Undecidability Results on Two-Variable Logics*

Erich Grädel‡ Martin Otto‡ Eric Rosen‡

Revised, May 1998

Abstract

It is a classical result of Mortimer that L^2, first-order logic with two variables, is decidable for satisfiability. We show that going beyond L^2 by adding any one of the following leads to an undecidable logic:

- very weak forms of recursion, viz.
  - (i) transitive closure operations
  - (ii) (restricted) monadic fixed-point operations
- weak access to cardinalities, through the Härtig (or equicardinality) quantifier
- a choice construct known as Hilbert's ε-operator.

In fact all these extensions of L^2 prove to be undecidable both for satisfiability, and for satisfiability in finite models. Moreover most of them are hard for Σ^1_1, the first level of the analytical hierachy, and thus have a much higher degree of undecidability than first-order logic.

---

* This research was partially supported by the German-Israeli Foundation for Scientific Research and Development.

‡ Lehrgebiet Mathematische Grundlagen der Informatik, RWTH Aachen, D-52056 Aachen, Germany, {graedel,otto,erosen}@informatik.rwth-aachen.de

Contents

1 Introduction

One of the most prominent problems of mathematical logic has been the *classical decision problem*: given a first-order sentence, determine whether it is satisfiable. It was an essential part of Hilbert's formalist programme for the foundations of mathematics. He called it *das Entscheidungsproblem*, literally "the decision problem" and considered it to be the central problem of mathematical logic.

After Church and Turing had proved that the classical decision problem is algorithmically unsolvable, the problem was transformed into a classification problem: which formula classes are decidable for satisfiablility and which are not? Traditionally most attention has been given to fragments of first-order logic that are determined by the quantifier prefix and the vocabulary of relation and function symbols of the given formulae. With respect to such fragments the classical decision problem is completely solved. A comprehensive account of this classification is given in [3].

Besides satisfiability a number of related problems have been studied by logicians. Two important variations are:

- Which classes are decidable for *finite satisfiability*?
- Which classes have the *finite model property*?

In computer science a number of logic problems have arisen (e.g., for knowledge representation systems, automatic verification, concurrent systems) that are closely related to the classical decision problem, but are not covered by the standard framework. Decidability problems for such logics are important for more practical reasons than the foundational issues associated with the classical decision problem. A knowledge representation system or a logic for reasoning about concurrent processes may be of limited use if the most basic questions are undecidable. Despite the differences between these logic problems and the classical decision problem, there are interesting connections that can be exploited. First, it should be noted that many other logic problems easily reduce to satisfiability. For instance the question whether a a sentence ψ is a consequence of the 'axioms' $\varphi_1, \dots, \varphi_n$ reduces to the question whether $\varphi_1 \wedge \dots \wedge \varphi_n \wedge \neg\psi$ is satisfiable. Further, many logics used in computer science applications can be seen as fragments of first-order logic. In most cases, though, these fragments are *not* those determined by quantifier prefix and vocabulary that were traditionally studied by logicians. However, a number of logics used in computer science can be embedded into L^2, the class of first-order formulae with only two distinct variables. In particular this is the case for *propositional modal logics* and for a number of *knowledge representation logics* (or *concept logics*) belonging to the so-called KL-ONE family [1]. Thus the classical result of Mortimer [34] that L^2 is decidable and indeed has the finite model property immediately implies corresponding results for these logics.

Of course, many logics of interest in computer science are not fragments of L^2. In fact they often contain constructors such as fixed-point operators, quantifiers over paths, counting constructs, etc. that are not even present in first-order logic. Nevertheless, many of them can still be seen as two-variable logics. For instance, the μ-calculus L_μ can easily be embedded in the extension of L^2 with a fixed-point operator and most of the more powerful variants of the KL-ONE knowledge representation logics are indeed sublogics of C^2, the extension of L^2 with counting quantifiers, or sublogics of C^2 with transitivity [4].

It is therefore interesting to determine the borderline between decidability and undecidability for two-variables logics. Of course, determining the decidability or undecidability of a logic is just a first step. If it is decidable one is interested in its complexity and in a practical algorithm. (Unfortunately, most logic problems are of rather high complexity.) If it is undecidable, one is theoretically interested in its degree of undecidability in terms of recursion-theoretic hierarchies. On the other hand, one may try to isolate a decidable fragment that is still expressive enough to describe an interesting portion of the problems the logic was designed for.

Here we provide a number of results from the undecidable side of that borderline. We show that going beyond L^2 by adding any one of the following leads to an undecidable logic:

- very weak forms of recursion, viz.
  - (i) transitive closure operations, or
  - (ii) (restricted) monadic fixed-point operations
- weak access to cardinalities, through the Härtig (or equicardinality) quantifier, or
- a choice construct known as Hilbert's ε-operator.

In fact all these extensions of L^2 prove to be undecidable both for satisfiability, and for satisfiability in finite models. Moreover most of them are hard for Σ^1_1, the first level of the analytical hierachy, and thus have a much higher degree of undecidability than first-order logic.

In a companion paper to this one, we show that C^2, the extension of two-variable first-order logic by counting quantifiers $\exists^{\geqslant i}$ and $\exists^{\leqslant i}$, is decidable [15].

Acknowledgment

We would like to acknowledge that the programme to investigate the borderline of decidability in the neighborhood of L^2 has been strongly advocated by Moshe Vardi, whose questions and conjectures have been an important incentive for our study.

2 Decidability, recursive inseparability and conservative classes

Let X be a class of formulae. We write

- $\mathit{sat}(X)$ for the set of $\psi \in X$ that are satisfiable;
- $\mathit{val}(X)$ for the set of logically valid $\psi \in X$;
- $\mathit{fin\text{-}sat}(X)$ for the set of $\psi \in X$ that have a finite model;
- $\mathit{inf\text{-}axioms}(X)$ for $\mathit{sat}(X) - \mathit{fin\text{-}sat}(X)$, the *infinity axioms* of X;
- $\mathit{non\text{-}sat}(X)$ for the set of unsatisfiable $\psi \in X$.

We say that a formula class X has the *finite model property* if every satisfiable formula in X also has a finite model, i.e., if $\mathit{sat}(X) = \mathit{fin\text{-}sat}(X)$.

In most cases, in particular if X is a class of first-order formulae, the finite model property of X immediately implies that the satisfiability problem, the finite satisfiability problem and the validity problem for X are decidable. Indeed, from the facts that *(i)* up to isomorphism, the finite structures of a given finite vocabulary are recursively enumerable, and *(ii)* the property that a given finite structure is a model of a given first-order sentence is decidable, it follows that $\mathit{sat}(X)$ is recursively enumerable. On the other side, the Completeness Theorem implies that also $\mathit{non\text{-}sat}(X)$ is recursively enumerable.

Further it should be noted that in most cases a proof of the finite model property actually gives an explicit bound on the size of the structures that have to be checked (as a function of, say, the length of the given formula). In this case one can also derive an upper bound on the decision complexity of the satisfiability problem.

A formula class without the finite model property contains *infinity axioms*, i.e., satisfiable formulae with only infinite models. Most of the formula classes that are decidable for satisfiability do in fact have the finite model property but there also are some important exceptions. For instance the first-order logic of one unary function is known to be decidable [7, 36] but the sentence $\forall x \forall y(fx = fy \rightarrow x = y) \wedge \exists x \forall y(fy \neq x)$ expresses that f is injective but not surjective and therefore has only infinite models.

A stronger variant of the unsolvability of the classical decision problem is Trakhtenbrot's *Inseparability Theorem* [37] which uses the concept of recursive inseparability.

Definition 2.1 Two disjoint sets X, Y are called *recursively inseparable* if there is no recursive set R such that $X \subseteq R$ and $R \cap Y = \varnothing$.

Theorem 2.2 (Trakhtenbrot) *The sets fin-sat*(FO), *inf-axioms*(FO) *and non-sat*(FO) *are pairwise recursively inseparable.*

Definition 2.3 Let X and Y be formula classes. A *conservative reduction* from X to Y is a recursive function $g : X \to Y$ such that for all $\psi \in X$:

(i) ψ is satisfiable if and only if $g(\psi)$ is satisfiable;

(ii) ψ is finitely satisfiable if and only if $g(\psi)$ is finitely satisfiable.

A formula class X is a *conservative reduction class* if there exists a conservative reduction from the set of all first-order formulae to X.

By Trakhtenbrot's Theorem, the sets *non-sat*(X), *fin-sat*(X), and *inf-axioms*(X) are pairwise recursively inseparable for every conservative reduction class X. Moreover it follows that *val*(X), *fin-sat*(X) and *non-sat*(X) are hard for the r.e. sets while *sat*(X) and *inf-axioms*(X) are hard for the co-r.e. sets. Note also that a conservative reduction class necessarily has infinity axioms.

3 First-order logic with two variables

We denote by L^k the restriction of first-order logic to formulae of relational vocabulary (i.e. without function symbols) that contain only the variables $x_1, \dots, x_k$. Note that interesting sentences in L^k are *not* in prenex normal form. Quite to the contrary, one extensively uses the possibility to re-use variables.

Example. To express that a graph $G = (V, E)$ contains a path of length 37, a sentence in prenex normal form needs 38 variables. By re-using variables, the same property is expressible in L^2, by a sentence of the form

$$\exists x \exists y (Exy \wedge \exists x (Eyx \wedge \exists y (Eyx \wedge \cdots) \cdots)).$$

It is well-known that the decision problem for L^k is unsolvable (even for formulae without equality) for all $k \geqslant 3$. This is an immediate consequence of the result by Kahr, Moore and Wang [28] that the satisfiablity problem of the prefix class $[\forall\exists\forall]$ is undecidable.

