Electronic Colloquium on Computational Complexity, Report No. 93 (2007)

The Complexity of the Counting Constraint Satisfaction Problem

Andrei A. Bulatov
School of Computimg Science, Simon Fraser University, Burnaby, Canada
abulatov@cs.sfu.ca

Abstract

The Counting Constraint Satisfaction Problem ($\#\mathrm{CSP}(\mathcal{H})$) over a finite relational structure $\mathcal{H}$ can be expressed as follows: given a relational structure $\mathcal{G}$ over the same vocabulary, determine the number of homomorphisms from $\mathcal{G}$ to $\mathcal{H}$. In this paper we characterize relational structures $\mathcal{H}$ for which $\#\mathrm{CSP}(\mathcal{H})$ can be solved in polynomial time and prove that for all other structures the problem is #P-complete.

1 Introduction

In the Counting Constraint Satisfaction Problem, $\#\mathrm{CSP}(\mathcal{H})$, over a finite relational structures $\mathcal{H}$ the objective is, given a finite relational structure $\mathcal{G}$, to compute the number of homomorphisms from $\mathcal{G}$ to $\mathcal{H}$. Various particular cases of the #CSP arise and have been extensively studied in a wide range of areas from logic, graph theory, and artificial intelligence [3, 18, 27, 34, 38, 47, 51, 53, 54, 57, 58], to statistical physics [2, 16, 45]. However, in different areas this problem often appears in different equivalent forms: (1) the problem of finding the number of models of a conjunctive formula, (2) the problem of computing the size (number of tuples) of the evaluation $Q(D)$ of a conjunctive query (without projection) Q on a database D and also (3) the problem of counting the number of assignments to a set of variables subject to specified constraints.

Since the seminal papers [55, 30], the complexity of the decision counterpart of #CSP, the Constraint Satisfaction Problem or CSP for short, has been an object of intensive study. The ultimate goal of that research direction is to classify finite relational structures with respect to the complexity of the corresponding CSP. We shall refer to this research problem as the *classification problem*. A number of significant results have been obtain, see e.g. [55, 30, 6, 8], but a full classification is far from being completed.

Although the classification problem for the general #CSP has been tackled for the first time very recently, a massive work has been done in the study of the complexity of various particular counting CSPs. These particular problems include classical combinatorial problems such as #CLIQUE, GRAPH RELIABILITY, ANTICHAIN, PERMANENT etc. [47, 53, 57, 58] expressible in the form of #CSP; the counting SATISFIABILITY and GENERALIZED SATISFIABILITY problems (in these problems the objective is to find the number of satisfying assignments to a propositional formula) [18, 54] which correspond to $\#\mathrm{CSP}(\mathcal{H})$ for 2-element structures $\mathcal{H}$, counting the number of solution of equations over finite semigroups [50, 44] and many others.

However, the real focus of research in this area has been #H-COLORING problem and its variants. In the #H-COLORING problem the aim is to find the number of homomorphisms from a given graph G to the fixed graph H. Thus, it is equivalent to $\#\mathrm{CSP}(\mathcal{H})$ where $\mathcal{H}$ is a graph. Dyer and Greenhill [27] have proved that, for every undirected graph H, its associated #H-COLORING problem is either in FP (we shall call such problems *tractable*) or #P-complete and they have also provided a complete characterization of the

ISSN 1433-8092

tractable problems. This result has been extended to the counting LIST #H-COLORING problem [24, 22], which allows additional restrictions on possible images of a node. Recently, Dyer, Goldberg, and Paterson [28, 29] obtained a similar classification for directed acyclic graphs. Furthermore, some other variants of the #H-COLORING problem for undirected graphs have been intensively studied during the last few years [20, 21]. Another direction in this area is the study of problems with restricted input, that is subproblems of the #H-COLORING problem in which the input graph G must be planar [38, 56], a partial k-tree [23], sparse or of low degree [34, 35], etc. Finally, we should mention the approach to counting problems using approximation and randomized algorithms, see e.g. [43, 26, 25].

In [4, 14] we started a systematic study of the classification problem for the general #CSP. The main approach chosen was the *algebraic approach* which has proved to be quite useful in the study of the decision CSP [40, 41, 6, 8]. This approach uses invariance properties of predicates definable in relational structures. Invariance properties are usually expressed as *polymorphisms* of the predicates, that is (multi-ary) operations on the universe of the relational structure compatible with the predicates.

In [4], we proved that if $\#\mathrm{CSP}(\mathcal{H})$ is tractable, then $\mathcal{H}$ has a *Mal'tsev* polymorphism, that is a ternary operation $m(x, y, z)$ satisfying the identities $m(x, y, y) = m(y, y, x) = x$. Another observation was that the *congruences*, i.e. the definable equivalence relations, of $\mathcal{H}$ play a very important role. In particular, these results have allowed us to come up with a nearly trivial prove of the result of [27]. In [5], another necessary condition for the tractability of $\#\mathrm{CSP}(\mathcal{H})$ has been identified. It imposes certain restrictions onto possible congruences of $\mathcal{H}$, in terms of sizes of their equivalence classes.

In this paper, after giving general definitions (Section 2.1) and introducing the basics of the algebraic approach (Sections 2.2 and 2.3), we go deeper into the structure of congruences of a relational structure (Section 3.1) and then identify several further necessary conditions for tractability (Section 3.2), again expressed in terms of properties of congruences. Then, in Section 4, we prove that, for every relational structure $\mathcal{H}$ satisfying all the conditions obtained, the problem $\#\mathrm{CSP}(\mathcal{H})$ can be solved in polynomial time. Thus, we completely solve the classification problem for the general counting CSP.

We intensively use methods and results from a number of areas of modern algebra: lattice theory, tame congruence theory, commutator theory and ring theory. To make the paper available for a wider audience we are avoiding the excessive use of algebraic terminology. In spite of that, some parts of the paper, Section 4 and especially proofs, are demanding: they require from the reader some familiarity with basic algebraic objects and ideas. The keen reader is referred to textbooks [15, 31, 33, 37]. The reader should be aware that to avoid yet another layer of objects we use algebraic terminology for relational structures, while in the algebraic literature the same concepts are used for "dual" objects, universal algebras.

2 Preliminaries

2.1 Relational structures and homomorphisms

Our notation concerning tuples and relational structures is fairly standard. Let $[n]$ denote the set $\{1, \ldots, n\}$. The set of all n-tuples of elements from a set H is denoted by H^n. We denotes tuples of elements in boldface, e.g. $\mathbf{a}$, and their components by $\mathbf{a}[1], \mathbf{a}[2], \ldots$. For a subset $I = \{i_1, \ldots, i_k\} \subseteq [n]$ and an n-tuple $\mathbf{a}$, by $\mathrm{pr}_I\mathbf{a}$ we denote the *projection of* $\mathbf{a}$ *onto* I, the k-tuple $(\mathbf{a}[i_1], \ldots, \mathbf{a}[i_k])$. For an n-ary relation $R \subseteq H^n$, its projection onto I is defined to be $\mathrm{pr}_I R = \{\mathrm{pr}_I\mathbf{a} \mid \mathbf{a} \in R\}$. If $D_i = \mathrm{pr}_i R$ for $i \in [n]$ we say that R is *subdirect product* of $D_1, \ldots, D_n$. If $D_1 = \ldots = D_n = H$ then R is said to be an *n-th subderect power* of H. For $\mathbf{a} = (\mathbf{a}[1], \ldots, \mathbf{a}[n])$ and $\mathbf{b} = (\mathbf{b}[1], \ldots, \mathbf{b}[m])$, $(\mathbf{a}, \mathbf{b})$ denotes the tuple $(\mathbf{a}[1], \ldots, \mathbf{a}[n], \mathbf{b}[1], \ldots, \mathbf{b}[m]$, while $\langle \mathbf{a}, \mathbf{b} \rangle$ denotes the pair of tuples.

A *vocabulary* is a finite set of relational symbols $R_1, \dots, R_n$ each of which has a fixed arity. A *relational structure* over the vocabulary $R_1, \dots, R_n$ is a tuple $\mathcal{H} = (H; R_1^{\mathcal{H}}, \dots, R_n^{\mathcal{H}})$ such that A is a non-empty set, called the *universe* of $\mathcal{H}$, and each $R_i^{\mathcal{H}}$ is a relation on H having the same arity as the symbol R_i. Let $\mathcal{G}, \mathcal{H}$ be relational structures over the same vocabulary $R_1, \dots, R_n$. A *homomorphism* from $\mathcal{G}$ to $\mathcal{H}$ is a mapping $\varphi\colon G \to H$ from the universe of $\mathcal{G}$ (the *instance*) to the universe H of $\mathcal{H}$ (the *template*) such that, for every relation $R^{\mathcal{G}}$ of $\mathcal{G}$ and every tuple $(\alpha_1, \dots, a_m) \in R^{\mathcal{G}}$, we have $(\varphi(\alpha_1), \dots, \varphi(\alpha_m)) \in R^{\mathcal{H}}$.

A relation R is said to be *primitive positive definable* (*pp-*) in $\mathcal{H}$, if it can be expressed using the predicates $R_i^{\mathcal{H}}$ of $\mathcal{H}$ together with the binary equality predicate on H (denoted Δ_H), conjunction, and existential quantification. We use $\mathsf{def}(\mathcal{H})$ to denote the set of all pp-definable relations.

2.2 Constraint Satisfaction Problem

The counting constraint satisfaction problem can be formulated in several ways (see Section 1). We use the model theoretic form of this problem.

Definition 1 *Let $\mathfrak{H}$ be a class of relational structures. In the* counting constraint satisfaction problem associated with $\mathfrak{H}$ ($\#\mathrm{CSP}(\mathfrak{H})$)*, the objective is, given a structure $\mathcal{H} \in \mathfrak{H}$ and a structure $\mathcal{G}$, to compute the number of homomorphisms from $\mathcal{G}$ to $\mathcal{H}$. We will refer to this problem as a* uniform *#CSP.*

If $\mathfrak{H}$ consists of a single structure $\mathcal{H}$, then we write $\#\mathrm{CSP}(\mathcal{H})$ instead of $\mathrm{CSP}(\{\mathcal{H}\})$ and refer to such a problem as a non-uniform homomorphism problem, *because the inputs are just source structures.*

Example 1 ($\#H$-COLORING, [27, 36, 46]) A graph $\mathcal{H}$ is a structure with a vocabulary consisting of one binary symbol R. Then $\#\mathrm{CSP}(\mathcal{H})$ is widely known as the $\#H$-COLORING Problem, in which the objective is to compute the number of homomorphisms from a given graph into $\mathcal{H}$.

Example 2 (#3-SAT, [18, 19, 57, 58]) An instance of the #3-SAT problem is specified by giving a propositional logic formula in CNF each clause of which contains 3 literals, and asking how many assignments satisfy it. Therefore, #3-SAT is equivalent to $\#\mathrm{CSP}(\mathcal{S}_3)$, where $\mathcal{S}_3$ is the 2-element relational structure with the universe $\{0, 1\}$ and the vocabulary $R_1, \dots, R_8$, the predicates $R_1^{\mathcal{S}_3}, \dots, R_8^{\mathcal{S}_3}$ are the 8 predicates expressible by 3-clauses.

Example 3 Let F be a finite field and #LINEAR EQUATIONS is the problem of finding the number of solutions to a system of linear equations over F. It is not hard to see that #LINEAR EQUATIONS is equivalent to $\#\mathrm{CSP}(\mathfrak{L})$, where $\mathfrak{L}$ is the class of relational structures with the universe F and the relations corresponding to hyperplanes of finite-dimensional vector spaces over F.

Example 4 (Equations over semigroups, [50, 44]) Let S be a finite semigroup, that is, a set with a binary associative operation. An equation over S is an expression of the form $x_1 \cdot x_2 \cdot \dots \cdot x_m = y_1 \cdot y_2 \cdot \dots \cdot y_m$ where $\cdot$ is the semigroup operation, and x_i, y_j are either indeterminants or constants. Then $\#\mathrm{EQN}_S^*$ stands for the problem of counting the number of solutions to a system of semigroup equations.

The problem $\#\mathrm{EQN}_S^*$ is equivalent to the problem $\#\mathrm{CSP}(\mathfrak{S})$ where $\mathfrak{S}$ is the class of structures with universe S and relations expressible as the set of solutions of a semigroup equation.

In the last two examples, as well as for many other uniform problems, there is a minor ambiguity concerning a representation of the input. We always assume that in uniform problems the relations of the template are represented explicitly, by a list of tuples in the relation. In Examples 3,4 such a representation is not the most natural one. However, the class of relations admitting a succinct representation is rather limited (see, e.g. [39]), and thus such representations are unsuitable for the study of the general problem. Morever, changing representation does not affect the complexity of non-uniform problems.

Every counting CSP belongs to the class #P. However, the exact complexity of $\#\mathrm{CSP}(\mathcal{H})$ strongly depends on the structure $\mathcal{H}$. We say that a relational structure $\mathcal{H}$ is *#-tractable* if $\#\mathrm{CSP}(\mathcal{H})$ is solvable in polynomial time; $\mathcal{H}$ is *#P-complete* if $\#\mathrm{CSP}(\mathcal{H})$ is #P-complete. Note that all reductions used in this paper are Turing reductions. The research problem we deal with in this paper is the following one.

Problem 1 (classification problem) *Characterize #-tractable and #P-complete relational structures.*

Example 5 (1) Dyer and Greenhill [27] proved that if H is an undirected graph then $\#H$-COLORING can be solved in polynomial time if and only if every connected component of H is either a complete bipartite graph, or a complete graph with all loops present, or a single vertex. Otherwise the problem is #P-complete.

(2) A 2-element relational structure $\mathcal{H}$ is #-tractable if and only if every predicate of $\mathcal{H}$ can be represented by a system of linear equations over the 2-element field [18, 19]. Otherwise, $\mathcal{H}$ is #P-complete.

(3) $\#\mathrm{CSP}(\mathfrak{L})$ is solvable in polynomial time.

(4) The problem $\#\mathrm{EQN}^*_S$ is solvable in polynomial time if and only if S is a direct product of a uniformly inflated Abelian group, and inflated left-zero semigroup, and an inflated right-zero semigroup. Otherwise $\#\mathrm{EQN}^*_S$ is #P-complete. For details see [44].

2.3 Polymorphisms, Algebras and Complexity

We have shown in [4] that polymorphisms of relational structures are a very powerful tool to study the complexity of counting problems. Any operation on a set H can be extended in a standard way to an operation on tuples over H, as follows. For any (m-ary) operation f, and any collection of tuples $\mathbf{a}_1, \ldots, \mathbf{a}_m \in H^n$, define $f(\mathbf{a}_1, \ldots, \mathbf{a}_m)$ to be $(f(\mathbf{a}_1[1], \ldots, \mathbf{a}_m[1]), \ldots, f(\mathbf{a}_1[n], \ldots, \mathbf{a}_m[n]))$. Then f *preserves* an n-ary relation R (or R is *invariant* under f, or f is a *polymorphism of* R) if for any $\mathbf{a}_1, \ldots, \mathbf{a}_m \in R$ the tuple $f(\mathbf{a}_1, \ldots, \mathbf{a}_m)$ belongs to R. For a given set of operations, C, the set of all relations invariant under every operation from C is denoted by $\mathsf{Inv}(C)$. For a relational structure $\mathcal{H}$ we use $\mathsf{Pol}(\mathcal{H})$ to denote the set of all operations preserving every relation of $\mathcal{H}$.

Example 6 *Let R be the solution space of a system of linear equations over a field F. Then the operation $m(x, y, z) = x - y + z$ is a polymorphism of R. Indeed, let $A \cdot \mathbf{x} = \mathbf{b}$ be the system defining R, and $\mathbf{x}, \mathbf{y}, \mathbf{z} \in R$. Then*

$$A \cdot m(\mathbf{x}, \mathbf{y}, \mathbf{z}) = A \cdot (\mathbf{x} - \mathbf{y} + \mathbf{z}) = A \cdot \mathbf{x} - A \cdot \mathbf{y} + A \cdot \mathbf{z} = \mathbf{b}.$$

In fact, the converse can also be shown: if R is invariant under m then it is the solution space of a certain system of linear equations.

The following propositions links together polymorphisms and pp-definability of relations.

Proposition 1 ([32, 1, 42]) *Let $\mathcal{H}$ be a finite structure, and let $R \subseteq H^r$ be a non-empty relation. Then R is preserved by all polymorphisms of $\mathcal{H}$ if and only if R is pp-definable in $\mathbf{A}$.*

The connection between polymorphisms and the complexity of counting CSPs is provided by the following result.

