



HAL
open science

Constraint Satisfaction Tractability from Semi-lattice Operations on Infinite Sets

Manuel Bodirsky, Dugald Macpherson, Johan Thapper

► **To cite this version:**

Manuel Bodirsky, Dugald Macpherson, Johan Thapper. Constraint Satisfaction Tractability from Semi-lattice Operations on Infinite Sets. *ACM Transactions on Computational Logic*, 2013, 14 (4), pp.30:1-19. 10.1145/2528933 . hal-00756929

HAL Id: hal-00756929

<https://hal.science/hal-00756929v1>

Submitted on 13 Mar 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Constraint Satisfaction Tractability from Semi-lattice Operations on Infinite Sets

Manuel Bodirsky, Laboratoire d'Informatique (LIX), CNRS UMR 7161, École Polytechnique
H. Dugald Macpherson, School of Mathematics, University of Leeds
Johan Thapper, Laboratoire de Recherche en Informatique (LRI), Université Paris-Sud 11

A famous result by Jeavons, Cohen, and Gyssens shows that every constraint satisfaction problem (CSP) where the constraints are preserved by a semi-lattice operation can be solved in polynomial time. This is one of the basic facts for the so-called *universal-algebraic approach* to a systematic theory of tractability and hardness in finite domain constraint satisfaction.

Not surprisingly, the theorem of Jeavons et al. fails for arbitrary infinite domain CSPs. Many CSPs of practical interest, though, and in particular those CSPs that are motivated by *qualitative reasoning calculi* from Artificial Intelligence, can be formulated with constraint languages that are rather well-behaved from a model-theoretic point of view. In particular, the automorphism group of these constraint languages tends to be *large* in the sense that the number of orbits of n -subsets of the automorphism group is bounded by some function in n .

In this paper we present a generalization of the theorem by Jeavons et al. to infinite domain CSPs where the number of orbits of n -subsets grows sub-exponentially in n , and prove that preservation under a semi-lattice operation for such CSPs implies polynomial-time tractability. Unlike the result of Jeavons et al., this includes CSPs that cannot be solved by Datalog.

Categories and Subject Descriptors: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Logic and constraint programming; Model theory

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Computational complexity, constraint satisfaction problems, semi-lattice operations, ω -categoricity

ACM Reference Format:

Bodirsky, M., Macpherson, H. D., Thapper, J. 2013. Constraint Satisfaction Tractability from Semi-lattice Operations on Infinite Sets. *ACM Trans. Comput. Logic* V, N, Article A (January YYYY), 19 pages. DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

1. INTRODUCTION

Constraint satisfaction problems are fundamental computational problems that arise in many areas of theoretical computer science. In recent years, a considerable amount of research has been concentrated on the classification of those constraint satisfaction problems that can be solved in polynomial time, and those that are computationally

The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 257039) and from the EPSRC grant EP/H00677X/1. The third author was supported by the LIX-Qualcomm Postdoctoral Fellowship.

Author's addresses: M. Bodirsky, Laboratoire d'Informatique, École Polytechnique, 91128 Palaiseau, France; H. D. Macpherson, School of Mathematics, University of Leeds, Leeds LS2 9JT, England; J. Thapper, Laboratoire de Recherche en Informatique, Université Paris-Sud 11, 91405 Orsay, France.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1529-3785/YYYY/01-ARTA \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

hard. In this paper, we contribute to this line of research and generalize an important tractability condition from finite domain constraint satisfaction to a broad class of infinite domain constraint satisfaction problems.

We work with the following definition of constraint satisfaction problems (CSPs), which is well-adapted to treat CSPs over infinite as well as finite domains. The definition is based on the concept of a *homomorphism* between relational structures, and equivalent to the standard definition for finite domain CSPs. A (*relational*) *structure* A consists of a (not necessarily finite) *domain* $D(A)$ (or simply A when no confusion can arise), and a set of relations on $D(A)$, each of a finite positive arity. Each relation is named by a *relation symbol* R ; the corresponding relation in A is denoted by R^A . The set of all relation symbols is called the *signature* of the structure. A *homomorphism* from a relational structure A to a relational structure B over the same signature is a mapping $f: D(A) \rightarrow D(B)$ such that for each relation symbol R , and tuple $t \in R^A$, it holds that $f(t) \in R^B$, where f is applied component-wise to t . The existence of a homomorphism from A to B is denoted by $A \rightarrow B$ and in this case A is said to be *homomorphic to* B . For a fixed structure, traditionally denoted by Γ , with finite relational signature τ the *constraint satisfaction problem for* Γ (denoted by $\text{CSP}(\Gamma)$) is the following problem.

CSP(Γ)

INSTANCE: A finite structure A over the signature τ .

QUESTION: Is there a homomorphism from A to Γ ?

To give an example, the three-colorability problem can be formulated as $\text{CSP}(K_3)$, where K_3 is the complete graph with three elements. A basic example of an infinite-domain CSP is $\text{CSP}(\langle \mathbb{Q}; < \rangle)$, where $(\mathbb{Q}; <)$ is the strict linear ordering of the rationals.

Jeavons et al. [1997] and Bulatov et al. [2005] made the ground-breaking observation that for finite structures Γ , the complexity of $\text{CSP}(\Gamma)$ is captured by the *polymorphisms* of Γ , defined as follows. When $f: D^k \rightarrow D$ is a k -ary function, and R is an n -ary relation over D , then we say that f *preserves* R if for all n -tuples $t_1, \dots, t_k \in R$, we have

$$(f(t_1[1], \dots, t_k[1]), \dots, f(t_1[n], \dots, t_k[n])) \in R.$$

A *polymorphism* of a relational structure Γ with domain $D = D(\Gamma)$ is a function from D^k to D that preserves all relations of Γ . In other words, a polymorphism is a homomorphism from Γ^k to Γ , for some k . The exploitation of polymorphisms for classifying the complexity of CSPs is sometimes referred to as the *universal-algebraic approach*. Indeed, very often tractability of $\text{CSP}(\Gamma)$ is linked to polymorphisms of Γ with certain ‘good properties’. For a finite domain $D = D(\Gamma)$, it is known that if $\text{CSP}(\Gamma)$ is not NP-hard, then Γ has for some $k \geq 2$ a polymorphism $f: D^k \rightarrow D$ satisfying

$$f(x_1, x_2, \dots, x_k) = f(x_2, x_3, \dots, x_k, x_1),$$

for all $x_1, \dots, x_k \in D$, and it is conjectured that $\text{CSP}(\Gamma)$ can be solved in polynomial-time whenever Γ has such a polymorphism [Bulatov et al. 2005; Barto and Kozik 2010]¹. This conjecture is known to hold in several special cases when f satisfies stronger identities. We will now look at one such case.

¹Barto and Kozik [2010] call f a *cyclic term* when it satisfies the additional requirement of being idempotent; $f(x, \dots, x) = x$ for all $x \in D$. They state the conditions for NP-hardness and the conjecture for polynomial-time tractability in terms of the absence or presence of such cyclic terms among the polymorphisms of the *core* of Γ , cf. Section 3.2. The two conditions are equivalent: let Δ be a finite structure and Γ its core. If f is a cyclic term of Γ and $g: \Delta \rightarrow \Gamma$ a homomorphism, then $(x_1, \dots, x_k) \mapsto f(g(x_1), \dots, g(x_k))$ is a polymorphism of Δ satisfying the condition in the text. Conversely, if f is such a polymorphism of Δ and $h: \Gamma \rightarrow \Delta$ a homomorphism, then $f'(x_1, \dots, x_k) = g(f(h(x_1), \dots, h(x_k)))$ is a polymorphism of Γ . Now $\alpha(x) = f'(x, \dots, x)$ is an automorphism of Γ , and $\alpha^{-1} \circ f'$ is a cyclic term of the core of Δ .

An operation $f: D^2 \rightarrow D$ is called

- *idempotent* if $f(x, x) = x$;
- *commutative* if $f(x, y) = f(y, x)$ for all $x, y \in D$;
- *associative* if $f(x, f(y, z)) = f(f(x, y), z)$ for all $x, y, z \in D$;
- a *semi-lattice operation* if it is commutative, associative, and idempotent.

Jeavons, Cohen, and Gyssens [1997] proved that for every finite structure Γ with a polymorphism that is a semi-lattice operation, the problem $\text{CSP}(\Gamma)$ can be solved in polynomial time. In this paper, we present a generalization of this result to a large class of infinite-domain CSPs. All infinite structures considered are assumed to be countably infinite. To state the result, we need the following definitions.

A bijective homomorphism with an inverse that is also a homomorphism is called an *isomorphism*. An *automorphism* of a relational structure Γ is an isomorphism between Γ and itself, and the set of all automorphisms of Γ is denoted by $\text{Aut}(\Gamma)$. For a subset S of the domain of Γ , the *orbit* of S in Γ is the set $\{\{\alpha(s) \mid s \in S\} \mid \alpha \in \text{Aut}(\Gamma)\}$. When S is of cardinality n , then we call the orbit of S in Γ an *orbit of n -subsets*. If the number of orbits of n -subsets of Γ is at least c^n for some $c > 1$ and all sufficiently large n , then we say that Γ has *exponential growth*. Otherwise, we say that Γ has *sub-exponential growth*, or that Γ is a *sub-exponential structure*. Note that every finite structure is a sub-exponential structure since for n greater than the domain size, there are no n -subsets at all, hence zero orbits. But also the structure $(\mathbb{Q}; <)$ (and all structures with domain \mathbb{Q} whose relations are first-order definable in $(\mathbb{Q}; <)$) is sub-exponential: it has only one orbit of n -subsets, for all n (see e.g. [Hodges 1997]). Our main result is the following.

THEOREM 1.1. *Let Γ be a sub-exponential structure with finite relational signature. If Γ has a semi-lattice polymorphism, then $\text{CSP}(\Gamma)$ can be solved in polynomial time.*

Finite domain structures with a semi-lattice polymorphism can be solved by a standard technique, known as *establishing arc-consistency* or *1-consistency* (see e.g. [Chen et al. 2011]). The situation is different for sub-exponential structures with a semi-lattice polymorphism. Consider for instance the structure $(\mathbb{Q}; \{(x, y, z) \mid x > y \vee x > z\})$. It has the same automorphism group as $(\mathbb{Q}; <)$ and hence is sub-exponential as well. This structure has the function $(x, y) \mapsto \min(x, y)$ on \mathbb{Q} as a polymorphism, but it has been shown that there is no k such that this problem can be solved by establishing k -consistency [Bodirsky and Kára 2010].

