

Homomorphism Problems for First-Order Definable Structures*

Bartek Klin¹, Sławomir Lasota¹, Joanna Ochremiak², and Szymon Toruńczyk¹

1 University of Warsaw

2 Universitat Politècnica de Catalunya

Abstract

We investigate several variants of the homomorphism problem: given two relational structures, is there a homomorphism from one to the other? The input structures are possibly infinite, but definable by first-order interpretations in a fixed structure. Their signatures can be either finite or infinite but definable. The homomorphisms can be either arbitrary, or definable with parameters, or definable without parameters. For each of these variants, we determine its decidability status.

1998 ACM Subject Classification F.4.1 Mathematical Logic–Model theory, F.4.3 Formal Languages–Decision problems

Keywords and phrases Sets with atoms, first-order interpretations, homomorphism problem

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2016.?

1 Introduction

First-order definable sets, although usually infinite, can be finitely described and are therefore amenable to algorithmic manipulation. Definable sets (we drop the qualifier *first-order* in what follows) are parametrized by a fixed underlying relational structure \mathcal{A} whose elements are called *atoms*. We shall assume that the first-order theory of \mathcal{A} is decidable. To simplify the presentation, unless stated otherwise, let \mathcal{A} be a countable set $\{\underline{1}, \underline{2}, \underline{3}, \dots\}$ equipped with the equality relation only; we shall call this the *pure set*.

► **Example 1.** Let

$$V = \{ \{a, b\} \mid a, b \in \mathcal{A}, a \neq b \},$$

$$E = \{ (\{a, b\}, \{c, d\}) \mid a, b, c, d \in \mathcal{A}, a \neq b \wedge a \neq c \wedge a \neq d \wedge b \neq c \wedge b \neq d \wedge c \neq d \}.$$

Both V and E are definable sets (over \mathcal{A}), as they are constructed from \mathcal{A} using (possibly nested) set-builder expressions with first-order guards ranging over \mathcal{A} . In general, we allow finite unions in the definitions, and finite tuples (as above) are allowed for notational convenience. Precise definitions are given in Section 2. The pair $G = (V, E)$ is also a definable set, in fact, a definable graph. It is an infinite Kneser graph (a generalization of the famous Petersen graph): its vertices are all two-element subsets of \mathcal{A} , and two such subsets are adjacent iff they are disjoint.

The graph G is \emptyset -definable: its definition does not refer to any particular elements of \mathcal{A} . In general, one may refer to a finite set of parameters $S \subseteq \mathcal{A}$ to describe an S -definable set. For instance, the set $\{a \mid a \in \mathcal{A}, a \neq \underline{1} \wedge a \neq \underline{2}\}$ is $\{\underline{1}, \underline{2}\}$ -definable. Definable sets are those which are S -definable for some finite $S \subseteq \mathcal{A}$. ◀

* This work is supported by Poland’s National Science Centre grant 2012/07/B/ST6/01497, and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement ERC-2014-CoG 648276 AUTAR).



Although definable relational structures correspond (up to isomorphism) to first-order interpretations well-known from logic and model theory [19], we prefer to use a different definition since standard set-theoretic notions directly translate into this setting. For example, a definable function $f : X \rightarrow Y$ is simply a function whose domain X , codomain Y , and graph $\Gamma(f) \subseteq X \times Y$ are definable sets. A relational structure is definable if its universe, signature, and interpretation function that maps each relation symbol to a relation on the universe, are definable. Finally, a definable homomorphism between definable structures over the same signature is a definable mapping between their universes that is a homomorphism, i.e., preserves every relation in the signature. All hereditarily finite sets (finite sets, whose elements are finite, and so on, recursively) are definable, and every finite relational structure over a finite signature is (isomorphic to) a definable one.

The classical *homomorphism problem* is the problem of determining whether there exists a homomorphism from a given finite source structure \mathbb{A} to a given finite target structure \mathbb{B} . This is also known as the Constraint Satisfaction Problem, and is clearly decidable (and NP-complete). The precise computational complexity has been thoroughly studied in the literature in many variants. The case when the target structure is fixed (and is called a *template*) is of particular interest, as it expresses many natural computational problems (such as k -colorability, 3-SAT, or solving systems of linear equations over a finite field). The famous Feder-Vardi conjecture states that for every fixed template \mathbb{B} , the corresponding constraint satisfaction problem $\text{CSP}(\mathbb{B})$ is either solvable in polynomial time or NP-complete [18].

In this paper, we consider the homomorphism problem for definable structures: given two definable structures \mathbb{A}, \mathbb{B} , does there exist a homomorphism from \mathbb{A} to \mathbb{B} ? Note that definable structures can be meaningfully considered as instances of a computational problem since they are finitely described with the set-builder notation and first-order formulas in the language of \mathcal{A} .

We remark that in the pure set \mathcal{A} with equality, every first-order formula is effectively equivalent to a quantifier-free formula. Thus, as long as complexity issues are ignored and decidability is the only concern, we can safely restrict to quantifier-free formulas.

► **Example 2.** The graph G from Example 1 does not map homomorphically to a clique of 3 vertices, which is another way of saying that G is not 3-colorable. In fact, G does not map homomorphically to any finite clique (the finite subgraph of G using only atoms $\underline{1}, \dots, \underline{2n}$ has chromatic number at least n , as it contains an n -clique). However, G maps homomorphically to the (easily definable) infinite clique on the set \mathcal{A} , by any injective mapping from V to \mathcal{A} . No such homomorphism is definable, as there is no definable injective function from V to \mathcal{A} , even with parameters. ◀

We consider several variants of the homomorphism problem:

- *Finite vs. infinite signature.* In the most general form, we allow the signature of both input structures to be infinite, but definable. In a restricted variant, the signature is required to be finite.
- *Finite vs. infinite structures.* In general, both input structures can be infinite, definable. In a restricted variant, one of the two input structures may be assumed to be finite.
- *Definability of homomorphisms.* In the general setting, we ask the question whether there exists an arbitrary homomorphism between the input structures. In other variants, the homomorphism is required to be definable, or to be \emptyset -definable.
- *Restrictions on homomorphisms.* Most often we ask about any homomorphism, but one may also ask about existence of a homomorphism that is injective, strong, or an embedding.

- *Fixing one structure.* In the uniform variant, both the source and the target structures are given on input. We also consider non-uniform variants, when one of the two structures is fixed.
- *Structured atoms.* In the basic setting, the underlying structure \mathcal{A} is the pure set, i.e., has no structure other than equality. One can also consider sets definable over other structures. For instance, if the underlying structure is (\mathbb{Q}, \leq) , the definitions of definable sets can refer to the relation \leq .