Proposition 2 ([4]) *Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be relational structures with the same universe. If $\mathsf{Pol}(\mathcal{H}_1) \subseteq \mathsf{Pol}(\mathcal{H}_2)$ then $\#\mathrm{CSP}(\mathcal{H}_2)$ is polynomial time reducible to $\#\mathrm{CSP}(\mathcal{H}_1)$.*

Theorem 2 amounts to say that all the information about the complexity of $\#\mathrm{CSP}(\mathcal{H})$ can be extracted from the family of polymorphisms of $\mathcal{H}$. Sets of polymorphisms often provide a more convenient and concise way of describing a class of constraint satisfaction problems. For example, in [4], we used polymorphisms to identify some conditions necessary for the #-tractability of a relational structure. A ternary operation $m(x, y, z)$ on a set H is said to be *Mal'tsev* if $m(x, y, y) = m(y, y, x) = x$ for all $x, y \in H$.

Proposition 3 ([4]) *If $\mathcal{H}$ is a relational structure which is invariant under no Mal'tsev operation then $\mathcal{H}$ is #P-complete.*

Notice that if $\mathcal{H}$ has a Mal'tsev polymorphism then the decision CSP corresponding to $\mathcal{H}$ can be solved in polynomial time [7, 13].

Example 7 A Mal'tsev operation $m(x, y, z)$ is a polymorphism of the graph H_1 shown in Fig. 1, where m is defined as

$$m(i_1j_1, i_2j_2, i_3j_3) = ij,$$

$i = i_1$ $[j = j_1]$ unless $i_1 = i_2$ $[j_1 = j_2]$, in this case $i = i_3$ $[j = j_3]$.

The graph H_2 has no Mal'tsev polymorphisms. Indeed, if some $f(x, y, z)$ is a Mal'tsev operation, then

$$f\left(\begin{pmatrix} a \\ c \end{pmatrix}, \begin{pmatrix} a \\ d \end{pmatrix}, \begin{pmatrix} b \\ d \end{pmatrix}\right) = \begin{pmatrix} b \\ d \end{pmatrix} \notin E(H_2).$$

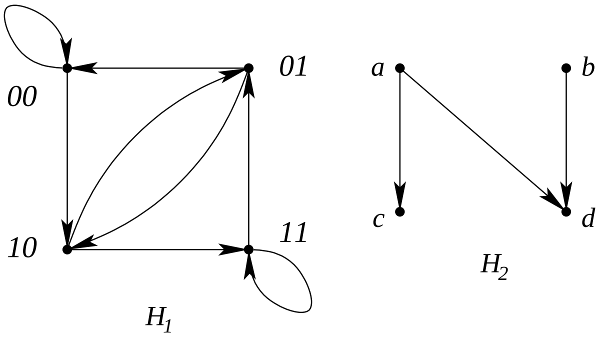

Figure 1:

In our algebraic definitions we follow [17, 49]. For algebraic notions and results concerning the decision CSP the reader is referred to [9, 11].

A (*universal*) *algebra* is an ordered pair $\mathbb{A} = (A, F)$ where A is a non-empty set and F is a family of finitary operations on A. The set A is called the *universe* of $\mathbb{A}$, the operations from F are called *basic*. An algebra with a finite universe is referred to as a *finite algebra*.

Any relational structure $\mathcal{H}$ with universe Hcan be converted into an algebra $\mathsf{Alg}(\mathcal{H}) = (H; \mathsf{Pol}(\mathcal{H}))$. Conversely, every algebra $\mathbb{A} = (A; F)$ corresponds to a class of structures $\mathsf{Str}(\mathbb{A})$ with universe A and relations from $\mathsf{Inv}(F)$. Using this correspondence we can define #-tractable algberas. An algebra $\mathbb{A}$ is said to be #-tractable if every structure $\mathcal{H} \in \mathsf{Str}(\mathbb{A})$ is #-tractable; it is said to be #P-complete if some $\mathcal{H} \in \mathsf{Str}(\mathbb{A})$ is #P-complete.

We shall express the complexity of $\#\mathrm{CSP}(\mathcal{H})$ in terms of $\mathsf{Alg}(\mathcal{H})$. For example, if an algebra has a Mal'tsev operation, it is called a *Mal'tsev algebra*. Proposition 3 implies that if $\#\mathrm{CSP}(\mathcal{H})$ is solvable in polynomial time then $\mathsf{Alg}(\mathcal{H})$ is Mal'tsev.

2.4 Subalgebras and congruences

We shall use various constructions on algebras, but two of these constructions, subalgebras and congruences, can be defined for relational structures, and are very useful and illustrative in this context.

A *subalgebra* of a structure $\mathcal{H} = (H; R_1^{\mathcal{H}}, \dots, R_k^{\mathcal{H}})$ is a unary relation definable in $\mathcal{H}$, and a *congruence* of $\mathcal{H}$ an equivalence relation definable in $\mathcal{H}$. For a subset $B \subseteq H$, the substructure of $\mathcal{H}$ *induced* by B is defined to be $\mathcal{H}|_B = (B; R_1^{\mathcal{H}}|_B, \dots, R_k^{\mathcal{H}}|_B)$, where $R_i|_B = R_i \cap B^{m_i}$, R_i is m_i-ary. For an equivalence relation α and $a \in H$, the class of α containing α is denoted by $a/_\alpha$ and the set of all classes of α by $H/_\alpha$. The *quotient structure* $\mathcal{H}/_\alpha$ is defined to be $\mathcal{H}/_\alpha = (H/_\alpha; R_1^{\mathcal{H}}/_\alpha, \dots, R_k^{\mathcal{H}}/_\alpha)$, where $R_i/_\alpha = \{(a_1/_\alpha, \dots, a_{m_i}/_\alpha) \mid (a_1, \dots, a_{m_i}) \in R_i\}$.

Example 8 Let $\mathcal{H}$ be a digraph without sources and sinks, i.e. the in-degree and out-degree of each vertex is non-zero. We define two binary relations on the vertex set H of $\mathcal{H}$: $(a,b) \in \theta$ if and only if a, b have a common out-neighbour and $(a,b) \in \eta$ if and only if a, b have a common in-neighbour; in other words, $\theta = \{(a,b) \mid (a,c), (b,c)$ for a certain $c \in H\}$, $\eta = \{(a,b) \mid (c,a), (c,b)$ for a certain $c \in H\}$. In general, θ, η are reflexive and symmetric relation. However, if $\mathcal{H}$ has a Mal'tsev polymorphism m, they are also transitive. Indeed, suppose that $(a, \beta) \in \theta$, $d \in H$ is their common out-neighbour and c is an out-neighbour of α. If c is not an out-neighbour of b, then $\mathcal{H}$ contains H_2 (see Fig. 1) as an induced subgraph, which contradicts the assumption that $\mathcal{H}$ has a Mal'tsev polymorphism. Therefore, the out-neighbourhoods of a, b are equal whenever $(a,b) \in \theta$, that implies transitivity. Thus, θ, η are congruences of $\mathcal{H}$.

For the graph H_3 shown in Fig. 2, the θ-classes are $\{a,b,c\}, \{d,e\}$ and the η-classes are $\{a,b,e\}, \{c,d\}$.

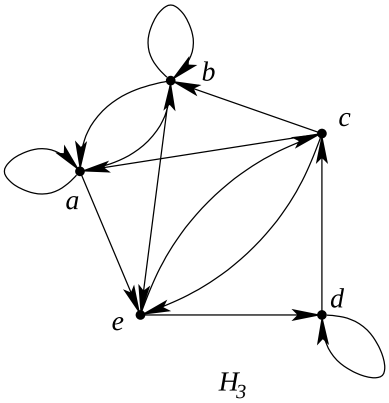

Figure 2:

Proposition 4 ([4]) *Let $\mathcal{H}$ be a relational structure, B and α its subalgebra and congruence respectively.*

(1) If $\mathcal{H}$ is #-tractable then so are $\mathcal{H}|_B$ and $\mathcal{H}/_\alpha$.

(2) If $\mathcal{H}|_B$ or $\mathcal{H}/_\alpha$ is #P-complete then $\mathcal{H}$ is #P-complete.

Let $R \in \mathsf{def}(\mathcal{H})$ be an n-ary relation. It can be viewed as a subalgebra of nth direct power of $\mathcal{H}$. A *congruence on* R is a $2n$-ary relation $Q \in \mathsf{def}(\mathcal{H})$ such that $\mathrm{pr}_{\{1,\dots,n\}} Q = \mathrm{pr}_{\{n+1,\dots,2n\}} Q = R$, and, if Q is treated as a binary relation on R, it is an equivalence relation.

The existence of a Mal'tsev polymorphism provides a necessary condition for the #-tractability of a relational structure. However, it is not a sufficient condition, as Example 9 shows. In the next section we prove two more necessary conditions, and a particular case of one of them is that proved [5].

Let α, β be congruences of a $\mathcal{H}$, where α, β are incomparable, that is, neither $\alpha \subseteq \beta$, nor $\beta \subseteq \alpha$. Let $A_1, \ldots, A_k$ and $B_1, \ldots, B_\ell$ be α- and β-classes respectively (see Fig.3). Then $M(\alpha, \beta)$ denotes the $k \times \ell$-matrix (m_{ij}), where $m_{ij} = |A_i \cap B_j|$.

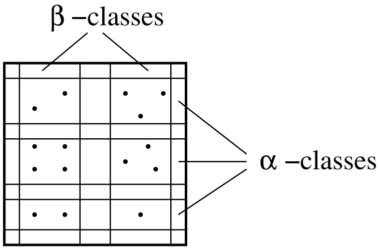

Figure 3:

Proposition 5 ([5]) *Let $\mathcal{H}$ be a relational structure, and let α, β be congruences of $\mathcal{H}$. If $\mathsf{rank}(M(\alpha,\beta)) > k$, where k is the number of classes in the smallest congruence containing both α and β, then $\#\mathrm{CSP}(\mathcal{H})$ is #P-complete.*

Example 9 Let $\mathcal{H}$ be the graph H_3 shown in Fig. 2, $\alpha = \theta_{H_3}$ and $\beta = \eta_{H_3}$. We have $A_1 = \{a, b, c\}, A_2 = \{e, d\}$, $B_1 = \{a, b, e\}, B_2 = \{c, d\}$ and

$$M(\alpha, \beta) = \left(\begin{array}{cc} 2 & 1 \\ 1 & 1 \end{array} \right).$$

By Proposition 5, the problem $\#\mathrm{CSP}(H_3)$ is #P-complete.

2.5 Varieties and Complexity

It will be convenient for us to jump back forth between model-theoretic and algebraic views to the CSP. The language of relational structures is more convenient when describing algorithms. On the other hand, standard algebraic constructions allow us to strengthen necessary conditions for #-tractability, and eventually formulate a criterion for #-tractability.

Definition 2 (1) *Let $\mathbb{A} = (A; F)$ be an algebra. The k-th* direct power *of $\mathbb{A}$ is the algebra $\mathbb{A}^k = (A^k; F)$ where we treat each (n-ary) operation $f \in F$ as acting on A^k component-wise.*

(2) *Let $\mathbb{A} = (A; F)$ be an algebra, and let B be a subset of A such that, for any (n-ary) $f \in F$, and for any $b_1, \ldots, b_n \in B$, we have $f(b_1, \ldots, b_n) \in B$. Then the algebra $\mathbb{B} = (B; F\big|_B)$, where $F\big|_B$ consists of restrictions of operations $f \in F$ to B, is called a* subalgebra *of $\mathbb{A}$.*

Note that a set B is a subalgebra of a structure $\mathcal{H}$ if and only if B is the universe of a subalgebra of $\mathsf{Alg}(\mathcal{H})$.

(3) *Let $\mathbb{A}_1 = (A_1; F_1)$ and $\mathbb{A}_2 = (A_2; F_2)$ such that $F_1 = \{f_i^1 \mid i \in I\}$, $F_2 = \{f_i^2 \mid i \in I\}$, and f_i^1, f_i^2 are of the same arity, $i \in I$. A mapping $\varphi : A_1 \to A_2$ is called a* homomorphism *from $\mathbb{A}_1$ to $\mathbb{A}_2$ if $\varphi f_i^1(a_1, \ldots, a_{n_i}) = f_i^2(\varphi(a_1), \ldots, \varphi(a_{n_i}))$ holds for all $i \in I$ and all $a_1, \ldots, a_{n_i} \in A_1$. If the mapping φ is onto then $\mathbb{A}_2$ is said to be a* homomorphic image *of $\mathbb{A}_1$.*

A common way of constructing homomorphic images is through congruences and quotient algebras. A *congruence* of an algebra $\mathbb{A} = (A; F)$ is an equivalence relation on A invariant under all operations from F. Let θ be a congruence of $\mathbb{A}$. The algebra $\mathbb{A}/_\theta = (A/_\theta; F/_\theta)$, where $F/_\theta = \{f/_\theta \mid f \in F\}$ and $f/_\theta$ is defined through the equality $f/_\theta(a_1/_\theta, \ldots, a_n/_\theta) = (f(a_1, \ldots, a_n))/_\theta$ is called a *quotient algebra*. Observe that an equivalence relation is a congruence of a structure $\mathcal{H}$ if and only if it is a congruence of $\mathsf{Alg}(\mathcal{H})$.

A property of algebras such that if an algebra enjoys the property then any its subalgebra, homomorphic image, and direct power also enjoys it, is said to be *hereditary*. Universal algebra mostly deals with hereditary properties [37, 49]. Therefore, the next theorem allows us to apply the methods of modern algebra to the study of the complexity of the counting CSP.

Theorem 1 ([4, 14]) *Let $\mathbb{A} = (A; F)$ be a finite algebra. Then*

(i) *if $\mathbb{A}$ is #-tractable then so is every subalgebra, homomorphic image, and direct power of $\mathbb{A}$.*

(ii) *if $\mathbb{A}$ has an #P-complete subalgebra, homomorphic image, or direct power, then $\mathbb{A}$ is #P-complete itself.*

For an algebra $\mathbb{A}$ the class of algebras that are homomorphic images of subalgebras of direct powers of $\mathbb{A}$ is called the *variety* generated by $\mathbb{A}$.

An operation f on the universe of an algebra $\mathbb{A} = (A; F)$ that preserves all relations invariant under F is called a *term* operation of $\mathbb{A}$. Every term operation of $\mathbb{A}$ can be obtained from operations of F by means of superposition.

An operation f on a set A is said to be *idempotent* if the equality $f(x, \ldots, x) = x$ holds for all x from A. Algebras whose basic operations are idempotent posess many useful properties that will assist in our investigation. The *full idempotent reduct* of an algebra $\mathbb{A} = (A; F)$ is the algebra $\mathsf{Id}(\mathbb{A}) = (A; F_{\rm id})$ where $F_{\rm id}$ consists of all idempotent term operations of $\mathbb{A}$. There is another way to characterize $F_{\rm id}$. If $\mathbb{A} = \mathsf{Alg}(\mathcal{H})$ for a certain relational structure $\mathcal{H}$, then $\mathsf{Id}(\mathbb{A}) = \mathsf{Alg}(\mathcal{H}_{\rm id})$, where $\mathcal{H}_{\rm id}$ is an expansion of $\mathcal{H}$ by unary relations C_h, $h \in \mathcal{H}$, and C_h is interpreted an a *constant relation* $\{(h)\}$, containing only one tuple, namely (h).

Theorem 2 ([4, 14]) *A finite algebra $\mathbb{A}$ is #-tractable* [*#P-complete*] *if and only if so is $\mathsf{Id}(\mathbb{A})$.*

If $\mathbb{A}$ is an idempotent algebra and the condition of Proposition 5 is true for every pair of congruences of $\mathbb{A}$ then $\mathbb{A}$ is said to be *congruence singular*. If every finite algebra in a variety is congruence singular then the variety is called congruence singular. We call a relational structure $\mathcal{H}$ congruence singular if $\mathsf{Alg}(\mathcal{H})$ generates a congruence singular variety. By Proposition 5 and Theorems 1, 2, every structure $\mathcal{H}$ that is not #P-complete is congruence singular. The main result of the paper is that this condition is sufficient for #-tractability.

Theorem 3 *A relational structure $\mathcal{H}$* [*an algebra $\mathbb{A}$*]*, is #-tractable if and only if $\mathcal{H}_{\rm id}$ is congruence singular* [*$\mathbb{A}$ generates a congruence singular variety*].

