

Homomorphism Problems for First-Order Definable Structures

Bartek Klin Sławomir Lasota

University of Warsaw
{klin,sl}@mimuw.edu.pl

Joanna Ochremiak Szymon Toruńczyk

University of Warsaw &
Universitat Politècnica de Catalunya
ochremiak@cs.upc.edu szymtor@mimuw.edu.pl

Abstract

We investigate several variants of the homomorphism problem: given two relational structures, do they admit a homomorphism? The input structures are possibly infinite, but definable by first-order interpretations in a fixed structure. The 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.

Categories and Subject Descriptors F.4.1 [Mathematical logic and formal languages]: Mathematical logic—Model theory; F.4.3 [Mathematical logic and formal languages]: Formal languages—Decision problems

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

1. Introduction

First-order definable sets, although often 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 most often assume that the first-order theory of \mathcal{A} is decidable. For simplicity, unless stated otherwise, let \mathcal{A} be a countable set $\{\underline{1}, \underline{2}, \underline{3}, \dots\}$ of atoms equipped with the equality relation only.

Example 1. Let

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

where ϕ is the quantifier-free formula expressing that a, b, c, d are pairwise distinct. 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 also allow finite unions in the definitions; 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 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 $S \subseteq \mathcal{A}$.

Although definable sets correspond to first-order interpretations well-known from logic and model theory, we prefer to use a different definition, as 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. Note that definable sets strictly generalize hereditarily finite sets (finite sets, whose elements are finite, and so on, recursively).

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 have finite descriptions, using set-builder notation and first-order formulas in the language of \mathcal{A} , and possibly finitely many parameters from \mathcal{A} . We remark that in the pure set, every first-order formula is effectively equivalent to a quantifier-free formula, so over this underlying structure, it is sufficient to consider such formulas, if we ignore complexity issues and care only for decidability.

Example 2. The graph G from Example 1 does not map homomorphically to the three-clique, which corresponds to the fact that G is not three-colorable. In fact, G does not map homomorphically to any finite clique (the finite subgraph of G using only atoms $\underline{1}, \dots, \underline{n}$ has chromatic number $n - 2$, by Lovasz' theorem [21]). However, G maps homomorphically to the infinite clique K on \mathcal{A} (which is a definable graph), by any injective mapping from V to \mathcal{A} . Note that no such homomorphism is definable, as there is no definable 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 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 other variants, one of the two input structures may be required 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 which are 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 .

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

Related work

Some of the variants considered in this paper are closely related to, or have been considered in previous work.

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 widely studied in the literature in many variants. The case when the target structure is fixed (and is called a *template*), is of particular interest to theoretical computer science, as it expresses many natural computational problems (such as k -colorability, 3-SAT, solving systems of linear equations over a finite field). The famous Feder-Vardi conjecture states that for every fixed template \mathbb{B} , the resulting constraint satisfaction problem is either solvable in polynomial time or NP-complete [15].

Bodirsky, Pinsker and coauthors [2, 6, 8] consider fixed infinite templates over finite signatures, and finite source structures given on input. Moreover, 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 the precise complexity bounds. Templates for which a complete complexity classification was obtained (modulo the Feder-Vardi conjecture) include all reducts of $(\mathbb{N}, =)$ [3], of

(\mathbb{Q}, \leq) [4], of the Rado graph [7], 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, which we use here. The construction of a canonical mapping relies on Ramsey theory, most conveniently applied through the use of the result of Kechris, Pestov, Todorćević concerning extremely amenable groups [18].

One of the results (cf. (3) of Theorem 3) stated in this paper says that the homomorphism problem is decidable in the case when both \mathbb{A} and \mathbb{B} are over a finite signature and are definable over the pure set or over (\mathbb{Q}, \leq) . The result is implicit in, and can be easily derived from [9]. Canonical functions play a crucial role in the proof. We explain this in greater detail in Section 6.

In a previous paper [19], the source structure is considered definable over the pure set, or more generally over (\mathbb{Q}, \leq) . In the case when the target structure is fixed and finite, it is shown that the complexity analysis can be reduced (with an exponential blowup) to the case of finite input structures. A more general decidability result concerns *locally finite templates*, i.e., templates which are definable, but in which every relation is required to contain only finitely many tuples. The decidability proof also employs Ramsey's theorem, most conveniently 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-Todorćević result. As we shall demonstrate, the local finiteness restriction is crucial and adding a single infinite definable relation leads to undecidability.

This paper, as well as [19], is part of a programme aimed at generalizing classical decision problems and computation models such as automata [11], Turing machines [12] and programming languages [10, 13], to sets with atoms. For other applications of sets with atoms (called there *nominal sets*) in computing, see [25].

Motivation

Testing existence of homomorphisms is at the core of many decidability problems related to combinatorics and logic. As shown in [9], 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 to 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 [20], 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 [22] 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 [22], see Section 7) to finite target structures. Finally in [19] 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.

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 introduce definable sets as in [19], and discuss their relevant properties.

2.1 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

$$\{e \mid a_1, \dots, a_n \in \mathcal{A}, \phi\}, \quad (1)$$

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 used in a definition of X , 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 4. 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. 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 (see below).

2.2 Closure properties

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

Lemma 5. *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

$$V, W \mapsto \{(v, w) \mid v \in V, w \in W, v \in w\},$$

$$V, W \mapsto \{(v, w) \mid v \in V, w \in W, v = w\},$$

$$V, W \mapsto \{(v, w) \mid v \in V, w \in W, v \subseteq w\}$$
 ($x \in y$ and $x \subseteq y$ are interpreted as false if y is an atom).

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 $=, \in, \subseteq$, relation symbols from the signature of \mathcal{A} , 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 \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 given by expressions, is also decidable. This demonstrates that definable sets are suitable for effectively performing set-theoretic manipulations and tests.

2.3 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 6. 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 . \square

Example 7. The graph G from Example 1 can be seen as a structure over the signature with a single binary predicate 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 4). 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 5.

Lemma 8. *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 .*

Proof: see Appendix A.1. \square

Remark 9. Definable structures over finite signatures correspond, up to definable isomorphism, to structures that *interpret with parameters* in \mathcal{A} , in the sense of model theory [17]. This can be deduced from Lemma 8 (see Appendix A.2).

2.4 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 $\underline{1}, \underline{2}, \dots$

In all decision problems defined in this paper, the input will consist of definable structures as described above.

3. Homomorphism problems

For simplicity, we prefer to drop the generality of the previous section. From now on, until the end of Section 5, let \mathcal{A} be the pure set. We postpone to Section 7 a discussion on 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 Σ .

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

It is not hard to prove that:

Theorem 10. \emptyset -DEFINABLE HOMOMORPHISM is decidable.

This implies (1) of Theorem 3. The proof is given e.g. in [19]. We sketch it here in order to illustrate 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 11. *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} .*

The following lemma is a consequence.

Lemma 12. *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 and such that $\psi \rightarrow \phi$, i.e., $\psi \vee \phi$ is equivalent to ϕ .

To find an \emptyset -definable homomorphism from \mathbb{A} to \mathbb{B} , apply Lemma 12 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 of the formula

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

ensuring that the function is a homomorphism. \square

In a similar vein one can decide the existence of homomorphisms that are injective, strong, or are embeddings, 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 10; 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 given $n \in \mathbb{N}$.

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 size of support:

Problem: DEFINABLE HOMOMORPHISM

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

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

Note the different nature of the two problems. On one hand, DEFINABLE HOMOMORPHISM is recursively enumerable, by an argument similar to the proof sketch of Theorem 10; on the other side HOMOMORPHISM is co-recursively enumerable, by a compactness argument. (Indeed, if every finite substructure of \mathbb{A} maps to \mathbb{B} , then \mathbb{A} maps to \mathbb{B} .)

Remark 13. 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. The decidability status of all these problems is the same. Theorems 14–17, to be stated below, apply to all these variants. \square

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

Theorem 14. DEFINABLE HOMOMORPHISM and HOMOMORPHISM are decidable if one of the input structures has a finite universe.

Proof: see Appendix B.1. \square

In contrast with Theorems 10 and 14, our first main result is that the general version of the homomorphism problem is undecidable:

Theorem 15. 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 15 is proved in Section 4.

More surprisingly, it turns out that DEFINABLE HOMOMORPHISM is even harder to decide:

Theorem 16. 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 16 is proved in Section 5, proving (2) of Theorem 3.

These two negative results are complemented by a positive one:

Theorem 17. HOMOMORPHISM is decidable for finite signatures.

Theorem 17 is implicit in [9], and is proved (although not stated explicitly) there in a special case when $\mathbb{A} = \mathbb{B}^n$, for $n \geq 1$, where \mathbb{B}^n denotes the product structure, and \mathbb{A} is a reduct of a finitely bounded Ramsey structure \mathcal{A} (see Section 7). For completeness, we give a self-contained proof of Theorem 17 in Section 6. This gives (3) of Theorem 3.