Contribution. For most combinations of these choices we determine the decidability status of the homomorphism problem. The decidability border turns out to be quite subtle and sometimes counterintuitive. The following theorem samples some of the opposing results proved in this paper:

► **Theorem 3.** *Let \mathcal{A} be the pure set. Given two definable structures \mathbb{A}, \mathbb{B} over a finite signature,*

1. *it is decidable whether there is a \emptyset -definable homomorphism from \mathbb{A} to \mathbb{B} ,*
2. *it is undecidable (but semidecidable) whether there is a definable homomorphism from \mathbb{A} to \mathbb{B} ,*
3. *it is decidable whether there is a homomorphism from \mathbb{A} to \mathbb{B} ,*
4. *it is undecidable (but co-semidecidable) whether a given \emptyset -definable partial mapping between the universes of \mathbb{A} and \mathbb{B} extends to a homomorphism.* ◀

In a previous paper [21], the constraint satisfaction problem is considered for source structures definable over the pure set, or more generally over (\mathbb{Q}, \leq) . We denote this problem by $\text{CSP}_{\text{def}}(\mathbb{B})$. The results from [21], together with the polynomial time reduction to the finite-template CSP which we provide here, imply complexity results for different variants of the constraint satisfaction problem:

► **Theorem 4.** *For any definable template \mathbb{B} over a finite signature:*

1. *the problem $\text{CSP}(\mathbb{B})$ is in NP,*
2. *the problem $\text{CSP}_{\text{def}}(\mathbb{B})$ is in NEXPTIME.* ◀

Related work. Some of the variants considered in this paper are closely related to previous work.

Bodirsky, Pinsker and coauthors [2, 8, 10] consider fixed infinite templates over finite signatures, and finite source structures given on input. They usually consider the template \mathbb{B} to be a reduct of a fixed structure \mathcal{A} with good properties, in particular, with a decidable first-order theory. Reducts are special cases of definable structures: a structure \mathbb{B} is a reduct of \mathcal{A} if \mathbb{B} is \emptyset -definable over \mathcal{A} and both have the same domains. In general, if the template \mathbb{B} is definable over a structure \mathcal{A} with decidable first-order theory, then \mathbb{B} itself has decidable first-order theory. It follows that the existence of a homomorphism from a given *finite* source structure \mathbb{A} is trivially decidable, as it can be expressed as an existential formula evaluated in \mathbb{B} . In this case, the interesting problem is to analyse precise complexity bounds. Templates for which a complete complexity classification was obtained (modulo the Feder-Vardi conjecture) include all reducts of countably infinite homogeneous graphs [3, 9, 12, 6], of (\mathbb{Q}, \leq) [4], and of the integers with the successor function $(\mathbb{Z}, +1)$ [5]. One of the key tools used in these results is the notion of a *canonical mapping*. The construction of a canonical mapping relies on Ramsey theory, most conveniently applied through the use of the result of KeCHRIS, Pestov, and Todorćević concerning extremely amenable groups [20].

For finite templates, it is shown in [21] that the complexity analysis of $\text{CSP}_{\text{def}}(\mathbb{B})$ can be reduced (with an exponential blowup) to the case of finite input structures. For example, 3-colorability of definable graphs is decidable and NEXPTIME-complete, because 3-colorability of finite graphs is NP-complete. A more general decidability result concerns *locally finite templates*, i.e., definable, possibly infinite templates (over definable, possibly infinite signatures) where every relation contains only finitely many tuples. The decidability proof also employs Ramsey theory, applied through the use of Pestov’s theorem concerning the topological dynamics of the group $\text{Aut}(\mathbb{Q}, \leq)$, which is a special case of the Kechris-Pestov-Todorcevic result. As we shall demonstrate here, for infinite signatures the local finiteness restriction is crucial and adding even a single infinite definable relation may lead to undecidability.

This paper, as well as [21], is part of a programme aimed at generalizing classical decision problems and computation models such as automata [15], Turing machines [16] and programming languages [14, 13, 22, 24], to sets with atoms. For other applications of sets with atoms (called there *nominal sets*) in computing, see [26].

Motivation. Testing existence of homomorphisms is at the core of many decision problems in combinatorics and logic. As shown in [11], decidability of pp-definability of a definable relation R in a definable structure \mathbb{A} can be reduced to deciding the existence of homomorphisms between definable structures. Another application is 0-1 laws, and deciding whether a sentence ϕ of the form $\exists R. \exists^* \forall^* \psi$ is satisfied with high probability in a finite random graph. In [23], after showing that the problem is equivalent to testing if ϕ holds in the infinite random graph, the authors give a complex Ramsey argument based on [25] to prove the decidability of the latter. The second step can be alternatively achieved by reducing to several instances of the homomorphism problem from structures definable over the ordered random graph (which is a Ramsey structure by [25], see Section 5) to finite target structures. Finally, in [21] the homomorphism problem for locally finite definable templates is used to test whether the logic IFP captures PTime over a certain class of finite structures, generalizing the Cai-Fürer-Immerman construction [17].

Acknowledgments. We are grateful to Albert Atserias, Manuel Bodirsky and Michael Pinsker for useful discussions.

2 Preliminaries

Throughout this section, fix a countable relational structure \mathcal{A} of *atoms*. We assume that the signature of \mathcal{A} is finite. We shall now introduce definable sets, following [21].

Definable sets. An *expression* is either a variable from some fixed infinite set, or a formal finite union (including the empty union \emptyset) of *set-builder expressions* of the form

$$\{ e \mid a_1, \dots, a_n \in \mathcal{A}, \phi \}, \tag{1}$$

where e is an expression, a_1, \dots, a_n are (bound) variables, and ϕ is a first-order formula over the signature of \mathcal{A} and over the set of variables. Free variables in (1) are those free variables of e and of ϕ which are not among a_1, \dots, a_n .

For an expression e with free variables V , any valuation $\text{val} : V \rightarrow \mathcal{A}$ defines in an obvious way a value $X = e[\text{val}]$, which is either an atom or a set, formally defined by induction on the structure of e . We then say that X is a *definable set with atoms*, and that it is *defined*

by e with val . Note that one set X can be defined by many different expressions. When we want to emphasize those atoms that are in the image of the valuation $\text{val} : V \rightarrow \mathcal{A}$, we say that the finite set $S = \text{val}(V) \subseteq \mathcal{A}$ *supports* X , or that X is *S -definable*.