Tractability of $\text{CSP}((\mathbb{Q}; \{(x, y, z) \mid x > y \vee x > z\}))$ is instead a consequence of the following more general result of Bodirsky and Kára [2009]: every structure Γ with domain \mathbb{Q} that has the same automorphism group as $(\mathbb{Q}; <)$ and that is preserved by the minimum function (or the maximum function) has a polynomial-time tractable CSP. Our result will be a proper generalization of this result and of the mentioned result of Jeavons, Cohen, and Gyssens.

Let us remark that for a general infinite structure, a semi-lattice polymorphism does not suffice to ensure tractability. For an arbitrary subset U of \mathbb{N} , let Γ_U be the structure $(\mathbb{N}; \{(x, y) \in \mathbb{N}^2 \mid x = y + 1\}, \{0\}, U)$. Every such structure has \min (and \max) as a semi-lattice polymorphism. We claim that $\text{CSP}(\Gamma_U)$ and $\text{CSP}(\Gamma_V)$ are different problems for distinct subsets U and V of \mathbb{N} . Let m be any element in $U \Delta V$, where Δ denotes the symmetric difference of the sets, and let A_m be the instance on variables $\{x_0, \dots, x_m\}$ containing the constraints $x_0 = 0$, $x_i = x_{i-1} + 1$ for $1 \leq i \leq m$, and $U(x_m)$. Then A_m is a satisfiable instance for precisely one of the problems $\text{CSP}(\Gamma_U)$ and $\text{CSP}(\Gamma_V)$. It follows that there are as many pairwise distinct CSPs of this type as there are subsets of \mathbb{N} ,

i.e., uncountably many. However, there are no more than countably many algorithms, hence $\text{CSP}(\Gamma_U)$ is undecidable for some U .

The example of the previous paragraph relies on the fact that all structures Γ_U have an *infinite* number of orbits of 2-subsets. When the number of orbits of n -subsets of Γ is finite for all n , then the structure is called ω -categorical in model theory [Hodges 1997]. For ω -categorical structures, polymorphisms still capture the computational complexity of $\text{CSP}(\Gamma)$ [Bodirsky and Nešetřil 2006]. Our main result can thus be seen as a contribution to the further extension of the universal-algebraic approach from finite to ω -categorical structures.

Overview

Our paper is structured as follows. In Section 2 we introduce a new algorithmic technique to solve infinite-domain constraint satisfaction problems, and present a reduction of $\text{CSP}(\Gamma)$ for structures Γ with a semi-lattice polymorphism to an *efficient sampling algorithm* for Γ . The basic idea is that when there is such an efficient sampling algorithm for Γ , then we can use the arc-consistency procedure for finite domain CSPs to solve $\text{CSP}(\Gamma)$ (we actually use the *uniform* version of the arc-consistency procedure where both A and a finite template B are part of the input). In fact, our technique works under a slightly more general assumption on Γ : instead of requiring the existence of a semi-lattice polymorphism, we only require that Γ has *totally symmetric* polymorphisms of all arities.

The next part of our paper, Section 3, is devoted to the proof that all sub-exponential structures Γ with totally symmetric polymorphisms of all arities admit such an efficient sampling algorithm. Here, our proof is based on a *classification* of those structures Γ . We would like to remark that the general algorithmic technique is applicable also for many structures with totally symmetric polymorphisms of all arities that are *not* sub-exponential, and this will be illustrated by some examples in Section 4. In fact, we make the conjecture that when Γ is an ω -categorical structure with a semi-lattice polymorphism, then $\text{CSP}(\Gamma)$ is in P. However, unlike the case of sub-exponential structures, we cannot provide a classification result like the one in Section 3 for this more general case, and so this remains an interesting open question.

2. ALGORITHM

One of the basic building blocks of our algorithm will be the arc-consistency procedure. Arc-consistency is sometimes called *hyperarc-consistency* when applied to structures with relations of arity greater than two. We start this section by recalling, in the case of finite relational structures, the connection between the applicability of this procedure, homomorphisms from the *set structure*, and the existence of *totally symmetric polymorphisms*.

Let B be a finite structure with a finite relational signature, and let A be an instance of $\text{CSP}(B)$. The *arc-consistency procedure* (AC) applied to the problem (A, B) works by reducing a set of possible images for each variable in A . If such a set becomes empty during the procedure, it follows that there can be no homomorphism, so AC rejects. Otherwise, AC accepts. Algorithm 1 gives the pseudocode for AC. We say that arc-consistency *solves* the problem $\text{CSP}(B)$ if, for every instance A , the procedure accepts *if and only if* $A \rightarrow B$. It can be implemented to run in time that is polynomial in $|A| + |B|$ – take e.g. AC-3 [Dechter 2003].

The *set structure* of B , denoted by $\mathcal{P}(B)$, has as vertices all non-empty subsets of the domain of B . For every k -ary relation R^B , we have $(U_1, \dots, U_k) \in R^{\mathcal{P}(B)}$ iff for every i and $u_i \in U_i$, there exists $u_1 \in U_1, \dots, u_{i-1} \in U_{i-1}, u_{i+1} \in U_{i+1}, \dots, u_k \in U_k$ such that $(u_1, \dots, u_k) \in R^B$.

ALGORITHM 1: $\text{AC}(A, B)$, closely following the pseudocode given in [Chen et al. 2011].

Input: Finite relational structures A and B over the same signature.

Output: Returns `accept` iff $A \rightarrow \mathcal{P}(B)$.

```

foreach  $x \in D(A)$  do
   $h(x) := D(B)$ 
end
repeat
  foreach  $R^A \in A$  do
    foreach  $(x_1, \dots, x_k) \in R^A$  do
      foreach  $i = 1, \dots, k$  do
         $h(x_i) := \pi_i(R^B \cap h(x_1) \times \dots \times h(x_k))$ ;
      end
    end
  end
until  $h(x)$  does not change for any  $x \in D(A)$ ;
if  $h(x) = \emptyset$  for some  $x \in D(A)$  then return reject;
else return accept;

```

A k -ary function $f: D^k \rightarrow D$ is called *totally symmetric* if for all $x_1, \dots, x_k, y_1, \dots, y_k \in D$ we have

$$f(x_1, \dots, x_k) = f(y_1, \dots, y_k) \text{ whenever } \{x_1, \dots, x_k\} = \{y_1, \dots, y_k\}.$$

We say that a structure B has *totally symmetric polymorphisms of all arities* if, for each $k \geq 1$, there is a k -ary polymorphism of Γ that is totally symmetric.

The following is well-known, cf. [Dalmau and Pearson 1999; Feder and Vardi 1999].

THEOREM 2.1. *Let B be a finite structure with a finite relational signature. The following are equivalent.*

- (1) *The arc-consistency procedure solves $\text{CSP}(B)$.*
- (2) *There is a homomorphism $\mathcal{P}(B) \rightarrow B$.*
- (3) *The structure B has totally symmetric polymorphisms of all arities.*

It is clear that when Γ has a semi-lattice operation f , then it also has a totally symmetric polymorphism f_n of arity n , for each $n \geq 2$:

$$f_n(x_1, \dots, x_n) := f(x_1, f(x_2, \dots, f(x_{n-1}, x_n) \dots)).$$

From Theorem 2.1, it thus follows that arc-consistency solves $\text{CSP}(B)$ whenever B is a finite relational structure with a semi-lattice polymorphism. In our arguments, we only need the weaker condition on totally symmetric polymorphisms of all arities; this gives us a stronger result.

The other component of our algorithm will be a procedure to efficiently “sample” appropriate finite structures homomorphic to an infinite structure Γ . Formally, we make the following definition.

Definition 2.2. Let Γ be a structure over a finite relational signature. We say that an algorithm is a *sampling algorithm* for Γ if, given a positive integer n , it computes a finite structure B that is homomorphic to Γ such that $A \rightarrow B$ if and only if $A \rightarrow \Gamma$, for every instance A with $|A| = n$. A sampling algorithm is called *efficient* if its running time is bounded by a polynomial in n .

We are now ready to describe Algorithm 2. Let Γ be a sub-exponential structure over a finite relational signature, and assume that Γ has totally symmetric polymorphisms of all arities.

ALGORITHM 2: CSP(Γ)**Input:** Finite relational structure A over the same signature as Γ .**Output:** Returns accept iff $A \rightarrow \Gamma$. $B := \text{Sample-}\Gamma(|A|);$ // Sample- Γ is a sampling algorithm for Γ **return** AC(A, B);

The main idea of our algorithm is to reduce CSP(Γ) to an appropriate *uniform* finite domain CSP. That is, when given an instance A of CSP(Γ), we reduce to the following problem: decide whether A maps homomorphically to B , where B is a finite structure returned by a sampling algorithm for Γ on input $|A|$, and B is considered as part of the input. We want to use arc-consistency for deciding whether A maps homomorphically to B . Hence, for this approach to work, we need to establish the following.

- (1) There should be an efficient sampling algorithm which samples some B from Γ .
- (2) The arc-consistency procedure applied to (A, B) should accept if and only if $A \rightarrow B$.

The first condition implies that the size of B is polynomial in the size of A , and since AC can be implemented to run in time that is polynomial in $|A| + |B|$, it follows that our algorithm will be polynomial in $|A|$. The second condition ensures that the algorithm gives the correct answer for every instance A of CSP(Γ).

It will be the purpose of Section 3 to prove that an efficient sampling algorithm for Γ exists. We state this result as follows.

THEOREM 2.3. *Let Γ be a sub-exponential structure with a finite relational signature and totally symmetric polymorphisms of all arities. Then there is an efficient sampling algorithm for Γ .*

We next prove one part of a generalization of the equivalence between the second and third item of Theorem 2.1 to infinite domains. This result has a converse for all ω -categorical structures, cf. Section 5.

