Homomorphism Problems for First-Order Definable Structures

Bartek Klin[1], Sławomir Lasota[1], Joanna Ochremiak[2], and Szymon Toruńczyk[1]

1 University of Warsaw
2 Universitat Politécnica 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...

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 $\{\underline{1}, \underline{2}, \underline{3}, \ldots\}$ 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 $S \subseteq \mathcal{A}$. ◀

Although definable relational structures correspond (up to isomorphism) to first-order interpretations well-known from logic and model theory [21], we prefer to use a different definition since standard set-theoretic notions directly translate into this setting. For example, a definable function

Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$f : X \to Y$ 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 CSP($\mathbb{B}$) is either solvable in polynomial time or NP-complete [19].

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 $\underline{1}, \ldots, \underline{2n}$ 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 [23], 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 $\mathrm{CSP}_{\mathrm{def}}(\mathbb{B})$. The results from [23], 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* $\mathrm{CSP}(\mathbb{B})$ *is in NP,*
2. *the problem* $\mathrm{CSP}_{\mathrm{def}}(\mathbb{B})$ *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 Todorcevic concerning extremely amenable groups [22].

For finite templates, it is shown in [23] that the complexity analysis of $\mathrm{CSP}_{\mathrm{def}}(\mathbb{B})$ 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 $\mathrm{Aut}(\mathbb{Q}, \leq)$, 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 [23], is part of a programme aimed at generalizing classical decision problems and computation models such as automata [14], Turing machines [15] and programming languages [13, 16, 24, 26], to sets with atoms. For other applications of sets with atoms (called there *nominal sets*) in computing, see [28].

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 [25], 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 [27] 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 [27], see Section 5) to finite target structures. Finally, in [23] 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 [18].

Acknowledgments.

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

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).

2 Preliminaries

Thorough 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 [23].

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

$$\{\, e \,|\, a_1, \ldots, a_n \in \mathcal{A}, \phi \,\}, \tag{1}$$

where e is an expression, $a_1, \ldots, a_n$ 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 $a_1, \ldots, a_n$.

For an expression e with free variables V, any valuation val : $V \to \mathcal{A}$ 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 val : $V \to \mathcal{A}$, 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 $\{a \,|\, a \in \mathcal{A}, a \neq \underline{1}\}$ is formally defined by the expression $\{a \,|\, a \in \mathcal{A}, a \neq b\}$, 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 $X \to Y$ 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 $\Sigma = \Sigma_1 \uplus \Sigma_2 \uplus \ldots \uplus \Sigma_l$ 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 assings 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 $\mathbb{A} = (A, I_1, \ldots, I_l)$ where $I_r = \{(\sigma, a_1, \ldots, a_r) \mid \sigma \in \Sigma_r, (a_1, \ldots, a_r) \in \sigma^{\mathbb{A}}\}$ is a definable set for $r = 1, \ldots, l$ (where l is the maximal arity in Σ).

▶ Remark 7. A definable Σ-structure $\mathbb{A} = (A, I_1, \ldots, I_l)$, for Σ finite or infinite, can be seen as a definable structure over a finite signature, denoted $\mathbb{A}_\Sigma$ and defined as follows. The universe of $\mathbb{A}_\Sigma$ is $A \uplus \Sigma$, and its relations are $I_1, \ldots, I_l$, of arity $2, \ldots, l{+}1$, respectively. The signature is finite, with just l symbols. Then homomorphisms between Σ-structures $\mathbb{A}$ and $\mathbb{B}$ correspond to those homomorphisms between $\mathbb{A}_\Sigma$ and $\mathbb{B}_\Sigma$ 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 $f : Y \to X$, 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 [21]. ◀

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 $\mathrm{val} : V \to S$ with $V = \{v_1, \ldots, v_n\}$ the free variables of e, then we also need to represent the tuple $\mathrm{val}(v_1), \ldots, \mathrm{val}(v_n)$ of elements of S. For the pure set $\mathcal{A}$, these elements can be represented as $\underline{1}, \underline{2}, \ldots$.

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** ([23]). *$\emptyset$-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 $\mathbb{A} = (A, I_1, \ldots, I_l)$ and $\mathbb{B} = (B, J_1, \ldots, J_l)$ 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 $\psi \to \phi$.