However, first-order logic with two variables is decidable.

Theorem 3.1 (Mortimer) *L^2 has the finite model property. Consequently, sat(L^2) is decidable.*

The bound on the model size implied by Mortimer's proof [34] is double exponential in the length of the given formulae. A single exponential bound on the model size and indeed a more transparent proof of the finite model property has recently been provided in [12]. By results of Fürer and Lewis (see [3, 9, 32]) this bound is essentially optimal.

It will be useful to introduce shorthand notation for the following basic predicate transformations that are quantifier-free L^2-definable. For binary E we denote the converse by E^{-1} and the symmetric reflexive closure by $E^{\leftrightarrow}$:

$$\begin{aligned} E^{-1} &:= \big\{(x,y) : Eyx\big\}, \\ E^{\leftrightarrow} &:= \big\{(x,y) : Exy \vee Eyx \vee x = y\big\}. \end{aligned}$$

We also speak of the domain of E, $\mathrm{dom}(E)$, and range of E, $\mathrm{range}(E)$, with the usual meanings:

$$\begin{aligned} \mathrm{dom}(E) &:= \big\{x : \exists y Exy\big\}, \\ \mathrm{range}(E) &:= \big\{x : \exists y Eyx\big\}. \end{aligned}$$

An important extension of L^2 is $L^2_{\infty\omega}$, infinitary logic with two variables. It may be thought of as the two-variable fragment of standard infinitary logic $L_{\infty\omega}$, or as the extension of L^2 that has infinitary disjunctions and conjunctions rather than just ordinary $\wedge$ and $\vee$: if Φ is any *set of formulae* of $L^2_{\infty\omega}$ then $\bigvee \Phi$ and $\bigwedge \Phi$ are formulae of $L^2_{\infty\omega}$. Since disjunctions and conjunctions can be formed over arbitrary sets, neither the syntax nor the semantics of $L^2_{\infty\omega}$ is recursive. Nevertheless infinitary logics are important since they contain many commonly studied logics as sublogics and admit simple and powerful descriptions of their expressive power via the Ehrenfeucht-Fraïssé method (see [6]).

4 Domino problems and local grids

Domino or tiling problems provide a simple and powerful method for proving undecidability results. They were introduced in the early 1960s by Wang as a tool to show the unsolvability of the $\forall\exists\forall$-prefix class in the pure predicate calculus. In the last thirty years they have been used to establish many undecidability results and lower complexity bounds for various systems of propositional logic, for subclasses of first-order logic and for decision problems in mathematical theories (see e.g. [3, 10, 16, 19, 28]).

The original, 'unconstrained' version of a domino problem is given by a finite set of dominoes or tiles, each of them an oriented unit square with coloured edges. The question is whether it is possible to cover the first quadrant in the Cartesian plane by copies of these tiles, without holes and overlaps, such that adjacent dominoes have matching colours on their common edge. The set of tiles is finite, but there are infinitely many copies of each tile available; rotation of the tiles is not allowed. Variants of this problem require that certain places (e.g. the origin, the bottom row or the diagonal) are tiled by specific tiles. A slightly more convenient definition is the following.

Definition 4.1 A *domino system* $\mathcal{D}$ is a triple (D, H, V) where D is a finite set of dominoes and $H, V \subseteq D \times D$ are two binary relations. Let S be any of the spaces $\mathbb{Z} \times \mathbb{Z}$, $\mathbb{N} \times \mathbb{N}$ or $\mathbb{Z}/s\mathbb{Z} \times \mathbb{Z}/t\mathbb{Z}$.[1] We say that $\mathcal{D}$ *tiles* S if there exists a *tiling* $\tau : S \to D$ such that for all $(x, y) \in S$:

(i) if $\tau(x, y) = d$ and $\tau(x + 1, y) = d'$ then $(d, d') \in H$;

(ii) if $\tau(x, y) = d$ and $\tau(x, y + 1) = d'$ then $(d, d') \in V$.

Proposition 4.2 *A domino system $\mathcal{D}$ admits a tiling of $\mathbb{Z} \times \mathbb{Z}$ if and only if it admits a tiling of $\mathbb{N} \times \mathbb{N}$.*

Proof. It is clear that a tiling of $\mathbb{Z} \times \mathbb{Z}$ also gives a tiling of $\mathbb{N} \times \mathbb{N}$. The converse is a nice application of König's Lemma. Suppose that τ is a tiling of $\mathbb{N} \times \mathbb{N}$ by $\mathcal{D}$. There exists at least one domino d such that for all n there exist $i, j > n$ with $\tau(i, j) = d$. Fix such a d. Further, for every $k \in \mathbb{N}$, let S_k be the square $\{-k, \dots, -1, 0, 1, \dots, k\} \times \{-k, \dots, -1, 0, 1, \dots, k\}$.

We define a finitely branching tree whose nodes are the correct tilings τ_k of S_k by $\mathcal{D}$ such that $\tau_k(0, 0) = d$. The root is the unique such tiling of S_0 and the children of a tiling τ_k are the possible extensions to tilings τ_{k+1} of S_{k+1}. This tree contains paths of any finite length. By König's Lemma it also contains an infinite path from the root, which means that $\mathcal{D}$ admits a tiling of $\mathbb{Z} \times \mathbb{Z}$. $\square$

In the context of conservative reductions we are also interested in periodic solutions of domino problems.

Definition 4.3 A domino system $\mathcal{D}$ is said to admit a *periodic tiling* if there is a tiling τ of $\mathbb{Z} \times \mathbb{Z}$ by $\mathcal{D}$ that has a horizontal and a vertical period $s, t > 0$ respectively. This means that for all points $(x, y) \in \mathbb{Z} \times \mathbb{Z}$ we have that

$$\tau(x, y) = \tau(x + s, y) = \tau(x, y + t).$$

A periodic tiling with periods s, t may be pictured as a tiling of a torus $\mathbb{Z}/s\mathbb{Z} \times \mathbb{Z}/t\mathbb{Z}$ obtained from gluing an $s \times t$ rectangle along the edges.

Berger [2] proved that the domino problem is undecidable. Gurevich and Koryakov [17] strengthened this to an inseparability result.

[1] $\mathbb{Z}/s\mathbb{Z}$ stands for $\{0, \dots, s-1\}$ with successor modulo s; this structure is isomorphic with the standard s-cycle.

Theorem 4.4 (Berger, Gurevich-Koryakov) *The set of domino systems that admit, respectively, no tiling and a periodic tiling are recursively inseparable.*

For a new proof of this theorem we refer to [3, Appendix A]. The proof shows that one can effectively associate with every first-order sentence ψ a domino system $\mathcal{D}$ which tiles $\mathbb{N} \times \mathbb{N}$ periodically if ψ has a finite model, and which admits no tiling of either $\mathbb{N} \times \mathbb{N}$ or $\mathbb{Z} \times \mathbb{Z}$ if ψ is unsatisfiable. It follows that a formula class X is a conservative reduction class if there exists a recursive function that associates with every domino system $\mathcal{D}$ a formula $\psi_{\mathcal{D}} \in X$ such that

(i) $\mathcal{D}$ tiles $\mathbb{N} \times \mathbb{N}$ or $\mathbb{Z} \times \mathbb{Z}$ if and only if $\psi_{\mathcal{D}}$ is satisfiable.

(ii) $\mathcal{D}$ admits a periodic tiling if and only if $\psi_{\mathcal{D}}$ has a finite model.

Let $\mathcal{D}$ be a domino system, $d \in D$ a designated domino. A tiling of $\mathbb{N} \times \mathbb{N}$ or $\mathbb{Z} \times \mathbb{Z}$ with $\mathcal{D}$ is called *recurrent* with respect to d if d occurs infinitely often in this tiling. The existence of recurrent tilings is not only undecidable, but undecidable of very high degree. Recall that Σ^1_1 and Π^1_1 form the first level of the analytic hierarchy in recursion theory. Both Σ^1_1 and Π^1_1 are located above the entire arithmetical hierarchy.

Theorem 4.5 (Harel) *The set of domino systems with a designated domino that admit a recurrent tiling of* $\mathbb{Z} \times \mathbb{Z}$ *is* Σ^1_1*-complete. The same applies to recurrent tilings of* $\mathbb{N} \times \mathbb{N}$.

For the proof of Σ^1_1-completeness (also of several useful variants of the recurrent domino problem) and for a number of applications in reductions see [19, 20].

Local grids. Two-dimensional grids form the basis of reductions from domino problems. Let $\mathcal{G}_{\mathbb{Z}}$ and $\mathcal{G}_{\mathbb{N}}$ be the infinite two-dimensional standard grids on $\mathbb{Z} \times \mathbb{Z}$ and on $\mathbb{N} \times \mathbb{N}$, respectively, with horizontal and vertical successor relations H and V:

$$\mathcal{G}_{\mathbb{Z}} = \Big(\mathbb{Z} \times \mathbb{Z}, H, V\Big) \quad \text{where} \quad \begin{array}{rcl} H & = & \big\{\big((x,y),(x+1,y)\big) : x,y \in \mathbb{Z}\big\} \\ V & = & \big\{\big((x,y),(x,y+1)\big) : x,y \in \mathbb{Z}\big\} \end{array}$$

Let $\mathcal{G}_{\mathbb{N}}$ be the substructure of $\mathcal{G}_{\mathbb{Z}}$ with universe $\mathbb{N} \times \mathbb{N}$. Similarly, $\mathcal{G}_m$ denotes the finite grid on $\mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$, which may be obtained from $\mathcal{G}_{\mathbb{Z}}$ by forming the quotient with respect to the congruence given by $(x,y) \sim (x',y')$ if $x \equiv x' \bmod m$ and $y \equiv y' \bmod m$.