Theorems 10–17 settle the decidability landscape for 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. We discuss this and other minor remaining problems in Sections 4 and 5.

3.3 Homomorphism extension problem

Theorem 17 may be a little surprising in light of Theorem 15. Indeed, Remark 6 allows viewing 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 15 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 18. HOMOMORPHISM EXTENSION is undecidable for finite signatures.

4. Homomorphism problem for infinite signatures

This section contains the proof of Theorem 15. We consider first the case when the source structure is fixed, and the target structure is the sole input. Thus, given a definable structure \mathbb{A} over a definable signature Σ , we consider the following problem:

Problem: $\text{HOM}(\mathbb{A}, -)$

Input: A definable structure \mathbb{B} over Σ

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

Example 19. 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:

$$\begin{aligned} A &= \mathcal{A} & B &= \mathcal{A} - \{a_0\} \\ R^{\mathbb{A}} &= \neq & R^{\mathbb{B}} &= \neq \end{aligned}$$

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 . 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 $i \in \mathbb{N}$. \square

Remark 20. An effect similar to Example 19 can be observed for \emptyset -definable \mathbb{A} and \mathbb{B} . Consider

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

(Here ab is simplified syntax for (a, b) .) Note that for any homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$, all atoms in \mathbb{A} are mapped to pairs that share the same first component (call it a_0). A reasoning similar to Example 19 shows that this determines a sequence of distinct atoms a_0, a_1, a_2, \dots such that $h(a_i) = a_0 a_{i+1}$ for each $i \in \mathbb{N}$.

Theorem 21. *There exists a \emptyset -definable structure \mathbb{A} for which the problem $\text{HOM}(\mathbb{A}, -)$ is undecidable.*

Proof. We reduce a quarter-plane tiling problem defined as follows. For a finite set $\mathcal{K} \ni K, L, \dots$ of colors, and for relations $\Gamma_H, \Gamma_V \subseteq \mathcal{K} \times \mathcal{K}$, a quarter-plane tiling is a function $\gamma : \mathbb{N}^2 \rightarrow \mathcal{K}$ such that

$$(\gamma(i, j), \gamma(i+1, j)) \in \Gamma_H \quad \text{and} \quad (\gamma(i, j), \gamma(i, j+1)) \in \Gamma_V$$

for $i, j \in \mathbb{N}$. By a well-known result of Berger [1], it is undecidable whether there exists a quarter-plane tiling for given \mathcal{K}, Γ_H and Γ_V .

Consider the (infinite but definable) signature Σ with:

- a unary predicate symbol P_a for each $a \in \mathcal{A}$, and
- binary relation symbols Π_1, Π_2, R and T .

Define a structure \mathbb{A} over Σ by:

$$\begin{aligned} A &= \mathcal{A} \cup \mathcal{A}^2 & P_a^{\mathbb{A}} &= \{a\} \quad \text{for } a \in \mathcal{A} \\ \Pi_1^{\mathbb{A}} &= \{(a, b), a \mid a, b \in \mathcal{A}\} & \Pi_2^{\mathbb{A}} &= \{(a, b), b \mid a, b \in \mathcal{A}\} \\ R^{\mathbb{A}} &= \{(a, b) \mid a, b \in \mathcal{A}, a \neq b\} & T^{\mathbb{A}} &= \mathcal{A}^2 \times \mathcal{A}^2 \end{aligned}$$

Note that $R^{\mathbb{A}}$ relates only atoms, $T^{\mathbb{A}}$ relates only (and all) pairs of atoms, and $\Pi_1^{\mathbb{A}}, \Pi_2^{\mathbb{A}}$ relate pairs of atoms to their components. Clearly, \mathbb{A} is \emptyset -definable.

Fix an atom $a_0 \in \mathcal{A}$. Denote

$$B_0 = \{ab \mid a, b \in \mathcal{A}, b \neq a_0\}.$$

Note that the set B_0 is $\{a_0\}$ -definable. Elements of B_0 are pairs of atoms, but we write ab instead of (a, b) , to distinguish them from pairs of atoms used in \mathbb{A} . The two kinds of pairs will serve different purposes in the encoding of the quarter-plane tiling problem. Intuitively, a pair (a, b) in \mathbb{A} will encode a point in the quarter-plane with coordinates a and b , while a pair ab in B_0 will model the fact that b encodes the successor of a in both axes of the quarter-plane.

Formally, given an instance \mathcal{K}, Γ_H and Γ_V of the quarter-plane tiling problem, define a $\{a_0\}$ -definable Σ -structure \mathbb{B} :

$$\begin{aligned} B &= B_0 \cup (B_0^2 \times \mathcal{K}) \\ P_a^{\mathbb{B}} &= \{ab \in B_0 \mid b \in \mathcal{A}\} \quad \text{for } a \in \mathcal{A} \\ \Pi_1^{\mathbb{B}} &= \{((ab, cd, K), ab) \mid ab, cd \in B_0, K \in \mathcal{K}\} \\ \Pi_2^{\mathbb{B}} &= \{((ab, cd, K), cd) \mid ab, cd \in B_0, K \in \mathcal{K}\} \\ R^{\mathbb{B}} &= \{(ab, cd) \mid ab, cd \in B_0, b \neq d\} \\ T^{\mathbb{B}} &= \{((ab, cd, K), (ef, gh, L)) \mid \\ &\quad ((b = e) \wedge (c = g) \rightarrow (K, L) \in \Gamma_H) \\ &\quad \wedge ((a = e) \wedge (d = g) \rightarrow (K, L) \in \Gamma_V)\} \end{aligned}$$

We shall now prove that \mathcal{K}, Γ_H and Γ_V admit a quarter-plane tiling if and only if there is a homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$.

For one direction, assume a quarter-plane tiling $\gamma : \mathbb{N}^2 \rightarrow \mathcal{K}$. Consider any enumeration of all atoms a_0, a_1, a_2, \dots with a_0 as the first element. Define $h : \mathbb{A} \rightarrow \mathbb{B}$ by:

$$\begin{aligned} h(a_i) &= a_i a_{i+1} \\ h(a_i, a_j) &= (a_i a_{i+1}, a_j a_{j+1}, \gamma(i, j)). \end{aligned}$$

It is easy to check that h is a homomorphism. Indeed, Π_1, Π_2 and all predicates P_a are preserved immediately. So is R , since $a_i \neq a_j$ implies $a_{i+1} \neq a_{j+1}$. For T to be preserved, for any $(a_i, a_j), (a_k, a_l) \in \mathcal{A}^2$, we need to check that

$$((a_i a_{i+1}, a_j a_{j+1}, \gamma(i, j)), (a_k a_{k+1}, a_l a_{l+1}, \gamma(k, l))) \in T^{\mathbb{B}}.$$

If $k = i + 1$ and $l = j$ then $(\gamma(i, j), \gamma(k, l)) \in \Gamma_H$ since γ is a tiling. If $k = i$ and $l = j + 1$ then $(\gamma(i, j), \gamma(k, l)) \in \Gamma_V$, for the same reason. In all other cases the condition holds trivially, by definition of $T^{\mathbb{B}}$.

For the other direction, consider any homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$. Interpretations of the predicates P_a in \mathbb{A} and \mathbb{B} ensure that for each $a \in \mathcal{A}$, necessarily $h(a) = ab$ for some $b \neq a_0$. Moreover, by the interpretations of Π_1 and Π_2 , for each $a, b \in \mathcal{A}$

$$h(a, b) = (h(a), h(b), K)$$

for some $K \in \mathcal{K}$.

Consider Σ, \mathbb{A} and \mathbb{B} restricted to the relation symbol R . The above implies that h restricts to a homomorphism from \mathcal{A} to B_0 that always returns its argument on the first component. This is essentially the same situation as in Example 19, and for reasons explained there, there must be an infinite sequence of distinct atoms a_0, a_1, a_2, \dots such that $h(a_i) = a_i a_{i+1}$ for each $i \in \mathbb{N}$.