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

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

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 $\mathbb{A} \to \mathbb{B}$.

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:

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

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 $h : \mathbb{A} \to \mathbb{B}$ 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 $h : \mathbb{A} \to \mathbb{B}$ determines an infinite sequence of distinct atoms $a_0, a_1, a_2, \ldots$ 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)):

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

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

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

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

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 $h : \mathbb{A} \to \mathbb{B}$, for some finite $S = \{a_1, \ldots, a_n\} \subseteq \mathcal{A}$. 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 $\{a_1, \ldots, a_n\}$. 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 $\{a_1, \ldots, a_n\}$, 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 $\{a_i a_j, a_j a_k, \ldots, a_m a_i\}$. ◂

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 [23]. 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. We discuss this and other minor remaining problems in Section 5, and in Appendices C and D.

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 $\mathbb{A}_\Sigma$ over a finite signature. Homomorphisms $\mathbb{A} \to \mathbb{B}$ correspond to those homomorphisms $\mathbb{A}_\Sigma \to \mathbb{B}_\Sigma$ that are the identity on the subset Σ of the universe of $\mathbb{A}_\Sigma$. 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 $f : A \to B$.
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:*

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

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 [23]. Moreover, Theorem 22 implies Theorem 4, since as shown in [23], for every finite template $\mathbb{B}'$, the complexity of $\mathrm{CSP}_{\mathrm{def}}(\mathbb{B}')$ is exponentially larger than the complexity of 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 $g : \bar{B} \to B$, 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 $\{1, \ldots, n\}$ by $[n]$. For a set C, numbers $m, n \in \mathbb{N}$ and an injective, monotone function $i : [m] \to [n]$, consider a projection mapping $\pi_i : C^n \to C^m$ onto m coordinates induced by i in the obvious way, i.e., $\pi_i(c_1, \ldots, c_n) = (c_{i(1)}, \ldots, c_{i(m)})$. 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 $C^{\leq r} = C \cup C^2 \cup \cdots \cup C^r$, 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 $i : [m] \to [n]$, there is a binary *projection relation* $\Pi_i \subseteq C^n \times C^m$ 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 $\mathrm{Aut}(\mathcal{A})$ 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 $\mathrm{Aut}(\mathcal{A})$ preserve the relations of $\mathbb{B}$. Reassuming, $\mathrm{Aut}(\mathcal{A})$ acts on the structure $\mathbb{B}$ by automorphisms. Similarly, $\mathrm{Aut}(\mathcal{A})$ acts on the structure $\mathbb{B}^{\leq r}$, inducing a quotient relational structure $\mathbb{B}^{\leq r}/\mathrm{Aut}(\mathcal{A})$ over the same signature. The elements of $\mathbb{B}^{\leq r}/\mathrm{Aut}(\mathcal{A})$ are orbits of $B^{\leq r}$ under the action of $\mathrm{Aut}(\mathcal{A})$; in other words, elements of $\mathbb{B}^{\leq r}/\mathrm{Aut}(\mathcal{A})$ are atomic types of k-tuples of atoms (an atomic type of a tuple of elements $(a_1, \ldots, a_k) \in \mathcal{A}^k$ specifies all equalities among the elements $a_1, \ldots, a_k$). Relation symbols are interpreted in $\mathbb{B}^{\leq r}/\mathrm{Aut}(\mathcal{A})$ existentially, as expected. A crucial but obvious observation is that the quotient structure $\mathbb{B}^{\leq r}/\mathrm{Aut}(\mathcal{A})$ is finite, by the following lemma.