In fact it will be more convenient to deal with structures that only locally resemble grids. This is made precise in the following definition.

Definition 4.6 A structure $\mathcal{G} = (G, H, V)$ is called a *local grid*, if H and V are the graphs of two injective functions h and v without fixed points, such that:

(i) $(H \cup H^{-1}) \cap (V \cup V^{-1}) = \varnothing$, i.e. h and v are independent everywhere.

(ii) h and v commute, i.e. $h \circ v = v \circ h$.

A local grid *without boundary* is one in which h and v are bijections.

Thus $\mathcal{G}_{\mathbb{Z}}$ and the $\mathcal{G}_m$ are local grids without boundary, $\mathcal{G}_{\mathbb{N}}$ is a local grid with boundary. Thinking in terms of the neighbourhood of a point a, condition (ii) for local grids means that those points reachable by traversing at most one H- and one V-edge, form a non-degenerate quadrangle that is closed at the far corner.

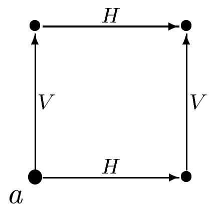

The following notion from classical model theory relaxes the conditions on (finite) axiomatizability in allowing an extended vocabulary for the purposes of the axiomatization.

Definition 4.7 A *projective characterization* of a class $\mathcal{C}$ of τ-structures is given by a sentence φ in some vocabulary $\tau' \supseteq \tau$, if $\mathcal{C}$ is the class of the τ-reducts of models of φ:

$$\mathcal{C} = \left\{ \mathfrak{A}|\tau : \mathfrak{A} \text{ a } \tau'\text{-structure}, \mathfrak{A} \models \varphi \right\}.$$

We say that a formula class L admits a projective characterization of $\mathcal{C}$ if there is a $\varphi \in L$ that projectively characterizes $\mathcal{C}$.

The following definition formalizes a natural logical closure requirement. It will be satisfied by all logics and fragments of logics we shall deal with.

Definition 4.8 A class L has the *relativization property* if for any sentence φ of L and unary predicate U, there is another sentence φ^U of L, whose models are those structures in which U carries as an induced substructure a model of φ. φ^U is called the *relativization* of φ to U. In formulae: $(\mathfrak{A}, U^{\mathfrak{A}}) \models \varphi^U$ if and only if $\mathfrak{A}|U^{\mathfrak{A}} \models \varphi$.

Recall that for first-order logic and infinitary logic, as well as for their bounded variable fragments, relativization is simply obtained by replacing all quantifications $\forall x \xi$ by $\forall x(Ux \rightarrow \xi)$ and $\exists x \xi$ by $\exists x(Ux \wedge \xi)$. Equally straightforward mechanisms apply to all other logics considered below. We will show that every extension L of L^2 that admits a projective characterization of local grids is a conservative reduction class. In fact we prove a stronger statement where the class of all local grids is replaced by any subclass $\mathcal{C}$, which is *rich* in the following sense:

(i) $\mathcal{C}$ contains at least one of the infinite standard grids, $\mathcal{G}_{\mathbb{Z}}$ or $\mathcal{G}_{\mathbb{N}}$.

(ii) For all $r \in \mathbb{N}$ there exists a $k \in \mathbb{N}$ such that $\mathcal{C}$ contains the grid $\mathcal{G}_{k \cdot r}$.

Theorem 4.9 *Let L be a formula class that contains L^2 and is closed under conjunction. If L admits a projective characterization of a rich class of local grids then L is a conservative reduction class. In particular, sat(L) and fin-sat(L) are co-r.e.-hard and r.e.-hard, respectively.*

If in addition L has the relativization property and admits a projective characterization of one of the infinite standard grids, $\mathcal{G}_{\mathbb{N}}$ or $\mathcal{G}_{\mathbb{Z}}$ (up to isomorphism of course), then moreover sat(L) is even Σ^1_1-hard.

Proof. Let $\alpha \in L$ be a formula providing a projective characterization of a rich class $\mathcal{C}$ of local grids. Given a domino system $\mathcal{D} = (D, H^{\mathcal{D}}, V^{\mathcal{D}})$, we construct a sentence $\psi_{\mathcal{D}} \in L^2$ whose vocabulary consists of the binary predicates H, V and monadic predicates P_d for $d \in D$ (we assume without loss of generality that the P_d are not used in α as auxiliary predicates). The conjuction of $\psi_{\mathcal{D}}$ with α then provides the desired conservative reduction from domino problems to L.

$$\begin{aligned}\psi_{\mathcal{D}} \quad := \quad & \forall x \bigwedge_{d \neq d'} \neg(P_d x \wedge P_{d'} x) \\ & \wedge \forall x \forall y \Big(Hxy \longrightarrow \bigvee_{(d,d') \in H^{\mathcal{D}}} (P_d x \wedge P_{d'} y) \Big) \\ & \wedge \forall x \forall y \Big(Vxy \longrightarrow \bigvee_{(d,d') \in V^{\mathcal{D}}} (P_d x \wedge P_{d'} y) \Big).\end{aligned}$$

We claim that $\mathcal{D}$ admits a tiling of $\mathbb{N} \times \mathbb{N}$ or of $\mathbb{Z} \times \mathbb{Z}$ if and only if $\alpha \wedge \psi_{\mathcal{D}}$ is satisfiable. Moreover, $\mathcal{D}$ admits a periodic tiling of $\mathbb{N} \times \mathbb{N}$ if and only if $\alpha \wedge \psi_{\mathcal{D}}$ has a finite model.

Suppose that $\tau : \mathbb{Z} \times \mathbb{Z} \to D$ is a correct tiling (such τ is implied by the existence of a tiling of $\mathbb{N} \times \mathbb{N}$ or of $\mathbb{Z} \times \mathbb{Z}$, compare Remark 4.2). α is satisfied in a suitable expansion of at least one of $\mathcal{G}_{\mathbb{Z}}$ or $\mathcal{G}_{\mathbb{N}}$. In either case a model for $\alpha \wedge \psi_{\mathcal{D}}$ is obtained by further expanding a corresponding expansion of $\mathcal{G}_{\mathbb{Z}}$ or $\mathcal{G}_{\mathbb{N}}$ with relations $P_d := \{(u,v) : \tau(u,v) = d\}$. If τ is periodic with horizontal and vertical periods s and t, we obtain a finite model as follows. Let r be the least common multiple of s and t. Since α projectively characterizes a rich class of local grids, there is some expansion of a grid $\mathcal{G}_{k \cdot r}$ that satisfies α. Obviously τ is also a correct tiling of $\mathbb{Z}/kr\mathbb{Z} \times \mathbb{Z}/kr\mathbb{Z}$ and thus provides a finite model for $\alpha \wedge \psi_{\mathcal{D}}$.

Conversely let $\mathfrak{A} \models \alpha \wedge \psi_{\mathcal{D}}$. Since $\mathfrak{A}$ is an expansion of a local grid we have injective functions h and v on the universe A such that $v \circ h = h \circ v$. Choose any element $a \in A$ and define a tiling $\tau : \mathbb{N} \times \mathbb{N} \to D$ as follows:

$$\tau(i,j) = d \text{ for the unique } d \text{ such that } \mathfrak{A} \models P_d h^i v^j a.$$

Since $\mathfrak{A} \models \psi_{\mathcal{D}}$ this mapping is well-defined and provides a correct tiling. If $\mathfrak{A}$ is finite the tiling is in fact periodic. Indeed, since h and v are injective there exist r, s such that $h^r(a) = a$ and $v^s(a) = a$. By the fact that h and v commute, this enforces that $h^r(b) = v^s(b) = b$ for all b that are reachable from a via applications of h and v. Thus $\tau(i,j) = \tau(i+r,j) = \tau(i,j+s)$ for all $(i,j) \in \mathbb{N} \times \mathbb{N}$. This finishes the proof of the first assertion concerning conservativity.

For the second assertion, let β be a sentence of L that projectively characterizes the isomorphism class of either $\mathcal{G}_{\mathbb{Z}}$ or of $\mathcal{G}_{\mathbb{N}}$. Let ξ be an infinity axiom of L that is satisfiable over countably infinite domains. Note that by assumption β itself is an example of such an infinity axiom. (Usually L will have much simpler infinity axioms with this property, though.) We choose ξ in a vocabulary that is disjoint from the vocabulary used in β and in $\psi_{\mathcal{D}}$. Of course we may also assume that the vocabularies of β and $\psi_{\mathcal{D}}$ only share H and V. It is then clear that the L-sentence $\beta \wedge \psi_{\mathcal{D}} \wedge \xi^{P_d}$ is satisfiable if and only if $(\mathcal{D}, d)$ admits a recurrent tiling of $\mathcal{G}_{\mathbb{Z}}$: the relativization of the infinity axiom ξ to P_d forces P_d to be infinite and this corresponds to an infinite occurrence of the domino d. The hardness claim now follows from Theorem 4.5. $\square$

5 L^2 with transitivity

The transitive closure of a binary relation E, denoted $\mathrm{TC}(E)$, is generated inductively as a limit

$$\mathrm{TC}(E) = E \cup E{\cdot}E \cup E{\cdot}E{\cdot}E \cup \dots ,$$

where $R{\cdot}P = \big\{(x,y) : \exists z(Rxz \wedge Pzy)\big\}$ denotes the composition of two binary relations R and P. Note that we choose to consider ordinary transitive closure rather than reflexive transitive closure as the basic notion here. Of course, one is definable from the other in quantifier-free first-order logic with equality (and without extra variables).