As syntactic sugar, we allow atoms to occur directly in set expressions. For example, what we write as the $\{\underline{1}\}$ -definable set $\{a \mid a \in \mathcal{A}, a \neq \underline{1}\}$ is formally defined by the expression $\{a \mid a \in \mathcal{A}, a \neq b\}$, together with a valuation mapping b to $\underline{1}$. Similarly, the set $\{\underline{1}, \underline{2}\}$ is $\{\underline{1}, \underline{2}\}$ -definable as a union of two singleton sets.

► **Remark 5.** To improve readability, it will be convenient to use standard set-theoretic encodings to allow a more flexible syntax. In particular, ordered pairs and tuples can be encoded e.g. by Kuratowski pairs, $(x, y) = \{\{x, y\}, \{x\}\}$. We will also consider as definable infinite families of symbols, such as $\{R_x : x \in X\}$, where R is a symbol and X is a definable set. Formally, such a family can be encoded as the set of ordered pairs $\{R\} \times X$, where the symbol R is represented by some \emptyset -definable set, e.g. \emptyset or $\{\emptyset\}$, etc. Here we use the fact that definable sets are closed under Cartesian products.

Closure properties. The following lemma is proved routinely by induction on the nesting of set-builder expressions.

► **Lemma 6.** *Definable sets are effectively closed under:*

- Boolean combinations $\cap, \cup, -$ and Cartesian products,
- images and inverse images under definable functions,
- quotients under definable equivalence relations,
- intersections and unions of definable families,
- the operations (below, $x \in y$ and $x \subseteq y$ are interpreted as false if y is an atom):
 $V, W \mapsto \{(v, w) \mid v \in V, w \in W, v \in w\},$
 $V, W \mapsto \{(v, w) \mid v \in V, w \in W, v \subseteq w\}.$ ◀

This implies that the set-builder notation (1) may be safely generalized by allowing bound variables to range not only over \mathcal{A} but also over other definable sets, and allowing in ϕ quantifiers of the form $\exists v \in V$ or $\forall v \in V$, for V a definable set presented by an expression. One may also use binary predicates $=, \in, \subseteq$ and binary operations $\cup, \cap, -, \times$. The resulting sets will still be definable. As an example, if V and W are definable sets, then so is

$$\{(v, w) \mid v \in V, w \in W, v \subseteq w \wedge \exists a \in \mathcal{A} \exists b \in \mathcal{A} (a, b) \in v\}.$$

Suppose that the first-order theory of \mathcal{A} is decidable (this applies in particular to the pure set). Then it is straightforward to prove that the validity of first-order sentences generalized as above, such as $\forall v \in V \exists w \in W v \subseteq w$ where V and W are definable sets presented by expressions, is also decidable. This demonstrates that definable sets are suitable for effectively performing set-theoretic manipulations and tests.

Definable relational structures. For any object in the set-theoretic universe (a relation, a function, a logical structure, etc.), it makes sense to ask whether it is definable. For example, a definable relation on X, Y is a relation $R \subseteq X \times Y$ which is a definable set of pairs, and a definable function $X \rightarrow Y$ is a function whose graph is definable. Along the same lines, a definable relational signature is a definable set of *symbols* Σ , together with a partition $\Sigma = \Sigma_1 \uplus \Sigma_2 \uplus \dots \uplus \Sigma_l$ into definable subsets, for $l \in \mathbb{N}$. We say that σ has *arity* r if $\sigma \in \Sigma_r$.

For a signature Σ , a definable Σ -structure \mathbb{A} consists of a definable universe A and a definable interpretation function which assigns a relation $\sigma^{\mathbb{A}} \subseteq A^r$ to each relation symbol $\sigma \in \Sigma$ of arity r . (We denote structures using blackboard font, and their universes using

the corresponding symbol in italics). More explicitly, such a structure can be represented by the tuple $\mathbb{A} = (A, I_1, \dots, I_l)$ where $I_r = \{(\sigma, a_1, \dots, a_r) \mid \sigma \in \Sigma_r, (a_1, \dots, a_r) \in \sigma^{\mathbb{A}}\}$ is a definable set for $r = 1, \dots, l$ (where l is the maximal arity in Σ).

► **Remark 7.** A definable Σ -structure $\mathbb{A} = (A, I_1, \dots, I_l)$, for Σ finite or infinite, can be seen as a definable structure over a finite signature, denoted \mathbb{A}_Σ and defined as follows. The universe of \mathbb{A}_Σ is $A \uplus \Sigma$, and its relations are I_1, \dots, I_l , of arity $2, \dots, l+1$, respectively. The signature is finite, with just l symbols. Then homomorphisms between Σ -structures \mathbb{A} and \mathbb{B} correspond to those homomorphisms between \mathbb{A}_Σ and \mathbb{B}_Σ that are the identity on Σ . ◀

► **Example 8.** The graph G from Example 1 can be seen as a structure over a signature with a single binary relation symbol E . To give an example of an infinite, definable signature, extend G to a structure \mathbb{A} by infinitely many unary predicates representing the neighborhoods of each vertex of G . To this end, define the signature $\Sigma = \{E\} \cup \{N_v \mid v \in V\}$, where V is the vertex set of G and N is a symbol (cf. Remark 5). The interpretation of N_v is specified by the set $I_1 = \{(N_v, w) \mid (v, w) \in E\}$ (where E is defined by the expression from Example 1), which is definable by Lemma 6. ◀

► **Lemma 9.** *For every S -definable set X there is an S -definable surjective function $f : Y \rightarrow X$, where Y is an S -definable subset of \mathcal{A}^k , for some $k \in \mathbb{N}$. Moreover, f and Y can be computed from X .* ◀

► **Remark 10.** By Lemma 9, definable structures over finite signatures coincide, up to definable isomorphism, with structures that admit a *first-order interpretation with parameters* in \mathcal{A} , in the sense of model theory [19]. ◀

Representing the input. Definable relational structures can be input to algorithms, as they are finitely presented by expressions defining the signature, the universe, and the interpretation function. If the input is an S -definable set X , defined by an expression e with valuation $\text{val} : V \rightarrow S$ with $V = \{v_1, \dots, v_n\}$ the free variables of e , then we also need to represent the tuple $\text{val}(v_1), \dots, \text{val}(v_n)$ of elements of S . For the pure set \mathcal{A} , these elements can be represented as $\underline{1}, \underline{2}, \dots$

3 Homomorphism problems

To simplify the presentation, we now drop some of the generality of the previous section. In this section let \mathcal{A} be the pure set. In Section 5 we shall discuss generalizations of our results to underlying structures other than the pure set.

3.1 \emptyset -definable homomorphism problem

Let's start with the following warm-up decision problem:

Problem: \emptyset -DEFINABLE HOMOMORPHISM

Input: \emptyset -definable structures \mathbb{A} and \mathbb{B} over a \emptyset -definable signature Σ .