Define $\gamma : \mathbb{N}^2 \rightarrow \mathcal{K}$ so that $\gamma(i, j)$ is the color K such that

$$h(a_i, a_j) = (a_i a_{i+1}, a_j a_{j+1}, K).$$

This is a quarter-plane tiling. Indeed, since

$$((a_i, a_j), (a_{i+1}, a_j)) \in T^{\mathbb{A}}$$

then necessarily

$((a_i a_{i+1}, a_j a_{j+1}, \gamma(i, j)), (a_{i+1} a_{i+2}, a_j a_{j+1}, \gamma(i+1, j))) \in T^{\mathbb{B}}$
 which, by definition of $T^{\mathbb{B}}$, implies that $(\gamma(i, j), \gamma(i+1, j)) \in \Gamma_H$.
 The condition for Γ_V follows analogously. \square

Remark 22. The problem $\text{HOM}(\mathbb{A}, -)$, for \mathbb{A} as in the proof of Theorem 21, remains undecidable even if one restricts input structures \mathbb{B} to be \emptyset -definable. To see this, Remark 20 is useful. Technically, in the proof above one replaces B_0 with

$$B_0^e = \{abc \mid a, b, c \in \mathcal{A}, a \neq c\},$$

redefines $P_a^{\mathbb{B}}, \Pi_1^{\mathbb{B}}, \Pi_2^{\mathbb{B}}$ and $T^{\mathbb{B}}$ so that they ignore the first components of triples from B_0^e , and changes $R^{\mathbb{B}}$ so that it only relates triples with the same first component:

$$R^{\mathbb{B}} = \{(abc, ade) \mid abc, ade \in B_0^e, c \neq e\}$$

The resulting structure \mathbb{B} is \emptyset -definable and, using Remark 20 instead of Example 19, the proof of Theorem 21 works analogously. \square

Another variant of the homomorphism problem keeps a target structure \mathbb{B} fixed, and treats the source structure as input:

Problem: $\text{HOM}(-, \mathbb{B})$

Input: A definable structure \mathbb{A} over Σ

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

It easily follows from Theorem 21 that this problem cannot be solvable in any practical sense: even if $\text{HOM}(-, \mathbb{B})$ were decidable for every \mathbb{B} , there could be no way to compute an algorithm to solve this problem, given a description of \mathbb{B} . In fact, a stronger negative result holds:

Theorem 23. *There exists a definable structure \mathbb{B} for which the problem $\text{HOM}(-, \mathbb{B})$ is undecidable.*

Proof. We proceed much as in the proof of Theorem 21, by a reduction from a seeded version of the quarter-plane tiling problem defined as follows. Given \mathcal{K}, Γ_H and Γ_V , for a finite sequence of colors $K_0, K_1, \dots, K_n \in \mathcal{K}$ (a seed), a legal tiling $\gamma : \mathbb{N}^2 \rightarrow \mathcal{K}$ is seeded if $\gamma(i, 0) = K_i$ for every $i \in \{0, 1, \dots, n\}$.

It is easy to see that there exist fixed \mathcal{K}, Γ_H and Γ_V such that it is undecidable whether a given seed admits a seeded tiling. Indeed, in Wang's proof of undecidability of the constrained tiling problem (see e.g. [14, App. A]), where tile sets encode Turing machines, it is enough to consider a set that encodes a universal Turing machine, and represent an input word for the machine as the seed.

Fix some atom $a_0 \in \mathcal{A}$ and consider \mathbb{B} defined as in the proof of Theorem 21, for the specific \mathcal{K}, Γ_H and Γ_V as above. Additionally, extend the signature with an infinite family of predicate symbols $\{Q_a \mid a \in \mathcal{A}\}$, and a finite family of predicate symbols $\{O_K \mid K \in \mathcal{K}\}$. Interpret these in \mathbb{B} as:

$$Q_a^{\mathbb{B}} = \{ba \in B_0 \mid b \in \mathcal{A}\} \quad \text{for } a \in \mathcal{A},$$

$$O_K^{\mathbb{B}} = \{(ab, cd, K) \mid ab, cd \in B_0\} \quad \text{for } K \in \mathcal{K}.$$

(In particular, $Q_{a_0}^{\mathbb{B}} = \emptyset$.) The structure \mathbb{B} is $\{a_0\}$ -definable.

Given a seed K_0, K_1, \dots, K_n , consider a structure \mathbb{A} over the extended signature as in the proof of Theorem 21. Pick any $n+2$ distinct atoms a_0, a_1, \dots, a_{n+1} starting with a_0 . Extend \mathbb{A} by:

$$Q_{a_{i+1}}^{\mathbb{A}} = \{a_i\} \quad \text{for } 0 \leq i \leq n,$$

$$Q_a^{\mathbb{A}} = \emptyset \quad \text{for } a \notin \{a_1, \dots, a_{n+1}\},$$

$$O_K^{\mathbb{A}} = \{(a_i, a_0) \mid i \leq n, K_i = K\} \quad \text{for } K \in \mathcal{K}.$$

The structure \mathbb{A} is $\{a_0, a_1, \dots, a_{n+1}\}$ -definable.

Homomorphisms from \mathbb{A} to \mathbb{B} then correspond to quarter-plane tilings for \mathcal{K}, Γ_H and Γ_V seeded by K_0, \dots, K_n . To see this, proceed as in the proof of Theorem 21, but note additionally that due to the interpretation of the P_a and Q_a in \mathbb{A} and \mathbb{B} , for any $h : \mathbb{A} \rightarrow \mathbb{B}$ there is holds that $h(a_i) = a_i a_{i+1}$ for $0 \leq i \leq n$. In other words, the infinite sequence of atoms determined by h as in the proof of Theorem 21, must begin with a_0, a_1, \dots, a_{n+1} . Finally, by the interpretation of O_K in \mathbb{A} and \mathbb{B} , the tiling γ derived from h satisfies $\gamma(i, 0) = K_i$ for $0 \leq i \leq n$, as requested. \square

Remark 24. The structure \mathbb{B} in Theorem 23 can be made \emptyset -definable, using the technique of Remarks 20 and 22. However, nonempty support of input structures \mathbb{A} used in the proof of Theorem 23 seems harder to avoid, as in the reduction, its size is unbounded. We leave open the question whether there exists a \mathbb{B} for which $\text{HOM}(-, \mathbb{B})$ is undecidable when restricted to \emptyset -definable input structures.

5. Definable homomorphism problem

This section contains the proof of Theorem 16. First, given a definable structure \mathbb{A} over a finite signature Σ , consider the problem:

Problem: $\text{DEF-HOM}(\mathbb{A}, -)$

Input: A definable structure \mathbb{B} over Σ

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

Example 19 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.

Example 25. Consider a signature with a single binary relation symbol R . Define structures \mathbb{A} and \mathbb{B} over this signature as follows:

$$A = \mathcal{A} \quad B = \{ab \mid a, b \in \mathcal{A}, a \neq b\}$$

$$R^{\mathbb{A}} \neq \emptyset \quad R^{\mathbb{B}} = \{(ab, cd) \mid a, b, c, d \in \mathcal{A}, a \neq b, c \neq d, a \neq c\}$$

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 atoms, one may put $h(a_n) = a_n a_{n+1}$ for each $n \in \mathbb{N}$.

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

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

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 $e_i = bc$ 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 $\{a_i a_j, a_j a_k, \dots, a_m a_i\}$.

This observation will be the core of the following undecidability theorem, much as Example 19 was the core of Theorem 21.

Theorem 26. *There exists an \emptyset -definable structure \mathbb{A} over a finite signature for which the problem $\text{DEF-HOM}(\mathbb{A}, -)$ is undecidable.*

Proof. The proof is similar to that of Theorem 21, with Example 25 replacing Example 19 as the core source of undecidability.

We reduce a periodic tiling problem defined as follows. For a finite set $\mathcal{K} \ni K, L, \dots$ of colors and relations $\Gamma_H, \Gamma_V \subseteq \mathcal{K} \times \mathcal{K}$, we say that a tiling $\gamma : \mathbb{N}^2 \rightarrow \mathcal{K}$ is *periodic* if there is a number $n \geq 1$ such that $\gamma(i, j) = \gamma(i+n, j) = \gamma(i, j+n)$ for $i, j \in \mathbb{N}$. It is well known [16] that it is undecidable whether a periodic tiling exists for given \mathcal{K}, Γ_H and Γ_V .

Consider a signature Σ with four binary relation symbols Π_1, Π_2, R and T . Define a structure \mathbb{A} over Σ as in the proof of Theorem 21, minus the interpretation of predicates P_a , which are now absent from the signature.

Given an instance \mathcal{K}, Γ_H and Γ_V of the periodic tiling problem, define a Σ -structure \mathbb{B} by:

$$\begin{aligned} B &= \{ab \mid a \neq b \in \mathcal{A}\} \\ &\cup \{(ab, cd, K) \mid a \neq b, c \neq d \in \mathcal{A}, K \in \mathcal{K}\} \\ \Pi_1^{\mathbb{B}} &= \{((ab, cd, K), ab) \mid a \neq b, c \neq d \in \mathcal{A}, K \in \mathcal{K}\} \\ \Pi_2^{\mathbb{B}} &= \{((ab, cd, K), cd) \mid a \neq b, c \neq d \in \mathcal{A}, K \in \mathcal{K}\} \\ R^{\mathbb{B}} &= \{(ab, cd) \mid a \neq b, c \neq d, a \neq c \in \mathcal{A}\} \\ T^{\mathbb{B}} &= \{((ab, cd, K), (ef, gh, L)) \mid \\ &\quad (e = b \wedge d = h \implies (K, L) \in \Gamma_H) \\ &\quad \wedge (d = g \wedge b = f \implies (K, L) \in \Gamma_V)\} \end{aligned}$$

We shall now prove that \mathcal{K}, Γ_H and Γ_V admit a periodic tiling if and only if there is a definable homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$.