LEMMA 2.4. *Let Γ be a structure over a finite relational signature. If Γ has totally symmetric polymorphisms of all arities, then $\mathcal{P}(S) \rightarrow \Gamma$ for all finite structures $S \rightarrow \Gamma$.*

PROOF. Let S be a finite structure, let $h: S \rightarrow \Gamma$ be a homomorphism, and let f be an m -ary totally symmetric polymorphism of Γ , where $m = k_{\max}|S|$ and k_{\max} is the maximum arity of any relation in Γ . Let $f': \mathcal{P}(S) \rightarrow \Gamma$ be the function defined on a non-empty set $X = \{x_1, \dots, x_i\} \subseteq S$ by $f'(X) = f(h(x_1), \dots, h(x_i), h(x_i), \dots, h(x_i))$, where the list of arguments of f is padded to length m by elements already occurring in X . We claim that f' is a homomorphism $\mathcal{P}(S) \rightarrow \Gamma$. We must show that $(f'(U_1), \dots, f'(U_k)) \in R^\Gamma$ for an arbitrary relation R , and tuples $(U_1, \dots, U_k) \in R^{\mathcal{P}(S)}$. For each $1 \leq i \leq k$, let $T_i \subseteq R^S \cap (U_1 \times \dots \times U_k)$ be a set of k -tuples such that $\{t[i] \mid t \in T_i\} = U_i$. Note that $\{t[i] \mid t \in T_j, 1 \leq j \leq k\} = U_i$, and that we can choose T_i to have size $l(i) := |U_i| \leq |S|$, for each i . Let $u_1^i, \dots, u_{l(i)}^i$ be an enumeration of U_i and $t_1^i, \dots, t_{l(i)}^i$ be an enumeration of T_i . We then have

$$\begin{aligned} f'(U_i) &= f'(\{u_1^i, \dots, u_{l(i)}^i\}) \\ &= f(h(t_1^1[i]), \dots, h(t_{l(1)}^1[i]), \dots, h(t_1^k[i]), \dots, h(t_{l(k)}^k[i]), \dots, h(t_{l(k)}^k[i])), \end{aligned}$$

for all $1 \leq i \leq k$. Since f is a polymorphism of Γ , and $h(t_j^i) \in R^\Gamma$ for all i and j , it follows that $(f'(U_1), \dots, f'(U_k)) \in R^\Gamma$. \square

The correctness and efficiency of our algorithm now follows from the previous lemma in conjunction with the existence of an efficient sampling algorithm for Γ .

THEOREM 2.5. *Let Γ be a structure over a finite relational signature with totally symmetric polymorphisms of all arities. If there exists an efficient sampling algorithm for Γ , then Algorithm 2 correctly solves $\text{CSP}(\Gamma)$ in polynomial time.*

PROOF. Let A be the input structure and let B be the structure returned by the sampling algorithm for Γ on input $|A|$. The sampling algorithm runs in polynomial time in $|A|$, so the size of B will be polynomial in $|A|$. Since $\text{AC}(A, B)$ can be implemented to run in time polynomial in $|A| + |B|$, it follows that the entire algorithm runs in polynomial time.

To show correctness, note that if $\text{AC}(A, B)$ rejects, then $A \not\rightarrow B$ which is equivalent to $A \not\rightarrow \Gamma$ since B was produced by the sampling algorithm for Γ on input $|A|$. We may therefore safely reject. Otherwise, $\text{AC}(A, B)$ terminates with a non-empty list $h(x) \subseteq B$ for each $x \in A$. Furthermore, for each k -ary relation R^A , tuple $(x_1, \dots, x_k) \in R^A$, index i , and element $d \in h(x_i)$, we know that $R^B \cap h(x_1) \times \dots \times h(x_k)$ contains a tuple (d_1, \dots, d_k) with $d_i = d$. In other words $(h(x_1), \dots, h(x_k)) \in R^{\mathcal{P}(B)}$, so the function $h: A \rightarrow \mathcal{P}(B)$ is a homomorphism. By assumption, Γ has totally symmetric polymorphisms of all arities, so Lemma 2.4 implies the existence of a homomorphism $g: \mathcal{P}(B) \rightarrow \Gamma$. In conclusion, we have a homomorphism $g \circ h: A \rightarrow \Gamma$. \square

As a direct corollary of Theorem 2.3 and Theorem 2.5, we get the following result, which also implies our main result, Theorem 1.1.

THEOREM 2.6. *Let Γ be a sub-exponential structure with a finite relational signature. If Γ has totally symmetric polymorphisms of all arities, then $\text{CSP}(\Gamma)$ is solvable in polynomial time.*

3. CLASSIFICATION

The purpose of this section is to prove Theorem 2.3, showing that every sub-exponential structure has an efficient sampling algorithm. Our approach is based on a *classification* of sub-exponential structures with totally symmetric polymorphisms of all arities. The general outline of this classification is as follows. We first present an argument that reduces the classification task to those sub-exponential structures Γ that are *model-complete cores* and have totally symmetric polymorphisms of all arities (Section 3.2). From there on, the classification follows a decomposition of the automorphism group of Γ . The next step is the reduction to those structures Γ having a *transitive* automorphism group (Section 3.3). We then use the fact that the automorphism group of Γ has only finitely many congruence relations to further reduce the classification task to the case that the automorphism group of Γ is *primitive*. Combining the central theorem from [Macpherson 1985] on primitive permutation groups with a sub-exponential number of orbits of n -subsets with Cameron's theorem on highly set-transitive permutation groups [Cameron 1976] we finish the classification in Section 3.4. Our main classification, Theorem 3.10, shows that every sub-exponential model-complete core that has totally symmetric polymorphisms of all arities has an *interpretation* in $(\mathbb{Q}; <)$.

We find it instructive to give a ‘top-down’ presentation of the classification proof, rather than starting from special cases and assembling more general sub-exponential structures from specific ones. We thus take a decomposition approach, and show first how to describe the most general case in terms of its components. This sometimes leads to forward references of results, but we believe that the reader will be compensated by a more accessible presentation.

3.1. Preliminaries

Before we start, we recall a basic fact which will frequently be used in the following arguments, and which explains the interaction between permutation group theory and

logic for ω -categorical structures. Let Γ be a relational structure. In this paper, we say that a relation R over the domain of Γ is *first-order definable* in Γ if there is a first-order formula $\phi(\bar{x})$ such that $\phi(\bar{a})$ is true in (Γ, \bar{a}) if and only if $\bar{a} \in R$. We say that a relational structure Γ' is *first-order definable in Γ* if Γ' and Γ have the same domain and every relation R in Γ' is first-order definable in Γ . We say that two relational structures are *first-order interdefinable* if one of them is first-order definable in the other, and vice versa. The following is a consequence of the theorem of Ryll-Nardzewski, Engeler, and Svenonius.

THEOREM 3.1 (SEE SECTION 6.3 [HODGES 1997]). *Let Γ be ω -categorical. Then a relation R is preserved by all automorphisms of Γ if and only if R has a first-order definition in Γ . In particular, two ω -categorical structures are first-order interdefinable if and only if they have the same automorphism group.*

This theorem makes possible a translation of terminology between logic and permutation groups. We illustrate its use with the following, which will be needed later on. A *congruence* of a permutation group is an equivalence relation that is preserved by all permutations in the group. A permutation group is called *primitive* if the only congruences are the equivalence relation with just one equivalence class, and the equivalence relation where all equivalence classes are of size one; it is called *imprimitive* otherwise. By Theorem 3.1, the congruences of the automorphism group of an ω -categorical structure Γ are precisely the first-order definable equivalence relations of Γ . We will say that Γ is *primitive* if its automorphism group is primitive.

A permutation group G on a countable set X is called *closed* if and only if it is the automorphism group of a relational structure with domain X . The topological explanations for this terminology can be found in [Gao 2008], Theorem 2.4.4. The corresponding topology is called the *topology of pointwise convergence* on $\text{Sym}(X)$, where $\text{Sym}(X)$ denotes the set of all permutations of X . In this topology, the open sets are unions of sets of the form $\{\alpha \in \text{Sym}(X) \mid \alpha(\bar{x}) = \bar{y}\}$, for n -tuples \bar{x}, \bar{y} of elements of X . A subset H of $G \subseteq \text{Sym}(X)$ is *dense (in G)* if G is the closure of H with respect to this topology.

3.2. Model-Complete Cores

An endomorphism of a relational structure Γ is a homomorphism from Γ to Γ ; we denote the set of all endomorphisms of Γ by $\text{End}(\Gamma)$. A relational structure is called a *core* if every endomorphism of Γ is an embedding². For a relational structure Δ , a *core of Δ* is a core structure Γ that is homomorphically equivalent to Δ , that is, there is a homomorphism from Γ to Δ and vice versa. A first-order formula is called *primitive positive* if it is of the form

$$\exists x_1, \dots, x_n (\psi_1 \wedge \dots \wedge \psi_n),$$

where ψ_1, \dots, ψ_n are atomic formulas. The importance of primitive positive definitions in this paper comes from the fact that relations with a primitive positive definition in a relational structure Γ are preserved by the polymorphisms of Γ .

The motivation of these definitions for constraint satisfaction with finite templates Δ comes from the following facts.

- Every finite relational structure Δ has a core Γ , and Γ is unique up to isomorphism.
- When Γ is a core of Δ , then Γ and Δ have the same CSP.
- In a finite core structure Γ , every orbit of n -tuples³ is primitive positive definable.

²An embedding of a relational structure Γ into a relational structure Δ is an isomorphism between Γ and an induced substructure of Δ .

³When (t_1, \dots, t_n) is an n -tuple of elements of Γ , then the *orbit of t (in Γ)* is the set $\{(\alpha(t_1), \dots, \alpha(t_n)) \mid \alpha \in \text{Aut}(\Gamma)\}$.

These properties have been generalized to ω -categorical structures. A structure Γ is *model-complete* if every embedding of Γ into itself preserves all first-order formulas. We later need the following result which follows by combining Theorem 18 in [Bodirsky 2007] and Lemma 10 in [Bodirsky and Pinsker].

LEMMA 3.2. *An ω -categorical relational structure Δ is a model-complete core if and only if the group of automorphisms of Δ is dense in the endomorphism monoid of Δ , that is, for every endomorphism f and finite subset U of $D(\Delta)$, there is an automorphism g of Δ agreeing with f on U .*

The following appears as Theorem 16 in [Bodirsky 2007]; see also [Bodirsky et al. 2012].