Decide: Is there an \emptyset -definable homomorphism from \mathbb{A} to \mathbb{B} ?

It is not hard to prove the following theorem, which gives (1) of Theorem 3:

► **Theorem 11** ([21]). \emptyset -DEFINABLE HOMOMORPHISM is decidable.

We sketch a proof here in order to illustrate the good algorithmic properties of definable sets, and to emphasize the contrast with later undecidability results.

Proof sketch. Our aim is to decide if two given \emptyset -definable Σ -structures $\mathbb{A} = (A, I_1, \dots, I_l)$ and $\mathbb{B} = (B, J_1, \dots, J_l)$ admit an \emptyset -definable homomorphism. The signature Σ is assumed to be part of the input (also, it can be computed from \mathbb{A} or from \mathbb{B}).

We will use the following facts that hold for the pure set \mathcal{A} , but also for many other structures with decidable first-order theories.

► **Lemma 12.** *For each number $n \in \mathbb{N}$, there are finitely (doubly exponentially) many first-order formulas with n free variables, up to equivalence in \mathcal{A} . Moreover, they can be computed from n .* ◀

The following lemma is a consequence.

► **Lemma 13.** *An \emptyset -definable set X has only finitely many \emptyset -definable subsets, and expressions defining these subsets can be enumerated from an expression defining X .*

Indeed, for each definable set X represented by a single set-builder expression of the form (1), replace ϕ by each (up to equivalence) quantifier-free formula ψ with the same free variables, such that $\psi \rightarrow \phi$.

To verify existence of an \emptyset -definable homomorphism from \mathbb{A} to \mathbb{B} , apply Lemma 13 to $X = A \times B$ and for every \emptyset -definable subset $R \subseteq A \times B$, test the validity of the first-order formula

$$\forall a \in A \exists! b \in B \ R(a, b)$$

ensuring that R is a graph of a function; and, for $i = 1 \dots l$, test the validity of the formula

$$\forall a_1, \dots, a_i \in A \forall b_1, \dots, b_i \in B \forall \rho \in \Sigma_i \bigwedge_{1 \leq j \leq i} R(a_j, b_j) \wedge I_i(\rho, a_1, \dots, a_i) \rightarrow J_i(\rho, b_1, \dots, b_i)$$

ensuring that the function is a homomorphism. ◀

In a similar vein one can decide the existence of homomorphisms that are injective, strong, or are embeddings (i.e. injective and strong), as all these properties are first-order definable.

The assumption that the structures \mathbb{A} and \mathbb{B} are \emptyset -definable is inessential in Theorem 11; the crucial assumption is that a homomorphism we ask for is required to be \emptyset -definable. In fact, a similar argument as above works even if the two given structures are definable instead of \emptyset -definable, and a homomorphism is allowed to be definable with n parameters, for a number $n \in \mathbb{N}$ given on input.

3.2 (Definable) homomorphism problem

In more relaxed versions of the homomorphism problem, we ask for a homomorphism that is definable without any bound on the number of parameters:

Problem: DEFINABLE HOMOMORPHISM

Input: Definable structures \mathbb{A} and \mathbb{B} over a definable signature Σ .

Decide: Is there a definable homomorphism from \mathbb{A} to \mathbb{B} ?

Or we may make no restriction on a homomorphism at all:

Problem: HOMOMORPHISM

Input: Definable structures \mathbb{A} and \mathbb{B} over a definable signature Σ .

Decide: Is there a homomorphism from \mathbb{A} to \mathbb{B} ?

These problems appear similar, but they are of rather different nature. On one hand, DEFINABLE HOMOMORPHISM is recursively enumerable, by an argument similar to the proof sketch of Theorem 11: if a definable homomorphism exists then one can find it by searching for homomorphisms definable with n parameters, for increasing values of n . On the other hand HOMOMORPHISM is co-recursively enumerable, by a compactness argument: if \mathbb{A} does not map homomorphically to \mathbb{B} then some finite substructure of \mathbb{A} does not map to \mathbb{B} either, and one can detect this by enumerating all finite substructures of \mathbb{A} and using Theorem 15 below.

► **Remark 14.** We might also consider natural variants of (DEFINABLE) HOMOMORPHISM, where one asks about existence of an injective homomorphism, or a strong homomorphism, or an embedding. Theorems 15–20, stated below, apply to all these variants as well. ◀

Below we show that both DEFINABLE HOMOMORPHISM and HOMOMORPHISM are undecidable in general. However, when one of the input structures has finite universe, both problems are decidable:

► **Theorem 15.** DEFINABLE HOMOMORPHISM and HOMOMORPHISM are decidable if one of the input structures has a finite universe. ◀

On the other hand, the general version of the homomorphism problem is undecidable:

► **Theorem 16.** HOMOMORPHISM is undecidable, even if one of the input structures is fixed. ◀

The fixed input structure is understood existentially; in particular, there exists a definable structure \mathbb{B} such that it is undecidable, for a given definable structure \mathbb{A} over the same signature, whether there is a homomorphism $\mathbb{A} \rightarrow \mathbb{B}$.

Theorem 16 is proved by a reduction from a classical quarter-plane tiling problem [1]. The following example illustrates a phenomenon used in the proof: a homomorphism can determine an infinite ordered sequence of atoms, and thus to enumerate coordinates within the quarter-plane.

► **Example 17.** Consider a signature with a single binary relation symbol R . For a chosen atom $a_0 \in \mathcal{A}$, define structures \mathbb{A} and \mathbb{B} over this signature as follows:

$$A = \mathcal{A} \quad R^{\mathbb{A}} = \neq \quad B = \mathcal{A} - \{a_0\} \quad R^{\mathbb{B}} = \neq$$

Note that \mathbb{A} is \emptyset -definable and \mathbb{B} is $\{a_0\}$ -definable. Considered as graphs, \mathbb{A} and \mathbb{B} are isomorphic to the countably infinite clique. However, no homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$ is definable. To see this, suppose towards contradiction that an S -definable homomorphism h actually exists for some finite S . We will exploit the fact that the S -definition of h is necessarily invariant under every bijection π of atoms such that $\pi(a) = a$ for all $a \in S$.