Observe that the condition of having a Mal'tsev polymorphism (term operation) is not included into the criterion. As we shall see later (Lemma 1) every congruence singular structure has a Mal'tsev polymorphism.

3 Congruence lattices and the structure of relations

3.1 Congruence lattices and types of prime quotients

In this section we look closer at the family of congruences of a relational structure $\mathcal{H}$. We shall assume that $\mathcal{H}$ has a Mal'tsev polymorphism $m(x, y, z)$. All definitions and results given here were originally introduced for algebras [15, 49]. As our algorithms are described in terms of relational structures, we reformulate them in terms of structures, replacing congruences of algebra with congruences of structures, and term operations of an algebra with polymorphisms of a structure. However, the notions we arrive to for a structure $\mathcal{H}$ are exactly the same as those defined for the algebra $\mathsf{Alg}(\mathcal{H})$.

The set of all congruences of $\mathcal{H}$ is denoted by $\mathsf{Con}(\mathcal{H})$. Let $\alpha, \beta \in \mathsf{Con}(\mathcal{H})$. The intersection of α and β is again a congruence of $\mathcal{H}$ is denoted $\alpha \wedge \beta$. As is well known, the smallest equivalence relation containing both α and β is the transitive closure of $\alpha \cup \beta$. It can be shown that this equivalence relation is a congruence of $\mathcal{H}$, denoted by $\alpha \vee \beta$. The set $\mathsf{Con}(\mathcal{H})$ together with the operations $\wedge$ (*meet*) and $\vee$ (*join*) is called the *congruence lattice* of $\mathcal{H}$. The set $\mathsf{Con}(\mathcal{H})$ is naturally ordered with respect to inclusion. The least element of $\mathsf{Con}(\mathcal{H})$ is the equality relation, denoted by Δ, and the greatest element is the full relation, denoted by $\bigtriangledown$.

If R is a relation pp-definable in$\mathcal{H}$, then $\mathsf{Con}(R)$ denotes the set of all congruences on R. This set depends on $\mathcal{H}$ as well as on R, but usually $\mathcal{H}$ is clear from the context. The set $\mathsf{Con}(R)$ is also a lattice.

Since $\mathcal{H}$ has a Mal'tsev polymorphism, the set $\mathsf{Con}(\mathcal{H})$ cannot be just an arbitrary collection of equivalence relation. In particular, every two members α, β of $\mathsf{Con}(\mathcal{H})$ must be *permutable*, that is $\alpha \circ \beta = \beta \circ \alpha$. This means that, for any α-class A and any β-class B belonging the same $\alpha \vee \beta$-class, $A \cap B$ is non-empty (see Fig.4).

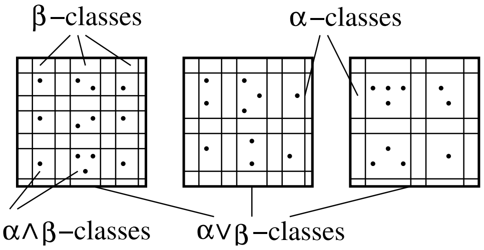

Figure 4:

Lemma 1 *If a relational structure $\mathcal{H}$ is congruence singular* [*an algebra $\mathbb{A}$ generates a congruence singular variety*], *then it has a Mal'tsev polymorphism* [*a Mal'tsev term operation*].

Proof: By a well known result of Mal'tsev [15], an algebra $\mathbb{A}$ has a Mal'tsev term operation if and only if any two congruences of any algebra in the variety generated by $\mathbb{A}$ are permutable. Therefore it suffices to prove that if the variety generated by $\mathsf{Alg}(\mathcal{H})$ for a structure $\mathcal{H}$ is congruence singular then it is congruence permutable.

As is easily seen, congruences α, β are permutable if and only if $\alpha \circ \beta = \beta \circ \alpha = \alpha \vee \beta$. Suppose $\mathcal{H}$ is congruence singular, $\mathbb{B} \in \mathsf{var}(\mathsf{Alg}(\mathcal{H}))$, and $\alpha, \beta \in \mathsf{Con}(\mathbb{B})$. If $\alpha \subseteq \beta$ or $\beta \subseteq \alpha$ then they are obviously permutable. If the congruences are incomparable then $\mathsf{rank}(M(\alpha, \beta)) = k$ where k is the number of $\alpha \vee \beta$-classes. This equality implies, in particular, that for any a, b from the same $\alpha \vee \beta$-class, say, a belongs to α-class A_1 and β-class B_1, and b belongs to α-class A_2 and β-class B_2, we have $A_1 \cap B_2 \neq \varnothing$

and $A_2 \cap B_1 \neq \varnothing$ (the corresponding entries of $M(\alpha,\beta)$ must be nonzero). Then $\langle a,b\rangle \in \alpha\circ\beta$, as any $c \in A_1 \cap B_2$ witnesses, and $\langle a,b\rangle \in \beta\circ\alpha$, as any $d \in A_2\cap B_1$ witnesses. Thus $\alpha\circ\beta = \beta\circ\alpha = \alpha\vee\beta$. □

A pair of congruences $\langle\alpha,\beta\rangle$ is said to be a *prime quotient* if $\alpha \le \beta$ and, for any γ such that $\alpha\le\gamma\le\beta$, either $\gamma = \alpha$ or $\gamma = \beta$.

We shall use some notions and results of tame congruence theory [37]. Tame congruence theory is a tool to study a local structure of universal algebras and relational structures through certain properties of prime quotients of the congruence lattice. In general, this theory identifies five possible types of such quotients defined in a rather sophisticated way. Fortunately, in our case of relational structures with a Mal'tsev polymorphism, only two of those types can occur, and the definition of these possible types can be significantly simplified.

If every polymorphism of a relational structure $\mathcal{H}$ is idempotent, then, for any congruence α of $\mathcal{H}$, every α-class A is a subalgebra. Indeed, for any $f(x_1,\dots,x_n) \in \mathsf{Pol}(\mathcal{H})$ and any $a_1,\dots,a_n \in A$, we have $(a_1,a_1),(a_2,a_1),\dots,(a_n,a_1) \in \alpha$, $f(a_1,\dots,a_1) = a_1$ and therefore

$$f\left(\begin{pmatrix}a_1\\a_1\end{pmatrix},\begin{pmatrix}a_1\\a_2\end{pmatrix},\dots,\begin{pmatrix}a_1\\a_n\end{pmatrix}\right) = \begin{pmatrix}f(a_1,\dots,a_n)\\a_1\end{pmatrix} \in \alpha.$$

Hence, $f(a_1,\dots,a_n)\in A$.

A prime quotient $\alpha \prec \beta$ is said to be of *affine* type, if, for any β-class B, there is a module M_B with the base set $B/_{\alpha}$ over a ring R_B such that for any $f(x_1,\dots,x_n,y_1,\dots,y_m) \in \mathsf{Pol}(\mathcal{H})$ and $a_1,\dots,a_m \in H$, if the operation $g(x_1,\dots,x_n) = f(x_1,\dots,x_n,a_1,\dots,a_m)$ preserves B, then it can be represented as an operation of the module M_B:

$$(g|_B(x_1,\dots,x_n))/_{\alpha} = c_1x_1 + \dots c_nx_n + a.$$

In all other cases, $\alpha\prec\beta$ has *Boolean* type.

Example 10 Let $\mathcal{L}_2$ be a 2-element relational structure whose relational symbols are interpreted as solution spaces to systems of linear equations. Then $\mathcal{L}_2$ has only two congruences: Δ_2, the equality relation, and ∇_2, the total binary relation. As Example 6 shows, the prime quotient $\Delta_2 \prec \nabla_2$ is of affine type. Thus, affine type corresponds to some kind of "linearity" in a broad sense.

Prime intervals $\alpha_1\prec\beta_1$ and $\alpha_2\prec\beta_2$ are said to be *projective* if $\beta_1\vee\alpha_2 = \beta_2$, $\beta_1\wedge\alpha_2 = \alpha_1$ or $\alpha_1\vee\beta_2 = \beta_1$, $\alpha_1\wedge\beta_2 = \alpha_1$. Thus projectivity is a binary relation on the set of prime intervals of $\mathsf{Con}(\mathcal{H})$. Two intervals that belong to the transitive closure of this relation are said to be *perspective* to each other.

Lemma 2 ([37], Lemma 6.2) *If $\alpha_1\prec\beta_1$ and $\alpha_2\prec\beta_2$ are perspective intervals in $\mathsf{Con}(\mathcal{H})$, then they have the same type.*

3.2 Congruence lattices of Mal'tsev Algebras

We will sometimes distinguish two cases: when the congruence lattice of our relational structure omits affine type, and when affine type occurs in this lattice.

3.2.1 Algebras omitting affine type.

If $\mathcal{H}$ omits affine type then, by Theorem 9.15 of [37], $\mathsf{Con}(\mathcal{H})$ is *distributive*, that is, for any $\alpha, \beta, \gamma \in \mathsf{Con}(\mathcal{H})$, the equality $\alpha \wedge (\beta \vee \gamma) = (\alpha \wedge \beta) \vee (\alpha \wedge \gamma)$ holds. Finite distributive lattices are exhaustively studied (see, e.g. [33]). In particular, there is a finite set, M, and a injective mapping $\pi\colon \mathsf{Con}(\mathcal{H}) \to 2^M$ (the set of all subsets) such that $\pi(\alpha \vee \beta) = \pi(\alpha) \cup \pi(\beta)$ and $\pi(\alpha \wedge \beta) = \pi(\alpha) \cap \pi(\beta)$. We use the following representation of a set M. Take a maximal chain C in $\mathsf{Con}(\mathcal{H})$, that is, a chain of congruences $\Delta = \theta_0 \prec \theta_1 \prec \ldots \prec \theta_\ell = \nabla$. The set M is defined to be the set of the prime quotients of the chain. Slightly abusing notion the quotient th_{i-1}, θ_i will be denoted by i. A congruence $\theta \in \mathsf{Con}(\mathcal{H})$ corresponds to the sets of quotients from M that are projective to quotients of the form $\gamma \prec \beta \leq \theta$. The bottom end of a prime quotient $\alpha \in \{1, \ldots, \ell\}$ will be denoted by α^-, and the top one by α^+.

Example 11 The lattice shown in Fig. 5(a) is distributive and its representation as a lattice of subsets is also shown.

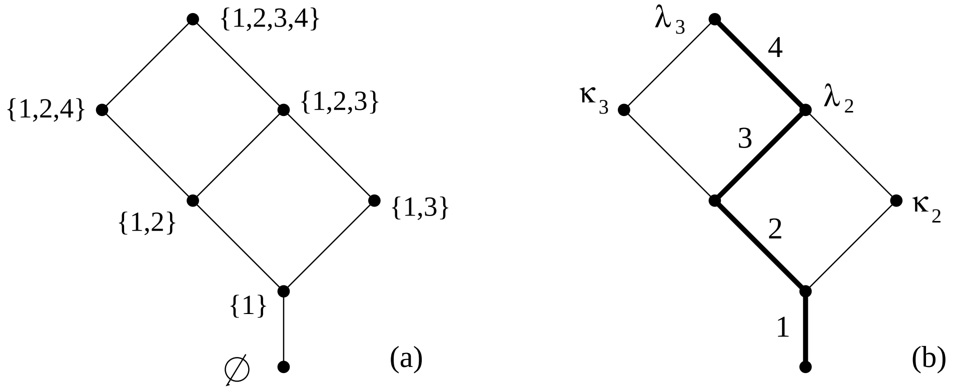

Figure 5:

The following proposition comprises properties of $\mathsf{Con}(\mathcal{H})$ that follow easily from the representation of this lattice as a lattice of subsets.

Proposition 6 *(1) Every prime interval in* $\mathsf{Con}(\mathcal{H})$ *is perspective to one and only one of the intervals of C.*

(2) For any $\alpha \in M$, that is, any prime interval in C, there is the greatest prime interval $\kappa_\alpha \prec \lambda_\alpha$ perspective to α; that is, for any $\beta \prec \gamma$ perspective to α we have $\beta \leq \kappa_\alpha$ and $\gamma \leq \lambda_\alpha$.

(3) For any $\alpha \in M$, the congruence κ_α is meet-irreducible*, that is, if $\kappa_\alpha = \beta \wedge \gamma$ than $\kappa_\alpha = \beta$ or $\kappa_\alpha = \gamma$ (see Fig.5(b).*

3.2.2 Algebras admitting affine type.

Let us again consider the congruence lattice $\mathsf{Con}(\mathcal{H})$. A congruence β is said to be *solvable* over α if there are $\alpha = \alpha_1 \prec \ldots \prec \alpha_k = \beta$ such that all the prime quotients $\alpha_i \prec \alpha_{i+1}$ have affine type. Then $\overset{s}{\sim}$ denotes a binary relation on $\mathsf{Con}(\mathcal{H})$ defined as follows: $\alpha \overset{s}{\sim} \beta$ if and only if $\alpha \vee \beta$ is solvable over $\alpha \wedge \beta$. If $\alpha \leq \beta$ then the set of all γ such that $\alpha \leq \gamma \leq \beta$ is said to be an *interval* in $\mathsf{Con}(\mathcal{H})$, denoted $[\alpha, \beta]$. The next proposition lists some properties of $\overset{s}{\sim}$ that follows from well known facts about modular lattices, Mal'tsev operations and Lemma 7.4, Theorem 7.7 from [37].

Proposition 7 *(1)* $\overset{s}{\sim}$ *is an equivalence relation and, moreover, a* congruence *of* $\mathsf{Con}(\mathcal{H})$*; that is, for any* $\alpha_1, \alpha_2, \beta_1, \beta_2 \in \mathsf{Con}(\mathcal{H})$ *such that* $\alpha_1 \overset{s}{\sim} \alpha_2, \beta_1 \overset{s}{\sim} \beta_2$*, we have* $(\alpha_1 \vee \beta_1) \overset{s}{\sim} (\alpha_2 \vee \beta_2), (\alpha_1 \wedge \beta_1) \overset{s}{\sim} (\alpha_2 \wedge \beta_2)$.

(2) Every class S *of* $\overset{s}{\sim}$ *has the greatest* η_S *and the least* θ_S *elements (with respect to* $\leq$*), and equals the interval* $[\theta_S, \eta_S]$*. Every prime quotient inside* S *has affine type.*

(3) The quotient lattice $\mathcal{L} = \mathsf{Con}(\mathcal{H})/\overset{s}{\sim}$ *is distributive (see Fig.6).*

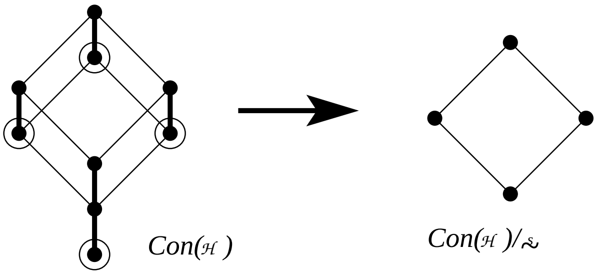

Figure 6: Congruence lattice and its quotient lattice modulo $\overset{s}{\sim}$. Prime quotients of affine type are shown by thick lines; the least elements in the classes of $\overset{s}{\sim}$ are encircled

Proposition 7(3) implies that $\mathcal{L}$ can be represented as a lattice of subsets of a finite set M. Similar to Subsection 3.2.1, M can be chosen to be the set of prime intervals of a maximal chain C in $\mathcal{L}$. Note that the endpoints of $\alpha \in M$ are sets S_1, S_2 of congruences from $\mathsf{Con}(\mathcal{H})$ (S_1 correspods to the bottom end of α). By α^- we denote the greatest element of S_1, and by α^+ the least element of S_2 such that $\alpha^- \leq \alpha^+$. Let $\beta \prec \gamma$ be the greatest interval in $\mathcal{L}$ perspective to α. Again, β and γ are sets T_1, T_2 of congruences from $\mathsf{Con}(\mathcal{H})$ (T_1 corresponds to β). By κ_α we denote the greatest element of T_1, and λ_α the least element in T_2 such that $\kappa_\alpha \leq \lambda_a$ (see Fig.7).