A binary relation E is transitive if $\mathrm{TC}(E) = E$. Note that L^2 cannot express transitivity of a binary relation. One way to see this involves Mortimer's result [34] that L^2 has the finite model property, i.e., that any satisfiable sentence of L^2 has a finite model. To show that transitivity cannot be expressed in L^2 it suffices to exhibit a sentence $\varphi(E) \in L^2$ such that $\varphi(E)$ has a model in which E is transitive, but no finite model in which E is transitive. Here is a simple example. Let φ_0 be the conjunction of $\forall x \neg Exx$ and $\forall x \exists y Exy$. It is easy to see that any model of φ_0 with transitive E must embed a copy of the natural numbers with their natural order relation, since for transitive E, the first conjunct forbids directed E-cycles. But actually a much stronger statement can be made.

Fact 5.1 *Transitivity is not expressible (even over finite structures) in two-variable infinitary logic* $L^2_{\infty\omega}$.

This follows from standard pebble game arguments. For instance, $L^2_{\infty\omega}$ cannot distinguish the undirected graphs consisting of a single hexagon (or even a square) and of two disjoint triangles, respectively. The edge relation of the latter is transitive while the former is not.

5.1 Characterizing local grids

Consider the weak extension of L^2 that — in addition to what can be expressed in L^2 itself — admits positive transitivity assertions about basic binary relations in the vocabulary. Formally, we may for the present purposes consider sentences that correspond to conjunctions $\varphi \wedge \theta$, where φ is an ordinary sentence of L^2 and θ is a statement $\mathrm{trans}[R_1, \dots, R_m]$ asserting the transitivity of some basic relations R_i in τ. Call the resulting "logic" $L^2 + \mathrm{trans}$. $L^2 + \mathrm{trans}$ is not a natural logic since it lacks the most basic closure properties. We merely consider it here since it forms the nucleus for the undecidability results to be considered. As a fragment of several more natural logics, we use $L^2 + \mathrm{trans}$ to show that these other logics must be undecidable, too. We shall prove next that this weak system suffices for the projective characterization of a rich class of local grids.

Fact 5.2 *$L^2 + \mathrm{trans}$ may be regarded as a fragment of first-order logic (with three variables): $L^2 \subsetneq L^2 + \mathrm{trans} \subsetneq L^3 \subsetneq \mathrm{FO}$.*

Matchings. As a preliminary step, we show how $L^2 + \mathrm{trans}$ can be used to characterize projectively those graphs whose edge relation E forms a matching. Note that this is equivalent with saying that E is the graph of an injective partial function whose domain and range are disjoint. For the projective characterization we use one extra binary predicate S whose intended interpretation will be the symmetric reflexive closure $E^{\leftrightarrow}$ of E. This intended connection between S and E is directly expressible in L^2 of course:

$$\forall x \forall y \Big(Sxy \leftrightarrow \big(Exy \vee Eyx \vee x = y \big) \Big).$$

For expressing the condition that the domain and range of E are disjoint, we may also use a straightforward sentence in L^2:

$$\forall x \neg \big(\exists y Exy \wedge \exists y Eyx \big).$$

It now suffices to assert in addition that S is transitive:

$$\mathrm{trans}[S].$$

Note that these three conditions are trivially satisfied by all structures of the form (V, E, S) for which E is a matching and S is the symmetric transitive closure of E. Conversely let (V, E, S) satisfy the above conditions. Then $S = E^{\leftrightarrow}$, and no vertex can simultaneously have E-successors and E-predecessors. Together with the transitivity requirement on S this implies further that no vertex can have more than one E-successor or -predecessor. For successors for instance, assume that a vertex v had two different E-successors. Then, by transitivity, there would have to be an E-edge between these, whence at least one of them would have to have an E-successor as well as an E-predecessor — which is ruled out.

Remark 5.3 *As far as* projective *characterizations or satisfiability in* $L^2 + \mathrm{trans}$ *are concerned, we may without loss of generality allow transitivity statements concerning arbitrary* L^2*-definable predicates rather than just basic relations.*

Consider a projective characterization involving a transitivity assertion of the form $\mathrm{trans}[\varphi]$ with $\varphi \in L^2$ — saying that the binary predicate $\{(x, y) | \varphi(x, y)\}$ is transitive. Its semantics remains unchanged if we adjoin an extra predicate R whose interpretation is fixed through the L^2-sentence $\forall x \forall y \big(Rxy \leftrightarrow \varphi(x, y)\big)$ and assert transitivity of this new basic relation R instead. We shall in particular from now on tacitly allow statements of the form $\mathrm{trans}[E^{\leftrightarrow}]$ for basic relations E directly (recall that $E^{\leftrightarrow}$, the symmetric reflexive closure of E, is L^2-definable from E).

Local grids in $L^2 + \mathrm{trans}$. We shall axiomatize in $L^2 + \mathrm{trans}$ a rich class of local grids without boundary (recall Definition 4.6). Those conditions on H and V in Definition 4.6 on local grids without boundary, which are not right away expressible in L^2 alone, are

(i) that they are functional and bijective.

(ii) that the induced functions commute.

The first point is taken care of by applying the considerations about matchings to predicates obtained from H and V by splitting them into two disjoint parts, each of which becomes a matching. In order to be able to achieve this in a straightforward and systematic manner we shall restrict attention to a rich subclass: it will here be useful to restrict the periodicities in components that have the topological type of a cylinder or torus to even numbers.

Definition 5.4 Call a local grid without boundary *even*, if the two functions h and v associated with H and V satisfy the following: whenever some vertex is a fixed point of some composed function $h^n \circ v^m$, then both n and m are even.

Obviously the class of even local grids without boundary is rich in the sense required for Theorem 4.9, so it suffices to show the following.

Proposition 5.5 $L^2 + \mathrm{trans}$ *admits a projective characterization of the class of even local grids without boundary.*

For the proof it remains to address points (i) and (ii) above. We introduce more auxiliary relations that will serve to reduce also the commutativity statement to a condition on transitivity. Altogether we shall work with the following auxiliary binary predicates, apart from the original H and V:

$$H_i, V_i \quad \text{and} \quad D_{ij}, M_{ij} \quad \text{where } i, j \in \{0, 1\}.$$

The intended interpretations are such that H becomes the disjoint union of H_0 and H_1 in such a way that horizontal edges indexed 0 alternate with those indexed 1 (and that vertically parallel horizontal edges are of the same kind); similarly for V and the V_i. The D_{ij} correspond to diagonal edges in such a way that $D_{ij} = H_i \cdot V_j$. M_{ij} will be the disjoint union of H_i, V_j, and D_{ij}. The standard grid $\mathcal{G}_{\mathbb{Z}}$ is expanded to a structure ${\mathcal{G}_{\mathbb{Z}}}^*$ according to the following (this expansion is also depicted in the sketch below):

$$\begin{aligned}
H_i &= \big\{\big((u,v),(u+1,v)\big) \in H : u \equiv i \bmod 2\big\},\\
V_i &= \big\{\big((u,v),(u,v+1)\big) \in V : v \equiv i \bmod 2\big\},\\
D_{ij} &= \big\{\big((u,v),(u+1,v+1)\big) \in V : u \equiv i \bmod 2 \text{ and } v \equiv j \bmod 2\big\},\\
M_{ij} &= H_i \cup V_j \cup D_{ij}.
\end{aligned}$$

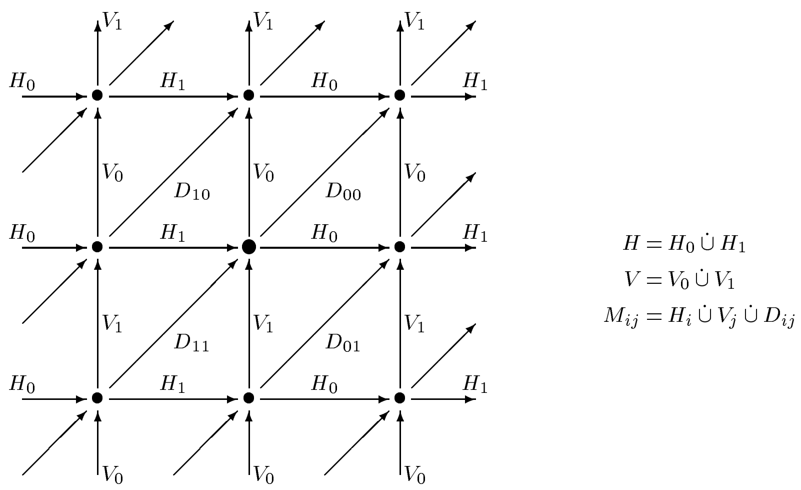

Any even local grid without boundary admits an expansion that locally looks like $\mathcal{G}_{\mathbb{Z}}{}^*$, since evenness of periods ensures that the grid may consistently be covered with adjacent 2×2-patches of the kind depicted above. (Note that each connected component of the given local grid may be treated separately.)