THEOREM 3.3. *Let Δ be an ω -categorical relational structure. Then*

- (1) Δ is homomorphically equivalent to a model-complete core Γ ;
- (2) the structure Γ is unique up to isomorphism, and ω -categorical or finite;
- (3) in Γ , every orbit of n -tuples is primitive positive definable.

For our classification project (and our algorithmic result), it therefore suffices to study the CSPs for model-complete cores of sub-exponential structures. Let us first show that the model-complete core of a sub-exponential structure is again sub-exponential. This follows from the following more general result.

PROPOSITION 3.4. *Let Δ be an ω -categorical relational structure, and let Γ be its model-complete core. Then for every n , the number of orbits of n -subsets in Γ is at most the number of orbits of n -subsets in Δ .*

PROOF. Let f be a homomorphism from Δ to Γ , and g be a homomorphism from Γ to Δ . Since Γ is a core, it follows that $f \circ g$ is an embedding, so g is injective. It now suffices to show that when two n -subsets t_1, t_2 of Γ are mapped by g to two n -subsets s_1, s_2 in the same orbit of n -subsets in Δ , then t_1 and t_2 lie in the same orbit of n -subsets in Γ . Let \bar{t}_1 be an n -tuple listing all the elements in t_1 , let α be an automorphism of Δ that maps s_1 to s_2 , and let $\bar{s}_2 = \alpha(g(\bar{t}_1))$. Since \bar{s}_2 lists all the elements of s_2 , we can arrange the elements of t_2 into an n -tuple \bar{t}_2 such that $g(\bar{t}_2) = \bar{s}_2$. By Theorem 3.3(3), there are primitive positive definitions ϕ_1 and ϕ_2 of the orbits of \bar{t}_1 and \bar{t}_2 . Since g, α , and f preserve primitive positive formulas, the tuple $\bar{t}_3 := f(\alpha(g(\bar{t}_1)))$ satisfies ϕ_1 . But $f(\alpha(g(\bar{t}_1))) = f(g(\bar{t}_2))$, and hence \bar{t}_3 also satisfies ϕ_2 . Therefore, ϕ_1 and ϕ_2 define the same orbit of n -tuples, and so \bar{t}_1 and \bar{t}_2 are in the same orbit. This implies that t_1 and t_2 are in the same orbit of n -subsets. \square

In general it might not be true that the model-complete core of a sub-exponential structure with a semi-lattice polymorphism has again a semi-lattice polymorphism. A finite example of this situation can be derived from Proposition 5.2 in [Larose and Zádori 2004]. This example shows a finite poset with a semi-lattice polymorphism which retracts to a poset without a semi-lattice polymorphism. By introducing constants, the latter structure can be turned into a core of the former. However, we always have the following.

PROPOSITION 3.5. *Let Δ be a finite or ω -categorical relational structure with an n -ary totally symmetric polymorphism. Then the model-complete core of Δ also has an n -ary totally symmetric polymorphism.*

PROOF. Let Γ be the model-complete core of Δ , and let $g: \Delta \rightarrow \Gamma$ and $h: \Gamma \rightarrow \Delta$ be homomorphisms. When f is an n -ary totally symmetric polymorphism of Δ , then $f': D(\Gamma)^n \rightarrow D(\Gamma)$ defined by $(x_1, \dots, x_n) \mapsto g(f(h(x_1), \dots, h(x_n)))$ is totally symmetric, and a polymorphism of Γ . \square

The reason that this proof does not show that the core of a structure with a semi-lattice polymorphism f has again a semi-lattice polymorphism is that f' need not be associative.

Note that a sampling algorithm for the core Γ of a structure Δ is also a sampling algorithm for Δ , since $A \rightarrow \Gamma$ if and only if $A \rightarrow \Delta$ for all structures A . It therefore suffices to show Theorem 2.3 for the special case of sub-exponential structures that are model-complete cores and have totally symmetric polymorphisms of all arities.

3.3. Reduction to the Transitive Case

Let Γ be a sub-exponential model-complete core. Since Γ is sub-exponential, it has in particular a finite number of orbits of 1-subsets, called *orbits* for short. A structure is called *transitive* if it has only one orbit.

PROPOSITION 3.6. *Let Γ be an ω -categorical model-complete core, and let Γ' be the expansion of Γ by all primitive positive definable relations. Let U be an orbit of Γ , and Δ' be the restriction of Γ' to U . Then Δ' is a transitive model-complete core.*

PROOF. First observe that every automorphism α of Γ is also an automorphism of Γ' and preserves U , and hence $\alpha|_U$ is an endomorphism of Δ' . Since the same also applies to the inverse of α , we have that also $\alpha|_U$ has an inverse in $\text{End}(\Delta')$, and therefore is an automorphism of Δ' . So, the restriction of an automorphism of Γ to U is an automorphism of Δ' , and therefore Δ' is transitive.

To show that Δ' is a model-complete core, let e be an endomorphism of Δ' , and let t be a k -tuple of elements from U . Then any primitive positive formula that holds on t in Γ also holds on $e(t)$, since Δ' is the restriction of an expansion of Γ by all primitive positive definable relations. Since Γ is a model-complete core, the orbits of k -tuples are primitive positive definable in Γ by Theorem 3.3(3), and hence there is an automorphism α of Γ that maps t to $e(t)$. Then $\alpha|_U$ is an automorphism of Δ' . This shows that $\text{Aut}(\Delta')$ is dense in $\text{End}(\Delta')$, and the statement follows from Lemma 3.2. \square

The following result is proved in Section 3.4.

THEOREM 3.7. *Let Δ be a transitive sub-exponential model-complete core with totally symmetric polymorphisms of all arities. Then Δ is either a structure of size 1, or it is isomorphic to a structure which is first-order interdefinable with $(\mathbb{Q}; <)$.*

We now analyze how the automorphism group G of Γ is built from its *transitive constituents*, that is, from the permutation groups of the form $\{\alpha|_U \mid \alpha \in \text{Aut}(\Gamma)\}$ on U where U is an orbit of Γ . In general, we only know that G is a subdirect product of its transitive constituents (see, e.g., [Cameron 1999]). In our case, we can make this decomposition more precise, since we have a good knowledge of the group $\text{Aut}((\mathbb{Q}; <))$.

LEMMA 3.8. *Let Γ be a sub-exponential structure with an $\text{Aut}(\Gamma)$ -invariant linear order $<$ defined on the union of two orbits U and V . Then U and V are convex with respect to $<$.*

PROOF. Assume to the contrary that there are elements $u_1, u_2 \in U$, $v_1 \in V$ such that $u_1 < v_1 < u_2$. Since u_1 and u_2 lie in the same orbit, there is an automorphism α of Γ such that $\alpha(u_1) = u_2$. Let $v_{i+1} = \alpha(v_i)$ for $i = 1, \dots, m-1$, and let $u_{i+1} = \alpha(u_i)$ for $i = 2, \dots, m-1$. Since α preserves the order $<$, we have $u_2 = \alpha(u_1) < \alpha(v_1) = v_2$, and $v_2 < \alpha(u_2) = u_3$. By repeated application of α , we obtain $u_i < v_i < u_{i+1}$ for all $i = 1, \dots, m-1$, and finally $u_m < v_m$. Hence, we can encode sequences $s \in \{0, 1\}^m$ in subsets $S \subseteq \Gamma$ of size m by letting $u_i \in S$ iff $s_i = 0$ and $v_i \in S$ iff $s_i = 1$. Different sequences of length m then correspond to different orbits of m -subsets in Γ . This contradicts the assumption that Γ is sub-exponential. \square

Let G and H be permutation groups on the sets X and Y , respectively. By an isomorphism between G and H (and in particular between automorphism groups of relational structures), we will always mean a group isomorphism induced by a bijection between X and Y . That is, G and H are isomorphic as permutation groups.

We now describe the automorphism groups of sub-exponential structures where the closures of the transitive constituents G_1, \dots, G_k for all orbits U_1, \dots, U_k of G are isomorphic to $\text{Aut}((\mathbb{Q}; <))$: the following theorem shows that in this case G is precisely what Cameron [1990] calls the *intransitive action* of the direct product $G_1 \times \dots \times G_k$ on $U_1 \cup \dots \cup U_k$. We will say that a permutation group $G \subseteq \text{Sym}(X)$ is *transitive on a subset* $Y \subseteq X$ if for all $x, y \in Y$, there exists a permutation $\alpha \in G$ such that $\alpha(x) = y$. Otherwise, we say that G is *intransitive on* Y .

LEMMA 3.9. *Let Γ be a sub-exponential structure with automorphism group G and orbits U_1, \dots, U_k , and for $1 \leq i \leq k$ let G_i be the transitive constituent of G on U_i . Assume that for each $i \leq k$ either U_i is of size one or the closure of G_i in $\text{Sym}(U_i)$ is isomorphic to $\text{Aut}((\mathbb{Q}; <))$. Then $\alpha \in G$ if and only if $\alpha|_{U_i} \in G_i$ for each $1 \leq i \leq k$.*

PROOF. Since the closure of G_i is isomorphic to $\text{Aut}((\mathbb{Q}; <))$ for each i , there is a G_i -invariant dense linear order $<_i$ on each orbit U_i . The result is trivial for $k = 1$, so assume that $k > 1$. Let G' be the intransitive action of $G_1 \times \dots \times G_k$ on $U_1 \cup \dots \cup U_k$. It suffices to show that G is dense in G' : since G is closed, they must then be equal. Assume for the sake of contradiction that G is not dense in G' . Then, there are finite sequences (\bar{u}, u') and (\bar{v}, v') in Γ , and an automorphism $\alpha \in G'$ such that $\alpha(\bar{u}, u') = (\bar{v}, v')$, but $\gamma(\bar{u}, u') \neq (\bar{v}, v')$ for all $\gamma \in G$. Furthermore we can choose (\bar{u}, u') and (\bar{v}, v') so that $\beta(\bar{u}) = \bar{v}$ for some $\beta \in G$. By applying β^{-1} to (\bar{v}, v') and α , we may then assume that $\bar{v} = \bar{u}$ and $\alpha(\bar{u}) = \bar{u}$. By a \bar{u} -interval, we will mean an inclusion-maximal convex subset of $U_i \setminus \bar{u}$, for some i , where convexity is evaluated with respect to $<_i$. Now u' lies in some \bar{u} -interval I , and since $\alpha(u') = v'$ and $\alpha(\bar{u}) = \bar{u}$, we have $v' \in I$ as well. By assumption we have $\gamma(u') \neq v'$ for all $\gamma \in \text{Aut}((\Gamma, \bar{u}))$, so $\text{Aut}((\Gamma, \bar{u}))$ is intransitive on I . On the other hand, if we let $\bar{w} := \bar{u} \cap U$, where U is the orbit of Γ containing I , then $\text{Aut}((\Gamma, \bar{w}))$ is clearly transitive on each \bar{w} -interval contained in U : each \bar{w} -interval contained in U is an orbit, and its corresponding transitive constituent of $\text{Aut}((\Gamma, \bar{w}))$ is isomorphic to a dense subgroup of $\text{Aut}((\mathbb{Q}; <))$. In particular, $\text{Aut}((\Gamma, \bar{w}))$ is transitive on I . It therefore follows that we can find a subsequence \bar{a} of \bar{u} containing \bar{w} , and an element $b \notin U$ such that $\text{Aut}((\Gamma, \bar{a}))$ is transitive on I but $\text{Aut}((\Gamma, \bar{a}, b))$ is not.