Since \mathbb{A} is a clique and \mathbb{B} has no self-loops, h must be injective. Pick the atom $a_1 = h(a_0)$. Clearly $a_1 \neq a_0$, since $a_0 \notin \mathbb{B}$. This means that $a_1 \in S$; indeed, if $a_1 \notin S$ then the S -definition of (the graph of) h would be invariant under a renaming π of atoms with $\pi(a_0) = a_0$ and $\pi(a_1) \neq a_1$, which cannot be since h is a function. Now consider $a_2 = h(a_1)$. Again, $a_2 \neq a_0$. Moreover we have $a_2 \neq a_1$, since $a_1 \neq a_0$ and h is injective. Moreover, $a_2 \in S$ by the same argument as for a_1 . This proceeds by induction, showing that infinitely many distinct atoms must belong to S , which contradicts the finiteness of S .

More importantly, each homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$ determines an infinite sequence of distinct atoms a_0, a_1, a_2, \dots such that $h(a_i) = a_{i+1}$ for each $i \in \mathbb{N}$. ◀

As it turns out, DEFINABLE HOMOMORPHISM is even harder to decide than HOMOMORPHISM:

► **Theorem 18.** DEFINABLE HOMOMORPHISM *is undecidable even if*

- (i) *a source structure \mathbb{A} over a finite signature is fixed; or*
- (ii) *a target structure \mathbb{B} is fixed.* ◀

Theorem 18 yields (2) of Theorem 3, and is proved by reduction from periodic and ultimately periodic variants of the tiling problem.

Example 17 shows a situation where definable homomorphisms do not exist, but non-definable ones do, and each of them induces an infinite sequence of atoms. In the following example definable homomorphisms do exist, and each of them determines a finite cycle of atoms. This observation is the core of the proof of Theorem 18, much as Example 17 is the core of Theorem 16.

► **Example 19.** Consider a signature with a single binary relation symbol R . Define structures \mathbb{A} and \mathbb{B} over this signature as follows (for readability we write ab to denote an ordered pair (a, b)):

$$\begin{aligned} A &= \mathcal{A} & B &= \{ab \mid a, b \in \mathcal{A}, a \neq b\} \\ R^{\mathbb{A}} &= \neq & R^{\mathbb{B}} &= \{(ab, cd) \mid a, b, c, d \in \mathcal{A}, a \neq b, c \neq d, a \neq c\} \end{aligned}$$

Note that there are many non-definable homomorphisms from \mathbb{A} to \mathbb{B} . For example, for any enumeration a_0, a_1, a_2, \dots of all atoms, one may put $h(a_n) = a_n a_{n+1}$ for each $n \in \mathbb{N}$.

However, definable homomorphisms $h : \mathbb{A} \rightarrow \mathbb{B}$ also exist. For example, there is an S -definable one for $S = \{\underline{1}, \underline{2}, \underline{3}\}$:

$$h(x) = x\underline{1} \quad h(\underline{1}) = \underline{1}\underline{2} \quad h(\underline{2}) = \underline{2}\underline{3} \quad h(\underline{3}) = \underline{3}\underline{1}$$

where $x \notin S$. Note how the values of h on S encode a cycle of atoms of length 3. This is a general phenomenon. Indeed, consider any S -definable homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$, for some finite $S = \{a_1, \dots, a_n\} \subseteq \mathcal{A}$. Denote $e_i = h(a_i)$ for $i = 1..n$. Each e_i is of the form $a_j a_k$ for some $1 \leq j \neq k \leq n$. Indeed, if some $e_i = bc$ (or $e_i = cb$) for some $b \notin S$, then the S -definition of (the graph of) h would be invariant under a renaming π of atoms with $\pi(a_i) = a_i$ and $\pi(b) \neq b$, which cannot be since h is a function.

One may view the e_i as edges of a directed graph with nodes $\{a_1, \dots, a_n\}$. This graph has n nodes, n edges, no self-loops, and, looking at the definition of $R^{\mathbb{B}}$, no two distinct edges have the same source. In other words, the graph is the graph of a function without fixpoints on $\{a_1, \dots, a_n\}$, therefore it contains a cycle of length at least 2. In other words, there is a subset of S of size at least 2 that is mapped to a set of the form $\{a_i a_j, a_j a_k, \dots, a_m a_i\}$. ◀

The two negative results in Theorems 16 and 18 are complemented by a positive one:

► **Theorem 20.** HOMOMORPHISM *is decidable for finite signatures.* ◀

This gives (3) of Theorem 3. Theorem 20 is implicit in the work of Bodirsky, Pinsker and Tsankov [11], where it is proved in a special case when $\mathbb{A} = \mathbb{B}^n$, for $n \geq 1$, and \mathbb{B} is a reduct of a finitely bounded Ramsey structure \mathcal{A} (cf. Section 5). Our self-contained proof of Theorem 20, given in Section 4, instead of using the machinery of *canonical mappings* goes by a direct reduction to the case when the target structure is finite, which is decidable as shown in [21]. Our reduction slightly generalizes a reduction due to Bodirsky and Mottet [7] in the special case of the target structure being a reduct of \mathcal{A} (both reductions need \mathcal{A} to be a finitely bounded homogeneous structure).

Theorems 11–20 settle the decidability landscape for the homomorphism problem almost entirely. One remaining open problem is the decidability status of DEFINABLE HOMOMORPHISM for a fixed target structure \mathbb{B} over a finite signature.

3.3 Homomorphism extension problem

Theorem 20 may be a little surprising in light of Theorem 16. Indeed, Remark 7 allows one to view an arbitrary definable Σ -structure as a definable structure \mathbb{A}_Σ over a finite signature. Homomorphisms $\mathbb{A} \rightarrow \mathbb{B}$ correspond to those homomorphisms $\mathbb{A}_\Sigma \rightarrow \mathbb{B}_\Sigma$ that are the identity on the subset Σ of the universe of \mathbb{A}_Σ . Thus by Theorem 16 we obtain undecidability, even for finite signatures, of the following slight generalization of HOMOMORPHISM:

Problem: HOMOMORPHISM EXTENSION

Input: Definable structures \mathbb{A} and \mathbb{B} over Σ and a definable partial mapping $f : A \rightarrow B$.

Decide: Is there a homomorphism from \mathbb{A} to \mathbb{B} extending f ?

The above remark proves (4) of Theorem 3:

► **Theorem 21.** HOMOMORPHISM EXTENSION is undecidable for finite signatures. ◀

4 Homomorphism problem for finite signatures

Throughout this section, we assume that Σ is a finite signature. For simplicity, assume that \mathcal{A} is the pure set; in Section 5 we discuss how the results generalize to other underlying structures. We consider the homomorphism problem for structures over Σ which are definable over \mathcal{A} . For simplicity, we assume that the input structures \mathbb{A} and \mathbb{B} are \emptyset -definable – the proof easily generalizes to arbitrary definable structures over a finite signature.