Note that the vertices are of four different types according to which index their outgoing V- and H-edges carry. Let $\eta_{ij}(x)$ for $i, j \in \{0, 1\}$ be L^2-formulae characterizing these four vertex classes. For instance

$$\begin{aligned}\eta_{01}(x) \quad &= \quad \exists y H_0 xy \wedge \exists y V_1 xy \wedge \exists y H_1 yx \wedge \exists y V_0 yx \\ & \qquad \wedge \neg \exists y H_1 xy \wedge \neg \exists y V_0 xy \wedge \neg \exists y H_0 yx \wedge \neg \exists y V_1 yx.\end{aligned}$$

Let us say that x is of type (i, j) if it satisfies η_{ij}. Let φ_0 be a sentence of L^2 which asserts that

1. V and H are irreflexive and antisymmetric. For instance: $\forall x \forall y \big(Vxy \rightarrow \neg Vyx\big)$.
2. H is the disjoint union of H_0 and H_1; V the disjoint union of V_0 and V_1.
3. all vertices are of exactly one of the above four types: $\forall x \dot{\bigvee} \eta_{ij}(x)$.
4. whenever x is of type (i, j) then its horizontal neighbours are of type $(1 - i, j)$ and its vertical neighbours are of type $(i, 1 - j)$:
   $\forall x \forall y \Big[\Big(\big((Hxy \vee Hyx) \wedge \eta_{01}(x)\big) \rightarrow \eta_{11}(y) \Big) \wedge \ldots \Big]$.
5. The domain of D_{ij} is the set of vertices of type (i, j), its range the set of points of type $(1 - i, 1 - j)$.
6. The D_{ij}, H_i, and V_j and their converses are all mutually disjoint.
7. M_{ij} is the disjoint union of H_i, V_j, and D_{ij}.

Let now φ be the conjunction of the L^2-sentence φ_0 with the following transitivity statements [2]:

$$\mathrm{trans}[H_0^{\leftrightarrow}, H_1^{\leftrightarrow}, V_0^{\leftrightarrow}, V_1^{\leftrightarrow}, D_{00}^{\leftrightarrow}, \ldots, D_{11}^{\leftrightarrow}, M_{00}, \ldots, M_{11}]. \tag{8}$$

It is readily checked that $\mathcal{G}_{\mathbb{Z}}{}^*$ satisfies φ. The crucial fact in connection with the transitivity requirements is the alternation of differently indexed variants of horizontal, vertical and diagonal successor relations which ensures that any of the predicates mentioned in (8) splits into disjoint connected components involving only the nearest neighbours of any given vertex.

[2]Note that we here use the relaxed variant of $L^2 + \mathrm{trans}$ in accordance with Remark 5.3 above.

Furthermore, as argued above, any even local grid without boundary can be expanded to become a model of φ through splitting H into H_0 and H_1, and and V into V_0 and V_1, in an alternating fashion consistent with conditions (3) and (4). Evenness is used in order to avoid clashes in components that fold back on themselves in forming cylinders or tori.

We claim that φ exactly provides a projective characterization of even local grids without boundary. It remains to show that any model of φ must be an even local grid without boundary. Evenness is in fact immediate, so that according to Definition 4.6 it remains to be shown that in any model of φ the predicates V and H are the graphs of completely disjoint, fixed-point free, bijective functions that commute. The non-trivial parts of this claim are functionality and commutativity. Note first that conditions (3), (4), and (5) in φ_0 imply in particular that the H_i, V_i, and D_{ij} each have disjoint domain and range.

Consider first functionality, of H say. Functionality of H_0 is enforced by transitivity of $H_0^{\leftrightarrow}$ (in (8)), since we already know that $\mathrm{dom}(H_0) \cap \mathrm{range}(H_0) = \varnothing$. H_0 in fact must be a matching. The same is true for H_1. Functionality carries over to the union $H = H_0 \dot{\cup} H_1$ because condition (3) in φ_0 also forces the domains of H_0 and H_1 to be disjoint. The same argument directly applies to V, and also to the D_{ij} if for the disjointness of domain and range we invoke (4) and (5) from φ_0. So we have that in any model of φ, H and V are functional, and that moreover the H_i, the V_i, and the D_{ij} must be matchings.

We turn to commutativity. Consider without loss of generality a vertex a of type $(0,0)$. The argument involves the predicate M_{00}, which by φ_0 is the union of H_0, V_0, and D_{00}. By (4) the neighbourhood of a must be as indicated in (a) of the following figure.

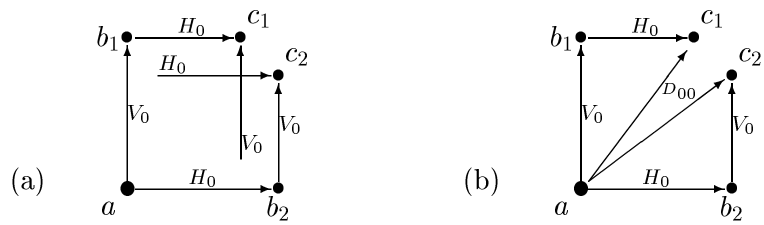

Here b_1 is of type $(0,1)$, b_2 of type $(1,0)$ and the c_i must both be of type $(1,1)$. We want to show that in fact $c_1 = c_2$. By transitivity of M_{00} as postulated in (8), c_1 and c_2 must both be linked to a by M_{00}-edges. The M_{00}-edge (a, c_1) for instance cannot be an H_0-edge (since already $(b_1, c_1) \in H_0$), or a V_0-edge (since already $(a, b_1) \in V_0$). The same arguments apply to c_2. Therefore by (7) both c_1 and c_2 are linked to a by D_{00}-edges as shown in (b). As D_{00} is a matching, c_1 and c_2 indeed coincide.

We have thus proved Proposition 5.5, which by Theorem 4.9 implies that $L^2 + \mathrm{trans}$ is a conservative reduction class. As $L^2 + \mathrm{trans}$ is contained in FO, we conclude:

Corollary 5.6 *sat*$(L^2 + \mathrm{trans})$ *and fin-sat*$(L^2 + \mathrm{trans})$ *are co-r.e.- and r.e.-complete, respectively.*

6 Two-variable transitive closure logics

First-order logic does not have the expressive means to model even very modest relational recursion. For instance it is well known and easily shown by means of Ehrenfeucht-Fraïssé games that connectivity is not a first-order definable property of (finite undirected) graphs. *Transitive closure logic* TC was introduced by Immerman [26] to address this weakness. If the transitive closure of the edge relation is made available, then connectivity becomes easily expressible even with the use of only two variables: it suffices to assert that any pair of distinct vertices is in the transitive closure of the edge relation.

6.1 TC^2 and DTC^2

TC augments the syntax and semantics of first-order logic by allowing the formation of new formulae that define the transitive closure of definable binary relations on k-tuples for all k. We here limit the application of the transitive closure operation to a framework which allows only two variables, so that only the very basic case of transitive closures of binary relations (on singletons) is taken into account. Syntactically there is a transitive closure operator denoted TC. It allows one to construct from a formula φ a new formula $\mathrm{TC}\varphi$ that is free in variables x and y. (Recall that in L^2 we only admit variable symbols x and y.) The semantics of the new formula is such that $\mathfrak{A} \models \mathrm{TC}\varphi[a, b]$ if and only if the pair (a, b) is in the transitive closure of the binary relation that is defined by φ over $\mathfrak{A}$.

$$\mathfrak{A} \models \mathrm{TC}\varphi[a, b] \qquad \text{iff} \qquad (a, b) \in \mathrm{TC}\Big(\{(c, d) \in A^2 : \mathfrak{A} \models \varphi[c, d]\}\Big).$$

Definition 6.1 Let *two-variable transitive closure logic* TC^2 be the simultaneous closure of L^2 under the usual first-order operations in L^2 and applications of the TC-operator.

Obviously $L^2 + \mathrm{trans} \subseteq \mathrm{TC}^2$ and this inclusion is strict as $L^2 + \mathrm{trans}$ is contained in first-order while TC^2 is not (consider connectivity: $\forall x \forall y \big(x \neq y \rightarrow \mathrm{TC}(Exy)\big)$).

An important variant of transitive closure logic TC is so-called *deterministic transitive closure logic*. It is denoted DTC and was also introduced by Immerman in [26]. The syntax is analogous to that for TC, but with the DTC-operator whose semantics is given in terms of deterministic transitive closures. The deterministic transitive closure DTC(E) of E by definition is the transitive closure of the relation

$$E_{\mathrm{det}} = \big\{(x,y) \in E : \forall y'\big(Exy' \rightarrow y = y'\big)\big\}.$$

More intuitively, $(a,b) \in \mathrm{DTC}(E)$ means that there is a deterministic E-path connecting a to b. The term deterministic transitive closure is somewhat misleading, however: $E \mapsto \mathrm{DTC}(E)$ is not a closure operator. In fact it is readily checked that $\mathrm{DTC}(\mathrm{DTC}(E)) = \mathrm{DTC}(E)$ if and only if E is the graph of a partial function whose domain is disjoint from its range.

Definition 6.2 Let *two-variable deterministic transitive closure logic* DTC^2 be the closure of L^2 under the first-order operations of L^2 and the DTC-operator.

The full logics TC and DTC (without restriction on the number of variables, and with TC- or DTC-operators applicable to binary predicates interpreted in higher arity) are important in descriptive complexity, since in expressive power over finite linearly ordered structures they exactly correspond to NLogspace and Logspace, respectively, [26]. It is clear that $\mathrm{DTC} \subseteq \mathrm{TC}$, simply because E_{det} is first-order definable from E. This inclusion is shown to be strict (over finite unordered structures) in [13]. In contrast, TC^2 and DTC^2 are incomparable in expressive power (even over finite structures).