Let $f_m(\Gamma)$ denote the number of orbits of $\text{Aut}(\Gamma)$ on m -sets and let d be an element of the domain of Γ . Then $f_m((\Gamma, d)) \leq m f_m(\Gamma) + (m + 1) f_{m+1}(\Gamma)$ (for there are at most $m f_m(\Gamma)$ orbits of $\text{Aut}((\Gamma, d))$ on m -sets containing d , and at most $(m + 1) f_{m+1}(\Gamma)$ orbits on m -sets omitting d). Since Γ is sub-exponential and \bar{a} and b are fixed, it follows that (Γ, \bar{a}, b) is sub-exponential. Let V_1 and V_2 be any two orbits of $\text{Aut}((\Gamma, \bar{a}, b))$ contained in $I \subseteq U$. Then the order $<_i$ on U , restricted to $V_1 \cup V_2$, is an $\text{Aut}((\Gamma, \bar{a}, b))$ -invariant linear order on $V_1 \cup V_2$. By Lemma 3.8, it follows that every orbit of $\text{Aut}((\Gamma, \bar{a}, b))$ contained in I is convex. Since $\text{Aut}((\Gamma, \bar{a}, b))$ is intransitive on I , it follows that there is an initial segment $I(b) \subsetneq I$ first-order definable in (Γ, \bar{a}, b) . Let $x \in I(b)$, $y \in I \setminus I(b)$, and pick an automorphism $\alpha \in \text{Aut}((\Gamma, \bar{a}))$ such that $\alpha(x) = y$. Now $\alpha(I(b)) \supseteq I(b)$ and $\alpha(I(b))$ is first-order definable in $(\Gamma, \bar{a}, \alpha(b))$. Let $b_1 = b$ and for $i \geq 1$, let $b_{i+1} = \alpha(b_i)$. By repeating this procedure, we get an increasing sequence of sets $I(b_1) \subsetneq I(b_2) \subsetneq \dots \subsetneq I(b_m)$ which are all definable in $(\Gamma, \bar{a}, \bar{b})$, where $\bar{b} = (b_1, \dots, b_m)$. For $i \geq 1$, pick $c_i \in I(b_{i+1}) \setminus I(b_i)$, so that each of the elements c_i lies in a different orbit of $(\Gamma, \bar{a}, \bar{b})$. We now encode a set $S \subseteq \{1, \dots, m\}$ as the subset T of Γ consisting of the elements in \bar{a}, \bar{b} , and $\{c_i \mid i \in S\}$. Let $S' \subseteq \{1, \dots, m\}$ be another set with encoding T' , and let $\alpha \in G$ be an automorphism such that $\alpha(T) = T'$. Then α has to fix \bar{a} and \bar{b} , i.e., α must be an automorphism of $(\Gamma, \bar{a}, \bar{b})$. Therefore α cannot map c_i to c_j for $i \neq j$, so S must be equal to S' . The set

T has size at most $2m + |\bar{a}|$, and since we can fix the size of \bar{a} , we conclude that the number of distinct orbits of m -subsets of Γ is at least $\Omega(2^{m/2})$. This contradicts the sub-exponentiality of Γ , so G must be dense in G' , and the result follows. \square

We also need the concept of *interpretations* from model theory, which we briefly recall in the following. Let σ and τ be relational signatures, Δ a σ -structure, and Γ a τ -structure. For an m -ary σ -formula ψ , we let $\psi(D(\Delta)^m)$ denote the set of tuples $\bar{a} \in D(\Delta)^m$ such that $\Delta \models \psi(\bar{a})$. A d -dimensional (first-order) interpretation I of Γ in Δ consists of (cf. [Hodges 1997])

- (i) a σ -formula $\partial_I(x_1, \dots, x_d)$;
- (ii) a surjective map $f_I: \partial_I(D(\Delta)^d) \rightarrow \Gamma$ (the *coordinate map*), and
- (iii) for each atomic τ -formula $\phi(y_1, \dots, y_m)$, a σ -formula $\phi_I(\bar{x}_1, \dots, \bar{x}_m)$, where $\bar{x}_1, \dots, \bar{x}_m$ are disjoint d -tuples of distinct variables,

such that for all atomic τ -formulas ϕ and all $\bar{a}_i \in \partial_I(D(\Delta)^d)$

$$\Gamma \models \phi(f_I(\bar{a}_1), \dots, f_I(\bar{a}_m)) \text{ if and only if } \Delta \models \phi_I(\bar{a}_1, \dots, \bar{a}_m). \quad (1)$$

THEOREM 3.10. *Let Γ be a sub-exponential model-complete core with totally symmetric polymorphisms of all arities. Then Γ has a first-order interpretation in $(\mathbb{Q}; <)$.*

PROOF. Let U_1, \dots, U_k be the orbits of Γ . Let Γ' be the expansion of Γ by all primitive positive definable relations, and let $\Gamma'_1, \dots, \Gamma'_k$ be the structures induced in Γ' by U_1, \dots, U_k . Since Γ is a model-complete core, those orbits are primitive positive definable in Γ by Theorem 3.3, and in particular preserved by the totally symmetric polymorphisms of Γ . By Proposition 3.6, Γ'_i is a transitive model-complete core, for all $i \leq k$. Since Γ'_i also has totally symmetric polymorphisms of all arities (obtained as the restrictions of the totally symmetric polymorphisms of Γ to U_i), we can apply Theorem 3.7, and conclude that each of the structures Γ'_i either has size 1, or is isomorphic to a structure which is first-order interdefinable with $(\mathbb{Q}; <)$. We write $<_i$ for the respective linear order that is first-order definable in Γ'_i .

Let G_i be the transitive constituent of $\text{Aut}(\Gamma)$ for orbit U_i . We claim that either U_i has size one, or the closure of G_i equals $\text{Aut}(\Gamma'_i)$. Clearly, every permutation from G_i and its closure preserve all the relations of Γ'_i . For the converse, let α be an automorphism of Γ'_i . It suffices to show that for every finite tuple t of elements from U_i there is a permutation $\beta \in G_i$ such that $\alpha(t) = \beta(t)$. Otherwise, t and $\alpha(t)$ would have been in different orbits of tuples in $\text{Aut}(\Gamma)$. Since orbits of tuples are primitive positive definable in model-complete cores, Γ' has a relation R such that $t \in R$ and $\alpha(t) \notin R$. Since Γ'_i is an induced substructure of Γ' , this contradicts the assumption that α is an automorphism of Γ'_i . We conclude that for each $i \leq k$ where U_i has more than one element, the closure of G_i is isomorphic to $\text{Aut}((\mathbb{Q}; <))$. We can apply Lemma 3.9, and deduce that the automorphism group G of Γ equals $\{\alpha \in \text{Sym}(U_1 \cup \dots \cup U_k) \mid \text{for all } i \leq k \text{ we have } \alpha|_{U_i} \in G_i\}$.

Our interpretation I of Γ in $(\mathbb{Q}; <)$ is $(k+1)$ -dimensional; the domain formula $\partial_I(x_0, x_1, \dots, x_k)$ is $\bigwedge_{1 \leq i < j \leq k} x_i \neq x_j \wedge \bigvee_{1 \leq i \leq k} x_0 = x_i$. The coordinate map f_I sends $\bar{a} \in \partial_I(\mathbb{Q}^{k+1})$ to U_i if $a_0 = a_i$. We choose f_I to be surjective, and if U_i is infinite we choose f_I such that for all $\bar{a}, \bar{b} \in \partial_I(\mathbb{Q}^{k+1})$ with $f_I(\bar{a}), f_I(\bar{b}) \in U_i$ we have $f_I(\bar{a}) <_i f_I(\bar{b})$ if and only if $a_0 < b_0$. In particular, if $f_I(\bar{a}), f_I(\bar{b}) \in U$, we have $f_I(\bar{a}) = f_I(\bar{b})$ if and only if $a_0 = b_0$. Now let $\phi(x_1, \dots, x_m)$ be an atomic τ -formula. We will show that the $m(k+1)$ -ary relation

$$R = \{(a_0^1, \dots, a_k^m) \mid \Gamma \models \phi(f_I(a_0^1, \dots, a_k^1), \dots, f_I(a_0^m, \dots, a_k^m))\}$$

is preserved by all automorphisms of $(\mathbb{Q}; <)$, and so by Theorem 3.1 has a first-order definition over $(\mathbb{Q}; <)$. This first-order definition becomes the formula ϕ_I of our interpretation I .

Let $\alpha \in \text{Aut}((\mathbb{Q}; <))$ be arbitrary. Write β for the mapping that sends an element b of Γ with pre-image c_0, \dots, c_k under f_I to $f_I(\alpha(c_0), \dots, \alpha(c_k))$. To see that this mapping is well-defined, let d_0, \dots, d_k be another pre-image of b under f_I . We have to show that $f_I(\alpha(d_0), \dots, \alpha(d_k)) = \beta(b)$. Since $f_I(\bar{c}) = f_I(\bar{d})$ implies $c_0 = d_0$, we also have $\alpha(c_0) = \alpha(d_0)$ and therefore $\beta(b) = f_I(\alpha(\bar{c})) = f_I(\alpha(\bar{d}))$. It is straightforward to verify that β preserves the relations U_i , and the relations $<_i$. Hence, by Lemma 3.9 we have that β is an automorphism of Γ .