▶ **Lemma 23.** *The group* $\mathrm{Aut}(\mathcal{A})$ *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 $\mathbb{B}^{\leq r}/\operatorname{Aut}(\mathcal{A})$, 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 $\mathbb{A} \to \mathbb{B}$ if, and only if there is a homomorphism $\mathbb{A}' \to \mathbb{B}'$.*

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

We now prove the "if" direction. Fix a homomorphism $f : \mathbb{A}^{\leq r} \to \mathbb{B}^{\leq r}/\operatorname{Aut}(\mathcal{A})$. Recall that $B \subseteq \mathcal{A}^k$. Consider the set $D = (A \times \{1, \ldots, k\})/\sim$, where the equivalence relation $\sim$ is defined as follows. Take $(a_1, \ldots, a_n) \in A^n$, for $n \leq r$. Then $f(a_1, \ldots, a_n) \in B^n/\operatorname{Aut}(\mathcal{A})$ corresponds to an atomic type of $(n \cdot k)$-tuples of atoms. In particular for $n = 2$, the atomic type concerns tuples $(x_1^1, \ldots, x_1^k, x_2^1, \ldots, x_2^k)$ 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 $A \times \{1, \ldots, k\}$ 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 $A \times \{1, \ldots, k\}$, 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 $e : D \to \mathcal{A}$. We define a function $h : A \to \mathcal{A}^k$, by composing the abstraction function $[\_]_\sim : A \times \{1, \ldots, k\} \to D$ with the function e:

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

Note that $h(a) \in B$ for every $a \in A$. It follows by construction that the function $h : A \to B$ 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 [21], is that $\mathcal{A}$ is countable and $\mathrm{Aut}(\mathcal{A})$ acts oligomorphically on $\mathcal{A}$.
- The number of orbits of $\mathcal{A}^n$ under the action of $\mathrm{Aut}(\mathcal{A})$ 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 [23] works are that $\mathrm{Aut}(\mathcal{A})$ is extremely amenable or, equivalently, that $\mathcal{A}$ is a *Ramsey structure* [22]. Examples of Ramsey structures include $(\mathbb{Q}, \leq)$ and the ordered random graph, by [27].

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 $\mathrm{Aut}(\mathbb{A})$.

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 àdiriger 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, B. Martin, M. Pinsker, and A. Pongrácz. Constraint satisfaction problems for reducts of homogeneous graphs. *CoRR*, abs/1602.05819, 2016.
7. M. Bodirsky and A. Mottet. Reducts of finitely bounded homogeneous structures, and lifting tractability from finite-domain constraint satisfaction. *CoRR*, abs/1601.04520, 2016.
8. M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. *J. Log. Comput.*, 16(3):359–373, 2006.
9. M. Bodirsky and M. Pinsker. Schaefer's theorem for graphs. In *Procs. STOC*, pages 655–664, 2011.
10. M. Bodirsky and M. Pinsker. Topological Birkhoff. *Transactions of the AMS*, 367:2527–2549, 2015.
11. M. Bodirsky, M. Pinsker, and T. Tsankov. Decidability of definability. *The Journal of Symbolic Logic*, 78(04):1036–1054, 2013.
12. M. Bodirsky and M. Wrona. Equivalence constraint satisfaction problems. In *Procs. CSL*, pages 122–136, 2012.
13. M. Bojańczyk, L. Braud, B. Klin, and S. Lasota. Towards nominal computation. In *Procs. POPL 2012*, pages 401–412, 2012.
14. M. Bojańczyk, B. Klin, and S. Lasota. Automata theory in nominal sets. *Log. Meth. Comp. Sci.*, 10, 2014.
15. M. Bojanczyk, B. Klin, S. Lasota, and S. Torunczyk. Turing machines with atoms. In *LICS*, pages 183–192, 2013.
16. M. Bojańczyk and S. Toruńczyk. Imperative programming in sets with atoms. In *Procs. FSTTCS 2012*, volume 18 of *LIPIcs*, 2012.
17. E. Börger, E. Grädel, and Y. Gurevich. *The classical decision problem*. Perspectives in Mathematical Logic. Springer, 1997.
18. J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identifications. *Combinatorica*, 12(4):389–410, 1992.
19. 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.
20. Y. Gurevich and I. Koryakov. Remarks on Berger's paper on the domino problem. *Siberian Mathematical Journal*, 13:319–321, 1972.
21. W. Hodges. *Model Theory*. Number 42 in Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993.
22. 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.
23. B. Klin, E. Kopczyński, J. Ochremiak, and S. Toruńczyk. Locally finite constraint satisfaction problems. In *Procs. of LICS'15*, pages 475–486, 2015.
24. B. Klin and M. Szynwelski. SMT solving for functional programming over infinite structures. In *Procs. MSFP'16*, pages 57–75, 2016.

