, and each Δ-predicate Q in ϕ by the corresponding Δ'-predicate Q', then we obtain a pp-definition ϕ' of R of length at most $k \cdot f(l) = k \cdot f(\lceil \frac{n}{k} \rceil)$. Note further that existentially quantifying all the additionally added dummy variables adds only constantly many symbols. Thus $\mathbf{A}$ has pp-definitions of length $O(k \cdot f(\lceil \frac{n}{k} \rceil)) = O(f(\lceil \frac{n}{k} \rceil))$.

Now let us prove the "if" direction. Let Γ be a relational basis of $\mathrm{Inv}(\mathbf{A})$. Let $e \colon A \to A^k$ be the map $x \mapsto (x, \dots, x)$. For every $Q \in \Gamma$, we define $Q' = \{(e(x_1), \dots, e(x_m)) \mid (x_1, \dots, x_m) \in Q\} \leq \mathbf{B}^m$, and we define the binary relations $P_i = \{((x_1, \dots, x_k), e(x_i)) \mid (x_1, \dots, x_k) \in B\} \leq \mathbf{B}^2$, for $i = 1, \dots, k$. We construct the relational basis of $\mathrm{Inv}(\mathbf{B})$ as $\Gamma' = \{Q' \mid Q \in \Gamma\} \cup \{P_1, \dots, P_k\}$.

Let $R' \leq \mathbf{B}^n$ and let $R \leq \mathbf{A}^{kn}$ be the relation R' considered as a kn-ary relation over A. By assumption, there exists a pp-definition $\phi(x_1, \dots, x_{nk})$ of $R \leq \mathbf{A}^{kn}$ over Γ of length in $O(f(\lceil \frac{kn}{k} \rceil)) = O(f(n))$. Let $z_1, \dots, z_\ell$ be its existentially quantified variables. Let us then define $\phi'(y_1, y_2, \dots, y_n)$ as a pp-formula over Γ' with existentially quantified variables $x'_1, \dots, x'_{nk}, z'_1, \dots, z'_\ell$ and predicates $P_i(y_j, x'_{(j-1)(k+i)})$ for all $i \in [k], j \in [n]$, as well as $Q'(u'_1, \dots, u'_m)$ for every predicate $Q(u_1, \dots, u_m)$ in ϕ with $u_i \in \{x_1, \dots, x_{nk}, z_1, \dots, z_\ell\}$. It is easy to check that ϕ' defines $R' \leq \mathbf{B}^n$ over Γ'. Clearly, the length of ϕ' is in $O(f(n))$. ◀

A.2 Proof of Lemma 15

Proof. Let Γ be a relational basis of $\mathrm{Inv}(\mathbf{A})$, and let $R \leq \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ with $\mathbf{A}_i \in \mathrm{HS}(\mathbf{A})$. So $\mathbf{A}_i = h_i(\mathbf{S}_i)$ for a subalgebra $\mathbf{S}_i \leq \mathbf{A}$ and a homomorphism $h_i \colon \mathbf{S}_i \to \mathbf{A}_i$. Note that the graph of this homomorphism $G_{h_i} = \{(a, h_i(a)) \colon a \in \mathbf{S}_i\} \leq \mathbf{S}_i \times \mathbf{A}_i$ is an invariant relation. We define Γ' to be the union of Γ and all binary relations G_h. It is not hard to see that $R' = \{(a_1, \dots, a_n) \in S_1 \times \dots \times S_n \mid (h_1(a_1), \dots, h_n(a_n)) \in R\}$ is invariant under $\mathbf{A}$. By assumption, R' has a pp-definition $\phi'(x_1, \dots, x_n)$ of length in $O(f(n))$. The relation R can then be defined by the pp-formula $\exists y_1, \dots, y_n \, (\bigwedge_{i=1}^n G_{h_i}(y_i, x_i) \wedge \phi(y_1, \dots, y_n))$, whose length is also in $O(f(n))$.

For the converse, let us consider relations $R \leq \mathbf{A}_1 \times \ldots \times \mathbf{A}_n$ with $\mathbf{A}_i \in \mathrm{HS}(\mathbf{A})$ from any relational basis Γ of $\mathrm{Inv}(\mathrm{HS}(\mathbf{A}))$. Then $\mathbf{A}_i = h_i(\mathbf{S}_i)$, for a subalgebra $\mathbf{S}_i \leq \mathbf{A}$ and a homomorphism $h_i \colon \mathbf{S}_i \to \mathbf{A}_i$. As above, $R' = \{(a_1, \ldots, a_n) \in A^n \mid (h_1(a_1), \ldots, h_n(a_n)) \in R\}$ is invariant under $\mathbf{A}$. It is not hard to see that $\Gamma' = \{R' \mid R \in \Gamma\}$ is a relational basis of $\mathrm{Inv}(\mathbf{A})$, and for any pp-definition ϕ of a relation $Q \leq \mathbf{A}^n$ over Γ, the formula ϕ' obtained by replacing every occurrence of the symbol R by R' defines Q over Γ'. ◀