For the ‘‘if’’ part, consider any S -definable homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$, for a finite set $S \subseteq \mathcal{A}$. Interpretations of Π_1 and Π_2 in \mathbb{A} and \mathbb{B} ensure that for each $a \in \mathcal{A}$, necessarily $h(a) = bc$ for some $b \neq c \in \mathcal{A}$. Moreover, for each $a, b \in \mathcal{A}$, there is $h(a, b) = (h(a), h(b), K)$ for some $K \in \mathcal{K}$.

Consider Σ, \mathbb{A} and \mathbb{B} restricted to the relation symbol R . The above implies that h restricts to an S -definable homomorphism from \mathcal{A} to $\{ab \mid a \neq b \in \mathcal{A}\}$. This is essentially as in Example 25, and for reasons explained there, there must be a sequence $(a_0, a_1, \dots, a_{n-1})$ of atoms from S , with $2 \leq n \leq |S|$, such that all pairs $a_0a_1, a_1a_2, \dots, a_{n-2}a_{n-1}, a_{n-1}a_0$ are values of h on some atoms from S . Denote those atoms $b_0, \dots, b_{n-1} \in S$, so that

$$h(b_0) = a_0a_1, h(b_1) = a_1a_2, \dots, h(b_{n-1}) = a_{n-1}a_0.$$

Note that we make no claims as to whether some b_i are equal to a_j , and to which ones. This is irrelevant for the following.

For $j \geq n$, define $a_j = a_i$, where i is the residue of j modulo n . Define $\gamma : \mathbb{N}^2 \rightarrow \mathcal{K}$ so that $\gamma(i, j)$ is the color K such that

$$h(b_i, b_j) = (a_i a_{i+1}, a_j a_{j+1}, K).$$

This is a legal periodic tiling. Indeed, since $T^{\mathbb{A}}$ is the full relation on pairs of atoms, for each $i, j \in \mathbb{N}$ we must have

$$\begin{aligned} ((a_i a_{i+1}, a_j a_{j+1}, \gamma(i, j)), (a_{i+1} a_{i+2}, a_j a_{j+1}, \gamma(i+1, j))) &\in T^{\mathbb{B}} \\ ((a_i a_{i+1}, a_j a_{j+1}, \gamma(i, j)), (a_i a_{i+1}, a_{j+1} a_{j+2}, \gamma(i, j+1))) &\in T^{\mathbb{B}} \end{aligned}$$

hence $(\gamma(i, j), \gamma(i+1, j)) \in \Gamma_H, (\gamma(i, j), \gamma(i, j+1)) \in \Gamma_V$ and the tiling is legal.

For the converse, let $\gamma : \mathbb{N}^2 \rightarrow \mathcal{K}$ be a periodic tiling with period n . Without loss of generality, $n \geq 2$. Pick any n atoms $a_0, \dots, a_{n-1} \in \mathcal{A}$. Define

$$\begin{aligned} h(x) &= xa_0 \\ h(a_i) &= a_i a_{i+1} \\ h(x, y) &= (xa_0, ya_0, \gamma(0, 0)) \\ h(a_i, y) &= (a_i a_{i+1}, ya_0, \gamma(i+1, 0)) \\ h(x, a_j) &= (xa_0, a_j a_{j+1}, \gamma(0, j+1)) \\ h(a_i, a_j) &= (a_i a_{i+1}, a_j a_{j+1}, \gamma(i+1, j+1)) \end{aligned}$$

where $x, y \notin \{a_0, \dots, a_{n-1}\}$. Let $S = \{a_0, \dots, a_{n-1}\}$.

The function h is clearly S -definable. Moreover, it is a homomorphism from \mathbb{A} to \mathbb{B} . Indeed, relations Π_1, Π_2 and R are preserved immediately by definition. To check that T is preserved, we need to demonstrate that h maps every pair of elements of \mathcal{A}^2 to a pair related by $T^{\mathbb{B}}$. Consider the value of h on some arbitrarily chosen pair of elements of \mathcal{A}^2 , say

$$(ab, cd, K) \quad \text{and} \quad (ef, gh, L).$$

We must show that the implications in the definition of $T^{\mathbb{B}}$ hold.

By definition of h , the value of h on elements of \mathcal{A}^2 is always of the form $(xa_i, ya_j, \gamma(i, j))$, for some $i, j \in \{0, \dots, n-1\}$, and for some atoms x, y that will be irrelevant for the present analysis. In particular, we know that $b, d, f, h \in \{a_0, \dots, a_{n-1}\}$. Choose $i, j \in \{0, \dots, n-1\}$ so that $b = a_i$ and $d = a_j$.

We only concentrate on the first implication in the definition of $T^{\mathbb{B}}$, as the other one is shown analogously. Suppose $b = e$ and $d = h$. Then $f = a_{i+1}$ (by the definition of h), and we obtain

$$(ab, cd) = (xa_i, ya_j) \quad (ef, gh) = (a_i a_{i+1}, za_j),$$

for some atoms x, y, z . We infer $K = \gamma(i, j)$ and $L = \gamma(i+1, j)$, hence (since γ is a tiling) $(K, L) \in \Gamma_H$ as required. \square

Remark 27. Note that the structure \mathbb{B} constructed in the proof above is always \emptyset -definable, so the problem $\text{DEF-HOM}(\mathbb{A}, -)$ remains undecidable on inputs restricted to \emptyset -definable structures.

As in the case of arbitrary homomorphisms, one can consider a dual variant of the definable homomorphism problem, for a fixed target Σ -structure \mathbb{B} :

Problem: $\text{DEF-HOM}(-, \mathbb{B})$

Input: A definable structure \mathbb{A} over Σ

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

At the price of considering infinite signatures, one can repeat the development of Section 4 to prove:

Theorem 28. *There exists an \emptyset -definable structure \mathbb{B} for which the problem $\text{DEF-HOM}(-, \mathbb{B})$ is undecidable.*

Proof: see Appendix C.1. \square

Remark 29. We do not know whether $\text{DEF-HOM}(-, \mathbb{B})$ remains undecidable for some structure \mathbb{B} over a finite signature, and/or when inputs are restricted to \emptyset -definable structures. Note, however, that by Theorem 26 there exists an \emptyset -definable structure \mathbb{A} over a finite signature for which $\text{DEF-HOM}(\mathbb{A}, -)$ is undecidable. \square

Remark 30. Homomorphisms constructed in the proofs of Theorems 21, 23, 26 and 28, are injective. Therefore the respective variants of the homomorphism problem remain undecidable when one asks about existence of an injective homomorphism. In Theorems 21 and 26, those homomorphisms are even embeddings, i.e., they reflect relations and predicates as well as preserve them. Therefore the existence of embeddings of fixed structures is undecidable. However, homomorphisms in the proofs of Theorems 23 and 28 are not embeddings, as they do not reflect predicates Q_a and O_K . Decidability of the existence of embeddings into fixed definable structures therefore remains open.

6. Homomorphism problem for finite signatures

This section contains the proof of Theorem 17:

Theorem 17. *HOMOMORPHISM is decidable for finite signatures.*

The result can be easily deduced from the proof in [9]. However, it is not explicitly stated there. Therefore, we present a self-contained proof, which follows the lines of [24].

Let \mathbb{A} and \mathbb{B} be definable structures over a finite signature Σ . Then there is a finite set $S \subseteq \mathcal{A}$ such that every relation of \mathbb{A} and of \mathbb{B} is S -definable. For simplicity, we assume that $S = \emptyset$, i.e. that every relation (and also the domains A, B) of \mathbb{A} and \mathbb{B} are \emptyset -definable. The proof generalizes to arbitrary definable structures over a finite signature as we discuss later. In the rest of this section fix \mathbb{A} and \mathbb{B} over a finite signature Σ , as described above.

One might try to prove Theorem 17 by proving that existence of a homomorphism from \mathbb{A} to \mathbb{B} implies existence of a definable one with a support of bounded size. Then one could use the approach of Section 3.1. This cannot work however, as shown in Example 19.

Instead, we first show that existence of a homomorphism from \mathbb{A} to \mathbb{B} implies existence of a *canonical* one (to be defined shortly); then we provide a way of representing canonical homomorphisms effectively, which yields a decision procedure to decide their existence. The key technical tool comes from topological dynamics, in the following theorem due to Pestov [23]:

Theorem 31. *Every continuous action of the topological group $\text{Aut}(\mathbb{Q}, \leq)$ on a compact space has a fixpoint.*

Since Pestov's theorem concerns the automorphism group of the rational numbers, it will be convenient to assume that \mathcal{A} is isomorphic to (\mathbb{Q}, \leq) . Therefore, \mathcal{A} has more structure than the pure set. Every set or structure definable over the pure set is also definable over \mathcal{A} (where definitions can refer to $=$ but also to \leq).