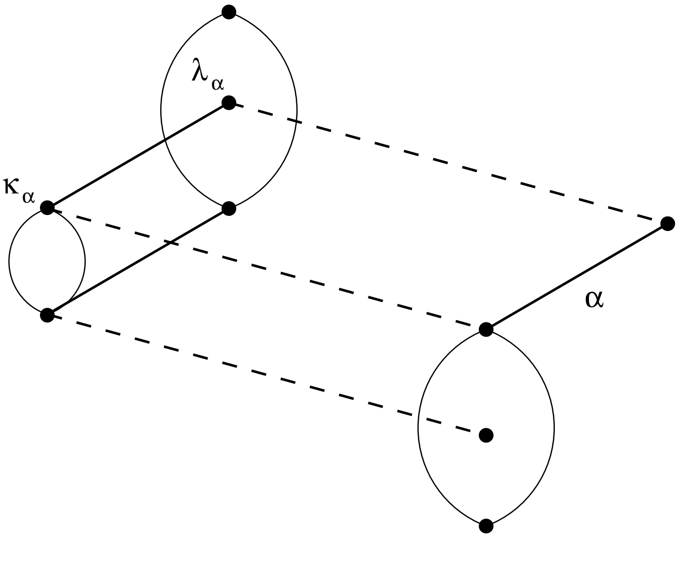

Figure 7: Congruence lattice and congruences $\kappa_\alpha, \lambda_\alpha$. Solid lines represent prime intervals of Boolean type, ovals represent $\overset{s}{\sim}$-classes

Proposition 8 *(1) The interval* $[\alpha^-, \alpha^+]$ *is perspective to* $[\kappa_\alpha, \lambda_\alpha]$.

(2) The intervals $[\alpha^-, \alpha^+]$ and $[\kappa_\alpha, \lambda_\alpha]$ are prime.
(3) The intervals $[\alpha^-, \alpha^+]$ and $[\kappa_\alpha, \lambda_\alpha]$ have Boolean type.
(4) The congruence κ_α is meet-irreducible.

3.3 Structure of relations invariant under a Maltsev operation

3.3.1 Basic properties

The following proposition contains some basic properties of Mal'tsev algebras and relations invariant under a Mal'tsev operation, that will be constantly used. Some of the results we cite below are traditionally stated in terms of algebras: a relation pp-definable in a structure $\mathcal{H}$ is treated as a subalgebra of a direct power of $\mathsf{Alg}(\mathcal{H})$. In order to keep the presentation uniform we formulate them in terms of relations and relational structures.

Proposition 9 *Let $\mathcal{H}$ be a structure with a Mal'tsev polymorphism and R an n-ary relation pp-definable in $\mathcal{H}$. Then for any $I \subseteq [n]$ the following properties hold*
1. R is rectangular, *that is if $\mathbf{a}, \mathbf{b} \in \mathrm{pr}_I R, \mathbf{c}, \mathbf{d} \in \mathrm{pr}_{[n]-I} R$ and $(\mathbf{a}, \mathbf{c}), (\mathbf{a}, \mathbf{d}), (\mathbf{b}, \mathbf{c}) \in R$, then $(\mathbf{b}, \mathbf{d}) \in R$.*
2. The relation $\theta_I = \{\langle \mathbf{a}, \mathbf{b}\rangle \in (\mathrm{pr}_I R)^2 \mid$ there is $\mathbf{c} \in \mathrm{pr}_{[n]-I} R$ such that $(\mathbf{a}, \mathbf{c}), (\mathbf{b}, \mathbf{c}) \in R\}$ is a congruence of $\mathrm{pr}_I R$.
3. R is a disjoint union of sets of the form $B \times C$ where B is a θ_I-class and C is a $\theta_{[n]-I}$-class.

Binary relations invariant with respect to a Mal'tsev operation have particularly simple form. Let B_1, B_2 be subalgebras of $\mathcal{H}$ and $\alpha_1 \in \mathsf{Con}(B_1)$, $\alpha_2 \in \mathsf{Con}(B_2)$. Let also φ be a mapping from $B_1/_{\alpha_1}$ to $B_2/_{\alpha_2}$. The *thick mapping* corresponding to φ is the binary relation $R = \{(a, b) \in B_1 \times B_2 \mid \varphi(a/_{\alpha_1}) = b/_{\alpha_2}\}$. Any congruence α is the thick mapping corresponding to the identity mapping on $H/_{\alpha}$.

Corollary 1 *Every binary relation compatible with $\mathbb{A}$ is a thick mapping.*

We shall intensively use thick mappings throughout the paper. Let $R \in \mathsf{def}(\mathcal{H})$ be a kth subdirect power of H. For $i, j \in [k]$ by $\psi_{i,j}$ we denote the thick mapping equal to $\mathrm{pr}_{i,j} R$. If it is a thick mapping corresponding to $\varphi\colon H/_{\alpha} \to H/_{\alpha}$ for some $\alpha \in \mathsf{Con}(H)$, we say that $\psi_{i,j}$ is a *thick mapping of level* α. Let $\beta \in \mathsf{Con}(\mathcal{H})$. By β^* we denote an equivalence relation on the set $[k]$ defined as follows: $\langle i, j\rangle \in \beta^*$ if and only if $\mathrm{pr}_{i,j} R$ is a thick mapping from $H/_{\gamma_1}$ to $H/_{\gamma_2}$ for some $\gamma_1, \gamma_2 \le \beta$. The following lemma follows from the definitions.

Lemma 3 *If $R \in \mathsf{def}(\mathcal{H})$ is a subdirect power of H then, for any $\alpha \in \mathsf{Con}(\mathcal{H})$, any α^*-class A, any $g, g' \in A$, and any sequence $g = g_1, \ldots, g_n = g'$ such that $\psi_{g_i, g_{i+1}}$ is a thick mapping of $H/_{\beta_i}$ to $H/_{\gamma_i}$ for some $\beta_i, \gamma_i \le \alpha$, $i \in [n-1]$, we have*

$$\psi_{g_1,g_2} \circ \ldots \circ \psi_{g_{n-1},g_n} \subseteq \psi_{g,g'}.$$

Lemma 3 implies that, for any congruence α and any α^*-class A, we can select a representative g_A and a family of mappings $\varphi_g\colon H/_{\alpha} \to H/_{\alpha}$, where $g \in A$, such that for any homomorphism $\psi\colon \mathcal{G} \to \mathcal{H}_g$, we have $\psi(g)/_{\alpha} = \varphi_g(\psi(g_A))$.

3.3.2 Boolean type and rectangularity properties

Let $\mathbb{A}$ be a finite algebra. The algebra $\mathbb{A}$ is called *subdirectly irreducible* if there is a congruence μ, the *monolith* of $\mathbb{A}$, such that $\Delta \prec \mu$ and, for any congruence $\gamma \neq \Delta$, we have $\mu \leq \gamma$. We call a relational structure $\mathcal{H}$ subdirectly irreducible if $\mathsf{Alg}(\mathcal{H})$ is subdirectly irreducible. The monolith of a subdirectly irreducible structure is defined as the monolith of $\mathsf{Alg}(\mathcal{H})$.

Let $R \in \mathsf{def}(\mathcal{H})$, where $\mathcal{H}$ is a subdirectly irreducible structure with a Mal'tsev polymorphism, be an k-ary subdirect power of $\mathcal{H}$. The equivalence relation μ^* is defined in the same way as before. In [12], we defined *coherent sets* of the subdirect power R satisfying these conditions, as classes of a certain partition of the set $[k]$[1]. We do not need here a precise definition of coherent sets, because if the interval $\Delta \prec \mu$ has Boolean type then it follows from Lemma 2.7 of [12] that the coherent sets are equal to the classes of μ^*.

Lemma 4 (Lemma 2.6, [12]) *Let R be a subdirect power of $\mathcal{H}$ and the structure $\mathcal{H}$ is subdirectly irreducible. Let also μ be its monolith and $B_1, \ldots, B_k$ μ-classes such that $R \cap (B_1 \times \ldots \times B_k) \neq \varnothing$. Let $I_1, \ldots, I_\ell$ be the coherent sets and*

$$B_{I_j} = \mathrm{pr}_{I_j} R \cap \prod_{i \in I_j} B_i.$$

Then $R \cap (B_1 \times \ldots \times B_k) = B_{I_1} \times \ldots \times B_{I_\ell}$.

For a congruence $\alpha \in \mathsf{Con}(\mathcal{H})$, let α^k denote the congruence of R consisting of pairs $\langle \mathbf{a}, \mathbf{b} \rangle$ of tuples such that $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \alpha$ for all $i \in [k]$ (it is an easy exercise to check that α^k is indeed a congruence).

Proposition 10 *Let $\mathcal{H}$ be a structure with a Mal'tsev polymorphism, let M be a maximal chain in $\mathsf{Con}(\mathcal{H})$, let R be a kth subdirect power of $\mathcal{H}$ and $\alpha \in M$. Let also $B_1, \ldots, B_k$ be classes of λ_α and $I_1, \ldots, I_\ell$ the classes of κ^*_α, $I_j = \{i_{j1}, \ldots, i_{jk_j}\}$. Then either $R \cap (B_1 \times \ldots B_k) = \varnothing$, or*

$$R/_{\kappa_\alpha^k} \cap (B_1/_{\kappa_\alpha} \times \ldots \times B_k/_{\kappa_\alpha}) = B_{I_1}/_{\kappa_\alpha^{|I_1|}} \times \ldots \times B_{I_\ell}/_{\kappa_\alpha^{|I_\ell|}},$$

where $R/_{\kappa_\alpha^k} = \{(\mathbf{a}[1]/_{\kappa_\alpha}, \ldots, \mathbf{a}[k]/_{\kappa_\alpha}) \mid \mathbf{a} \in R\}$ and $B_{I_j} = \mathrm{pr}_{I_j} R \cap \prod_{i \in I_j} B_i$, and

$$B_{I_j} = \{(a, \psi_{i_{j1}, i_{j2}}(a), \ldots, \psi_{i_{j1}, i_{jk_j}}(a)) \mid a \in B_j/_{\kappa_\alpha}\}.$$

Proof: The relation $R/_{\kappa_\alpha^k}$ can be treated as a subdirect power of $\mathcal{H}/_{\kappa_a l}$. Then the proposition follows straightforwardly from Lemmas 2.6 and 2.7 of [12], and also from Proposition 8(3),(4). □

If a structure $\mathcal{H}$ with a Mal'tsev polymorphism omits affine type, then we can obtain even stronger rectangularity-type condition. Recall that in this case the congruence lattice of $\mathcal{H}$ is distributive. A Mal'tsev algebra ($\mathsf{Alg}(\mathcal{H})$ in our case) generating a variety, in which every algebra has a distributive congruence lattice is called *arithmetical*. Arithmetical algebras are exhaustively studied. We will use the following result [52] describing the structure of relations invariant with respect to such algebras.

Proposition 11 *Let $\mathbb{D}$ be a subdirect product of $\mathbb{A}_1, \ldots, \mathbb{A}_k$. Then $\mathbb{D}$ can be uniquely determined by the thick mappings ψ_{ij} for $i, j \in [k]$. More precisely, $\mathbf{a} \in \mathbb{D}$ if and only if $(\mathbf{a}[i], \mathbf{a}[j]) \in \psi_{ij}$ for all $i, j \in [k]$.*

Corollary 2 *If $\mathcal{H}$ is a structure with a Mal'tsev polymorphism omitting affine type then $\#\mathrm{CSP}(\mathcal{H})$ is polynomial time equivalent to $\#\mathrm{CSP}(\mathcal{H}')$, where $\mathcal{H}'$ is a relational structure with the same universe as $\mathcal{H}$, all relational symbols of which are binary and interpreted as thick mappings of $\mathcal{H}$.*

[1] In [12], we used the algebraic terminology: R is a subdirect product of subdirectly irreducible Mal'tsev algebras.

4 Necessary condition for tractability

In this section we prove two more necessary conditions for #-tractability. Both of them follow from Proposition 5, but they allow us to design an algorithm for #CSP.

If the algebra corresponding to the structure $\mathcal{H}$ does not omit the affine type, then we have a stronger necessary condition for the tractability of $\#\mathrm{CSP}(\mathcal{H})$.

Proposition 12 *If $\mathcal{H}$ is congruence singular then for any congruences $\delta \leq \alpha \prec \beta \in \mathsf{Con}(\mathcal{H})$ such that $\alpha \prec \beta$ has affine type, any n-ary relation $R \in \mathsf{def}(\mathcal{H})$ and any sequences $A_1, \ldots, A_n$ and $B_1, \ldots, B_n$ of α-classes such that A_i, B_i belong to the same β-class ($i \in [n]$), if $R_1 = R \cap (A_1 \times \ldots \times A_n) \neq \varnothing$, $R_2 = R \cap (B_1 \times \ldots \times B_n) \neq \varnothing$, then $|R_1/_{\delta^n}| = |R_2/_{\delta^n}|$.*

We make use of some basics of commutator theory in congruence modular varieties (see the seminal book [31]). Let $\mathbb{A}$ be a Mal'tsev algebra and $\alpha, \beta, \gamma \in \mathsf{Con}(\mathbb{A})$. The congruence α *centralizes* β *modulo* γ, denoted $C(\alpha, \beta; \gamma)$, if, for any (n-ary) term operation f, any $\langle u, v\rangle \in \alpha$ and any $\langle a_1, b_1\rangle, \ldots, \langle a_{n-1}, b_{n-1}\rangle \in \beta$,

$$\begin{aligned} &\langle f(u, a_1, \ldots, a_{n-1}), f(u, b_1, \ldots, b_{n-1})\rangle \in \gamma \\ \Longleftrightarrow \quad &\langle f(v, a_1, \ldots, a_{n-1}), f(v, b_1, \ldots, b_{n-1})\rangle \in \gamma. \end{aligned}$$

The smallest congruence γ such that $C(\alpha, \beta; \gamma)$ is called the *commutator* of α, β, denoted $[\alpha, \beta]$.

Proposition 13 ([31]) *Let $\mathbb{A}$ be a Mal'tsev algebra and $\alpha, \beta, \gamma \in \mathsf{Con}(\mathbb{A})$. Then*

(1) $[\alpha, \beta] = [\beta, \alpha]$;

(2) if $\alpha \prec \beta$ and this interval has affine type if and only if $[\beta, \beta] \leq \alpha$;

(3) if $\alpha \leq \beta$ and $[\beta, \beta] \leq \alpha$, there is a congruence θ of β (which is considered as a subalgebra of $\mathbb{A}^2$) such that the set $\{\langle (a, b), (c, d)\rangle \mid \langle a, b\rangle, \langle c, d\rangle \in \alpha\}$ is a class of θ.

Proof: (of Proposition 12.) By switching to the quatient structure $\mathcal{H}/_{\delta}$ we may assume that δ is the equality relation. To prove Proposition 12 we consider the universal algebra $\mathbb{A} = (H; \mathsf{Pol}(\mathcal{H}))$ and the subalgebra $\mathbb{B}$ of $\mathbb{A}^n$ with the universe R. Thus we consider R as a subalgebra of $\mathbb{A}^n$.

CLAIM 1. For the algebra $\mathbb{B}$, $[\beta^n, \beta^n] \leq \alpha^n$.

Let f be a (k-ary) term operation of $\mathbb{A}$, and let $\langle \mathbf{u}, \mathbf{v}\rangle \in \beta^n$ and $\langle \mathbf{a}_1, \mathbf{b}_1\rangle, \ldots, \langle \mathbf{a}_{k-1}, \mathbf{b}_{k-1}\rangle \in \beta^n$. If $\langle f(\mathbf{u}, \mathbf{a}_1, \ldots, \mathbf{a}_{k-1}), f(\mathbf{u}, \mathbf{b}_1, \ldots, \mathbf{b}_{k-1})\rangle \in \alpha^n$ then $\langle f(\mathbf{u}[i], \mathbf{a}_1[i], \ldots, \mathbf{a}_{k-1}[i]), f(\mathbf{u}[i], \mathbf{b}_1[i], \ldots, \mathbf{b}_{k-1}[i])\rangle \in \alpha$ for each $i \in [n]$. Since $C(\beta, \beta; \alpha)$, this implies $\langle f(\mathbf{v}[i], \mathbf{a}_1[i], \ldots, \mathbf{a}_{k-1}[i]), f(\mathbf{v}[i], \mathbf{b}_1[i], \ldots, \mathbf{b}_{k-1}[i])\rangle \in \alpha$ for each $i \in [n]$. Thus $\langle f(\mathbf{v}, \mathbf{a}_1, \ldots, \mathbf{a}_{k-1}), f(\mathbf{v}, \mathbf{b}_1, \ldots, \mathbf{b}_{k-1})\rangle \in \alpha^n$.