For $\mathrm{DTC}^2 \not\subseteq \mathrm{TC}^2$ consider graphs whose automorphism group acts transitively on the set of vertices (any vertex can be mapped to any other vertex by some automorphism). Any two vertices in such a graph are indistinguishable in L^2, whence all L^2-definable predicates are quantifier-free definable. (The latter statement is based on the observation that any L^2-type in two variables x and y is fully determined by the individual L^2-types of the components x and y together with the quantifier-free type of the pair (x,y).) Now compare two such graphs: $\mathfrak{C} = (C,E)$ a finite directed cycle, and $\mathfrak{T} = (T,E)$ a finite two-dimensional torus. By a torus we here mean the structure obtained from a two-dimensional rectangular grid, with edge predicate E corresponding to the union of the horizontal and vertical successor relations, by gluing corresponding horizontal and vertical borders. For both these graphs all TC-applications to quantifier-free definable predicates are trivial in the sense that the results of TC-applications to quantifier-free definable predicates are again quantifier-free definable, and moreover the translation is the same for the two graphs. It follows that TC^2 cannot distinguish them. DTC^2, however does distinguish them, as $\mathfrak{C} \models \forall x \forall y \mathrm{DTC}(Exy)$ while $\mathfrak{T} \models \forall x \forall y \neg \mathrm{DTC}(Exy)$. A similar argument also shows that $\mathrm{TC}^2 \not\subseteq \mathrm{DTC}^2$: over structures that consist of one or several disjoint copies of two-dimensional tori $\mathfrak{T}$, DTC^2 collapses to quantifier-free L^2

and in particular cannot express connectivity for these, while this is possible using TC^2. Hence DTC^2 and TC^2 are incomparable over finite graphs.

Fact 5.1 and in particular the example given for its illustration, directly apply to DTC^2 as well as TC^2, so that we have the following.

Fact 6.3 *In expressive power (even over finite structures) both* DTC^2 *and* TC^2 *are not contained in* $L^2_{\infty\omega}$.

Infinity axioms. It may be instructive to give simple examples of sentences in TC^2 and DTC^2 that only have infinite models (the existence of such sentences will follow again from the undecidability results to be proved below). For $L^2 + \mathrm{trans} \subseteq \mathrm{TC}^2$ we have already exhibited such a sentence above: the conjunction of $\forall x \neg Exx$, $\forall x \exists y Exy$, and $\mathrm{trans}[E]$ rendered in TC^2 as $\forall x \forall y \big(\mathrm{TC}(Exy) \leftrightarrow Exy\big)$. Any model of this sentence embeds a copy of the natural numbers with the natural order relation.

In DTC^2 we may projectively characterize the standard successor structure over the natural numbers, using a unary predicate U to mark 0 as the initial element. The conjunction of L^2-sentences saying that there is a unique element in U, that this element of U has no E-predecessor, and the DTC^2-sentence $\forall x \forall y \big(x \neq y \leftrightarrow (\mathrm{DTC}(Exy) \vee \mathrm{DTC}(Eyx)\big)$ is as desired.

It is also interesting to note that the conjunction of $\forall x \neg Exx$, $\forall x \forall y (Exy \vee Eyx)$, and $\mathrm{trans}[E]$ axiomatizes in TC^2 the class of all linear orderings.

Theorem 6.4 *Both* TC^2 *and* DTC^2 *are conservative reduction classes. Both logics therefore are undecidable for satisfiability as well as for satisfiability in finite models.*

Since $L^2 + \mathrm{trans} \subseteq \mathrm{TC}^2$ we already know that TC^2 is a conservative reduction class. In direct analogy with the treatment for $L^2 + \mathrm{trans}$ the claim for DTC^2 is proved by showing the following.

Proposition 6.5 DTC^2 *admits a projective characterization of even local grids without boundary.*

Proof. We merely indicate those variations in the argument for $L^2 + \mathrm{trans}$ that yield the analogous facts for DTC^2. Firstly, matchings are even easier to characterize in DTC^2 than in TC^2 as E is a matching if and only if $\mathrm{DTC}(E) = E$ and $\mathrm{DTC}(E^{-1}) = E^{-1}$. Returning to the expansion $\mathcal{G}_{\mathbb{Z}}{}^*$ of the standard grid we used above, we may thus express that the V_i, H_i and D_{ij} are matchings.

Once more it remains to enforce commutativity of the V_i and H_i. This time we do not make use of the M_{ij} because these admit no deterministic paths. Instead, looking for instance at a vertex of type $(0,0)$ as an initial point, we now consider derived predicates defined by the formulae

$$\begin{aligned}\varphi_{00}^{\mathrm{VH}}(x,y) &= \big(\eta_{00}(x) \wedge V_0xy\big) \vee \big(\eta_{01}(x) \wedge H_0xy\big),\\ \varphi_{00}^{\mathrm{HV}}(x,y) &= \big(\eta_{00}(x) \wedge H_0xy\big) \vee \big(\eta_{10}(x) \wedge V_0xy\big).\end{aligned}$$

For $\mathcal{G}_{\mathbb{Z}}{}^*$ we have that $(x,y) \in \mathrm{DTC}\big(\varphi_{00}^{\mathrm{VH}}\big) \iff \varphi_{00}^{\mathrm{VH}}(x,y) \vee D_{00}xy$, and
$(x,y) \in \mathrm{DTC}\big(\varphi_{00}^{\mathrm{HV}}\big) \iff \varphi_{00}^{\mathrm{HV}}(x,y) \vee D_{00}xy$.

Again, these equalities are local in $\mathcal{G}_{\mathbb{Z}}{}^*$ in the sense of involving only the immediate neighbourhoods of points. The same therefore holds similarly in any of the intended expansions of even local grids without boundary. The above equalities may be formalized in DTC^2 in a straightforward manner. Together with the condition that D_{00} is functional, they imply that H_0 and V_0 commute at vertices of type $(0,0)$. Adding analogous conditions for vertices of the other three types we thus once more obtain a projective characterization of the even local grids, this time in DTC^2. □

6.2 Σ_1^1-completeness

Proposition 6.6 *There are sentences in* TC^2 *and in* DTC^2 *that projectively characterize* $\mathcal{G}_{\mathbb{Z}}$ *up to isomorphism.*

Before proving the proposition, we observe that TC^2 and DTC^2 have the relativization property, whence Theorem 4.9 applies to show Σ_1^1-hardness. It is not hard to see, on the other hand, that TC and DTC both satisfy a Löwenheim-Skolem theorem, which guarantees models of at most countable cardinality for all satisfiable sentences – in fact this is well known even for least fixed-point logic LFP, which extends both TC and DTC. Applying a canonical Skolemization procedure to models that are realized as expansions of $(\mathbb{N},<)$, and using extra predicates to render transitive closures first-order definable in this scenario, one obtains a reduction of the satisfiability problems for TC and DTC to the existential second-order theory of $(\mathbb{N},<)$. Towards a projective characterization of, for instance, $\mathrm{TC}\varphi(x,y)$ one may introduce a new ternary predicate $\mathrm{d}_\varphi xyn$ for reachability of y from x on a φ-path whose length is at most n. This predicate has an obvious first-order implicit definition in expansions of $(\mathbb{N},<)$ based on the usual inductive definition w.r.t. n. With this apparatus, $\mathrm{TC}\varphi(x,y)$ becomes equivalent to $\exists n d_\varphi xyn$. Modulo suitable Skolemization, then, satisfiability of TC-sentences reduces to the existential second-order assertion that there are interpretations of the basic predicates and of the auxiliary predicates d_φ over $(\mathbb{N},<)$ that satisfy certain first-order conditions. This implies that *sat*(TC) and *sat*(DTC) are in Σ_1^1.

Corollary 6.7 *sat*(TC^2) *and sat*(DTC^2) *are* Σ_1^1*-complete.*

Proof of the proposition. The claim of the proposition is immediate for TC^2. We need merely add to the projective characterizations of even local grids without boundary conjuncts asserting that the predicates $H \dot{\cup} V$ and $H \dot{\cup} V^{-1}$ (and hence also $H^{-1} \dot{\cup} V$ and $H^{-1} \dot{\cup} V^{-1}$) are free of directed cycles, and that the overall structure is connected with respect to the joint relation $H \cup V$.

It is a little more complicated in DTC^2 to enforce two-dimensional infinite extension and a single component in local grids. Using for convenience three more monadic predicates O, B, and C, it is easy to assert in DTC^2 the following. In the standard model $\mathcal{G}_{\mathbb{Z}}$, an interpretation of B, respectively C, as the set of points on one particular horizontal, respectively vertical, axis clearly satisfies these assertions.

(i) $O = B \cap C$ consists of a single vertex.

(ii) $B - O$ consists of exactly those vertices that are connected to O by a (deterministic) H-path or by a (deterministic) H^{-1}-path; similarly for $C - O$ and V.

(iii) There is no cyclic (deterministic) H-path within B and similarly no cyclic (deterministic) V-path in C.

(iv) Every vertex that does not itself belong to C is connected to a vertex in C by a (deterministic) H-path or H^{-1}-path.