6.1 Canonical functions

We recall from [9] canonical functions and the main result about them. Our presentation is inspired by [24]. In this section, there is no mention of relational structures.

We use standard notions of (topological) groups, (continuous) group actions and orbits. In particular, we consider the *pointwise convergence topology* on the set of all functions from A to B , where a basic open neighborhood of a function $f : A \rightarrow B$ is specified by a finite set $S \subseteq A$, and consists of those functions $g : A \rightarrow B$ which agree with f on S .

The connection of continuous group actions with definable sets comes from the following lemma.

Lemma 32. *Let \mathcal{D} denote the set of all definable sets over \mathcal{A} . Then $\text{Aut}(\mathcal{A})$ acts on \mathcal{D} via $\pi \cdot e[\text{val}] = e[\pi \circ \text{val}]$, for $\pi \in \text{Aut}(\mathcal{A})$, an expression e and valuation val of the free variables of e . This action is continuous, for the discrete topology on \mathcal{D} and the pointwise convergence topology on $\text{Aut}(\mathcal{A})$.*

It follows that if x, X are definable sets such that $x \in X$ then $\pi \cdot x \in \pi \cdot X$ for all $\pi \in \text{Aut}(\mathcal{A})$. In particular, if X is \emptyset -definable, then X is an $\text{Aut}(\mathcal{A})$ -invariant subset of \mathcal{D} .

To avoid confusion later, and also for greater generality, let us assume that two, possibly different groups G, H act on A and on B , respectively. The following definition is crucial.

Definition 33. Let G act on A and H act on B . A function $f : A \rightarrow B$ is *canonical* if for every $n \geq 1$, every G -orbit of A^n is mapped (componentwise) by f to a single H -orbit of B^n .

Below we state the main result about canonical functions, Theorem 34. Roughly, it says that under some topological assumptions, every function from A to B induces a "similar" canonical function. To state this precisely, we need some terminology from topology.

A topological group G is *extremely amenable* if every continuous action of G on a compact space X has a fixpoint. By Pestov's theorem [23], the group $\text{Aut}(\mathbb{Q}, \leq)$, with the topology of pointwise convergence, is extremely amenable.

We say that H acts *oligomorphically* on B if for every $n \in \mathbb{N}$, B^n has finitely many orbits under the componentwise action of H .

In the statement below, A and B are discrete, and the closure is with respect to the pointwise convergence topology on $A \rightarrow B$.

Theorem 34 ([9]). *Assume that G is extremely amenable and acts continuously on A , and that H acts oligomorphically on B . Then, for every function $h : A \rightarrow B$ there exists a canonical function $f : A \rightarrow B$ such that $f \in \overline{H \cdot h \cdot G}$.*

We introduce some terminology which we find useful in proving Theorem 34, and also in its applications.

For $n \in \mathbb{N}$, denote $\{1, \dots, n\}$ by $[n]$. Let \mathcal{K} be the category whose objects are sets $[n]$, for $n \geq 1$, and whose morphisms are order-preserving inclusions. A *projection system*, also known as a simplicial set, is a contravariant functor P from \mathcal{K} to the category of sets. More explicitly, P is a family of sets P_n , one for each $n \in \mathbb{N}$, and for each $i : [m] \rightarrow [n]$, a mapping $\pi_i : P_n \rightarrow P_m$ called a *projection*, such that for $[k] \xrightarrow{i} [m] \xrightarrow{j} [n]$, the composition $P_n \xrightarrow{\pi_j} P_m \xrightarrow{\pi_i} P_k$ is equal to $\pi_{j \circ i}$. A mapping of projection systems P, Q is a natural transformation $\alpha : P \rightarrow Q$, i.e., a family of functions $\alpha_n : P_n \rightarrow Q_n$, one for each $n \in \mathbb{N}$, which commutes with projections: for each increasing $i : [m] \rightarrow [n]$,

$$\alpha_m \circ \pi_i^P = \pi_i^Q \circ \alpha_n. \quad (2)$$

For a set A , denote by A^* the projection system whose components are A^n , for $n \geq 1$, and with $\pi_i : A^n \rightarrow A^m$ being the obvious projection onto m coordinates, for an increasing $i : [m] \rightarrow [n]$. A function $f : A \rightarrow B$ naturally induces a mapping $f^* : A^* \rightarrow B^*$ of projection systems.

The componentwise action of G on A^* is an action by projection system automorphisms, i.e., an element $\rho \in G$ induces a mapping of the projection system A^* to itself, which has an inverse mapping. This action induces the quotient projection system, denoted A^*/G , whose components are A^n/G , the orbits of A^n under the action of G , and the projections are the natural ones. Similarly we define the action of H on B^* and the quotient B^*/H .

The groups G and H act on the set of mappings $\alpha : A^* \rightarrow B^*$:

- H acts from the left, by $(\rho \cdot \alpha)(\bar{a}) = \rho(\alpha(\bar{a}))$ for $\rho \in H$.
- G acts from the right, by $(\alpha \cdot \sigma)(\bar{a}) = \alpha(\sigma(\bar{a}))$ for $\sigma \in G$.

Moreover, the two actions commute, i.e., $(\rho \cdot f) \cdot \sigma = \rho \cdot (f \cdot \sigma)$.

Let $\kappa : (A \rightarrow B) \rightarrow (A^* \rightarrow B^*/H)$ be defined so that $\kappa(f)$ is the composition of f^* with the quotient mapping $B^* \rightarrow B^*/H$. The following lemma is simply a reformulation of the definition of canonicity.

Lemma 35. *$f : A \rightarrow B$ is canonical iff $\kappa(f) : A^* \rightarrow B^*/H$ is invariant under the (right) action of G .*

We now proceed to proving Theorem 34. First, two lemmas:

Lemma 36. *Any mapping of projection systems $u : A^* \rightarrow B^*/H$ can be lifted to a mapping $f : A \rightarrow B$, such that $\kappa(f) = u$.*

Proof. Choose an enumeration a_1, a_2, \dots of the elements of A . We lift u to a mapping $f^* : A^* \rightarrow B^*$ inductively, so that $H \cdot f^*(a_1 \dots a_n) = u(a_1 \dots a_n)$. In each step, this can be achieved because u is a mapping of projection systems. \square

Lemma 37. *If H acts oligomorphically on B , then the set of projection system mappings $A^* \rightarrow B^*/H$, equipped with the topology of pointwise convergence is a compact space.*

Proof: see Appendix D.1. \square

Finally, we prove Theorem 34.

Proof. Let $h : A \rightarrow B$. Then $\kappa(h) : A^* \rightarrow B^*/H$; let $K = \overline{G \cdot \kappa(h)}$ denote the closure of the orbit of $\kappa(h)$ in the space of projection system mappings $A^* \rightarrow B^*/H$. It follows from Lemma 37

that K is compact. It is clear that G acts continuously on the set of mappings $A^* \rightarrow B^*/H$, and hence also on K , the closure of a G -orbit. By extreme amenability, there is a fixpoint $u \in K$, where u is a mapping $u : A^* \rightarrow B^*/H$. Let f be as in Lemma 36, so that $\kappa(f) = u$. It follows that $f \in \overline{H \cdot h \cdot G}$ (see Appendix D.2.). \square

6.2 Back to relational structures

We now interpret Theorem 34 in terms of relational structures.

Let \mathcal{A} be isomorphic to (\mathbb{Q}, \leq) . Let \mathbb{A}, \mathbb{B} be structures over finite signatures, both definable over \mathcal{A} , and such that each relation of \mathbb{A} and \mathbb{B} is \emptyset -definable. Let $G = H = \text{Aut}(\mathcal{A})$; this group acts continuously on A and on B , and preserves each relation of \mathbb{A} and \mathbb{B} (by Lemma 32). In particular, every $\rho \in \text{Aut}(\mathcal{A})$ induces an automorphism of \mathbb{A} and of \mathbb{B} . If \mathcal{A} is isomorphic to (\mathbb{Q}, \leq) , then the two assumptions of Theorem 34 are satisfied:

- G is extremely amenable, thanks to Pestov's result (Theorem 31), since \mathcal{A} is isomorphic to (\mathbb{Q}, \leq) .
- H acts oligomorphically on B . This is because clearly, $\text{Aut}(\mathcal{A})$ acts oligomorphically on \mathcal{A} , and so it acts oligomorphically on every \emptyset -definable subset $X \subseteq \mathcal{A}^n$, for $n \geq 1$, and, by Lemma 8, on every \emptyset -definable set; in particular, on B .