We treat the congruence β^n as a subalgebra of $\mathbb{B}^2$; let us denote it by $\mathbb{C}$. Let $A_1, \ldots, A_k$ be the α^n-classes of $\mathbb{B}$ and $|A_i| = \ell_i$. By Proposition 13 there is a congruence γ of $\mathbb{C}$ such that the set D of pairs of the form $(\mathbf{a}, \mathbf{b})$, $\mathbf{a}, \mathbf{b} \in \mathbb{B}$ and $\langle \mathbf{a}, \mathbf{b}\rangle \in \alpha^n$. Let $\gamma' = \gamma \vee \alpha^{2n}$. The set D is a class of γ'.

CLAIM 2. Every class E of γ' is the union $(A_1 \times A_{\varphi(1)}) \cup \ldots \cup (A_k \times A_{\varphi(k)})$ for a certain bijective mapping $\varphi : [k] \to [k]$; and for one of the classes φ is the identity mapping.

Since $\alpha^{2n} \subseteq \gamma'$, if $A_i \times A_j \cap E \neq \varnothing$ then $A_i \times A_j \subseteq E$. Suppose that there are $(\mathbf{a}, \mathbf{b}), (\mathbf{c}, \mathbf{d}) \in E$ such that $\langle \mathbf{a}, \mathbf{c} \rangle \in \alpha^n$, but $\langle \mathbf{b}, \mathbf{d} \rangle \notin \alpha^n$. As $\alpha^{2n} \subseteq \gamma'$, we may assume $\mathbf{a} = \mathbf{c}$. Let us consider γ' as a 4-ary relation on $\mathbb{B}$. Let also f be a Mal'tsev operation of $\mathbb{A}$. Then we have

$$f\begin{pmatrix} \mathbf{a} & \mathbf{a} & \mathbf{b} \\ \mathbf{a} & \mathbf{b} & \mathbf{b} \\ \mathbf{a} & \mathbf{a} & \mathbf{d} \\ \mathbf{a} & \mathbf{d} & \mathbf{d} \end{pmatrix} = \begin{pmatrix} \mathbf{b} \\ \mathbf{a} \\ \mathbf{d} \\ \mathbf{a} \end{pmatrix} \in \gamma' \quad \text{and} \quad f\begin{pmatrix} \mathbf{a} & \mathbf{a} & \mathbf{b} \\ \mathbf{b} & \mathbf{a} & \mathbf{a} \\ \mathbf{a} & \mathbf{a} & \mathbf{d} \\ \mathbf{b} & \mathbf{a} & \mathbf{a} \end{pmatrix} = \begin{pmatrix} \mathbf{b} \\ \mathbf{b} \\ \mathbf{d} \\ \mathbf{b} \end{pmatrix} \in \gamma',$$

which implies that $\langle \mathbf{b}, \mathbf{d} \rangle \in \alpha^n$, a contradiction.

Let also β' denote the congruence $\alpha^n \times \beta^n$. It is not hard to see that $\gamma' \vee \beta' = \beta^n \times \beta^n$ and $\gamma' \wedge \beta' = \alpha^n \times \alpha^n$.

Clearly, every class of $\alpha^n \times \alpha^n$ is the direct product of two classes R_1, R_2 of α^n. Therefore, its size is $|R_1| \cdot |R_2|$. Thus, the first two rows of the matrix $M(\gamma', \beta')$ look as follows

$$\begin{pmatrix} \ell_1^2 & \ell_2^2 & \cdots & \ell_k^2 \\ \ell_1 \ell_{\varphi(1)} & \ell_2 \ell_{\varphi(2)} & \cdots & \ell_k \ell_{\varphi(k)} \end{pmatrix}.$$

If $\#\mathrm{CSP}(R) \subseteq \#\mathrm{CSP}(\mathcal{H})$ is not #P-complete, then these two rows are proportional, that is

$$\frac{\ell_1}{\ell_{\varphi(1)}} = \frac{\ell_2}{\ell_{\varphi(2)}} = \ldots = \frac{\ell_k}{\ell_{\varphi(k)}}.$$

For any $i \in \{1, \ldots, k\}$, let m be such that $\varphi^m(i) = i$. Since

$$\frac{\ell_i}{\ell_{\varphi(i)}} = \frac{\ell_\varphi(i)}{\ell_{\varphi^2(i)}} = \ldots = \frac{\ell_\varphi^{m-1}(i)}{\ell_{\varphi^m(i)}},$$

we have $\ell_i^2 = \ell_{\varphi(i)} \ell_{\varphi^{-1}(i)}$. As this holds for every i and $\varphi(i) = i$ for no i, we conclude that $\ell_i = \ell_j$ for any pair i, j from the same orbit of φ. Finally, for each pair $i, j \in \{1, \ldots, k\}$, there is a row in $M(\gamma', \beta')$ of the form $\begin{pmatrix} \ell_1 \ell_{\psi(1)} & \ell_2 \ell_{\psi(2)} & \cdots & \ell_k \ell_{\psi(k)} \end{pmatrix}$ such that $\psi(i) = j$. By what was proved above, $\ell_i = \ell_j$. □

We will also need another corollary from Proposition 5. Let T be a k-dimensional array, that is a collection of numbers $T[i_1, \ldots, i_k]$ indexed by tuples $(i_1, \ldots, i_k)$, where $1 \leq i_k \leq m_k$. The array T has rank 1, denoted $\mathsf{rank}(T) = 1$, if for each $\ell \in [k]$, and any $i_1, \ldots, i_{\ell-1}, i_{\ell+1}, \ldots, i_k, j_1, \ldots, j_{\ell-1}, j_{\ell+1}, \ldots, j_k$ with $i_u, j_u \in [m_u]$, we have

$$\frac{T[i_1, \ldots, i_{\ell-1}, 1, i_{\ell+1}, \ldots, i_k]}{T[j_1, \ldots, j_{\ell-1}, 1, j_{\ell+1}, \ldots, j_k]} = \ldots = \frac{T[i_1, \ldots, i_{\ell-1}, m_\ell, i_{\ell+1}, \ldots, i_k]}{T[j_1, \ldots, j_{\ell-1}, m_\ell, j_{\ell+1}, \ldots, j_k]}.$$

It is not hard to see that this condition can equivalently be expressed as follows: for each $\ell \in [k]$ there are numbers $t_1^\ell, \ldots, t_{m_k}^\ell$ such that

$$T[i_1, \ldots, i_k] = t_{i_1}^1 \cdot \ldots \cdot t_{i_k}^\ell.$$

Now let R be a relation pp-definable in a structure $\mathcal{H}$ with a Mal'tsev polymorphism, and let $\gamma_1, \ldots, \gamma_k$ be congruences on R such that for each $i \in [k]$

$$\gamma_i \vee (\gamma_1 \wedge \ldots \wedge \gamma_{i-1} \wedge \alpha_{i+1} \wedge \ldots \wedge \gamma_k) = \gamma_1 \vee \ldots \vee \gamma_k \qquad (1)$$

Let also C be a class of $\gamma = \gamma_1 \vee \ldots \vee \gamma_k$, and let $A^i_1, \ldots, A^i_{m_i}$ be the classes of γ_i from C. The condition (1) means that for any $i_1, \ldots, i_k$ the set $A^1_{i_1} \cap \ldots \cap A^k_{i_k}$ is a nonempty class of $\beta = \gamma_1 \wedge \ldots \wedge \gamma_k$, and any two classes of this form are different. We consider a k-dimensional array $M(C; \gamma_1, \ldots, \gamma_k)$, where

$$M(C; \gamma_1, \ldots, \gamma_k)[i_1, \ldots, i_k] = |A^1_{i_1} \cap \ldots \cap A^k_{i_k}|.$$

Proposition 14 *Let $\gamma_1, \ldots, \gamma_k$ be congruences of a structure $\mathcal{H}$ that has a Mal'tsev polymorphism, let them satisfy the condition (1), and let C be a class of $\gamma_1 \vee \ldots \vee \gamma_k$. Then,* $\mathsf{rank}(M(C; \alpha_1, \ldots, \alpha_k) = 1$ *or* $\#\mathrm{CSP}(\mathcal{H})$ *is #P-complete.*

Proof: We consider the congruences γ_i and $\beta_i = \gamma_1 \wedge \ldots \wedge \gamma_{i-1} \wedge \gamma_{i+1} \wedge \ldots \wedge \gamma_k$. To simplify the notation we assume $i = k$. If $\#\mathrm{CSP}(\mathcal{H})$ is not #P-complete then $\mathsf{rank}(M(C; \gamma_k, \beta_k)) = 1$. Let $A^i_1, \ldots, A^i_{m_i}$ be the classes of γ_i from C. The classes of β_k have the form $A^1_{i_1} \cap \ldots \cap A^{k-1}_{i_{k-1}}$, the classes of $\gamma_k \wedge \beta_k$ are the classes of $\gamma_1 \wedge \ldots \wedge \gamma_k$. Therefore every row of $M(C; \gamma_k, \beta_k)$ is equal to

$$(M(C; \gamma_1, \ldots, \gamma_k)[i_1, \ldots, i_{k-1}, 1], \ldots, M(C; \gamma_1, \ldots, \gamma_k)[i_1, \ldots, i_{k-1}, m_k])$$

for some $i_1, \ldots, i_{k-1}$. Since $\mathsf{rank}(M(C; \gamma_k, \beta_k)) = 1$, we get

$$\frac{M(C; \gamma_1, \ldots, \gamma_k)[i_1, \ldots, i_{k-1}, 1]}{M(C; \gamma_1, \ldots, \gamma_k)[j_1, \ldots, j_{k-1}, 1]} = \ldots = \frac{M(C; \gamma_1, \ldots, \gamma_k)[i_1, \ldots, i_{k-1}, m_k]}{M(C; \gamma_1, \ldots, \gamma_k)[j_1, \ldots, j_{k-1}, m_k]}.$$

The corollary is proved. □

An important example of a collection of congruences satisfying the condition (1) is the following. Let $\alpha \in M$, and let $I_1, \ldots, I_k$ be the classes of κ^*_α. A congruence γ_j is defined as follows: $\langle \mathbf{a}, \mathbf{b} \rangle \in \gamma_j$ if and only if $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \alpha^-$ for $i \in I_j$ and $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \alpha^+$ otherwise.

5 Algorithms: prerequisites

5.1 Decision CSPs over a Mal'tsev algebra.

If a relational structure $\mathcal{H}$ has a Mal'tsev polymorphism, then the decision CSP with the template $\mathcal{H}$ can be solved in polynomial time [7, 13]. Here we shall use the algorithm presented in [13]. This algorithm builts a sort of a succinct (polynomial size) representation for the set of all solutions.

Let n be a positive integer, let H be a finite set, let $\mathbf{a}$, $\mathbf{b}$ be n-ary tuples and let (i, a, b) be any element in $[n] \times H^2$. We say that $(\mathbf{a}, \mathbf{b})$ *witnesses* (i, a, b) if $\mathrm{pr}_{[i-1]}\mathbf{a} = \mathrm{pr}_{[i-1]}\mathbf{b}$, $\mathbf{a}[i] = a$, and $\mathbf{b}[i] = b$. We also say that $\mathbf{a}$ and $\mathbf{b}$ witness (i, a, b) meaning that $(\mathbf{a}, \mathbf{b})$ witnesses (i, a, b).

Let R be any n-ary relation on H. The *signature* of R, $\mathsf{Sig}_R \subseteq [n] \times H^2$, is defined to be the set containing all those $(i, a, b) \in [n] \times H^2$ witnessed by tuples in R, that is

$$\mathsf{Sig}_R = \{(i, a, b) \in [n] \times H^2 : \exists \mathbf{a}, \mathbf{b} \in R \text{ such that } (\mathbf{a}, \mathbf{b}) \text{ witnesses } (i, a, b)\}.$$

Note that in our notation $(i, a, b) \in \mathsf{Sig}_R$ if and only if $\langle a, b \rangle$ belongs to the relation θ_i computed for the relation $\mathrm{pr}_{[i]}R$ (see Section 3.3.1). In particular, for any $(i, a, b) \in \mathsf{Sig}_R$ and any $\mathbf{a} \in \mathrm{pr}_{[i]}R$ with $\mathbf{a}[i] = a$ the tuple $\mathbf{b}$ such that $\mathrm{pr}_{[i-1]}\mathbf{b} = \mathrm{pr}_{[i-1]}\mathbf{a}$ and $\mathbf{b}[i] = 1$ also belongs to $\mathrm{pr}_{[i]}R$.

A subset R' of R is called a *representation* of R if $\mathsf{Sig}_{R'} = \mathsf{Sig}_R$. If furthermore, $|R'| \leq 2|\mathsf{Sig}_R|$ then R is called a *compact* representation of R. Observe that every relation R has compact representations.

Let $\mathcal{H}$ be a relational structure and $R' \subseteq H^n$ for some n. By $\langle R' \rangle_{\mathcal{H}}$ we denote the relation *generated* by R', that is, the smallest relation R definable in $\mathcal{H}$ and such that $R' \subseteq R$. Since $\mathcal{H}$ is usually clear from the context we shall omit this subscript. The key lemma proved in [13] states that if R is a relation definable in a relational structure with a Mal'tsev polymorphism, and R' is a representation of R, then $\langle R' \rangle = R$. Given an instance $\mathcal{G}$ of the constraint satisfaction problem $\mathrm{CSP}(\mathcal{H})$, $m = |\mathcal{G}|$, the set of all solutions $\Phi(\mathcal{G}, \mathcal{H})$ to this problem can be thought of as an m-ary definable relation in $\mathcal{H}$. The algorithm presented in [13] finds a compact representation of this set.

We will need to know unary and binary projections of the relation $\Phi(\mathcal{G}, \mathcal{H})$, that is, sets of the form $\psi_g = \{\varphi(g) \mid \varphi \in \Phi(\mathcal{G}, \mathcal{H})\}$ for $g \in \mathcal{G}$ and $\psi_{g,h} = \{(\varphi(g), \varphi(h)) \mid \varphi \in \Phi(\mathcal{G}, \mathcal{H})\}$ for $g, h \in \mathcal{G}$. It is not hard to see (see also [13]) that if R' is a compact representation of $\Phi(\mathcal{G}, \mathcal{H})$, then $\psi_g, \psi_{g,h}$ are equal to $\langle \mathrm{pr}_g R' \rangle$ and $\langle \mathrm{pr}_{g,h} R' \rangle$. Therefore, we may assume that we have a precomputed table that for each subset of $\mathcal{H}$, and for each subset of $\mathcal{H} \times \mathcal{H}$ shows the unary or binary relation generated by this subset; and every time we need to find ψ_g or $\psi_{g,h}$ using a compact representation R', we just find the corresponding projection of R' and look up the table.

If there is no complexity restriction imposed, as in the case of precomputation, the relation generated by some set $Q \subseteq \mathcal{H}^n$ can be computed by employing a standard method. Let $Q = \{\mathbf{a}_1, \ldots, \mathbf{a}_m\}$. First, find all m-ary polymorphisms of $\mathcal{H}$. This can be done using the *indicator problem* [42]. Next, include into $\langle Q \rangle$ all tuples that can be represented as $f(\mathbf{a}_1, \ldots, \mathbf{a}_m)$ for an m-ary polymorphism f.

5.2 Reduction to subdirect powers.

In general, for an instance $\mathcal{G}$ of $\#\mathrm{CSP}(\mathcal{H})$ the sets ψ_g, $g \in \mathcal{G}$, are subalgebras of $\mathcal{H}$ that are not necessarily equal to $\mathcal{H}$. For us, however, it is much more convenient to deal with the case when $\Phi(\mathcal{G}, \mathcal{H})$ is a subdirect power of $\mathcal{H}$, that is $\psi_g = \mathcal{H}$ for all $g \in \mathcal{G}$. We show how to transform the problem so that ψ_g be $\mathcal{H}$ for all $g \in \mathcal{G}$. To do this we borrow some methods from the multi-sorted CSP, see, e.g. [10].

Let $D_1, \ldots, D_n$ be the subalgebras of $\mathcal{H}$ (including H itself). We shall assume that along with every (n-ary) relational symbol R and any $D_{i_1}, \ldots, D_{i_n}$ the vocabulary of $\mathcal{H}$ contains a symbol R' such that $R'^{\mathcal{H}} = R \cap (D_{i_1} \times \ldots \times D_{i_n})$. Then we define a relational structure $\chi(\mathcal{H})$ as follows. The universe of $\chi(\mathcal{H})$ is $D = D_1 \times \ldots \times D_n$; the ith component of an element $\overline{a} \in D$ is denoted by $\overline{a}[i]$. For any (n-ary) relation R pp-definable in $\mathcal{H}$ we set $(\overline{a}_1, \ldots, \overline{a}_n) \in \chi(R)$ if and only if $(\overline{a}_1[i_1], \ldots, \overline{a}_n[i_n]) \in R$, where $D_{i_j} = \mathrm{pr}_j R$. In particular, each unary relation of $\chi(\mathcal{H})$ contains all elements of D and, therefore, can be thrown out. For any coordinate position i of any non-unary relation R, the set $\mathrm{pr}_i \chi(R)$ equals D. Finally, to define $\chi(\mathcal{H})$ formally we for each relational symbol R we interpret it as $R^{\chi(\mathcal{H})} = \chi(R)$.