Any local grid without boundary satisfying these conditions is isomorphic to $\mathcal{G}_{\mathbb{Z}}$. Overall connectedness follows since all vertices are connected with the single member o of O. B and C are forced to form horizontal and vertical axes in this connected local grid without boundary; both extend infinitely because they cannot form cycles by (iii). They also cannot wrap in spirals about a grid of the type of a cylinder (or a torus), since they meet in exactly one vertex by (i): if $n, m \in \mathbb{Z}$ and a are such that $v^n \circ h^m(a) = a$, then $v^n \circ h^m(o) = o$ and $h^m(o) = v^{-n}(o)$, whence (i) implies that $n = m = 0$. $\square$

7 Two-variable fixed-point logics

Fixed-point operations are the most natural way to augment first-order logic with recursive mechanisms. Consider a formula $\varphi(X, \overline{x})$ with a second-order variable X and a distinguished tuple $\overline{x}$ of pairwise distinct first-order variables whose number matches the arity r of X. If $\mathfrak{A}$ provides an interpretation for all symbols in φ apart from X and the $\overline{x}$, then $\varphi(X, \overline{x})$ gives rise to an operation $F^{\mathfrak{A}}_{\varphi}$ on r-ary predicates over A:

$$F^{\mathfrak{A}}_{\varphi} \colon P \longmapsto \{\overline{a} \in A^r : \mathfrak{A} \models \varphi[P, \overline{a}]\}.$$

If φ is monotone in X (i.e., if $P \subseteq P'$ implies $F^{\mathfrak{A}}_{\varphi}(P) \subseteq F^{\mathfrak{A}}_{\varphi}(P')$) then $F^{\mathfrak{A}}_{\varphi}$ possesses a *least fixed-point* $\mathrm{LFP}(F^{\mathfrak{A}}_{\varphi})$. By definition, $\mathrm{LFP}(F^{\mathfrak{A}}_{\varphi})$ is the $\subseteq$-minimal $P \subseteq A^r$ such that $F_\varphi^{\mathfrak{A}}(P) = P$. In more procedural terms, $\mathrm{LFP}(F_\varphi^{\mathfrak{A}})$ may be defined as the limit or saturation point of the sequence P_α of predicates obtained by iteration of F_φ on the empty predicate:

$$\mathrm{LFP}(F_\varphi^{\mathfrak{A}}) = \bigcup_\alpha P_\alpha \quad \text{where} \quad \begin{array}{lcl} P_0 & = & \emptyset, \\ P_{\alpha+1} & = & F_\varphi^{\mathfrak{A}}(P_\alpha), \\ P_\lambda & = & \bigcup_{\alpha<\lambda} P_\alpha, \text{ for limit ordinals } \lambda. \end{array}$$

By monotonicity, the P_α form an increasing sequence. It is clear therefore that over finite $\mathfrak{A}$ of size n saturation of this sequence occurs after at most n^r steps. Over infinite $\mathfrak{A}$, $\mathrm{card}(\mathfrak{A})^+$ (the cardinal successor of $\mathrm{card}(\mathfrak{A})$) provides a strict upper bound on the ordinal length of the sequence before saturation. If φ is monotone over all $\mathfrak{A}$, then the formula

$$\psi = [\mathrm{LFP}_{X,\overline{x}}\varphi]\,\overline{x}$$

is taken to describe the least fixed-point associated with F_φ. Its semantics is defined according to

$$\mathfrak{A} \models \psi[\overline{a}] \quad \text{iff} \quad \overline{a} \in \mathrm{LFP}\big(F_\varphi^{\mathfrak{A}}\big).$$

Some comments about free first-order variables in φ other than the $\overline{x}$ is in order. If φ is of the form $\varphi(X, \overline{x}, \overline{z})$ with additional free $\overline{z}$, then the $\overline{z}$ are treated as fixed parameters for the fixed-point process concerning X and $\overline{x}$; the free first-order variables of $[\mathrm{LFP}_{X,\overline{x}}\varphi]\overline{x}$ are those free in φ plus the $\overline{x}$.

Least fixed-point logic LFP is the extension of first-order logic obtained by closing with respect to the formation of least fixed-points for $\varphi(X, \overline{x})$ that are *positive* (and thus monotone) in X. It is clear that LFP strictly extends TC. Its importance in descriptive complexity is due to the fact that over linearly ordered finite structures it precisely captures PTIME. This well-known result is due to Immerman [25] and Vardi [38].

7.1 FP^2

There are several natural ways of restricting LFP to the context of two-variable logic. The straightforward notion would permit the application of LFP-operators $\mathrm{LFP}_{X,x}$ and $\mathrm{LFP}_{X,y}$ for monadic X and $\mathrm{LFP}_{X,xy}$ (and $\mathrm{LFP}_{X,yx}$) for binary X.

It is interesting to note that TC^2 is a sublogic of this liberal variant of two-variable fixed-point logic. Consider $\varphi = \varphi(x, y)$ and let $\psi(x, y) := \mathrm{LFP}_{X,y}\big(\varphi \vee \exists x(Xx \wedge \varphi)\big)$ with monadic X. Note that x is a parameter in this fixed-point! It is readily checked that $\mathrm{TC}\varphi$ is equivalent with ψ. The above liberal notion of two-variable fixed-point logic therefore is not a sublogic of infinitary two-variable logic $L^2_{\infty\omega}$, even over finite structures (compare Fact 5.1). The reason for this lies in the presence of first-order parameters in

LFP-processes. Standard arguments show that $L^2_{\infty\omega}$ is on the other hand closed, over finite structures (or indeed over any class of structures of bounded cardinality), under

- LFP-applications involving binary predicate variables, and
- LFP-applications involving monadic predicate variables as long as the first-order variable that is not directly used in the fixed-point is not free in φ.

The first type in fact subsumes the second, in the sense that $\mathrm{LFP}_{X,x}\varphi(X,x)$ for monadic X may be simulated using $\mathrm{LFP}_{X',xy}\big(x = y \wedge \varphi'(X',x,y)\big)$ with binary X'. For $\varphi'(X',x,y)$ take the result of replacing atoms Xx or Xy in φ by $X'xx$ or $X'yy$, respectively. Note that this translation is within L^2.

The following form of adjoining a least fixed-point operation to L^2 seems to be the weakest possible one, and in particular is weak enough to keep the resulting logic within $L^2_{\infty\omega}$ over any class of structures of bounded cardinality.

Definition 7.1 Let FP^2 be the simultaneous closure of L^2 under L^2-constructors and monadic least fixed-points of the form $\mathrm{LFP}_{X,x}\varphi$ where y is not free in φ or $\mathrm{LFP}_{X,y}\varphi$ where x is not free in φ. X stands for any unary second-order variable, and in both cases we require φ to be positive in X.

Fact 7.2 *Owing to the fact that no first-order parameters are allowed in fixed-point processes,* FP^2 *is a (proper) fragment of infinitary two-variable logic* $L^2_{\infty\omega}$ *over finite structures.*

This fact distinguishes FP^2 from all the other extensions of L^2 considered in this paper. In particular $\mathrm{TC}^2 \not\subseteq \mathrm{FP}^2$ since $\mathrm{TC}^2 \not\subseteq L^2_{\infty\omega}$ even over finite structures (see Fact 6.3). In fact FP^2 and TC^2 are incomparable over finite structures. For $\mathrm{FP}^2 \not\subseteq \mathrm{TC}^2$ one may use the pebble games appropriate for TC as introduced in [11] to show that for instance no sentence of TC^2 distinguishes all finite successor structures $\big(\{0,\dots,n\},<,S\big)$ with $S = \{(m,m+1) : 0 \leqslant m < n-1\}$ from their variants where S has an additional forward edge, say from $\lfloor n/2 \rfloor - 1$ to $\lfloor n/2 \rfloor + 1$. We shall see below in Lemma 7.5 (ii) that FP^2 does distinguish these classes.

An example to the effect that $\mathrm{FP}^2 \not\subseteq L^2_{\infty\omega}$ over infinite structures is obtained from the fact that the class of all well-orderings is axiomatizable in FP^2 (as we shall see in Lemma 7.5 (i)). By a well-known known result of López-Escobar [33], however, well-orderings are not even projectively axiomatizable in $L_{\infty\omega}$, let alone axiomatizable in $L^2_{\infty\omega}$. This observation should be contrasted with the fact that the class of all well-orderings is closed under $L^2_{\infty\omega}$-equivalence. In fact each individual well-ordering is characterized up to isomorphism by a single $L^2_{\infty\omega}$-sentence. The class of all well-orderings consists, however, of a class-size union of $L^2_{\infty\omega}$-equivalence classes (and thus, in particular, cannot be axiomatized by a disjunction indexed by all these classes). The appropriate extension of Fact 7.2 to arbitrary structures is the following: any FP^2-definable class of structures is closed under $L^2_{\infty\omega}$-equivalence. This observation in turn implies that $\mathrm{FP}^2 \subseteq L^2_{\infty\omega}$ in restriction to any class of structures of bounded cardinality, and in particular over finite structures.

Aside: least versus inductive fixed-points. An important variation of fixed-point logic LFP was introduced by Gurevich and Shelah in [18]. For *inflationary* or *inductive* fixed-points, the positivity (and monotonicity) requirement on formulae $\varphi(X, \overline{x})$ in fixed-points is waived. Limits are enforced instead by passage to a modified iteration scheme: rather than passing from P_α to $F_\varphi(P_\alpha) = \{\overline{a} : \varphi(P_\alpha, \overline{a})\}$ in the successor step, a cumulative evolution is built in according to $P_{\alpha+1} := P_\alpha \cup F_\varphi(P_\alpha)$. The saturation point reached in this sequence of iterations (starting from $P_0 = \emptyset$ and taking unions at limit ordinals, as before) is called the *inductive* fixed-point of φ, written $\mathrm{IFP}_{X,\overline{x}}\varphi$. The resulting extension of first-order logic IFP is termed *inductive fixed-point logic* in contrast with least fixed-point logic. Since the modified iteration scheme coincides with the original one if F_φ happens to be monotone, it is clear that LFP $\subseteq$ IFP. It is shown in [18] that over finite structures IFP and LFP are equally expressive.