Since homomorphisms $\mathbb{A} \rightarrow \mathbb{B}$ form a closed subset of $A \rightarrow B$, and G, H induce automorphisms of \mathbb{A}, \mathbb{B} , Theorem 34 yields:

Corollary 38. *If there is a homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$, then there is a canonical one.*

If $\mathbb{A}, \mathbb{B}, G, H$ are as above, then a mapping of projection systems $f : A^*/G \rightarrow B^*/H$ is *homomorphic* if for every relation symbol R , if n is its arity, then for every $O \in A^n/G$ such that $O \subseteq R^{\mathbb{A}}$, it holds that $f(O) \subseteq R^{\mathbb{B}}$.

By definition, a canonical mapping $f : A \rightarrow B$ induces a mapping of the quotient projection systems, which we denote $f^*/GH : A^*/G \rightarrow B^*/H$. Note that a canonical function $f : A \rightarrow B$ is a homomorphism iff $f^*/GH : A^*/G \rightarrow B^*/H$ is homomorphic.

Lemma 39. *There is a homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$ iff there is a homomorphic mapping $k : A^*/G \rightarrow B^*/H$ of projective systems.*

Proof: see Appendix D.3. \square

6.3 Small representations

We shall now see that under certain assumptions, the set of all mappings $A^*/G \rightarrow B^*/H$ can be finitely represented. This will lead to an effective procedure testing the condition in Lemma 39.

A projection system Q is *m-simple* if for any $x, y \in Q_n$ with $m \leq n$, if $\pi_i(x) = \pi_i(y)$ for all $i : [m] \rightarrow [n]$, then $x = y$.

Lemma 40. *If B is an \emptyset -definable subset of \mathcal{A}^n for some n and $H = \text{Aut}(\mathcal{A})$, then the system B^*/H is 2-simple.*

Proof: see Appendix D.4. \square

Corollary 41. *If B is an \emptyset -definable subset of \mathcal{A}^n for some n and $H = \text{Aut}(\mathcal{A})$, then any $h : A^*/G \rightarrow B^*/H$ is determined by its component $h_2 : A^2/G \rightarrow B^2/G$. Moreover, $h_m : A^m/G \rightarrow B^m/G$ can be effectively computed from h_2 , for every $m \geq 1$.*

For $n \geq 1$ define an *n-projection system* as a projection system P whose components P_m are empty for $m > n$. For any projection system P , let $F_n(P)$ denote the n -projection system obtained by truncation; similarly for a mapping $f : P \rightarrow Q$ let $F_n(f) : F_n(P) \rightarrow F_n(Q)$ be the truncated mapping. In what follows, for a set A we denote $F_n(A^*)$ by $A^{\leq n}$, $F_n(A^*/G)$ by $A^{\leq n}/G$, and for a canonical mapping $h : A \rightarrow B$, we denote $F_n(h^*/GH)$ by $h^{\leq n}/GH : A^{\leq n}/G \rightarrow B^{\leq n}/H$.

Lemma 42. *If B is an \emptyset -definable subset of \mathcal{A}^n for some n and $H = \text{Aut}(\mathcal{A})$, then every mapping $k : A^{\leq 3}/G \rightarrow B^{\leq 3}/H$ of 3-projection systems lifts to a canonical mapping $h : A \rightarrow B$ such that $h^{\leq 3}/GH = k$.*

Proof: see Appendix D.5. \square

Corollary 43. *If B is an \emptyset -definable subset of \mathcal{A}^n for some n and $H = \text{Aut}(\mathcal{A})$, then every mapping $k : A^{\leq 3}/G \rightarrow B^{\leq 3}/H$ of 3-projection systems has a unique extension $f : A^*/G \rightarrow B^*/H$.*

Proof of Theorem 17. We can assume that the universe of \mathbb{B} is a set of tuples of atoms of a fixed length. Indeed, given a \emptyset -definable structure \mathbb{B} over a finite signature Σ , apply Lemma 8 to obtain $B' \subseteq \mathcal{A}^n$ and a surjection $g : B' \rightarrow B$. Then compute a definable structure \mathbb{B}' with universe B' over the signature Σ . A relation symbol ρ in Σ is interpreted in \mathbb{B}' as the inverse image under g of its interpretation in \mathbb{B} . It is easy to see that there is a homomorphism from \mathbb{A} to \mathbb{B} if and only if there is a homomorphism from \mathbb{A} to \mathbb{B}' . Therefore, we can assume that B is an \emptyset -definable subset of \mathcal{A}^n .

Lemma 39 yields a decision procedure for deciding whether \mathbb{A} maps homomorphically to \mathbb{B} : scan through all mappings of 3-projection systems $f : A^{\leq 3}/G \rightarrow B^{\leq 3}/H$, and for each of them check if its unique extension (whose existence is guaranteed by Corollary 43) is homomorphic. This can be checked by Corollary 41 applied to m the maximal arity of relations in Σ . \square

7. 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. Few cases are decidable, which is quite unexpected.

Our proofs work, or can be easily adapted to the variant of the problem when one asks about the existence of an injective homomorphism, or a strong homomorphism, or an embedding.

7.1 Underlying structure \mathcal{A}

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

Preliminaries. 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 is 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.

Undecidability results. Theorems 15, 16 and 18 hold for every infinite structure \mathcal{A} . For Theorems 15 and 18 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 16 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} .

\emptyset -definable homomorphism. The \emptyset -definable homomorphism problem considered in Theorem 10 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 [17], 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 10 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} .

Finite source or target structure. 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 [19] work are that $\text{Aut}(\mathcal{A})$ is extremely amenable or, equivalently, that \mathcal{A} is a *Ramsey structure*[18]. Examples of Ramsey structures include (\mathbb{Q}, \leq) and the ordered random graph, by [22].

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

Finite signatures. Finally, let us briefly discuss deciding the existence of homomorphisms between structures over a finite signature. In the most general form, the proof of Theorem 17 presented in Section 6 works under the following assumptions:

- The structure \mathbb{A} is definable over a decidable Ramsey structure \mathcal{A} . It is shown in [9] 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 6 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, (\mathbb{Q}, \leq) is finitely bounded, as witnessed by the family \mathcal{F} consisting of two structures, the directed 3-cycle and the 2-vertex graph with no edges, explaining why the value 3 appears in the statement of Lemma 42. It is straightforward to generalize this lemma to a finitely bounded homogeneous structure (see [9]). Any finitely bounded homogeneous structure is effectively ω -categorical, and thus decidable.

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

7.2 Remaining problems

Besides the open problems listed above, we leave a few open questions (of limited significance) relating to the decidability border:

- Is there a definable structure \mathbb{B} , for which $(\text{DEF-})\text{HOM}(-, \mathbb{B})$ is undecidable on inputs restricted to \emptyset -definable structures?
- Is there a definable structure \mathbb{B} such that the question whether a given \mathbb{A} embeds in \mathbb{B} is undecidable?
- Is there a definable structure \mathbb{B} over a finite signature, for which $\text{DEF-HOM}(-, \mathbb{B})$ remains undecidable?

We leave open the decidability of the *isomorphism problem*: decide whether two definable structures \mathbb{A}, \mathbb{B} (say, over the pure set) are isomorphic. An equivalent problem 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 [9]: 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 $\text{Aut}(\mathbb{A})$.