25 P. Kolaitis and M. Vardi. The decision problem for the probabilities of higher-order properties. In *Procs. STOC'87*, pages 425–435, 1987.
26 E. Kopczyński and S. Toruńczyk. LOIS: an application of SMT solvers. In *Proc. SMT Workshop*, 2016. To appear.
27 J. Nešetřil and V. Rödl. Partitions of finite relational and set systems. *Journal of Combinatorial Theory, Series A*, 22(3):289 – 312, 1977.
28 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 9

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 $(a_1, \ldots, a_n)$, where $a_1, \ldots, a_n$ are the free variables of e; then Y is a definable subset of $\mathcal{A}^n$. Let $f : Y \to X$ be the function whose graph is $\{ ((a_1, \ldots, a_n), e) \mid \phi \}$, which is clearly definable and surjective.

If $X = X_1 \cup \cdots \cup X_r$ is a union of set-builder expressions, then for each X_i construct a definable surjective function $f_i : Y_i \to X_i$ 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 $U_1, \ldots, U_r$. Finally, take $Y = \bigcup_{i=1}^r Y_i \times U_i$ and $f = \bigcup_{i=1}^r g_i$, where $g_i : Y_i \times U_i \to X_i$ first projects onto Y_i, and then applies f_i. Then $Y \subseteq \mathcal{A}^{n+m}$ and $f : Y \to X$ is surjective and definable by Lemma 6. ◂

A.2 Proof of Remark 10

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 $f\colon B \to A$ be a surjective definable mapping obtained from Lemma 9, 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:

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

This is a definable set, by Lemma 6. 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 6, $\sim$ 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 $\phi_{\mathrm{dom}}, \phi_{=}, (\phi_\sigma)_{\sigma\in\Sigma}$ 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 6. 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. ◂

B Proofs from Section 3

B.1 Proof of Theorem 15

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

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

which in turn reduces to validity of a first-order formula in $\mathcal{A}$ as observed in Section 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 [23]. 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 11 (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 $\bot$; 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. ◀

C Homomorphism problem for infinite signatures

This section contains the proof of Theorem 16. 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: HOM($\mathbb{A}$,-)
Input: A definable structure $\mathbb{B}$ over Σ
Decide: Is there a homomorphism from $\mathbb{A}$ to $\mathbb{B}$?

▶ Remark 25. Recall Example 17. A similar effect can be observed for $\emptyset$-definable $\mathbb{A}$ and $\mathbb{B}$. Consider

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

(Here ab is simplified syntax for (a, b).) Note that for any homomorphism $h : \mathbb{A} \to \mathbb{B}$, all atoms in $\mathbb{A}$ are mapped to pairs that share the same first component (call it a_0). A reasoning similar to Example 17 shows that this determines a sequence of distinct atoms $a_0, a_1, a_2, \ldots$ such that $h(a_i) = a_0 a_{i+1}$ for each $i \in \mathbb{N}$.

▶ **Theorem 26.** *There exists a $\emptyset$-definable structure $\mathbb{A}$ for which the problem* HOM($\mathbb{A}$,-) *is undecidable.*

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

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

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 $\mathbb{B}$:

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

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 $h : \mathbb{A} \to \mathbb{B}$.

For one direction, assume a quarter-plane tiling $\gamma : \mathbb{N}^2 \to \mathcal{K}$. Consider any enumeration of all atoms $a_0, a_1, a_2, \ldots$ with a_0 as the first element. Define $h : \mathbb{A} \to \mathbb{B}$ 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 $h : \mathbb{A} \to \mathbb{B}$. 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 17, and for reasons explained there, there must be an infinite sequence of distinct atoms $a_0, a_1, a_2, \ldots$ such that $h(a_i) = a_i a_{i+1}$ for each $i \in \mathbb{N}$.