Here is the main result of this section:

► **Theorem 22.** Given a \emptyset -definable structure \mathbb{B} over a finite signature, one can compute a finite structure \mathbb{B}' such that:

- $\text{CSP}(\mathbb{B})$ is polynomial-time reducible to $\text{CSP}(\mathbb{B}')$,
- $\text{CSP}_{\text{def}}(\mathbb{B})$ is polynomial-time reducible to $\text{CSP}_{\text{def}}(\mathbb{B}')$.

Note that Theorem 22 implies Theorem 20, as finite structures \mathbb{B}' are a special case of *locally finite* ones, and decidability of the homomorphism problem for locally finite target structures has been shown in [21]. Moreover, Theorem 22 implies Theorem 4, since as shown in [21], for every finite template \mathbb{B}' , the complexity of $\text{CSP}_{\text{def}}(\mathbb{B}')$ is exponentially larger than the complexity of $\text{CSP}(\mathbb{B}')$.

Theorem 22 is a slight extension of results implicit in the work of Bodirsky, Pinsker and Tsankov [11] that provided a decision procedure for testing the existence of a homomorphism from $\mathbb{A} = \mathbb{B}^n$ to \mathbb{B} , where \mathbb{B} is assumed to be a reduct of \mathcal{A} (with further assumptions about \mathcal{A} , which apply also in our case, as discussed in Section 5). Instead, we allow both \mathbb{A} and \mathbb{B} to be arbitrary definable structures over \mathcal{A} (in particular, they need not be reducts). Our reduction in Theorem 22 is based on a reduction due to Bodirsky and Mottet [7], generalized to the case of definable structures rather than reducts (recall from Remark 10 that according to our definition, definable structures correspond to structures which interpret in \mathcal{A} via first-order interpretations).

Proof of Theorem 22. The remaining part of this section is devoted to demonstrating Theorem 22. In the sequel we fix a \emptyset -definable structure \mathbb{B} over the signature Σ (assumed to be finite). We show how to effectively construct a finite structure \mathbb{B}' as described in the theorem.

First we observe that without loss of generality we may assume that the universe B of \mathbb{B} is a subset of \mathcal{A}^k , for some k . To see this, apply Lemma 9 to obtain $\bar{B} \subseteq \mathcal{A}^k$ and a surjection $g : \bar{B} \rightarrow B$, both definable and computable from \mathbb{B} . Then compute a definable structure $\bar{\mathbb{B}}$ with universe \bar{B} over the same signature as \mathbb{B} , where every relation symbol is interpreted in $\bar{\mathbb{B}}$ as the inverse image under g of its interpretation in \mathbb{B} . Finally, observe that there is a homomorphism from \mathbb{A} to \mathbb{B} if, and only if there is a homomorphism from \mathbb{A} to $\bar{\mathbb{B}}$. Thus from now on we assume that $B \subseteq \mathcal{A}^k$, for some k .

We now define some notation. For $n \in \mathbb{N}$, denote $\{1, \dots, n\}$ by $[n]$. For a set C , numbers $m, n \in \mathbb{N}$ and an injective, monotone function $i : [m] \rightarrow [n]$, consider a projection mapping $\pi_i : C^m \rightarrow C^n$ onto m coordinates induced by i in the obvious way, i.e., $\pi_i(c_1, \dots, c_m) = (c_{i(1)}, \dots, c_{i(m)})$. Let \mathbb{C} be a structure with universe C and let $r \geq 2$ be an integer at least as large as the maximal arity of the relations in \mathbb{C} . We define a structure $\mathbb{C}^{\leq r}$ with universe $C^{\leq r} = C \cup C^2 \cup \dots \cup C^r$, as follows. If R is a relation symbol of arity k in the signature of \mathbb{C} , then the signature of $\mathbb{C}^{\leq r}$ contains a *unary* symbol U_R . If $S \subseteq C^k$ is the interpretation of R in \mathbb{C} , then the interpretation of U_R in $\mathbb{C}^{\leq r}$ is the set $S \subseteq C^k \subseteq C^{\leq r}$, treated as a unary relation. Moreover, for $m \leq n \leq r$ and each monotone injection $i : [m] \rightarrow [n]$, there is a binary *projection relation* $\Pi_i \subseteq C^m \times C^n$ in $\mathbb{C}^{\leq r}$ which is the graph of the projection π_i .

We use standard notions of group actions and orbits. The group $\text{Aut}(\mathcal{A})$ acts on \mathcal{A}^k , where an automorphism of \mathcal{A} acts coordinatewisely on elements of \mathcal{A}^k . Note that this action preserves $B \subseteq \mathcal{A}^k$, since B is \emptyset -definable. For the same reason, automorphisms of $\text{Aut}(\mathcal{A})$ preserve the relations of \mathbb{B} . Reassuming, $\text{Aut}(\mathcal{A})$ acts on the structure \mathbb{B} by automorphisms. Similarly, $\text{Aut}(\mathcal{A})$ acts on the structure $\mathbb{B}^{\leq r}$, inducing a quotient relational structure $\mathbb{B}^{\leq r}/\text{Aut}(\mathcal{A})$ over the same signature. The elements of $\mathbb{B}^{\leq r}/\text{Aut}(\mathcal{A})$ are orbits of $B^{\leq r}$ under the action of $\text{Aut}(\mathcal{A})$; in other words, elements of $\mathbb{B}^{\leq r}/\text{Aut}(\mathcal{A})$ are atomic types of k -tuples of atoms (an atomic type of a tuple of elements $(a_1, \dots, a_k) \in \mathcal{A}^k$ specifies all equalities among the elements a_1, \dots, a_k). Relation symbols are interpreted in $\mathbb{B}^{\leq r}/\text{Aut}(\mathcal{A})$ existentially, as expected. A crucial but obvious observation is that the quotient structure $\mathbb{B}^{\leq r}/\text{Aut}(\mathcal{A})$ is finite, by the following lemma.