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

References

- [1] R. Berger. The undecidability of the domino problem. *Mem. Amer. Math. Soc.*, 66, 1966.
- [2] M. 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] M. Bodirsky and J. Kára. The complexity of equality constraint languages. *Theory Comput. Syst.*, 43(2):136–158, 2008.
- [4] M. Bodirsky and J. Kára. The complexity of temporal constraint satisfaction problems. *J. ACM*, 57(2), 2010.
- [5] M. Bodirsky, B. Martin, and A. Mottet. Constraint satisfaction problems over the integers with successor. In *Procs. ICALP*, volume 9134 of *LNCS*, pages 256–267, 2015.
- [6] M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. *J. Log. Comput.*, 16(3):359–373, 2006.
- [7] M. Bodirsky and M. Pinsker. Schaefer's theorem for graphs. In *Procs. STOC*, pages 655–664, 2011.
- [8] M. Bodirsky and M. Pinsker. Topological Birkhoff. *Transactions of the AMS*, 367:2527–2549, 2015.
- [9] M. Bodirsky, M. Pinsker, and T. Tsankov. Decidability of definability. *The Journal of Symbolic Logic*, 78(04):1036–1054, 2013.
- [10] M. Bojańczyk, L. Braud, B. Klin, and S. Lasota. Towards nominal computation. In *Procs. POPL 2012*, pages 401–412, 2012.
- [11] M. Bojańczyk, B. Klin, and S. Lasota. Automata theory in nominal sets. *Log. Meth. Comp. Sci.*, 10, 2014.
- [12] M. Bojańczyk, B. Klin, S. Lasota, and S. Toruńczyk. Turing machines with atoms. In *LICS*, pages 183–192, 2013.
- [13] M. Bojańczyk and S. Toruńczyk. Imperative programming in sets with atoms. In *Procs. FSTTCS 2012*, volume 18 of *LIPICs*, 2012.
- [14] E. Börger, E. Grädel, and Y. Gurevich. *The classical decision problem*. Perspectives in Mathematical Logic. Springer, 1997.
- [15] T. Feder and M. Y. Vardi. The computational structure of monotone monadic snp and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, Feb. 1999.
- [16] Y. Gurevich and I. Koryakov. Remarks on Berger's paper on the domino problem. *Siberian Mathematical Journal*, 13:319–321, 1972.
- [17] W. Hodges. *Model Theory*. Number 42 in Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993.
- [18] A. Kechris, V. Pestov, and S. Todorcevic. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geometric & Functional Analysis GAFA*, 15(1):106–189, 2005.
- [19] B. Klin, E. Kopczyński, J. Ochremiak, and S. Toruńczyk. Locally finite constraint satisfaction problems. In *Procs. of LICS'15*, pages 475–486, 2015.
- [20] P. Kolaitis and M. Vardi. The decision problem for the probabilities of higher-order properties. In *Procs. STOC'87*, pages 425–435, 1987.
- [21] L. Lovász. Kneser's conjecture, chromatic number, and homotopy. *Journal of Combinatorial Theory, Series A*, 25(3):319 – 324, 1978.
- [22] J. Nešetřil and V. Rödl. Partitions of finite relational and set systems. *Journal of Combinatorial Theory, Series A*, 22(3):289 – 312, 1977.
- [23] V. Pestov. On free actions, minimal flows, and a problem by Ellis. *Trans. Amer. Math. Soc.*, 350:4149–4165, 1998.
- [24] M. Pinsker. Making the infinite finite: Polymorphisms on Ramsey structures. Workshop on Algebra and CSPs, Fields Inst. Toronto, 2011.
- [25] A. M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge University Press, 2013.

A. Proofs from Section 2

A.1 Proof of Lemma 8

If X is described by a single set-builder expression of the form (1), then take Y to be the set defined by the same expression, with e replaced by (a_1, \dots, a_n) , where a_1, \dots, a_n are the free variables of e ; then Y is a definable subset of \mathcal{A}^n . Let $f : Y \rightarrow X$ be the function whose graph is $\{((a_1, \dots, a_n), e) \mid \phi\}$, which is clearly definable and surjective.

If $X = X_1 \cup \dots \cup X_r$ is a union of set-builder expressions, then for each X_i construct a definable surjective function $f_i : Y_i \rightarrow X_i$ as above. By embedding \mathcal{A}^m into \mathcal{A}^n for $m \leq n$, we can assume that there is a single exponent $n \in \mathbb{N}$ such that each Y_i is a subset of \mathcal{A}^n . The last step is to replace the disjoint union of the Y_i 's by a single subset Y of \mathcal{A}^k , for some k . This can be done by taking m large enough, so that \mathcal{A}^m partitions into r disjoint nonempty, \emptyset -definable subsets U_1, \dots, U_r . Finally, take $Y = \bigcup_{i=1}^r Y_i \times U_i$ and $f = \bigcup_{i=1}^r g_i$, where $g_i : Y_i \times U_i \rightarrow X_i$ first projects onto Y_i , and then applies f_i . Then $Y \subseteq \mathcal{A}^{n+m}$ and $f : Y \rightarrow X$ is surjective and definable by Lemma 5. \square

A.2 Proof of Remark 9

We sketch one direction: if a relational structure \mathbb{A} over a finite signature is definable over \mathcal{A} , then it interprets in \mathcal{A} . Indeed, let $f : B \rightarrow A$ be a surjective definable mapping obtained from Lemma 8, with $B \subseteq \mathcal{A}^k$. Lift the structure of \mathbb{A} to a structure \mathbb{B} with universe B , by taking the inverse images of the relations:

$$\sigma^{\mathbb{B}} = \{(x_1, \dots, x_k) : (f(x_1), \dots, f(x_k)) \in \sigma^{\mathbb{A}}\}.$$

This is a definable set, by Lemma 5. Moreover, f is a homomorphism from \mathbb{B} to \mathbb{A} . As a result, \mathbb{A} is isomorphic to \mathbb{B}/\sim , where \sim is the kernel of f , i.e., $x \sim y$ iff $f(x) = f(y)$. Again by Lemma 5, \sim is a definable subset of $B \times B \subseteq \mathcal{A}^{2k}$. Since $B \subseteq \mathcal{A}^k$ and $\sim \subseteq \mathcal{A}^{2k}$ are definable, there are formulas ϕ_{dom} and $\phi_{=}$ which define them. Similarly, for each symbol $\sigma \in \Sigma$, $\sigma^{\mathbb{B}} \subseteq B^l$, where l is the arity of σ , so there is a formula ϕ_{σ} defining $\sigma^{\mathbb{B}}$. The formulas $\phi_{\text{dom}}, \phi_{=}, (\phi_{\sigma})_{\sigma \in \Sigma}$ define an interpretation of \mathbb{B}/\sim in \mathcal{A} , and, as noted above, \mathbb{B}/\sim is isomorphic to \mathbb{A} .

The opposite direction (every structure \mathbb{A} which interprets in \mathcal{A} is definable) is straightforward, since the usual expressions defining the universe of the structure \mathbb{A} and its relations, are allowed by Lemma 5. In particular, the universe of \mathbb{A} is defined as the quotient of \mathcal{A}^k under a definable equivalence relation, where k is the dimension of the interpretation. \square

B. Proofs from Section 3

B.1 Proof of Theorem 14

[Proof sketch] If the source structure \mathbb{A} has a finite universe, say $A = \{a_1, \dots, a_n\}$, then the two problems coincide, as every homomorphism $\mathbb{A} \rightarrow \mathbb{B}$ is definable. Both problems reduce to validity of the following generalized first-order formula in \mathbb{B}

$$\exists x_1, \dots, x_n \bigwedge_{1 \leq k \leq l, 1 \leq i_1, \dots, i_k \leq n} \forall \rho \in \Sigma_k \rho(a_{i_1}, \dots, a_{i_k}) \rightarrow \rho(x_{i_1}, \dots, x_{i_k}),$$

which in turn reduces to validity of a first-order formula in \mathcal{A} as observed in Section 2.2, and is thus decidable.

Both problems are also decidable if the target structure \mathbb{B} has a finite universe. Structures with finite universe is a special case of a *locally finite* structures, and thus decidability of HOMOMORPHISM follows from [19]. To see decidability of DEFINABLE HOMOMORPHISM, for simplicity assume that \mathbb{A} and \mathbb{B} are \emptyset -definable. We claim that the problem reduces to \emptyset -DEFINABLE HOMOMORPHISM and hence is decidable by Theorem 10 (the general case of a definable \mathbb{B} with a finite universe is shown analogously). Indeed,

suppose f is an S -definable homomorphism from \mathbb{A} to \mathbb{B} . Consider the following modification of the defining expression of f : change all occurrences of terms $v = a$, for $a \in S$, to \perp ; and all occurrences of terms $v \neq a$, for $a \in S$, to \top . As \mathbb{B} is assumed to be finite and \emptyset -definable, the modified expression still defines a function from A to B , and the function is clearly \emptyset -definable. As \mathbb{A} is \emptyset -definable, the function is a homomorphism as required. \square

C. Proofs from Section 5

C.1 Proof of Theorem 28

We apply the technique used in the proof of Theorem 23 to modify the proof of Theorem 26, with appropriate changes.

This time, the reduction is from a seeded ultimately-periodic tiling problem. For a finite set $\mathcal{K} \ni K, L, \dots$ of colors and relations $\Gamma_H, \Gamma_V \subseteq \mathcal{K} \times \mathcal{K}$, an ultimately periodic tiling is a function $\gamma : \mathbb{N}^2 \rightarrow \mathcal{K}$ such that for all $0 \leq i, j$,

$$\begin{aligned} (\gamma(i, j), \gamma(i+1, j)) &\in \Gamma_H \quad \text{and} \\ (\gamma(i, j), \gamma(i, j+1)) &\in \Gamma_V, \end{aligned}$$

and such that for some numbers n (the *head*) and m (the *period*),

$$\begin{aligned} \gamma(i, j) &= \gamma(i+n, j) && \text{for all } i \geq m, j \in \mathbb{N} \\ \gamma(i, j) &= \gamma(i, j+n) && \text{for all } j \geq m, i \in \mathbb{N}. \end{aligned}$$

Additionally, a tiling is seeded by $K_0, K_1, \dots, K_k \in \mathcal{K}$ if $\gamma(i, 0) = K_i$ for every $i \in \{0, 1, \dots, k\}$.