To show that R is preserved by α , let $\bar{a} \in R$ be arbitrary. We have to show that $\alpha(\bar{a}) \in R$, which is the case if and only if $\Gamma \models \phi(f_I(\alpha(a_0^1), \dots, \alpha(a_k^1)), \dots, f_I(\alpha(a_0^m), \dots, \alpha(a_k^m)))$. But

$$\begin{aligned} & (f_I(\alpha(a_0^1), \dots, \alpha(a_k^1)), \dots, f_I(\alpha(a_0^m), \dots, \alpha(a_k^m))) \\ &= (\beta(f_I(a_0^1, \dots, a_k^1)), \dots, \beta(f_I(a_0^m, \dots, a_k^m))) \end{aligned}$$

satisfies ϕ if and only if $(f_I(a_0^1, \dots, a_k^1), \dots, f_I(a_0^m, \dots, a_k^m))$ satisfies ϕ since $\beta \in \text{Aut}(\Gamma)$. And $\Gamma \models \phi(f_I(a_0^1, \dots, a_k^1), \dots, f_I(a_0^m, \dots, a_k^m))$ since $\bar{a} \in R$, which concludes the proof.

To verify that I satisfies the condition (1) from the definition of interpretations, let $\bar{a}_1, \dots, \bar{a}_m \in \partial_I(\mathbb{Q}^{(k+1)})$.

$$\begin{aligned} \Gamma \models \phi(f_I(\bar{a}_1), \dots, f_I(\bar{a}_m)) &\Leftrightarrow (\bar{a}_1, \dots, \bar{a}_m) \in R \\ &\Leftrightarrow (\mathbb{Q}; <) \models \phi_I(\bar{a}_1, \dots, \bar{a}_m) \end{aligned}$$

□

A relational structure Γ is *homogeneous* (or *ultrahomogeneous*) if every isomorphism between finite induced substructures of Γ can be extended to an automorphism of Γ . The structure $(\mathbb{Q}; <)$ is a well-known example of a homogeneous structure. A structure Γ with signature τ *admits quantifier-elimination* if every first-order τ -formula is over Γ equivalent to a quantifier-free τ -formula. A homogeneous structure with finite relational signature admits quantifier-elimination [Hodges 1997]. It follows that all the formulas that appear in the interpretation given in Theorem 3.10 can be assumed to be quantifier-free.

Our interest in interpretations also stems from the following result, which together with the remarks of Section 3.2, implies Theorem 2.3.

LEMMA 3.11. *Every structure Γ with a finite relational signature τ and a d -dimensional interpretation in $(\mathbb{Q}; <)$ has an efficient sampling algorithm. On input n , the output of the algorithm is a structure of size at most $(dn)^d$.*

PROOF. Let I be a d -dimensional interpretation of Γ in $(\mathbb{Q}; <)$. On input n , we compute a finite τ -structure B which is the induced substructure of Γ on the domain $f_I(\partial_I([dn]^d))$. For an m -ary $R \in \tau$, we do the following. Let ϕ_I be the interpretation in $(\mathbb{Q}; <)$ of the atomic formula $R(y_1, \dots, y_m)$, with ϕ_I given in quantifier-free conjunctive normal form. We now evaluate ϕ_I on each sequence of d -tuples, $\bar{a}_1, \dots, \bar{a}_m \in \partial_I([dn]^d)$. The tuple $(f_I(\bar{a}_1), \dots, f_I(\bar{a}_m))$ is in R^B iff the formula $\phi_I(\bar{a}_1, \dots, \bar{a}_m)$ is true. Since ϕ is of constant size in n , it follows that we can compute the relation R^B in $O((dn)^{dm})$ time. The signature τ is finite, so there is a relation of highest arity, independent of n , which provides the upper bound on the time complexity of the algorithm.

Next, let A be a finite τ -structure with n elements, and assume that $s: A \rightarrow \Gamma$ is a homomorphism. This implies $A \rightarrow B$ as well: the image of A under s has at most n

elements, $b_1, \dots, b_n \in \Gamma$. Let $\bar{a}_1, \dots, \bar{a}_n$ be tuples in $\partial_I(\mathbb{Q}^d)$ such that $f_I(\bar{a}_i) = b_i$, and let $g: s(A) \rightarrow \partial_I(\mathbb{Q}^d)$ be the function such that $g(b_i) = \bar{a}_i$ for all i . Note that $\bar{a}_1, \dots, \bar{a}_n$ contains at most dn distinct values of \mathbb{Q} . By homogeneity of $(\mathbb{Q}; <)$, it follows that there is an automorphism α of $(\mathbb{Q}; <)$ such that $\alpha(a_{ij}) \in [dn]$ for each i and j , where $\bar{a}_i = (a_{i1}, \dots, a_{id})$. Now for each relation symbol $R \in \tau$, we have $(x_1, \dots, x_m) \in R^\Gamma$ iff $\phi_I(g(x_1), \dots, g(x_m))$ holds in $(\mathbb{Q}; <)$ iff $\phi_I(\alpha(g(x_1)), \dots, \alpha(g(x_m)))$ holds in $(\mathbb{Q}; <)$, and this in turn is true if and only if $(f_I(\alpha(g(x_1))), \dots, f_I(\alpha(g(x_m)))) \in R^B$. It follows that $f_I \circ \alpha \circ g \circ s$ is a homomorphism from A to B . \square

We are left with the task to prove Theorem 3.7.

Remark. It follows from Lemma 3.11 that all ω -stable ω -categorical structures with finite signature and indiscernible strictly minimal sets in M^{eq} (for definitions of those concepts, see e.g. [Hodges 1997]) have an efficient sampling algorithm, since Lachlan (Theorem 3.1 in [Lachlan 1987]) showed that all those structure have a first-order interpretation in $(\mathbb{Q}; <)$.

3.4. The Transitive Case

In this section, we prove Theorem 3.7 using known results about primitive sub-exponential structures. A permutation group G is *highly set-transitive* if it has exactly one orbit of n -subsets, for every $n > 0$.

THEOREM 3.12 ([MACPHERSON 1985]). *Let G be a primitive but not highly set-transitive permutation group on an infinite set X . If c is a real number with $1 < c < 2^{1/5}$, then G has more than c^n orbits of n -subsets of X , for all sufficiently large n .*

THEOREM 3.13 (SEE SECTION 3.4 IN [CAMERON 1990]). *A permutation group G on a countably infinite set is highly set-transitive iff its closure is isomorphic to a permutation group that contains $\text{Aut}((\mathbb{Q}; <))$.*

To prove Theorem 3.7, we first show that when Γ is a transitive sub-exponential model-complete core with totally symmetric polymorphisms of all arities, then it is also primitive. We will need two more lemmas. Let R and S be two binary relations. An *alternating closed walk on R and S of length $2n$* is a sequence of elements $(x_0, x_1, \dots, x_{2n})$, with $x_{2n} = x_0$, and such that $(x_{2i}, x_{2i+1}) \in R$ and $(x_{2i+1}, x_{2i+2}) \in S$, for $0 \leq i < n$.

LEMMA 3.14. *Let R and S be two binary relations that are preserved by a totally symmetric function f_n of arity $n \geq 1$. If there is an alternating closed walk on R and S of length $2n$, then $R \cap S^{-1} \neq \emptyset$.*

PROOF. Since $(x_{2i}, x_{2i+1}) \in R$ for $0 \leq i < n$, we have $(y, z) \in R$, for $y = f_n(x_0, x_2, \dots, x_{2n-2})$ and $z = f_n(x_1, x_3, \dots, x_{2n-1})$. Similarly, since $(x_{2i+1}, x_{2i+2}) \in S$ for $0 \leq i < n$, we have $(z, z') \in S$, for $z' = f_n(x_2, x_4, \dots, x_{2n})$. Note that

$$y = f_n(x_0, x_2, \dots, x_{2n-2}) = f_n(x_2, x_4, \dots, x_{2n} = x_0) = z'.$$

Therefore, $(y, z) \in R$ and $(z, y) \in S$, hence $(y, z) \in R \cap S^{-1}$. \square

An *orbital* of Γ is an orbit of $\text{Aut}(\Gamma)$ acting component-wise on ordered pairs of elements of Γ . Every transitive structure always has the *trivial orbital* $\{(x, x) \mid x \in \Gamma\}$.

LEMMA 3.15. *Let Γ be a model-complete core with finitely many orbitals and a totally symmetric polymorphism f_n of arity n for all $n \geq 1$. Let X be an equivalence class of a first-order definable equivalence relation on Γ . If $\alpha^m(X) = X$ for some $\alpha \in \text{Aut}(\Gamma)$ and $m > 0$, then $\alpha(X) = X$.*

PROOF. By assumption, there exists a smallest integer $r \geq 1$ such that $\alpha^r(X) = X$. Let $x \in X$, and for $k \in \mathbb{Z}$, let $O(k)$ be the orbital of Γ containing the tuple $(x, \alpha^k(x))$. Then, we have the inclusion $\{(\alpha^n(x), \alpha^{n+k}(x)) \mid n \in \mathbb{Z}\} \subseteq O(k)$.