For an instance $\mathcal{G}$ of $\#\mathrm{CSP}(\mathcal{H})$, the following algorithm constructs an instance $\mathcal{G}'$ of $\#\mathrm{CSP}(\chi(\mathcal{H}))$.

```
Algorithm Subdirect
INPUT: an instance $\mathcal{G}$ of $\#\mathrm{CSP}(\mathcal{H})$
OUTPUT: an instance $\mathcal{G}'$ of $\#\mathrm{CSP}(\chi(\mathcal{H}))$ with the same universe as $\mathcal{G}$

Step 1      find a compact representation of $\Phi(\mathcal{G}, \mathcal{H})$
Step 2      for each $g \in \mathcal{G}$ find $\psi_g$
Step 3      for each ($n$-ary) relational symbol $R$ do
Step 3.1        for each tuple $(g_1, \ldots, g_n) \in R^{\mathcal{G}}$ do
Step 3.1.1          let $R'$ be the relational symbol such that
                    $R'^{\mathcal{H}} = R^{\mathcal{H}} \cap (\psi_{g_1} \times \ldots \times \psi_{g_n})$
```

```
Step 3.1.2        include $(g_1,\dots,g_n)$ into $R'^{\mathcal{G}'}$
              endfor
          endfor
Step 4    output $\mathcal{G}'$
```

The next easy lemma completes the reduction.

Lemma 5 *Let $\mathcal{G}$ is an instance of $\#\mathrm{CSP}(\mathcal{H})$ and $\mathcal{G}'$ an instance of $\#\mathrm{CSP}(\chi(\mathcal{H}))$ consructed by algorithm* `Subdirect`. *Let also $\psi_g = \mathrm{pr}_g\Phi(\mathcal{G},\mathcal{H})$ for $g \in \mathcal{G}$. Then $\Phi(\mathcal{G}',\chi(\mathcal{H}))$ is a subdirect power of $\chi(\mathcal{H})$ and*

$$|\Phi(\mathcal{G}',\chi(\mathcal{H}))| = |\Phi(\mathcal{G},\mathcal{H})| \cdot \prod_{g\in G} \frac{|D|}{|\psi_g|}.$$

Proof: Let $\varphi \in \Phi(\mathcal{G}',\chi(\mathcal{H}))$ be a homomorphism from $\mathcal{G}'$ to $\chi(\mathcal{H})$. Let us define a mapping $\chi^{-1}(\varphi)$ from $\mathcal{G}$ to $\mathcal{H}$ as follows. (Note that $\mathcal{G}$ and $\mathcal{G}'$ have a common universe.) For $g \in \mathcal{G}$ if $\varphi(g) = \overline{a}$ and $\psi_g = D_i$ then set $\chi^{-1}(\varphi)(g) = \overline{a}[i]$. By the construction of $\chi(\mathcal{H})$ and $\mathcal{G}'$, if we change the value $\varphi(g)$ for some $g \in \mathcal{G}$ with $\psi_g = D_i$ to any $\overline{b}$ such that $\overline{b}[i] = \overline{a}[i]$, then the resulting mapping φ' is still a homomorphism from $\mathcal{G}'$ to $\chi(\mathcal{H})$ and $\chi^{-1}(\varphi') = \chi^{-1}(\varphi)$. Conversely, for any homomorphism $\psi \in \Phi(\mathcal{G},\mathcal{H})$, any mapping $\varphi\colon \mathcal{G}' \to \chi(\mathcal{H})$ such that $\chi^{-1}(\varphi) = \psi$ is a homomorphism of $\mathcal{G}'$ to $\chi(\mathcal{H})$. This straightforwardly implies the result. □

5.3 Structure of Mal'tsev instances

Let $\mathcal{G}$ be a $\#\mathrm{CSP}(\mathcal{H})$ instance and $|\mathcal{G}| = m$. We shall asuume that the universe G of $\mathcal{G}$ equals to $[m]$. Clearly, the set $\Phi(\mathcal{G},\mathcal{H})$ can be thought of as an m-ary relation definable in $\mathcal{H}$, or as a subalgebra of the mth direct power of $\mathbb{A} = \mathsf{Alg}(\mathcal{H})$. By the results of the previous subsection we may assume that $R = \Phi(\mathcal{G},\mathcal{H})$ is a subdirect power of $\mathcal{H}$. Recall that for a congruence $\theta \in \mathsf{Con}(\mathcal{H})$ by θ^m we denote the congruence of R such that $\langle \mathbf{a},\mathbf{b}\rangle \in \theta^m$ if and only if $\langle \mathbf{a}[g],\mathbf{b}[g]\rangle \in \theta$ for all $g \in \mathcal{G}$. For congruences $\beta \le \gamma \in \mathsf{Con}(\mathcal{H})$ and a mapping $\pi\colon \mathcal{G} \to \mathcal{H}/_{\beta}$, by $\pi/_{\gamma}$ we denote a mapping from $\mathcal{G}$ to $\mathcal{H}\gamma$ defined by $\pi/_{\gamma}(g) = \pi(g)/_{\gamma}$. If $\psi_{g,h}$ is a thick mapping of level β then we treat $\psi_{g,h}/_{\beta^2}$ as a mapping that maps classes of β to classes of β. We need some structural properties of R.

Let $\alpha \in M$. Let also $A_1,\dots,A_k$ be the κ^*_α-classes and $g_1,\dots,g_k$ representatives of these classes. Let π be an element of $R/_{(\alpha^+)^m}$; such an element can be thought of as a homomorphism from $\mathcal{G}$ to $\mathcal{H}/_{\alpha^+}$, but not all such homomorphisms are elements of $R/_{(\alpha^+)^m}$. By $C^u_1,\dots,C^u_{s_u}$ we denote the κ_α-classes from $\pi(g_u)/_{\lambda_\alpha}$ for $u \in [\ell]$.

Lemma 6 *Every prime quotient in the interval $[\kappa^m_\alpha,\lambda^m_\alpha]$ in the congruence lattice $\mathsf{Con}(R)$ has Boolean type, the interval $[\kappa^m_\alpha,\lambda^m_\alpha]$ is a distributive lattice isomorphic to the lattice $2^{[k]}$ of subsets of a k-element set, and every congruence in this interval can be represented as η_J, $J \subseteq [k]$, defined as follows: $\langle \mathbf{a},\mathbf{b}\rangle \in \eta_J$ if and only if $\langle \mathbf{a}[i],\mathbf{b}[i]\rangle \in \kappa_\alpha$ whenever $i \notin J$ and $\langle \mathbf{a}[i],\mathbf{b}[i]\rangle \in \lambda_\alpha$ when $i \in J$.*

Proof: We consider the chain of congruences $\kappa^m_\alpha < \eta_1 < \eta_{\{1,2\}} < \dots < \eta_{[k]} = \lambda_\alpha$. First, we show that this chain is maximal. Note that as the quotient $\eta_{\{1,\dots,v\}} < \eta_{\{1,\dots,v,v+1\}}$ is projective to $\kappa_\alpha < \eta_{v+1}$. Therefore it suffices to show that the quotients of the form $\kappa^m_\alpha < \eta_v$ are prime. To simplify the notation we assume $v = 1$. By replacing R with $R/_{\kappa^m_\alpha}$ we also assume that $\kappa_\alpha = \Delta$.

Let $\kappa_\alpha^m < \theta \leq \eta_1$. Then, (a) for any $\langle \mathbf{a}, \mathbf{b} \rangle \in \theta$, $\mathbf{a}[g] = \mathbf{b}[g]$ for all $g \notin A_1$, and, (b) by Proposition 10, $\mathbf{a}[g_1] = \mathbf{b}[g_1]$ if and only if $\mathbf{a} = \mathbf{b}$. Since $\kappa_\alpha^m < \theta \leq \eta_1$, there are $\langle \mathbf{a}, \mathbf{b} \rangle \in \theta$ such that $\mathbf{a} \neq \mathbf{b}$. This means that $\mathbf{a}[g_1] \neq \mathbf{b}[g_1]$. As $\kappa_\alpha \prec \lambda_\alpha$, for any $\langle a, b \rangle \in \lambda_\alpha$ there are $\langle \mathbf{a}', \mathbf{b}' \rangle \in \theta$ with $\mathbf{a}'[g_1] = a, \mathbf{b}'[g_1] = b$. By (b) this implies $\theta = \eta_1$. It is also easy to check that θ does not centralize itself modulo κ_α^m that implies that the quotient $\kappa_\alpha^m \prec \eta_1$ has Boolean type.

We have proved that the chain $\kappa_\alpha^m < \eta_1 < \eta_{\{1,2\}} < \ldots < \eta_{[k]} = \lambda_\alpha$ is maximal, and, by Lemma 2, each of its prime quotients has Boolean type. Now, in a modular lattice every prime quotient is perspective to one of quotients of any maximal chain. Thus we conclude that every prime quotient from the interval has Boolean type.

Finally, by Lemma 6.6 of [37], this implies that this interval is a distributive lattice. Since the congruences $\eta_1, \ldots, \eta_\ell$ are *atoms* of this lattice, and $\eta_1 \vee \ldots \vee \eta_\ell = \lambda_\alpha^m$, every element θ of this interval can be represented in the form

$$\theta = \bigvee_{u \in J} \eta_u = \eta_J$$

for some $J \subseteq [k]$. □

Lemma 7 *For any choice of $i_u \in [s_u]$, $u \in [k]$, there is an element $\varphi \in R/_{\kappa_\alpha^m}$ such that for each $u \in [k]$, and each $g \in A_u$*

$$\varphi(g) = \psi_{g_u,g}/_{\kappa_\alpha^2}(C_{i_u}^u).$$

Proof: If we choose $B_g = \pi(g)/_{\lambda_\alpha}$ then π witnesses that $R \cap (B_1 \times \ldots B_m) \neq \varnothing$. As the coherent sets of R are equal to $A_1, \ldots, A_k$, by Proposition 10, we have that

$$R/_{\kappa_\alpha^k} \cap (B_1/_{\kappa_\alpha} \times \ldots \times B_k/_{\kappa_\alpha}) = B_{A_1}/_{\kappa_\alpha^{|A_1|}} \times \ldots \times B_{A_k}/_{\kappa_\alpha^{|A_k|}},$$

where $B_{A_u} = \mathrm{pr}_{A_u} R \cap \prod_{g \in A_u} B_g$, and for any $g, h \in A_u$ we have $\mathrm{pr}_{g,h} B_{A_u} = \psi_{g,h} \cap (B_g \times B_h)$. The result follows. □

Lemma 8 *There is $J \subseteq [k]$ such that for any π, an element from $R/_{(\alpha^+)^m}$, there are i_u, $u \in [k] - J$, with $i_u \in [s_u]$ satisfying the following conditions. Every homomorphism $\varrho \in R/_{(\alpha^-)^m}$ with $\varrho/_{(\alpha^+)^m} = \pi$ can be represented as follows: there are i_u for $u \in J$ with $i_u \in [s_u]$ such that $\varrho(g_u) \in C_{i_u}^u$ for $u \in [k]$ and, for any $g \in A_u$, $u \in [k]$, we have*

$$\varrho(g) = \pi(g) \cap \psi_{g_u,g}/_{\kappa_\alpha^2}(C_{i_u}^u).$$

Conversely, for any choice of $C_{i_u}^1, \ldots, C_{i_k}^k$ the mapping ϱ defined in this way is an element of $R/_{(\alpha^-)^m}$, and $\varrho/_{\alpha^+} = \pi$.

Proof: Observe that in the congruence lattice $\mathsf{Con}(R)$ we have $\kappa_a l^m \wedge (\alpha^+)^m = (\alpha^-)^m$ and $\kappa_a l^m \leq \kappa_a l^m \vee (\alpha^+)^m \leq \lambda_\alpha^m$. By Lemma 6, $\kappa_a l^m \vee (\alpha^+)^m = \eta_J$ for some $J \subseteq [k]$. This means that there are i_u, $u \in [k] - J$, with $i_u \in [s_u]$, such that for any $\varrho \in R/_{(\alpha^-)^m}$, with $\varrho/_{\alpha^+} = \pi$, we have $\varrho(g_u) \in C_{i_u}^u$ for $u \in [k] - J$.

Take $\varrho \in R/_{(\alpha^-)^m}$ with $\varrho/_{\alpha^+} = \pi$. Clearly, $\varrho/_{\kappa_\alpha}$ belongs to $R/_{\kappa_a l^m}$, and by what we showed above $\varrho(g_u) \in C_{i_u}^u$ for $u \in [k] - J$. Then the first part of the lemma follows from Lemma 7.

To prove the converse statement, let us denote the η_J-class containing π by D. Since κ_α^m and $(\alpha^+)^m$ permute, for any κ_α^m-class $C \subseteq D$ and any $(\alpha^+)^m$-class C', the intersection $C \cap C'$ is nonempty. Therefore, for any $\varphi \in R/_{\kappa_\alpha^m}$ such that $\varphi(g_u) = C^u_{i_u}$ for $u \in [k] - J$, there is $\varrho \in R/_{(\alpha^-)^m}$ such that $\varrho/_{\kappa_\alpha} = \varphi$ and $\varrho/_{\alpha^+} = \pi$; that is $\varrho(g) = \varphi(g) \cap \pi(g)$. Together with Lemma 6 this implies the result. □

Let $J \subseteq [k]$ be the set defined in Lemma 8 for $\alpha \in M$ and the κ_α^*-classes $A_1, \ldots, A_k$. A congruence γ_u, $u \in J$ is defined as follows: $\langle \mathbf{a}, \mathbf{b} \rangle \in \gamma_u$ if and only if $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \alpha^-$ for $i \in A_u \cup \bigcup_{v \in [k]-J} A_v$, and $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \alpha^+$ otherwise.

Lemma 9 *The congruences γ_u, $u \in J$, satisfy the condition (1).*

Proof: Without loss of generality we assume $J = \{1, \ldots, q\}$. First, observe that $\gamma_1 \wedge \ldots \wedge \gamma_q = (\alpha^-)^m$ and $\gamma_1 \vee \ldots \vee \gamma_q = (\alpha^+)^m$. Since $(\alpha^+)^m \vee (\kappa_\alpha)^m = \eta_J$, $(\alpha^+)^m \wedge (\kappa_\alpha)^m = (\alpha^-)^m$, and the lattice $\mathsf{Con}(R)$ is modular, the intervals $[(\alpha^-)^m, (\alpha^+)^m]$ and $[(\kappa_\alpha)^m, \eta_J]$ are isomorpic, where an isomorphism can be defined by $\varphi(x) = x \vee (\kappa_\alpha)^m$. Therefore we may consider $\beta_1, \ldots, \beta_q$ instead of $\gamma_1, \ldots, \gamma_q$, where $\beta_u = \gamma_u \vee (\kappa_\alpha)^m$ and $\langle \mathbf{a}, \mathbf{b} \rangle \in \beta_u$ if and only if $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \kappa_\alpha$ for $i \in A_u \cup \bigcup_{v \in [k]-J} A_v$ and $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \lambda_\alpha$ otherwise. We also may assume that $\kappa_\alpha = \Delta$. To simplify the notation we prove the condition (1) for $i = 1$.

By Lemma 8, $\langle \mathbf{a}, \mathbf{b} \rangle \in$
$beta_1$ if and only if $\mathrm{pr}_{A_1 \cup A_{q+1} \cup \ldots \cup A_k} \mathbf{a} = \mathrm{pr}_{A_1 \cup A_{q+1} \cup \ldots \cup A_k} \mathbf{b}$, $\mathrm{pr}_{A_2 \cup \ldots A_q} \mathbf{a}, \mathrm{pr}_{A_2 \cup \ldots A_q} \mathbf{b} \in \mathrm{pr}_{A_2 \cup \ldots A_q} \mathbb{D}$, and $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \lambda_\alpha$ for $i \in A_2 \cup \ldots A_q$. Similarly, $\langle \mathbf{a}, \mathbf{b} \rangle \in \beta_2 \wedge \ldots \wedge \beta_q$ if and only if $\mathrm{pr}_{A_1} \mathbf{a}, \mathrm{pr}_{A_1} \mathbf{b} \in \mathrm{pr}_{A_1} \mathbb{D}$, $\langle \mathbf{a}[i], \mathbf{b}[i] \rangle \in \lambda_\alpha$ for $i \in A_1$, and $\mathrm{pr}_{A_2 \cup \ldots A_k} \mathbf{a} = \mathrm{pr}_{A_2 \cup \ldots A_k} \mathbf{b} \in \mathrm{pr}_{A_2 \cup \ldots A_k} \mathbb{D}$.