It is not difficult to see that there exist fixed \mathcal{K}, Γ_H and Γ_V such that it is undecidable whether a given seed admits a seeded ultimately periodic tiling. The argument is similar to the one in the proof of Theorem 23, with the additional observation that while in Wang's encoding of Turing machines (see [14, App. A]), arbitrary tilings correspond to infinite runs, it is easy to modify the encoding so that ultimately periodic tilings correspond to finite accepting runs.

Consider \mathbb{B} defined as in the proof of Theorem 26 for the specific \mathcal{K}, Γ_H and Γ_V for which the seeded ultimately periodic tiling problem is undecidable. Additionally, extend the signature with an infinite family of predicate symbols $\{P_a, Q_a \mid a \in \mathcal{A}\}$, and a finite family of predicates $\{O_K \mid K \in \mathcal{K}\}$. Interpret these in \mathbb{B} as in the proof of Theorems 21 and 23:

$$\begin{aligned} P_a^{\mathbb{B}} &= \{ab \in B_0 \mid b \in \mathcal{A}\} && \text{for } a \in \mathcal{A}, \\ Q_a^{\mathbb{B}} &= \{ba \in B_0 \mid b \in \mathcal{A}\} && \text{for } a \in \mathcal{A}, \\ O_K^{\mathbb{B}} &= \{(ab, cd, K) \mid ab, cd \in B_0\} && \text{for } K \in \mathcal{K}. \end{aligned}$$

The structure \mathbb{B} is \emptyset -definable.

Given a seed K_0, K_1, \dots, K_k , consider a structure \mathbb{A} over the extended signature as in the proof of Theorem 21. Pick any sequence of $n+2$ distinct atoms a_0, a_1, \dots, a_{n+1} and extend \mathbb{A} as in the proof of Theorem 23. The structure \mathbb{A} is $\{a_0, a_1, \dots, a_{n+1}\}$ -definable.

\mathcal{K}, Γ_H and Γ_V admit an ultimately periodic tiling seeded by K_0, \dots, K_n if and only if there is a homomorphism from \mathbb{A} to \mathbb{B} . To see this, proceed as in the proof of Theorem 26, but note additionally that due to the interpretation of the Q_a in \mathbb{A} and \mathbb{B} , for any $h : \mathbb{A} \rightarrow \mathbb{B}$ there must be

$$h(a_i) = a_i a_{i+1} \quad \text{for } 0 \leq i \leq n$$

Moreover, all a_0, \dots, a_{n+1} must be in every support S of h . Looking back at Example 25, notice that not only the graph considered there must contain a cycle, but every node in the graph determines a unique directed path that starts from it, and ultimately ends in a cycle. This means that the sequence $a_0 a_1, a_1 a_2, \dots, a_n a_{n+1}$ must extend to a sequence of edges that ends in a cycle of length $n \geq 2$,

and every edge (pair of atoms) in that sequence is a value of h on some atom from S .

Using this, proceed as in the proofs of Theorems 23 and 26. In particular, a tiling with head n and period m determines a homomorphism supported by $n + m + 1$ atoms. Note also that every homomorphism $h : \mathbb{A} \rightarrow \mathbb{B}$ does determine a *periodic* tiling, but that tiling is not necessarily seeded by K_0, \dots, K_n . For a seeded tiling one needs to resort to an ultimately periodic tiling, since there is no guarantee that edges a_0a_1, a_1a_2, \dots for the atoms a_0, a_1, \dots fixed in the definition of \mathbb{A} , lie on the cycle determined by h . \square

D. Proofs from Section 6

D.1 Proof of Lemma 37

From oligomorphicity of the action of H on B , the system B^*/H is *finitary*, i.e., each of its components is finite. In general, if P and Q are projective systems and Q is finitary, then the set of projection system mappings $P \rightarrow Q$ with the topology of pointwise convergence is compact. Indeed, $P \rightarrow Q$ embeds into $\prod_{n \geq 1} Q_n^{P_n}$ in the natural way. Moreover, if the latter is equipped with the product topology, then this embedding is a homeomorphism of $P \rightarrow Q$ onto a subset of $\prod_{n \geq 1} Q_n^{P_n}$ which is closed, as follows from the form of the consistency conditions (2). From Tychonoff's theorem, $\prod_{n \geq 1} Q_n^{P_n}$ is compact, hence so is $P \rightarrow Q$. \square

D.2 Proof of Theorem 34 cont.

Let $h : A \rightarrow B$. Then $\kappa(h) : A^* \rightarrow B^*/H$; let $K = \overline{G \cdot \kappa(h)}$ denote the closure of the orbit of $\kappa(h)$ in the space of projection system mappings $A^* \rightarrow B^*/H$. It follows from Lemma 37 that K is compact. It is clear that G acts continuously on the set of mappings $A^* \rightarrow B^*/H$, and hence also on K , the closure of a G -orbit. By extreme amenability, there is a fixpoint $u \in K$, where u is a mapping $u : A^* \rightarrow B^*/H$.

Let f be obtained from Lemma 36, so that $\kappa(f) = u$. We show that $f \in \overline{H \cdot h \cdot G}$. To this end, choose any finite set $S \subseteq A$; we exhibit $\rho \in H$ and $\sigma \in G$ such that $\rho \cdot f \cdot \sigma$ and h agree on S .

Let $\bar{a} \in A^*$ be a tuple enumerating S . Since $u \in K = \overline{\kappa(h) \cdot G}$, it follows that $u(\bar{a})$ is of the form $(\kappa(h) \cdot \sigma)(\bar{a})$, for some $\sigma \in G$. From $\kappa(f) = u$ we have:

$$H \cdot f^*(\bar{a}) = u(\bar{a}) = (\kappa(h) \cdot \sigma)(\bar{a}) = H \cdot (h^* \cdot \sigma)(\bar{a})$$

Therefore, $f^*(\bar{a}) = (\rho \cdot h^* \cdot \sigma)(\bar{a})$, for some $\rho \in H$, so f and $\rho \cdot h \cdot \sigma$ agree on S . Hence, every basic open neighborhood of f contains an element of $H \cdot h \cdot G$, proving that $f \in \overline{H \cdot h \cdot G}$. \square

D.3 Proof of Lemma 39

Suppose that $h : \mathbb{A} \rightarrow \mathbb{B}$ is a homomorphism; by Corollary 38 we can assume that it is canonical. Then f^*/GH is homomorphic.

Conversely, suppose that $k : A^*/G \rightarrow B^*/H$ is homomorphic. Let $u : A^* \rightarrow B^*/H$ be the mapping obtained by composing k with the quotient mapping $A^* \rightarrow A^*/G$. By Lemma 36, there is an $f : A \rightarrow B$ such that $\kappa(f) = u$. By construction, $\kappa(f) = u$ is G -invariant; in particular, f is a canonical mapping, and $f/GH = k$ is homomorphic. Hence f is a homomorphism. \square

D.4 Proof of Lemma 40

Indeed, if B is an \emptyset -definable subset of \mathcal{A}^n and $H = \text{Aut}(\mathcal{A})$, then every orbit $O \in B^k/H$ corresponds to a \leq -type of kn -tuples of atoms, and thus is uniquely determined by its projections to 2-element subsets of coordinates. \square

D.5 Proof of Lemma 42

We construct an auxiliary structure $\mathcal{D} = (D, \leq)$, where

$$D = (A \times \{1, \dots, n\})/\sim.$$

The equivalence relation \sim and the relation \leq are defined as follows.

Take $(a_1, a_2, a_3) \in A^3$ and let O be its orbit. Then $g(O) \in B^{\leq 3}/H$ corresponds to a \leq -type of $3n$ -tuples of atoms. The \leq -type concerns tuples $(x_1^1, \dots, x_1^n, x_2^1, \dots, x_2^n, x_3^1, \dots, x_3^n)$ and specifies a relation

$$x_k^i \leq x_l^j \quad (3)$$

for certain $1 \leq i, j \leq n$ and $1 \leq k, l \leq 3$. Put $(a_k, i) \leq (a_l, j)$ in $A \times \{1, \dots, n\}$ if the \leq -type specifies the relation (3). This defines a transitive relation on $A \times \{1, \dots, n\}$. Define \sim to be its symmetric part, $\sim = \leq \cap \leq^{-1}$, to obtain a partially ordered set $\mathcal{D} = (D, \leq)$.

Since \mathcal{D} is a countable total order, there is an embedding $e : \mathcal{D} \rightarrow \mathcal{A}$. We define a function $h : A \rightarrow \mathcal{A}^n$, by essentially composing the abstraction function $[-]_{\sim} : A \times \{1, \dots, n\} \rightarrow \mathcal{D}$ with the embedding e :

$$h(a) = (e([(a, 1)]_{\sim}), \dots, e([(a, n)]_{\sim})).$$

It follows from the construction that h is a canonical function from A to B , and that the induced function h^*/GH extends k , as required. \square