Define $\gamma : \mathbb{N}^2 \to \mathcal{K}$ 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. ◀

▶ Remark 27. The problem HOM($\mathbb{A}$,-), for $\mathbb{A}$ as in the proof of Theorem 26, remains undecidable even if one restricts input structures $\mathbb{B}$ to be $\emptyset$-definable. To see this, Remark 25 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 25 instead of Example 17, the proof of Theorem 26 works analogously. ◀

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

Problem: 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 26 that this problem cannot be solvable in any practical sense: even if 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 28.** *There exists a definable structure $\mathbb{B}$ for which the problem* HOM(-, $\mathbb{B}$) *is undecidable.*

Proof. We proceed much as in the proof of Theorem 26, 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 $K_0, K_1, \ldots, K_n \in \mathcal{K}$ (a *seed*), a legal tiling $\gamma : \mathbb{N}^2 \to \mathcal{K}$ is *seeded* if $\gamma(i, 0) = K_i$ for every $i \in \{0, 1, \ldots, n\}$.

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. [17, 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 26, 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:

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

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

Given a seed $K_0, K_1, \ldots, K_n$, consider a structure $\mathbb{A}$ over the extended signature as in the proof of Theorem 26. Pick any $n + 2$ distinct atoms $a_0, a_1, \ldots, a_{n+1}$ starting with a_0. Extend $\mathbb{A}$ by:

$$\begin{aligned} Q^{\mathbb{A}}_{a_{i+1}} &= \{a_i\} && \text{for } 0 \le i \le n, \\ Q^{\mathbb{A}}_{a} &= \emptyset && \text{for } a \notin \{a_1, \ldots, a_{n+1}\}, \\ O^{\mathbb{A}}_{K} &= \{(a_i, a_0) \mid i \le n,\ K_i = K\} && \text{for } K \in \mathcal{K}. \end{aligned}$$

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

Homomorphisms from $\mathbb{A}$ to $\mathbb{B}$ then correspond to quarter-plane tilings for $\mathcal{K}$, Γ_H and Γ_V seeded by $K_0, \ldots, K_n$. To see this, proceed as in the proof of Theorem 26, but note additionally that due to the interpretation of the P_a and Q_a in $\mathbb{A}$ and $\mathbb{B}$, for any $h : \mathbb{A} \to \mathbb{B}$ it holds that $h(a_i) = a_i a_{i+1}$ for $0 \le i \le n$. In other words, the infinite sequence of atoms determined by h as in the proof of Theorem 26, must begin with $a_0, a_1, \ldots, a_{n+1}$. 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 $0 \le i \le n$, as requested. ◀

▶ Remark 29. The structure $\mathbb{B}$ in Theorem 28 can be made $\emptyset$-definable, using the technique of Remarks 25 and 27. However, nonempty support of input structures $\mathbb{A}$ used in the proof of Theorem 28 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 Hom(-,$\mathbb{B}$) is undecidable when restricted to $\emptyset$-definable input structures.

D Definable homomorphism problem

This section contains the proof of Theorem 18. First, given a definable structure $\mathbb{A}$ over a finite signature Σ, consider the problem:

Problem: 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 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 Example 19 definable homomorphisms do exist, and each of them determines a finite cycle of atoms. This observation is the core of the following undecidability theorem, much as Example 17 was the core of Theorem 26.

▶ **Theorem 30.** *There exists an $\emptyset$-definable structure $\mathbb{A}$ over a finite signature for which the problem* Def-Hom($\mathbb{A}$,-) *is undecidable.*

Proof. The proof is similar to that of Theorem 26, with Example 19 replacing Example 17 as the core source of undecidability.

We reduce a periodic tiling problem defined as follows. For a finite set $\mathcal{K} \ni K, L, \ldots$ of colors and relations $\Gamma_H, \Gamma_V \subseteq \mathcal{K} \times \mathcal{K}$, we say that a tiling $\gamma : \mathbb{N}^2 \to \mathcal{K}$ is *periodic* if there is a number $n \ge 1$ such that $\gamma(i, j) = \gamma(i + n, j) = \gamma(i, j + n)$ for $i, j \in \mathbb{N}$. It is well known [20] 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 26, 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:

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

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 $h : \mathbb{A} \to \mathbb{B}$.