Since Γ has finitely many orbitals, we can find integers $0 < l < k$ such that $O(k) = O(l)$. In fact, we can do this while ensuring that $l \equiv r - 1 \pmod{r}$. Note that $(\alpha^l(x), \alpha^{l+l+1}(x)) \in O(l+1)$, and that $(\alpha^{i+l+1}(x), \alpha^{i+1}(x)) \in O(-l)$, for all i . In particular, the following sequence is an alternating closed walk on $O(l+1)$ and $O(-l) = O(l)^{-1}$ of length $2(k-l)$.

$$(x, \alpha^{l+1}(x), \alpha^1(x), \alpha^{l+2}(x), \alpha^2(x), \dots, \alpha^{k-l-1}(x), \alpha^k(x), x)$$

As Γ is a model-complete core, each orbital is primitive positive definable in Γ , and hence preserved by f_n for each $n \geq 1$. From Lemma 3.14, it now follows that the orbitals $O(l+1)$ and $O(l)$ intersect, and therefore they must be equal. This implies that the tuples $(x, \alpha^{l+1}(x))$ and $(x, \alpha^l(x))$ are in the same orbital, so there exists an automorphism β of Γ which fixes x and maps $\alpha^{l+1}(x)$ to $\alpha^l(x)$. Since β fixes x , we have $\beta(X) = X$, and from the choice of l , we have $\alpha^{l+1}(X) = X$. It follows that $\alpha^l(x) = \beta(\alpha^{l+1}(x)) \in X$, so $\alpha^l(X) = X$, and hence $\alpha^{r-1}(X) = X$ as well. Due to the choice of r , this is only possible if $r = 1$, so we conclude that α preserves X . \square

PROPOSITION 3.16. *Let Γ be a transitive sub-exponential model-complete core with totally symmetric polymorphisms of all arities. Then Γ has size 1, or it is infinite and primitive.*

PROOF. Assume that Γ has size at least 2. Let E_0 and E_1 denote the congruence of $\text{Aut}(\Gamma)$ with equivalence classes of size 1, and the congruence with a single equivalence class, respectively. Let E be an inclusion-maximal congruence from the set of all congruences different from E_1 . Existence of E follows from the existence of E_0 , and Γ having finitely many first-order definable equivalence relations. This follows from Theorem 3.1, since each first-order definable binary relation over Γ consists of a union of orbitals, and there are finitely many such orbitals since Γ is sub-exponential. We want to show that E must in fact be E_0 from which it follows that $\text{Aut}(\Gamma)$ is primitive.

Let $D = D(\Gamma)$. By D/E we will denote the set of equivalence classes of E . For $x \in D$, let $x[E]$ denote the equivalence class of E containing x , and for $\alpha \in \text{Aut}(\Gamma)$, let α_E denote the function on D/E which maps $x[E]$ to $\alpha(x)[E]$ for each equivalence class $x[E]$ of E . (It follows from E being a congruence that α_E is well-defined.) We then have that $\text{Aut}(\Gamma)/E := \{\alpha_E \mid \alpha \in \text{Aut}(\Gamma)\}$ is a permutation group on D/E .

Let $H = \text{Aut}(\Gamma)/E$. For any n -subsets $\{X_1, \dots, X_n\}$ and $\{Y_1, \dots, Y_n\}$ of D/E , pick $x_i \in X_i$ and $y_i \in Y_i$ for each i . If $\alpha \in \text{Aut}(\Gamma)$ satisfies $\alpha(\{x_1, \dots, x_n\}) = \{y_1, \dots, y_n\}$ then $\alpha_E(\{X_1, \dots, X_n\}) = \{Y_1, \dots, Y_n\}$. Hence the number of orbits on n -subsets of H is sub-exponential whenever Γ is sub-exponential.

Assume that H is finite. Pick two distinct equivalence classes X and Y of E . Since Γ is transitive, we can find $\alpha \in \text{Aut}(\Gamma)$ such that $\alpha(X) = Y$. But if H is finite, then $\alpha^m(X) = X$ for some $m > 0$, so $\alpha(X) = X$ by Lemma 3.15, a contradiction. Therefore H must be infinite. Congruences of H are in one-to-one correspondence with the congruences of $\text{Aut}(\Gamma)$ containing E . As the latter are precisely E and E_1 , it follows that H is primitive.

It now follows from Theorem 3.12 that H is highly set-transitive, and so from Theorem 3.13 that the closure of H is isomorphic to a closed permutation group H' that contains $\text{Aut}(\langle \mathbb{Q}; < \rangle)$. The group H' either has one or two non-trivial orbitals. If it only has one non-trivial orbital, then so does H , hence for any two distinct equivalence classes X and Y of E , there is an automorphism α of Γ such that $\alpha(X) = Y$ and $\alpha(Y) = X$. Again by Lemma 3.15, it follows that $\alpha(X) = X$, a contradiction. So H' has two non-trivial orbitals, one of which is the order $<$ on \mathbb{Q} . Via the isomorphism, H

thus has a non-trivial orbital $<$ which is a linear order on the equivalence classes of E . Let R be a binary relation on Γ defined by $R(x, y)$ iff $x[E] < y[E]$. For $\alpha \in \text{Aut}(\Gamma)$, we have that α_E preserves $<$, so α preserves R . Assume now that the equivalence classes of E have size greater than 1. We can then encode a sequence in $\{0, 1\}^n$ as a set $\{x_1, y_1, \dots, x_n, y_n\} \subseteq \Gamma$: choose x_i, y_i so that $R(y_i, x_{i+1})$ for $1 \leq i < n$, encode a value 0 in position i by enforcing $E(x_i, y_i)$, $x_i \neq y_i$, and encode a value 1 by enforcing $R(x_i, y_i)$. The relations E, \neq , and R are all preserved by $\text{Aut}(\Gamma)$, so if two $2n$ -subsets encodes distinct sequences, then they must be contained in distinct orbits of $2n$ -subsets. Hence, the number of orbits of $2n$ -subsets is greater than or equal to 2^n , which contradicts Γ being sub-exponential. So the equivalence classes of E are of size 1, i.e., $E = E_0$, and Γ is primitive. \square

PROOF OF THEOREM 3.7. Let Γ be a transitive sub-exponential model-complete core with totally symmetric polymorphisms of all arities. If Γ is of size greater than 1, then it is infinite and primitive (Proposition 3.16), so has a highly set-transitive automorphism group G (Theorem 3.12). Hence, the closed group G is isomorphic to a group that contains $\text{Aut}((\mathbb{Q}; <))$ (Theorem 3.13). Via this isomorphism, we identify the domain of G and the domain of Γ with \mathbb{Q} in the following. By Theorem 3.1, all relations of Γ are first-order definable in $(\mathbb{Q}; <)$.

If the binary relation $<$ is an orbital of Γ , then $<$ is first-order definable in Γ by Theorem 3.1, and interdefinability with $(\mathbb{Q}; <)$ follows. Otherwise, the smallest orbital of Γ that contains $<$ also contains a pair (x, y) such that $x > y$. It follows that Γ has an automorphism α such that $\alpha(x) = y$ and $\alpha(y) = x$, i.e., $\alpha^2(x) = x$. Since Γ has totally symmetric polymorphisms of all arities, we can apply Lemma 3.15 to deduce $\alpha(x) = x$, a contradiction. \square

4. BEYOND SUB-EXPONENTIAL GROWTH: EXAMPLES

In this section we illustrate the possibility of extending our tractability result beyond sub-exponential structures by providing a number of examples of structures with exponential growth that are interpretable in $(\mathbb{Q}; <)$, and have semi-lattice operations. The tractability of these examples then follows from Theorem 2.5 and Lemma 3.11.

Example 4.1. Let ' $<$ ', ' $=$ ', and ' $>$ ' denote the usual inequality and equality relations on \mathbb{Q} . Let Γ_1 be the relational structure over \mathbb{Q}^2 with binary relations $R_{\rho, \sigma}$ for $\rho, \sigma \in \{<, =, >\}$, where $R_{\rho, \sigma}((x_1, y_1), (x_2, y_2))$ is defined by $\rho(x_1, x_2) \wedge \sigma(y_1, y_2)$. This structure is exponential. To obtain a lower bound on the growth rate of Γ_1 , pick an n -subset $A = \{a_1, \dots, a_n\} \subseteq \mathbb{Q}^2$, and assume that the projection of A on the second component contains exactly k distinct values. For $1 \leq i \leq k$, let A_i denote the number of elements that have the i th largest second component value. Then, $A_1 + \dots + A_k = n$ determines a *composition of n* , i.e., an expression for n as an ordered sum of positive integers. Two sets $A, B \subseteq \mathbb{Q}^2$ which determine different compositions must be in different orbits. Thus, the number of orbits of n -subsets of \mathbb{Q}^2 is at least 2^{n-1} , the number of compositions of n .

Hence, Theorem 2.6 does not apply to Γ_1 . Instead, tractability can be inferred as follows. The structure Γ_1 has a two-dimensional interpretation I in $(\mathbb{Q}; <)$. The formula ∂_I is always true, f_I is the identity on \mathbb{Q}^2 , and the interpretations of $R_{\rho, \sigma}$ are as given above. Furthermore, it is easy to verify that Γ_1 is invariant under the semi-lattice operations given by (component-wise) \min and \max . Tractability of $\text{CSP}(\Gamma_1)$ now follows from Theorem 2.5 and Lemma 3.11.

For any $d > 2$, this example can be generalized to a structure with 3^d relations of arity d , and with a d -dimensional interpretation in $(\mathbb{Q}; <)$. The CSP of each such structure is polynomial-time solvable. \square




Fig. 1. Four relational structures with signature (R, S) ; the relations R and S are given by the solid and dashed arrows, respectively.

Example 4.2. The *age* of a structure Γ is defined as the class of all finite structures isomorphic to a substructure of Γ . Let (R, S) be a signature with two binary relation symbols, and let $T = \{T_1, T_2, T_3, T_4\}$ be the set of structures in Fig. 4, where the tuples of R (S) are given by the solid (dashed) arrows. Let \mathcal{C} be the class of all finite structures with signature (R, S) for which every three-element substructure is isomorphic to a structure in T . It can be shown that \mathcal{C} is an *amalgamation class* so that its *Fraïssé limit* exists (cf. Theorem 6.1.2 in [Hodges 1997]). This is the up to isomorphism unique countable homogeneous structure with age \mathcal{C} .

The following describes a relational structure Γ_2 with age \mathcal{C} which can be verified to be homogeneous. It follows that Γ_2 is isomorphic to the Fraïssé limit of \mathcal{C} . Let $\Gamma_2 = (\mathbb{Q}^2; R, S)$, where $R((x_1, y_1), (x_2, y_2))$ is the relation $x_1 = x_2 \wedge y_1 < y_2$, and $S((x_1, y_1), (x_2, y_2))$ is the relation $x_1 < x_2$. The growth rate of Γ_2 can be bounded as in the previous example, and here it turns out that the number of orbits of n -subsets is precisely 2^{n-1} . The structure Γ_2 also has a two-dimensional interpretation in $(\mathbb{Q}; <)$ and semi-lattice polymorphisms given by (component-wise) \min and \max , so tractability follows once again from Theorem 2.5 and Lemma 3.11. \square

Example 4.3. Let $\Gamma_3 := (U \cup V; M, <)$ be the following relational structure. The domain $U \cup V$ is the disjoint union of two copies of \mathbb{Q} . The binary relation M defines a perfect matching between the elements of U and the elements of V , and the binary relation $<$ defines a dense linear order on $U \cup V$ such that $u < v$ for all $u \in U$ and $v \in V$, and for $v_1, v_2 \in V$, we have $v_1 < v_2$ iff $u_1 < u_2$ for the elements $u_1, u_2 \in U$ with $(u_1, v_1), (u_2, v_2) \in M$.