Take $\mathbf{a}, \mathbf{b} \in \mathbb{D}$ such that $\langle \mathbf{a}, \mathbf{b} \rangle \in \lambda_\alpha^m$ and $\mathbf{a}[i] = \mathbf{b}[i]$ for $i \in A_{q+1} \cup \ldots \cup A_k$, and define $\mathbf{c}$ to be the tuple with $\mathbf{c}[i] = \mathbf{a}[i]$ if $i \in A_1$ and $\mathbf{c}[i] = \mathbf{b}[i]$ if $i \in A_2 \cup \ldots \cup A_k$. By Lemma 8, $\mathbf{c} \in \mathbb{D}$ and $\langle \mathbf{a}, \mathbf{c} \rangle \in \beta_1$, $\langle \mathbf{c}, \mathbf{b} \rangle \in \beta_2 \wedge \ldots \wedge \beta_q$. Thus $\langle \mathbf{c}, \mathbf{b} \rangle \in \gamma_1 \vee (\beta_2 \wedge \ldots \wedge \beta_q)$. □

6 Algorithms: computing the number of solutions

6.1 The algorithm

Suppose that $\mathcal{H}$ is congruence singular. Let $\mathcal{G}$ be an instance of $\#\mathrm{CSP}(\mathcal{H})$. A mapping $\pi\colon \mathcal{G} \to \mathcal{H}/_\theta$ for $\theta \in \mathsf{Con}(\mathcal{H})$ will be called a *mapping of level θ*. For a mapping π of level θ, by $\Phi(\mathcal{G}, \mathcal{H}, \pi)$ we denote the set of all homomorphisms $\varrho \in \Phi(\mathcal{G}, \mathcal{H})$ with $\varrho/_\theta = \pi$. We recursively compute numbers of the form $|\Phi(\mathcal{G}, \mathcal{H}, \pi)|$ for the instance $\mathcal{G}$ and mappings π of level α^+, $\alpha \in M$. We assume that the universe G of $\mathcal{G}$ is $[m]$. If π is a mapping of level ℓ then $|\Phi(\mathcal{G}, \mathcal{H}, \pi)| = |\Phi(\mathcal{G}, \mathcal{H})|$, and if π is a mapping of level 0 then $|\Phi(\mathcal{G}, \mathcal{H}, \pi)| = 1$. Let $\alpha \in M$ and let π be a mapping from $\mathcal{G}$ to $\mathcal{H}/_{\alpha^+}$. We show how to reduce computing the number $|\Phi(\mathcal{G}, \mathcal{H}, \pi)|$ to computing numbers $|\Phi(\mathcal{G}, \mathcal{H}, \varrho)|$ for certain ϱ, mappings from $\mathcal{G}$ to $\mathcal{H}/_{(\alpha-1)^+}$.

Let $A_1, \ldots, A_k$ be the κ_α^*-classes and $g_1, \ldots, g_k$ their representatives. Let $C_u^1, \ldots, C_u^{s_u}$ be the κ_α-classes from $\pi(g_u)/_{\lambda_\alpha}$, the λ_α-class containing elements from $\pi(g_u)$, for $u \in [k]$. Let $J \subseteq [k]$ and i_u, $u \in [k] - J$, with $i_u \in [s_u]$ be the set corresponding to α^+, and κ_α-classes corresponding to π as in Lemma 8. Without loss of generality we assume $J = [q]$. The next statement follows straightforwardly from Lemma 8.

Proposition 15 *For any q-tuple $\mathbf{s}$ such that $\mathbf{s}[u] \in [m_u]$, the mapping $\varrho_{\mathbf{s}} \colon \mathcal{G} \to \mathcal{H}/_{\alpha^-}$, where for each $g \in A_u$, $u \in [q]$*

$$\varrho_{\mathbf{s}}(g) = \begin{cases} \psi_{g_u,g}(B^u_{\mathbf{s}[u]}) \cap \pi(g), & \text{if } u \le q \\ \psi_{g_u,g}(B^u_{i_u}) \cap \pi(g), & \text{if } u > q. \end{cases}$$

is a homomorphism from $\mathcal{G}$ to $\mathcal{H}/_{\alpha^-}$.

It is not hard to see that sets $\Phi(\mathcal{G},\mathcal{H},\varrho_{\mathbf{s}})$ are the classes of the congruence $(\alpha^-)^m$ on the relation $\Phi(\mathcal{G},\mathcal{H},\pi)$. Clearly, $(\alpha^-)^m = \gamma_1 \cap \ldots \cap \gamma_q$, where $\langle \mathbf{a},\mathbf{b}\rangle \in \gamma_u$ if and only if $\langle \mathbf{a}[g],\mathbf{b}[g]\rangle \in \alpha^-$ if $g \in A_u$ or $g \in A_{q+1} \cup \ldots \cup A_k$, and $\langle \mathbf{a}[g],\mathbf{b}[g]\rangle \in \alpha^+$ otherwise. By Lemma 9 the congruences $\gamma_1,\ldots,\gamma_q$ satisfy condition (1).

Let $T(\pi)$ denote a q-dimensional $s_1 \times \ldots \times s_q$ array such that its entry indexed by $\mathbf{s}$ is equal to $|\Phi(\mathcal{G},\mathcal{H},\varrho_{\mathbf{s}})|$. By Proposition 14, $T(\pi)$ has rank 1, that is, there are numbers $t^u_1,\ldots,t^u_{s_u}$ such that

$$|\Phi(\mathcal{G},\mathcal{H},\varrho_{\mathbf{s}})| = t^1_{\mathbf{s}[1]} \cdot \ldots \cdot t^q_{\mathbf{s}[q]}.$$

These numbers t^i_j can be found as follows. Fix a tuple $\mathbf{s}$. By $\mathbf{s}^i_v$ we denote the tuple, all entries of which are equal to the corresponding entries of $\mathbf{s}$, except for the ith entry that is equal to v. Then set

$$t^1_j = |\Phi(\mathcal{G},\mathcal{H},\varrho_{\mathbf{s}^1_j})| \quad \text{and} \quad t^i_j = \frac{|\Phi(\mathcal{G},\mathcal{H},\varrho_{\mathbf{s}^i_j})|}{|\Phi(\mathcal{G},\mathcal{H},\varrho_{\mathbf{s}^1_j})|} \qquad \text{for } i \in \{2,\ldots,q\}.$$

Now, as the numbers of the form t^i_j are known, we have

$$\begin{aligned}\Phi(\mathcal{G},\mathcal{H},\pi) &= \sum_{\mathbf{s}} \Phi(\mathcal{G},\mathcal{H},\varrho_{\mathbf{s}}) = \sum_{\mathbf{s}} t^1_{\mathbf{s}[1]} \cdot \ldots \cdot t^q_{\mathbf{s}[q]} \\ &= t^1_1 \left(\sum_{\mathbf{s}[2],\ldots,\mathbf{s}[q]} t^2_{\mathbf{s}[2]} \cdot \ldots \cdot t^q_{\mathbf{s}[q]} \right) + \ldots + t^1_{s_1} \left(\sum_{\mathbf{s}[2],\ldots,\mathbf{s}[q]} t^2_{\mathbf{s}[2]} \cdot \ldots \cdot t^q_{\mathbf{s}[q]} \right) \\ &= \ldots = \prod_{j=1}^{q} \sum_{i=1}^{m_j} t^j_i,\end{aligned}$$

that can be computed easily.

Finally, we make use of the following implication of Proposition 12.

Corollary 3 *Let $\varrho \in \Phi(\mathcal{G},\mathcal{H},\pi)/_{(\alpha^-)^m}$ and $\varrho_1,\varrho_2 \in \Phi(\mathcal{G},\mathcal{H},\varrho)$. Then $|\Phi(\mathcal{G},\mathcal{H},\varrho_1)| = |\Phi(\mathcal{G},\mathcal{H},\varrho_2)|$.*

Then for any mapping $\varrho \in \Phi(\mathcal{G},\mathcal{H},\pi)/_{(\alpha^-)^m}$ and any mapping $\varrho' \in \Phi(\mathcal{G},\mathcal{H},\varrho)/_{((\alpha-1)^+)^m}$, we have $|\Phi(\mathcal{G},\mathcal{H},\varrho)| = |\Phi(\mathcal{G},\mathcal{H},\varrho')| \cdot |\Phi(\mathcal{G},\mathcal{H},\varrho)/_{((\alpha-1)^+)^m}|$. The number $|\Phi(\mathcal{G},\mathcal{H},\varrho)/_{(\alpha-1)^+}|$ can be found using algorithm `Uniform` from the next subsection.

6.2 Uniform counting CSPs

Let $\alpha \in M$, π be a mapping of level α^-, and ϱ be a mapping of level $(\alpha-1)^+$. We need a method to find the number

$$\frac{|\Phi(\mathcal{G},\mathcal{H},\pi)|}{|\Phi(\mathcal{G},\mathcal{H},\varrho)|}.$$

We consider first the case when $(\alpha - 1)^+$ is the equality relation. In this case the required number can be found by algorithm UNIFORM using a compact representation R'' of $\Phi(\mathcal{G}, \mathcal{H}, \pi)$. Note also that such a representation can be found by the same algorithm from [13] applied to the instance $\mathcal{G}'$ with the same universe as $\mathcal{G}$ and additional unary constraints $\pi(g)$ imposed on each $g \in \mathcal{G}$. We shall assume that for each (n-ary) relational symbol R from the vocabulary of $\mathcal{H}$, and any set $\{i_1, \ldots, i_k\} \in [n]$, the vocabulary of $\mathcal{H}$ also contains a k-ary relational symbol $\mathrm{pr}_{\{i_1,\ldots,i_k\}} R$ interpreted as $\mathrm{pr}_{\{i_1,\ldots,i_k\}} R^{\mathcal{H}}$. For an instance $\mathcal{G}$ of $\#\mathrm{CSP}(\mathcal{H})$ and $g \in \mathcal{G}$ we denote by $\mathcal{G}_g$ the relational structure with universe $G - \{g\}$ and such that $(g_1, \ldots, g_n) \in R^{\mathcal{G}}$ for some relational symbol R and $g_{i_1} = \ldots = g_{i_\ell} = g$ and the rest of its entries are different from g we exclude this tuple from $R^{\mathcal{G}_g}$, and include the tuple $\mathrm{pr}_{[n]-\{i_1,\ldots,i_\ell\}}(g_1, \ldots, g_n)$ into $\mathrm{pr}_{[n]-\{i_1,\ldots,i_\ell\}} R^{\mathcal{G}_g}$. Recall that we assume $G = [m]$.

```
Algorithm Uniform
INPUT: an compact representation $R''$ of $\Phi(\mathcal{G},\mathcal{H},\pi)$, $\pi \in \Phi(\mathcal{G},\mathcal{H})/_{\beta^m}$, $\beta \overset{s}{\sim} \Delta$
OUTPUT: the cardinality of $\Phi(\mathcal{G},\mathcal{H},\pi)$
Step 1    set $N := 1$, $S := R''$, and $\overline{\mathcal{G}} := \mathcal{G}$
Step 2    for $g = m$ to 1 do
Step 2.1      let $\theta$ be a congruence of $\mathcal{H}$ such that $\langle a,b\rangle \in \theta$ if and only if
              $(g,a,b) \in \mathsf{Sig}_S$; since $\Delta \le \theta \le \beta$,$\theta$ is uniform over $\Delta$; let $w$ be the
              size of its classes
Step 2.2      set $N := N \cdot w$
Step 2.3      set $S := \mathrm{pr}_{[g-1]} S$ and $\overline{\mathcal{G}} := \overline{\mathcal{G}}_g$
          endfor
Step 3    output $N$
```

The correctness of algorithm UNIFORM follows from the rectangularity of $\langle S\rangle$, and the observation that the congruence θ constructed on Step 2.1 can be defined as follows: $\langle a, b\rangle \in \theta$ if and only if there is $\mathbf{a} \in \mathrm{pr}_{[g-1]}\langle S\rangle$ such that $(\mathbf{a}, a) \in \langle S\rangle$ and $(\mathbf{a}, b) \in \langle S\rangle$, that is w is the number of possible extensions of a tuple from $\mathrm{pr}_{[g-1]}\langle S\rangle$.

Observe that if we know the signature of the relation $\Phi(\mathcal{G}, \mathcal{H}, \pi)/_{\alpha^m}$ we still can use algorithm `Uniform`, for we can consider $\Phi(\mathcal{G}, \mathcal{H}, \pi)/_{\alpha^m}$ as a relation on $\mathcal{H}/_{\alpha}$. Therefore the problem we are facing now is to find the signature of this relation. Unfortunately, it is not clear at all how to obtain this signature using the signature or a compact representation of $\Phi(\mathcal{G}, \mathcal{H}, \pi)$, nor we can use the algorithm from [13] to compute the signature of $\Phi(\mathcal{G}, \mathcal{H}/_{\alpha}, \pi)$, since in general $\Phi(\mathcal{G}, \mathcal{H}/_{\alpha}, \pi) \neq \Phi(\mathcal{G}, \mathcal{H}, \pi)/_{\alpha^m}$. Instead, to compute each member of the required signature we find a compact representation of a certain modified problem using the algorithm from [13].

More specifically, we first find the *θ-signature* of the relation $\Phi(\mathcal{G}, \mathcal{H}, \pi)$. Let n be a positive integer, let H be a finite set, let θ be an equivalence relation on H, let $\mathbf{a}$, $\mathbf{b}$ be n-ary tuples and let (i, a, b) be any element in $[n] \times H^2$. We say that $(\mathbf{a}, \mathbf{b})$ *θ-witnesses* (i, a, b) if $\langle \mathbf{a}[j], \mathbf{b}[j]\rangle \in \theta$ for each $j < i$, $\mathbf{a}[i] = a$, and $\mathbf{a}'[i] = b$. Let R be any n-ary relation on H. The θ-signature of R, $\theta\mathsf{Sig}_R \subseteq [n] \times H^2$, is defined to be the set containing all those $(i, a, b) \in [n] \times H^2$ θ-witnessed by tuples in R, that is

$$\theta\mathsf{Sig}_R = \{(i, a, b) \in [n] \times H^2 : \exists \mathbf{a}, \mathbf{b} \in R \text{ such that } (\mathbf{a}, \mathbf{b})\ \theta\text{-witnesses } (i, a, b)\}.$$

We shall assume that for each subalgebra B of $\mathcal{H}$ the vocabulary of $\mathcal{H}$ contains a unary relational symbol R_B such that $R_B^{\mathcal{H}} = B$. Let $\mathcal{G}$ be an instance of $\#\mathrm{CSP}(\mathcal{H})$, let $g_1, \ldots, g_k \in \mathcal{G}$, and let $B_1, \ldots, B_k$

```
Algorithm θ-Signature
INPUT: an instance $\mathcal{G}$ of $\#\mathrm{CSP}(\mathcal{H})$, and a congruence
    $\alpha \in \mathsf{Con}(\mathcal{H})$
OUTPUT: the $\theta$-signature of $\Phi(\mathcal{G},\mathcal{H},\pi)$
Step 1      find a compact representation of $\Phi(\mathcal{G},\mathcal{H},\pi)$
Step 2      set $S := \varnothing$ (the $\theta$-signature of $\Phi(\mathcal{G},\mathcal{H},\pi)$)
Step 3      for each $(i,a,b) \in \{1,\ldots,m\} \times H^2$ do
Step 3.1        if there is $\mathbf{a} \in R'$ such that $\mathbf{a}[i] = a$ then do
Step 3.1.1          find a compact representation $R''$ of $\Phi(\mathcal{G}',\mathcal{H},\pi)$
                    where
            $\mathcal{G}' = \mathcal{G} \cup \{\langle g_1, (\mathbf{a}[1])/_\theta\rangle, \ldots, \langle g_{i-1}, (\mathbf{a}[i-1])/_\theta\rangle\})$
Step 3.1.2          if $b \in \langle \mathrm{pr}_i R''\rangle$ then $S := S \cup \{(i,a,b)\}$
                endif
            endfor
Step 5      return $S$
```

be subalgebras of $\mathcal{H}$. By $\mathcal{G} \cup \{\langle g_1, B_1\rangle, \ldots, \langle g_k, B_k\rangle\}$ we denote the relational structure with the same universe as $\mathcal{G}$, and such that the interpretation of every relational symbol $R \not\in \{R_{B_1}, \ldots, R_{B_k}\}$ equals $R^{\mathcal{G}}$ while the interpretation of R_{B_j} equals $R^{\mathcal{G}}_{B_j} \cup \{g_j\}$. Thus, the elements $g_1, \ldots, g_k$ are forced to be mapped to $B_1, \ldots, B_k$ respectively. It is easy to see that the algorithm θ-SIGNATURE finds the θ-signature of $\Phi(\mathcal{G},\mathcal{H},\pi)$. The signature of $\Phi(\mathcal{G},\mathcal{H},\pi)/_{\theta^m}$ can then be found by replacing each $(i,a,b) \in S$ by $(i, a/_\theta, b/_\theta)$.