► **Lemma 23.** *The group $\text{Aut}(\mathcal{A})$ acts oligomorphically on \mathbb{B} , i.e., the action splits B^n into finitely many orbits, for every $n \geq 1$.*

We now define the structure \mathbb{B}' promised in Theorem 22 as $\mathbb{B}^{\leq r}/\text{Aut}(\mathcal{A})$, where $r \geq 3$ is a fixed number at least as large as the maximal arity of the relations in Σ . As required, the structure \mathbb{B}' is finite. It remains to prove the two items of Theorem 22. Both reductions are shown in the same way. Let \mathbb{A} be given, where \mathbb{A} is either finite or \emptyset -definable. Define \mathbb{A}' as $\mathbb{A}^{\leq r}$. Note that if \mathbb{A} is \emptyset -definable, then so is \mathbb{A}' . Moreover, (the definition of) \mathbb{A}' is computable from \mathbb{A} in polynomial time, for a fixed signature Σ . To complete the reductions, it remains to prove the following:

► **Claim 24.** *There is a homomorphism $\mathbb{A} \rightarrow \mathbb{B}$ if, and only if, there is a homomorphism $\mathbb{A}' \rightarrow \mathbb{B}'$.*

The “only if” direction is immediate; from a given homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$, a homomorphism $h' : \mathbb{A}' \rightarrow \mathbb{B}'$ is obtained by taking the pointwise extension $h^{\leq r} : \mathbb{A}^{\leq r} \rightarrow \mathbb{B}^{\leq r}$ of h (also a homomorphism), and then post-composing $h^{\leq r}$ with the quotient homomorphism from $\mathbb{B}^{\leq r}$ to $\mathbb{B}^{\leq r}/\text{Aut}(\mathcal{A})$.

We now prove the “if” direction. Fix a homomorphism $f : \mathbb{A}^{\leq r} \rightarrow \mathbb{B}^{\leq r} / \text{Aut}(\mathcal{A})$. Recall that $B \subseteq \mathcal{A}^k$. Consider the set $D = (A \times \{1, \dots, k\}) / \sim$, where the equivalence relation \sim is defined as follows. Take $(a_1, \dots, a_n) \in A^n$, for $n \leq r$. Then $f(a_1, \dots, a_n) \in B^n / \text{Aut}(\mathcal{A})$ corresponds to an atomic type of $(n \cdot k)$ -tuples of atoms. In particular for $n = 2$, the atomic type concerns tuples $(x_1^1, \dots, x_1^k, x_2^1, \dots, x_2^k)$ and for each $1 \leq i, j \leq k$ and $1 \leq l, m \leq 2$, specifies a relation $x_l^i = x_m^j$ or $x_l^i \neq x_m^j$. Put $(a_1, i) \sim (a_2, j)$ in $A \times \{1, \dots, k\}$ if $(a_1, i) = (a_2, j)$ or the atomic type specifies the relation $x_1^i = x_2^j$. This defines an equivalence relation on $A \times \{1, \dots, k\}$, where $r \geq 3$ is essential for transitivity; it is also important here that f is a homomorphism and hence preserves projections. Since the set D is at most countable, there is an injective function $e : D \rightarrow \mathcal{A}$. We define a function $h : A \rightarrow \mathcal{A}^k$, by composing the abstraction function $[_]\sim : A \times \{1, \dots, k\} \rightarrow D$ with the function e :

$$h(a) = (e([(a, 1)]\sim), \dots, e([(a, k)]\sim)).$$

Note that $h(a) \in B$ for every $a \in A$. It follows by construction that the function $h : A \rightarrow B$ is a homomorphism from \mathbb{A} to \mathbb{B} . ◀

5 Concluding remarks

We investigated the homomorphism problem for definable relational structures. Our contribution is a detailed decidability border in the landscape of different variants of the problem.

Most of our proofs work, or can be easily adapted to the variant of the problem where one asks about the existence of an injective homomorphism, or a strong homomorphism, or an embedding. The only exceptions are Theorems 16 and 18 for the case where the target structure \mathbb{B} is fixed. Our proofs there work for the case of injective homomorphisms, but not for strong homomorphisms or embeddings, and the decidability of these cases remain open.

Underlying structure \mathcal{A} . We briefly describe the assumptions on the structure \mathcal{A} for which the results presented in this paper still hold.

The definitions and lemmas in Section 2 hold for an arbitrary structure \mathcal{A} . However, one needs to specify how inputs are represented, specifically, the parameters involved in the input. To represent all definable sets over \mathcal{A} , we should assume that there is an effective enumeration of its universe. Furthermore, to effectively perform tests on definable sets one needs to assume that the structure is *decidable*: given any first-order formula ϕ over the signature of \mathcal{A} with n free variables, and an n tuple \bar{a} of elements of \mathcal{A} , it must be decidable if $\phi, \bar{a} \models \mathcal{A}$. For simplicity we assume that the signature of \mathcal{A} is finite, to avoid questions concerning the encoding of relation symbols.

Theorems 16, 18 and 21 hold for every infinite structure \mathcal{A} . For Theorems 16 and 21 this is clear, as every structure definable over the pure set is also definable over arbitrary infinite \mathcal{A} , and existence of a homomorphism does not depend on \mathcal{A} . For Theorem 18 this is less clear, since the existence of definable homomorphisms depends on \mathcal{A} . However, an inspection of the proof shows that the result holds for arbitrary \mathcal{A} .

The \emptyset -definable homomorphism problem considered in Theorem 11 is decidable (with the same proof) as long as the following conditions hold:

- \mathcal{A} is ω -categorical, i.e., it is the only countable model of its first-order theory. An equivalent condition, due to the Ryll-Nardzewski-Engeler-Svenonius theorem [19], is that \mathcal{A} is countable and $\text{Aut}(\mathcal{A})$ acts oligomorphically on \mathcal{A} .
- The number of orbits of \mathcal{A}^n under the action of $\text{Aut}(\mathcal{A})$ is computable from a given $n \in \mathbb{N}$.

We call such structures *effectively ω -categorical*. Any effectively ω -categorical structure is (isomorphic to) a decidable structure, so every definable set can be represented. Theorem 11 can be easily generalized so that arbitrary definable structures \mathbb{A}, \mathbb{B} are given on input, as well as a finite set $S \subseteq \mathcal{A}$, and the algorithm determines whether there exists an S -definable homomorphism from \mathbb{A} to \mathbb{B} .

Regarding Theorem 15, in the case when the source structure \mathbb{A} is assumed to be finite, it is sufficient that \mathcal{A} is a decidable structure. In the case when the target structure \mathbb{B} is finite, and arbitrary homomorphisms are considered, the assumptions under which the proof from [21] works are that $\text{Aut}(\mathcal{A})$ is extremely amenable or, equivalently, that \mathcal{A} is a *Ramsey structure* [20]. Examples of Ramsey structures include (\mathbb{Q}, \leq) and the ordered random graph, by [25].

We do not know how to generalize to other atoms the case where only definable homomorphisms to a finite \mathbb{B} are considered.