For the "if" part, consider any S-definable homomorphism $h : \mathbb{A} \to \mathbb{B}$, 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 19, and for reasons explained there, there must be a sequence $(a_0, a_1, \ldots, a_{n-1})$ of atoms from S, with $2 \leq n \leq |S|$, such that all pairs $a_0a_1,\ a_1a_2, \ldots, a_{n-2}a_{n-1},\ a_{n-1}a_0$ are values of h on some atoms from S. Denote those atoms $b_0, \ldots, b_{n-1} \in S$, so that

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

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

$$h(b_i, b_j) = (a_ia_{i+1}, a_ja_{j+1}, K).$$

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

$$((a_ia_{i+1}, a_ja_{j+1}, \gamma(i, j)), (a_{i+1}a_{i+2}, a_ja_{j+1}, \gamma(i{+}1, j))) \in T^{\mathbb{B}}$$
$$((a_ia_{i+1}, a_ja_{j+1}, \gamma(i, j)), (a_ia_{i+1}, a_{j+1}a_{j+2}, \gamma(i, j{+}1))) \in T^{\mathbb{B}}$$

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

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

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

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 $i, j \in \{0, \ldots, n-1\}$, and for some atoms x, y that will be irrelevant for the present analysis. In particular, we know that $b, d, f, h \in \{a_0, \ldots, a_{n-1}\}$. Choose $i, j \in \{0, \ldots, n-1\}$ 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

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

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. ◀

▶ Remark 31. Note that the structure $\mathbb{B}$ constructed in the proof above is always $\emptyset$-definable, so the problem 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: 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 C to prove:

▶ **Theorem 32.** *There exists an $\emptyset$-definable structure $\mathbb{B}$ for which the problem* Def-Hom(-,$\mathbb{B}$) *is undecidable.*

Proof. We apply the technique used in the proof of Theorem 28 to modify the proof of Theorem 30, with appropriate changes.

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

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

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

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

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

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 28, with the additional observation that while in Wang's encoding of Turing machines (see [17, 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 30 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 26 and 28:

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

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

Given a seed $K_0, K_1, \ldots, K_k$, consider a structure $\mathbb{A}$ over the extended signature as in the proof of Theorem 26. Pick any sequence of $n+2$ distinct atoms $a_0, a_1, \ldots, a_{n+1}$ and extend $\mathbb{A}$ as in the proof of Theorem 28. The structure $\mathbb{A}$ is $\{a_0, a_1, \ldots, a_{n+1}\}$-definable.

$\mathcal{K}$, Γ_H and Γ_V admit an ultimately periodic tiling seeded by $K_0, \ldots, K_n$ if and only if there is a homomorphisms from $\mathbb{A}$ to $\mathbb{B}$. To see this, proceed as in the proof of Theorem 30, but note additionally that due to the interpretation of the Q_a in $\mathbb{A}$ and $\mathbb{B}$, for any $h : \mathbb{A} \to \mathbb{B}$ there must be

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

Moreover, all $a_0, \ldots, a_{n+1}$ must be in every support S of h. Looking back at Example 19, 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 $a_0a_1, a_1a_2, \ldots a_na_{n+1}$ must extend to a sequence of edges that ends in a cycle of length $n \ge 2$, 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 28 and 30. In particular, a tiling with head n and period m determines a homomorphism supported by $n + m + 1$ atoms. Note also that every homorphism $h : \mathbb{A} \to \mathbb{B}$ does determine a *periodic* tiling, but that tiling is not necessarily seeded by $K_0, \ldots, K_n$. For a seeded tiling one needs to resort to an ultimately periodic tiling, since there is no guarantee that edges $a_0a_1, a_1a_2, \ldots$ for the atoms $a_0, a_1, \ldots$ fixed in the definition of $\mathbb{A}$, lie on the cycle determined by h. ◀

▶ Remark 33. We do not know whether 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 30 there exists an $\emptyset$-definable structure $\mathbb{A}$ over a finite signature for which Def-Hom($\mathbb{A}$,-) is undecidable. ◀

▶ Remark 34. Homomorphisms constructed in the proofs of Theorems 26, 28, 30 and 32, are injective. Therefore the respective variants of the homomorphism problem remain undecidable when one asks about existence of an injective homomorphism. In Theorems 26 and 30, 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 28 and 32 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.