Complexity. Observe that the problem of finding the number $|\Phi(\mathcal{G},\mathcal{H},\pi)|$ reduces to finding $s_1 + \ldots + s_k$ numbers of the form $|\Phi(\mathcal{G},\mathcal{H},\varrho)|$, where $\varrho\colon \mathcal{G} \to \mathcal{H}/_{(\alpha\,-\,1)^+}$, and solving the same number of uniform problems. Clearly, $k \le |\mathcal{G}| = m$, $s_i \le |\mathcal{H}| = a$, and $|M| \le a^2$. If the uniform problem can be solved in time $p(m)$ then the overall time complexity of the algorithm is $(amp(m))^{a^2}$.

7 #H-COLORING

Theorem 3 yields a complete classification of #P-complete and polynomial time solvable #H-COLORING problems. However, it is difficult to express the criterion stated in the theorem in terms of (di)graphs. By [27], an (undirected) graph H gives rise to a polynomial time solvable #H-COLORING problem if and only if every connected component of H is either trivial, or a complete bipartite graph, or a complete graph with loops at all vertices. In [14], we observed that an undirected graph satisfies this condition if and only if it is invariant under a Mal'tsev operation.

In this section we compare the classification result from [28, 29] for directed acyclic graphs (DAGs for short) with Theorem 3. We show that every congruence singular DAG satisfies the *Lovász-goodness* condition introduced in [28, 29]. The two conditions must be equivalent, however, the converse implication probably uses some nontrivial properties of DAGs and is more difficult to prove.

A DAG $H = (V,E)$ is called *layered* V can be partitioned into subsets $V_1, \ldots, V_\ell$ such that for any $(v,w) \in E$ we have $v \in V_i$, $w \in V_{i+1}$ for a ceratin $i \le \ell$. Let $v \in V_i$, $w \in V_j$, $i < j$. Then H_{v*} denotes the subgraph of H induced by the vertices u such that there is a path from v to u; similarly, H_{*w}

denotes the subgraph of H induced by the vertices u such that there is a path and from u to w; and $H_{vw} = H_{v*} \cap H_{*w}$. The vertex set of the graph $H_{xy}H_{x'y'}$, where $H_{xy} = (V, E)$ and $H_{x'y'} = (V', E')$, is the set $((V \cap V_i) \times (V' \cap V_i)) \cup \ldots \cup ((V \cap V_j) \times (V' \cap V_j))$, a pair $((v, v'), (w, w'))$ is an edge if and only if $(v, w) \in E$ and $(v', w') \in E'$. It is proved in [28] that $H_{xy}H_{x'y'}$ has only one connected component that spans all layers from i to j. If such main connected components of graphs $H_{xy}H_{x'y'}$ and $H_{zt}H_{z't'}$ are isomorphic then we write $H_{xy}H_{x'y'} \equiv H_{zt}H_{z't'}$. Finally a layered graph is said to be *Lovász-good* if for any $0 \le i < j \le \ell + 1$ and any $x, x' \in V_i$, $y, y' \in V_j$ we have $H_{xy}H_{x'y'} \equiv H_{xy'}H_{x'y}$.

The key lemma for this result is a special case of the result by Lovász [48] that we cite in our notation.

Lemma 10 *If $|\Phi(G, H_1)| = |\Phi(G, H_2)|$ for all graphs G then graphs H_1, H_2 are isomorphic.*

We show that if H is congruence singular then $|\Phi(G, H_{xy}H_{x'y'})| = |\Phi(G, H_{xy'}H_{x'y})|$ for any $x, x' \in V_i$, $y, y' \in V_j$, and any graph G. This implies that $H_{xy}H_{x'y'}$ and $H_{xy'}H_{x'y}$ are isomorphic, and so $H_{xy}H_{x'y'} \equiv H_{xy'}H_{x'y}$. We use an observation made in [28] that $|\Phi(G, H_1H_2)| = |\Phi(G, H_1)| \cdot |\Phi(G, H_2)|$. If $G = (W, F)$ is not layered then $|\Phi(G, H_{xy}H_{x'y'})| = |\Phi(G, H_{xy'}H_{x'y})| = 0$. Let W_1, W_2 denote the set of vertices on the highest and on the lowest layers of G, respectively. As we know, $\Phi(G, H)$ is a relation pp-definable in H. Now let η_1, η_2 be congruences of $\Phi(G, H)$ such that $\langle \varphi, \varphi' \rangle \in \eta_i$, $i = 1, 2$, iff $\varphi(v) = \varphi'(v)$ for all $v \in W_i$. It is not hard to see that sets of the form H_{u*} are classes of η_1, sets of the form H_{*w} are classes of η_2, and sets of the form H_{uw} are classes of $\eta_1 \wedge \eta_2$. Since H is congruence singular, we have $\mathsf{rank}(M(\eta_1, \eta_2)) = k$ where k is the number of classes in $\eta_1 \vee \eta_2$. Hence

$$\begin{vmatrix} |\Phi(G, H_{xy})| & |\Phi(G, H_{xy'})| \\ |\Phi(G, H_{x'y})| & |\Phi(G, H_{x'y'})| \end{vmatrix} = 0,$$

or $\Phi(G, H_{xy}), \Phi(G, H_{x'y'})$ or $\Phi(G, H_{xy'}), \Phi(G, H_{x'y})$ are in different classes of $\eta_1 \vee \eta_2$. In the latter case either $|\Phi(G, H_{x'y})| = |\Phi(G, H_{xy'})| = 0$ or $|\Phi(G, H_{xy})| = |\Phi(G, H_{x'y'})| = 0$. The result follows.

Observe that in this argument congruence singularity is used in a very restricted way: Only projective congruences of only those subalgebras of direct powers of H that are representable in the form $\Phi(G, H)$.

8 Concluding remarks

The result obtained in the paper is rather general. It includes as particular case the results of [18, 27, 22, 28, 29, 44]. However, those results are stated in terms of particular problems, and deriving them from Theorem 3 requires extra research. We also should note that in some cases, e.g., [27], the #P-completeness results obtained for particular problems are stronger than those which follow from our result. For instance, #P-complete #H-COLORING problems in the case of undirected graphs remain #P-complete even when restricted to inputs of bounded degree.

A major question left unanswered is how to check if a given relational structure is congruence singular. This problem may turn out to be even undecidable.

References

[1] V.G. Bodnarchuk, L.A. Kaluzhnin, V.N. Kotov, and B.A. Romov. Galois theory for post algebras. i. *Kibernetika*, 3:1–10, 1969.

[2] G.R. Brightwell and P. Winkler. Graph homomorphisms and phase transitions. *Journal of Combinatorial Theory, Ser. B*, 77:221–262, 1999.

[3] R. Bubley, M. Dyer, C. Greenhill, and M. Jerrum. On approximately counting colourings of small degree graphs. *SIAM Journal of Computing*, 29:387–400, 1999.

[4] A. Bulatov and V. Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. In *Proceedings of the 44th IEEE Symposium on Foundations of Computer Science, FOCS'03*, pages 562–571, Cambridge, MA, USA, October 2003. IEEE Computer Society.

[5] A. Bulatov and M. Grohe. The complexity of partition functions. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming, ICALP'04*, pages 294–306, Turku, Finland, July 2004.

[6] A.A. Bulatov. A dichotomy theorem for constraints on a three-element set. In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science, FOCS'02*, pages 649–658, Vancouver, Canada, November 2002. IEEE Computer Society.

[7] A.A. Bulatov. Mal'tsev constraints are tractable. Technical Report PRG-RR-02-05, Computing Laboratory, University of Oxford, Oxford, UK, 2002.

[8] A.A. Bulatov. Tractable conservative constraint satisfaction problems. In *Proceedings of the 18th Annual IEEE Simposium on Logic in Computer Science*, pages 321–330, Ottawa, Canada, June 2003. IEEE Computer Society.

[9] A.A. Bulatov and P.G. Jeavons. Algebraic approach to multi-sorted constraints. Technical Report PRG-RR-01-18, Computing Laboratory, University of Oxford, Oxford, UK, 2001.

[10] A.A. Bulatov and P.G. Jeavons. An algebraic approach to multi-sorted constraits. In *Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming (CP'03),* pages 197–202, Kinsale, Ireland, 2003.

[11] A.A. Bulatov, P.G. Jeavons, and A.A. Krokhin. Constraint satisfaction problems and finite algebras. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming—ICALP'00*, volume 1853 of *LNCS*, pages 272–282. Springer-Verlag, 2000.

[12] Andrei A. Bulatov. Three-element mal'tsev algebras. *Acta Sci. Math. (Szeged)*, 72:519–550, 2006.

[13] Andrei A. Bulatov and Víctor Dalmau. A simple algorithm for mal'tsev constraints. *SIAM J. Comput.*, 36(1):16–27, 2006.

[14] Andrei A. Bulatov and Víctor Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. *Information and Computation*, 205(5):651–678, 2007.

[15] S. Burris and H.P. Sankappanavar. *A course in universal algebra*, volume 78 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1981.

[16] R. Burton and J. Steif. Nonuniqueness of measures of maximal entropy for subshifts of finite type. *Ergodic Theory and Dynamical Systems*, 14:213–236, 1994.

[17] P.M. Cohn. *Universal Algebra*. Harper & Row, 1965.

[18] N. Creignou and M. Hermann. Complexity of generalized satisfiability counting problems. *Information and Computation*, 125(1):1–12, 1996.

[19] N. Creignou, S. Khanna, and M. Sudan. *Complexity Classifications of Boolean Constraint Satisfaction Problems*, volume 7 of *SIAM Monographs on Discrete Mathematics and Applications*. SIAM, 2001.

[20] J. Diaz, M. Serna, and D.M. Thilikos. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science, DIMACS/DIMATIA Workshop on Graphs, Morphism and Statistical Physics. American Mathematical Society*, chapter Recent results on parameterized H-coloring. To appear.

[21] J. Diaz, M. Serna, and D.M. Thilikos. The complexity of restrictive H-coloring. *To appear in Discrete Applied Mathematics*.

[22] J. Diaz, M. Serna, and D.M. Thilikos. Counting list H-colorings and variants. Technical Report LSI-01-27-R, Departament LSI, Universitat Politècnica de Catalunya, 2001.

[23] J. Diaz, M. Serna, and D.M. Thilikos. Counting h-colorings of partial k-trees. *Theoretical Computer Science*, 281:291–309, 2002.

[24] Q. Donner. On the number of list h-colorings. *J. Graph Theory*, 16(3):239–245, 1992.

[25] M. Dyer, A. Frieze, and M. Jerrum. On counting independent sets in sparse graphs. *SIAM J. on Computing*, 31:1527–1541, 2002.

[26] M. Dyer, L.A. Goldberg, C. Greenhill, and M. Jerrum. On the relative complexity of approximate counting problems. In *Proccedings of Approximation Algorithms for Combinatorial Optimization 3rd International Workshop (APPROX00)*, volume 1913 of *LNCS*, pages 108–119. Springer-Verlag, 2000.

[27] M. Dyer and C. Greenhill. The complexity of counting graph homomorphisms. *Random Structures and Algorithms*, 17:260–289, 2000.

[28] Martin E. Dyer, Leslie Ann Goldberg, and Mike Paterson. On counting homomorphisms to directed acyclic graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, (121), 2005.

[29] Martin E. Dyer, Leslie Ann Goldberg, and Mike Paterson. On counting homomorphisms to directed acyclic graphs. In *ICALP (1)*, pages 38–49, 2006.

[30] T. Feder and M.Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM Journal of Computing*, 28:57–104, 1998.

[31] R. Freese and R. McKenzie. *Commutator theory for congruence modular varieties*, volume 125 of *London Math. Soc. Lecture Notes*. London, 1987.

[32] D. Geiger. Closed systems of function and predicates. *Pacific Journal of Mathematics*, pages 95–100, 1968.

[33] G. Gratzer. *General Lattice Theory*. Birkhäuser Verlag, Basel, 1998.

[34] C. Greenhill. The complexity of counting colourings and independent sets in sparse graphs and hypergraphs. *Computational Complexity*, 9:52–73, 2000.

[35] P. Hell and J. Nešetřil. Counting list homomorphisms for graphs with bounded degrees. *Discrete Mathematics*. to appear.

[36] P. Hell and J. Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Ser.B*, 48:92–110, 1990.

[37] D. Hobby and R.N. McKenzie. *The Structure of Finite Algebras*, volume 76 of *Contemporary Mathematics*. American Mathematical Society, Providence, R.I., 1988.

[38] H.B. Hunt III, M.V. Marathe, V. Radhakrishnan, and R.E. Stearns. The complexity of planar counting problems. *SIAM Journal on Computing*, 27:1142–1167, 1998.

[39] P. Idziak, P. Markovic, R. McKenzie, M. Valeriote, and R. Willard. Tractability and learnability arising from algebras with few subpowers. In *Proceedings of the 22th Annual IEEE Simposium on Logic in Computer Science*. IEEE Computer Society, 2007.

[40] P.G. Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200:185–204, 1998.

[41] P.G. Jeavons, D.A. Cohen, and M.C. Cooper. Constraints, consistency and closure. *Artificial Intelligence*, 101(1-2):251–265, 1998.

[42] P.G. Jeavons, D.A. Cohen, and M. Gyssens. How to determine the expressive power of constraints. *Constraints*, 4:113–131, 1999.

[43] M. Jerrum and A. Sinclair. The Markov chain Monte Carlo method: an approach to approximate counting and integration. In *Approximation Algorithms for NP-hard Problems*, pages 482–520. PSW, 1996.

[44] Ondrej Klíma, Benoit Larose, and Pascal Tesson. Systems of equations over finite semigroups and the #csp dichotomy conjecture. In *MFCS*, pages 584–595, 2006.

[45] J.L. Lebowitz and G. Gallavotti. Phase transitions in binary lattice gases. *Journal of Math. Physics*, 12:1129–1133, 1971.

[46] L.A. Levin. Universal enumeration problems. *Problems on Information Transmission*, 9:265–266, 1973.

[47] N. Linial. Hard enumeration problems in geometry and combinatorics. *SIAM Journal on Algebraic and Discrete Methods*, 7(2):331–335, 1986.

[48] L. Lovász. Operations with structures. *Acta. Math. Acad. Sci. Hung.*, 18:321–328, 1967.

[49] R.N. McKenzie, G.F. McNulty, and W.F. Taylor. *Algebras, Lattices and Varieties*, volume I. Wadsworth and Brooks, California, 1987.

[50] Gustav Nordh and Peter Jonsson. The complexity of counting solutions to systems of equations over finite semigroups. In *COCOON*, pages 370–379, 2004.

[51] P. Orponen. Dempster's rule of combination is #-complete. *Artificial Intelligence*, 44:245–253, 1990.

[52] A.F. Pixley. Characterizations of arithmetical varieties. *Algebra Universalis*, 9(1):87–98, 1979.

[53] J.S. Provan and M.O. Ball. The complexity of counting cuts and of computing the probability that a graph is connected. *SIAM Journal on Computing*, 12(4):777–788, 1983.

[54] D. Roth. On the hardness of approximate reasonning. *Artificial Intelligence*, 82:273–302, 1996.

[55] T.J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th ACM Symposium on Theory of Computing (STOC'78)*, pages 216–226, 1978.

[56] S.P. Vadhan. The complexity of counting in sparse, regular and planar graphs. *SIAM Journal on Computing*, 31(2):398–427, 2001.

[57] L. Valiant. The complexity of computing the permanent. *Theoretical Computing Science*, 8:189–201, 1979.

[58] L. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de/