Regarding Theorems 4, 20 and 22, the proofs presented in Section 4 work under the following assumptions:

- The structure \mathbb{A} is definable over a decidable Ramsey structure \mathcal{A} . For the second item of Theorem 4, we need some additional mild complexity assumptions about \mathcal{A} , e.g. that its first-order theory is decidable in NEXPTIME (for most reasonable structures it is in PSPACE). It is shown in [11] that if \mathcal{A} is a Ramsey structure, then extending \mathcal{A} by finitely many constants still yields a Ramsey structure. Clearly, this preserves decidability of the structure. From this it follows that the assumption made in Section 4 that the relations of \mathbb{A} and \mathbb{B} are \emptyset -definable is not relevant, since if they are S -definable over \mathcal{A} for some finite $S \subseteq \mathcal{A}$, then they are \emptyset -definable over \mathcal{A} extended by elements of S as constants.
- The structure \mathbb{B} is definable over a structure \mathcal{B} which is homogeneous and *finitely bounded*. We say that a structure \mathcal{B} over a signature Γ is finitely bounded if there is a finite set \mathcal{F} of finite Γ -structures such that for every finite Γ -structure \mathbb{A} , \mathbb{A} embeds into \mathcal{B} iff no structure from \mathcal{F} embeds into \mathbb{A} . For example, the pure set is finitely bounded, as witnessed by an empty family \mathcal{F} . This property is crucial for the proof of Claim 24. It is straightforward to generalize this claim to a finitely bounded homogeneous structure (see [11]). Any finitely bounded homogeneous structure is effectively ω -categorical, and thus decidable. Moreover, any expansion of a finitely bounded homogeneous structure by a constant is homogeneous and finitely bounded [7].

We do not know whether the finite boundedness condition can be dropped, while assuming that \mathcal{B} is effectively ω -categorical.

Open problems. Perhaps the most significant open question that remains is the decidability of the *isomorphism problem*: decide whether two definable structures \mathbb{A}, \mathbb{B} (say, over the pure set) are isomorphic, or whether there is a definable isomorphism between them. An equivalent formulation of the former question is the *orbit problem*: given a definable structure \mathbb{A} and two elements $x, y \in A$, decide whether there is an automorphism of \mathbb{A} which maps x to y .

This is related to an open problem from [11]: decide whether a given relation R is first-order definable in a given structure \mathbb{A} . Indeed, a unary predicate $R \subseteq A$ is first-order definable in \mathbb{A} iff it is preserved by all automorphisms of \mathbb{A} , iff no $x \in R$ and $y \in A - R$ lie in the same orbit of $\text{Aut}(\mathbb{A})$.

References

- 1 Robert Berger. The undecidability of the domino problem. *Mem. Amer. Math. Soc.*, 66, 1966.
- 2 Manuel Bodirsky. Complexity classification in infinite-domain constraint satisfaction. *Mémoire d'habilitation à diriger des recherches*, Université Diderot – Paris 7, Available at arXiv:1201.0856, 2012.
- 3 Manuel Bodirsky and Jan Kára. The complexity of equality constraint languages. *Theory Comput. Syst.*, 43(2):136–158, 2008.
- 4 Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *J. ACM*, 57(2), 2010.
- 5 Manuel Bodirsky, Barnaby Martin, and Antoine Mottet. Constraint satisfaction problems over the integers with successor. In *Procs. ICALP*, volume 9134 of *LNCS*, pages 256–267, 2015.
- 6 Manuel Bodirsky, Barnaby Martin, Michael Pinsker, and András Pongrácz. Constraint satisfaction problems for reducts of homogeneous graphs. In *Procs. of ICALP'16*, 2016. To appear.
- 7 Manuel Bodirsky and Antoine Mottet. Reducts of finitely bounded homogeneous structures, and lifting tractability from finite-domain constraint satisfaction. In *Procs. of LICS'16*, 2016. To appear.
- 8 Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *J. Log. Comput.*, 16(3):359–373, 2006.
- 9 Manuel Bodirsky and Michael Pinsker. Schaefer's theorem for graphs. In *Procs. STOC*, pages 655–664, 2011.
- 10 Manuel Bodirsky and Michael Pinsker. Topological Birkhoff. *Transactions of the AMS*, 367:2527–2549, 2015.
- 11 Manuel Bodirsky, Michael Pinsker, and Todor Tsankov. Decidability of definability. *The Journal of Symbolic Logic*, 78(04):1036–1054, 2013.
- 12 Manuel Bodirsky and Michal Wrona. Equivalence constraint satisfaction problems. In *Procs. CSL*, pages 122–136, 2012.
- 13 M. Bojańczyk and S. Toruńczyk. Imperative programming in sets with atoms. In *Procs. FSTTCS 2012*, volume 18 of *LIPIcs*, 2012.
- 14 Mikołaj Bojańczyk, Laurent Braud, Bartek Klin, and Sławomir Lasota. Towards nominal computation. In *Procs. POPL 2012*, pages 401–412, 2012.
- 15 Mikołaj Bojańczyk, Bartek Klin, and Sławomir Lasota. Automata theory in nominal sets. *Log. Meth. Comp. Sci.*, 10, 2014.
- 16 Mikołaj Bojańczyk, Bartek Klin, Sławomir Lasota, and Szymon Toruńczyk. Turing machines with atoms. In *LICS*, pages 183–192, 2013.
- 17 Jin-yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for graph identifications. *Combinatorica*, 12(4):389–410, 1992.
- 18 Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic snc and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, February 1999. URL: <http://dx.doi.org/10.1137/S0097539794266766>, doi:10.1137/S0097539794266766.
- 19 Wilfrid Hodges. *Model Theory*. Number 42 in *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1993.
- 20 A.S. Kechris, V.G. Pestov, and S. Todorćevic. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geometric & Functional Analysis GAFA*, 15(1):106–189, 2005.
- 21 Bartek Klin, Eryk Kopczyński, Joanna Ochremiak, and Szymon Toruńczyk. Locally finite constraint satisfaction problems. In *Procs. of LICS'15*, pages 475–486, 2015.

- 22 Bartek Klin and Michal Szynwelski. SMT solving for functional programming over infinite structures. In *Procs. MSFP'16*, pages 57–75, 2016.
- 23 P. Kolaitis and M. Vardi. The decision problem for the probabilities of higher-order properties. In *Procs. STOC'87*, pages 425–435, 1987.
- 24 Eryk Kopczyński and Szymon Toruńczyk. LOIS: an application of SMT solvers. In *Procs. SMT Workshop*, volume 1716 of *CEUR Proceedings*, pages 51–60, 2016.
- 25 Jaroslav Nešetřil and Vojtech Rödl. Partitions of finite relational and set systems. *Journal of Combinatorial Theory, Series A*, 22(3):289 – 312, 1977.
- 26 A. M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge University Press, 2013.