The structure Γ_3 is invariant under the semi-lattice operations given by \min and \max defined with respect to the order $<$ on $U \cup V$. It has two orbits and $\text{Aut}(\Gamma_3)$ is isomorphic (as an abstract group) to $\text{Aut}((\mathbb{Q}; <))$. By Lemma 3.9, this implies that Γ_3 does not have sub-exponential growth. But Γ_3 has a 2-dimensional interpretation I in $(\mathbb{Q}; <)$, so $\text{CSP}(\Gamma_3)$ is polynomial-time solvable: let $\partial_I(x, y)$ be the formula $x \neq y$, and let $f_I(x, y)$ be the copy of x in U if $x < y$ and the copy of x in V if $x > y$. The matching M on (x_1, y_1) and (x_2, y_2) is interpreted by the formula $x_1 = x_2 \wedge x_1 < y_1 \wedge x_2 > y_2$ and the order $<$ on (x_1, y_1) and (x_2, y_2) is interpreted by the formula $(x_1 < y_1 \wedge x_2 > y_2) \vee (x_1 < y_1 \wedge x_2 < y_2 \wedge x_1 < x_2) \vee (x_1 > y_1 \wedge x_2 > y_2 \wedge x_1 < x_2)$. \square

5. CONCLUDING REMARKS

In this article we prove that constraint satisfaction problems for templates Γ where the number of orbits of n -subsets of Γ grows sub-exponentially in n can be solved in polynomial time when Γ has a semi-lattice polymorphism. In fact, we showed the stronger result which only requires the existence of totally symmetric polymorphisms of all arities, instead of requiring the existence of a semi-lattice polymorphism. This algorithmic result can be showed in two stages:

- (1) In the first stage, we reduce $\text{CSP}(\Gamma)$, for structures Γ with totally symmetric polymorphisms of all arities, to solving certain uniform finite domain CSPs, and to the task to find an efficient sampling algorithm for Γ .
- (2) In the second stage, we *classify* sub-exponential structures Γ with totally symmetric polymorphisms of all arities, and use the classification to verify that there always exists an efficient sampling algorithm for Γ .

The reduction presented in the first stage crucially relies on the fact that when Γ has totally symmetric polymorphisms of all arities, then for all finite induced substructures S of Γ the set structure of S homomorphically maps to Γ (Lemma 2.4). We want to remark that this connection has a converse when Γ is ω -categorical.

LEMMA 5.1. *Let Γ be an ω -categorical structure over a finite relational signature. Then Γ has totally symmetric polymorphisms of all arities if and only if $\mathcal{P}(S) \rightarrow \Gamma$ for all finite structures S with $S \rightarrow \Gamma$.*

PROOF. The forward direction was proved in Lemma 2.4. The remaining direction can be proved using a common technique for constructing homomorphisms to ω -categorical structures. Given an arbitrary positive integer n , we want to produce an n -ary totally symmetric polymorphism f of Γ . The idea of the proof is as follows: let l_1, l_2, \dots be an enumeration of the elements of $D = D(\Gamma)$, and let $L_k = \{l_1, \dots, l_k\}$. For each $k \geq 1$, let F_k be the set of homomorphisms from $\mathcal{P}(\Gamma[L_k])$ to Γ . Introduce an equivalence relation \sim on F_k by defining $f \sim g$ iff $f = \alpha \circ g$ for some automorphism α of Γ . Let \tilde{F}_k denote the set of equivalence classes of F_k under \sim .

Arrange the elements of $\bigcup_{k \geq 1} \tilde{F}_k$ into a forest containing at least one infinite tree: each $\tilde{f}_1 \in \tilde{F}_1$ is defined to be the root of a separate tree, and for each $k > 1$, and $f_k \in F_k$, define the parent of $\tilde{f}_k \in \tilde{F}_k$ to be the equivalence class containing the restriction of f_k to the non-empty subsets of L_{k-1} . This definition is independent of the choice of representative in \tilde{f}_k , so each equivalence class in \tilde{F}_k , $k > 1$, has precisely one parent. Since Γ is ω -categorical, it follows that there are finitely many equivalence classes for a fixed k . Hence, there are finitely many trees and each tree is finitely branching in each node. By assumption, \tilde{F}_k is non-empty for each $k \geq 1$, so some tree has unbounded height. Now, König's tree lemma implies the existence of an infinite path $\tilde{f}_1, \tilde{f}_2, \dots$ in some tree. Assume that there are representatives $f_1 \in \tilde{f}_1, f_2 \in \tilde{f}_2, \dots, f_k \in \tilde{f}_k$, such that f_{k-1} is the restriction of f_k to the non-empty subsets of L_{k-1} . We show that this path can be extended indefinitely: choose $g_{k+1} \in \tilde{f}_{k+1}$ arbitrarily and let g_k be its restriction to the non-empty subsets of L_k . Then, there exists an automorphism α such that $f_k = \alpha \circ g_k$. It follows that f_k is the restriction of $\alpha \circ g_{k+1}$ to the non-empty subsets of L_k , hence we can define $f_{k+1} := \alpha \circ g_{k+1}$.

Now, for any n -tuple (x_1, \dots, x_n) over D , define $f(x_1, \dots, x_n) = f_m(\{x_1, \dots, x_n\})$, where m is an any integer such that $\{x_1, \dots, x_n\} \subseteq L_m$. By the construction of the sequence f_1, f_2, \dots , the function f is a well-defined totally symmetric n -ary function on D . To verify that f is a polymorphism of Γ , let R^Γ be an r -ary relation, and $t_1, \dots, t_n \in R^\Gamma$. Let $U_i = \{t_1[i], \dots, t_n[i]\}$, for $1 \leq i \leq r$. Assume without loss of generality that m has been chosen large enough so that $\bigcup_{i=1}^r U_i \subseteq L_m$. Then $f(t_1[i], \dots, t_n[i]) = f_m(U_i)$ for all i . By definition of the set structure, we have $(U_1, \dots, U_r) \in R^{\mathcal{P}(\Gamma[L_m])}$. Since f_m is a homomorphism, we conclude that $(f_m(U_1), \dots, f_m(U_r)) \in R^\Gamma$, so f is indeed a polymorphism. \square

Because of the general applicability of the algorithmic approach in Section 2, and because many structures of exponential growth that have totally symmetric polymorphisms seem to be well-behaved (see Section 4), we make the following conjecture.

CONJECTURE 5.2. *Let Γ be an ω -categorical structure with finite relational signature and totally symmetric polymorphisms of all arities. Then $\text{CSP}(\Gamma)$ can be solved in polynomial time.*

Acknowledgements

We would like to thank the referees for their very helpful comments.

REFERENCES

- BARTO, L. AND KOZIK, M. 2010. New conditions for Taylor varieties and CSP. In *Proceedings of LICS*. 100–109.
- BODIRSKY, M. 2007. Cores of countably categorical structures. *Logical Methods in Computer Science* 3, 1, 1–16.
- BODIRSKY, M., HILS, M., AND MARTIN, B. 2012. On the scope of the universal-algebraic approach to constraint satisfaction. *Logical Methods in Computer Science (LMCS)* 8, 3:13. Available at arXiv:0909.5097v3. An extended abstract that announced some of the results appeared in the proceedings of Logic in Computer Science (LICS'10).
- BODIRSKY, M. AND KÁRA, J. 2009. The complexity of temporal constraint satisfaction problems. *Journal of the ACM* 57, 2, 1–41. An extended abstract appeared in the Proceedings of the Symposium on Theory of Computing (STOC'08).
- BODIRSKY, M. AND KÁRA, J. 2010. A fast algorithm and Datalog inexpressibility for temporal reasoning. *ACM Transactions on Computational Logic* 11, 3.
- BODIRSKY, M. AND NEŠETŘIL, J. 2006. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation* 16, 3, 359–373.
- BODIRSKY, M. AND PINSKER, M. Minimal functions on the random graph. *Israel Journal of Mathematics*.
- BULATOV, A. A., KROKHIN, A. A., AND JEAVONS, P. G. 2005. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing* 34, 720–742.
- CAMERON, P. J. 1976. Transitivity of permutation groups on unordered sets. *Mathematische Zeitschrift* 148, 127–139.
- CAMERON, P. J. 1990. *Oligomorphic Permutation Groups*. Cambridge University Press, Cambridge.
- CAMERON, P. J. 1999. *Permutation Groups*. LMS Student Text 45. Cambridge University Press, Cambridge.
- CHEN, H., DALMAU, V., AND GRUSSIEN, B. 2011. Arc consistency and friends. *Accepted for publication in the Journal on Logic and Computation*. CoRR abs/1104.4993.
- DALMAU, V. AND PEARSON, J. 1999. Closure functions and width 1 problems. In *Proceedings of CP*. 159–173.
- DECHTER, R. 2003. *Constraint Processing*. Morgan Kaufmann.
- FEDER, T. AND VARDI, M. Y. 1999. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM Journal on Computing* 28, 57–104.
- GAO, S. 2008. *Invariant Descriptive Set Theory*. Pure and applied mathematics. Taylor and Francis.
- HODGES, W. 1997. *A shorter model theory*. Cambridge University Press, Cambridge.
- JEAVONS, P., COHEN, D., AND GYSSENS, M. 1997. Closure properties of constraints. *Journal of the ACM* 44, 4, 527–548.
- LACHLAN, A. H. 1987. Structures coordinatized by indiscernible sets. *Annals of Pure and Applied Logic* 34, 245–273.
- LAROSE, B. AND ZÁDORI, L. 2004. Finite posets and topological spaces in locally finite varieties. *Algebra Universalis* 52, 2-3, 119–136.
- MACPHERSON, H. D. 1985. Orbits of infinite permutation groups. *Proceedings of the LMS* 51, 3, 246–284.