It is straigtforward to modify Definition 7.1 to obtain a corresponding weak two variable fragment IFP^2. Again, it is clear that

$$L^2 \subsetneq \mathrm{FP}^2 \subseteq \mathrm{IFP}^2 \subsetneq L^2_{\infty\omega}.$$

The simulation of inductive fixed-points through least fixed-points over positive formulae in [18] involves an increase in the arity of fixed-points and correspondingly the use of more first-order variables. It is not too surprising therefore, that for our variable- and arity-restricted fragments we find $\mathrm{FP}^2 \subsetneq \mathrm{IFP}^2$ even for finite structures.

Fact 7.3 FP^2 *is contained in monadic second-order logic (with only two first-order variables).* IFP^2 *on the other hand is not contained in monadic second-order logic even for finite structures. It follows that* $\mathrm{FP}^2 \subsetneq \mathrm{IFP}^2$ *(for finite structures).*

Sketch of proof. Definability of monadic least fixed-points in monadic second-order logic is clear: in fact the very definition of $\subseteq$-minimal fixed-points is directly expressible. For an example of an IFP^2-definable class of finite structures that is not definable in monadic second-order logic, one may use a class of word models corresponding to a language that is not regular: let $\mathcal{C}$ be the class of finite structures of even cardinality (call it $2n$), linearly ordered by $<$ and with a monadic predicate P containing the first n elements. It follows from the Büchi-Elgot-Trakhtenbrot characterization of the regular languages in terms of monadic second-order logic (see e.g. [6, Chap. 5]) that $\mathcal{C}$ is not definable in monadic second-order logic within the class of finite structures $(A, <, P)$ that are linearly ordered by $<$. A definition in IFP^2, however, is easily given within this class. In L^2 we may assert that P forms a proper initial segment of $(A, <)$. The rest is taken care of by asserting that the inductive fixed-point over the following formula $\xi(X,x)$ evaluates to the full set A. Let $\varphi^U_{\min}(x)$ be the obvious L^2-formula saying that x is the $<$-least element of U, $\varphi^U_{\max}(x)$ the corresponding one for the maximal element.

$$\begin{aligned}\xi(X,x) \quad := \quad & \varphi^P_{\min}(x) \vee \varphi^{\overline{P}}_{\max}(x) \\ & \vee \Big(\exists x \exists y \big(Px \wedge \neg Py \wedge \neg Xx \wedge \neg Xy\big) \wedge \big(\varphi^{P\setminus X}_{\min}(x) \vee \varphi^{\overline{P}\setminus X}_{\max}(x)\big)\Big).\end{aligned}$$

($\overline{P}$ stands for the complement of P.) For correctness observe that in each iteration exactly one point from P and one from $\overline{P}$ will enter the inductive fixed-point. ☐

7.2 Well-foundedness, order, and successors

A typical example of the expressive power of FP^2 is the definability of the well-founded support of a binary predicate. If E is a binary predicate, then its *well-founded support* $\mathrm{WF}(E)$ is the set of those points that have no infinite descending E-chain below them: $x \in \mathrm{WF}(E)$ if and only if there is no infinite sequence $(x_i)_{i \geqslant 0}$ such that $x_0 = x$ and $Ex_{i+1}x_i$ for all $i \geqslant 0$. It is easy to check that the following formula defines the well-founded support of the predicate defined by $\varphi(x,y)$:

$$\psi(x) := \big[\mathrm{LFP}_{X,x}\big(\forall y(\varphi(y,x) \rightarrow Xy)\big)\big]x\,.$$

Thus FP^2 is closed under the WF-operation. It is instructive to note in passing that over finite structures (!) well-founded supports are also definable in TC^2.

Remark 7.4 TC^2 *is closed under the* WF*-operation over finite structures, but not over arbitrary structures, in fact not even over countably infinite structures.*

First observe that, if $\mathrm{dom}(E)$ is finite, then $\mathrm{WF}(E)$ consists of all those x that are not E-reachable from any y that is part of an E-cycle (here a loop is also considered a cycle). This is expressible in TC^2 since $\{y : (y,y) \in \mathrm{TC}(E)\}$ precisely defines the set of points on E-cycles, and $\neg\exists y\big(\exists x(x = y \wedge \mathrm{TC}(Exy)) \wedge \mathrm{TC}(Exy)\big)$ thus defines $\mathrm{WF}(E)$. ($\exists x(x = y \wedge \mathrm{TC}(Exy))$ says that $(y,y) \in \mathrm{TC}(E)$.)

Well-foundedness is not definable over arbitrary structures in TC^2, as may be inferred from the undefinability of well-order in $L_{\infty\omega}$ [33], since $\mathrm{TC}^2 \subseteq L_{\infty\omega}$ (actually transitive closures of binary relations are obviously definable in three variables using countable disjunctions). It is instructive to see, however, that even over countably infinite structures, TC^2 cannot express well-foundedness (whereas FP^2 and $L^2_{\infty\omega}$ do, as we shall see in Lemma 7.5 below). For instance, the well-founded linear ordering $(\mathbb{N},<)$ of the naturals is TC^2-equivalent with that on the naturals with a copy of the integers appended $(\mathbb{N},<) + (\mathbb{Z},<)$. Indeed, TC^2 collapses to L^2 over these two structures. To see this one can show that the following are true over $(\mathbb{N},<)$ and $(\mathbb{N},<) + (\mathbb{Z},<)$:

(i) the L^2-definable monadic predicates are just the boolean combinations of intervals $[0, m]$, for $m \in \mathbb{N}$;

(ii) the L^2-definable binary predicates are precisely the boolean combinations of $<$, equality, and monadic conditions concerning x and y separately, of the kind described in (i);

(iii) applications of the TC-operator do not lead out of the class of predicates described in (ii).

(iii) follows from the observation that for any partition of the universe into finitely many intervals $I_1 + I_2 + \cdots + I_k$, where only I_k is infinite, the following holds: if R is quantifier-free definable from $<$ in restriction to each $I_i \times I_j$, then the same is also true of $R \cup R \circ R$, and hence inductively of $\mathrm{TC}(R)$. Together (i) – (iii) imply that TC^2 can no more distinguish $(\mathbb{N}, <)$ from $(\mathbb{N}, <) + (\mathbb{Z}, <)$ than L^2 or even FO can.

An infinity axiom. Definability of $\mathrm{WF}(E)$ also immediately serves to obtain an infinity axiom in FP^2. For instance, the conjunction of $\forall x \exists y Exy$ with $\forall x \mathrm{WF}(E)x$ forces the existence of an infinite cycle-free E-path.

The real strength of FP^2 is brought out in the following lemma, which will also form the cornerstone in our proof of undecidability. It really seems intriguing that *well-orderings* are axiomatizable by a sentence of FP^2. Orderings as such, on the other hand, are not axiomatizable in FP^2, in fact the class of all orderings is not even closed under $L^2_{\infty\omega}$-equivalence (and it is transitivity that seems to be responsible).

Lemma 7.5 *(i) There is an* FP^2*-sentence* $\varphi_{wo}(<)$ *whose models are exactly all well-orderings* $(A, <)$.

(ii) There is an FP^2*-sentence* $\varphi_{wo,s}(<, S)$ *whose models are exactly all well-orderings with successor relation, i.e., all* $(A, <, S)$ *where* $(A, <)$ *is a well-ordering and* S *the associated successor relation.*

Over finite structures FP^2 therefore admits characterizations of linear orderings and linear orderings with successor, as well as a projective characterization of pure successor structures.

Proof. Let φ_0 be the L^2-sentence $\varphi_0 := \forall x \forall y \big(x = y \vee x < y \vee y < x\big)$. Let φ_{wo} be the conjunction of φ_0 with the sentence

$$\varphi_1 := \forall x \big[\mathrm{LFP}_{X,x}\big(\forall y (y < x \rightarrow Xy)\big)\big]x\,,$$

which says that $<$ is well-founded. φ_{wo} provides the desired characterization of well-orderings. φ_1 forces $<$ to be irreflexive and antisymmetric in forbidding loops and

2-cycles. The point is that $\varphi_0 \wedge \varphi_1$ forces transitivity as well: if $a_1 < a_2 < a_3$ in a model of φ_{wo}, then $a_1 \neq a_3$ since $<$ is antisymmetric, and $a_3 \not< a_1$ since $<$ may not contain a 3-cycle. Therefore $a_1 < a_3$ follows by φ_0.

Towards characterizing the successor relation S of a well-ordering $<$ consider the following formula:

$$\xi(X, x) := \forall y \big(Sxy \to \big(\forall x \big(x < y \to \forall y \big(y < x \to Xy \big) \big) \big) \big).$$

Note that ξ is positive in X. Suppose now that $(A, <)$ is a well-ordering and consider the fixed-point generation for $\mathrm{LFP}_{X,x}\xi(X, x)$ over some expansion $(A, <, S)$. We assume that $(A, <, S)$ satisfies the L^2-sentence $\varphi_2(<, S)$ which says that $S \subseteq <$, and that all elements apart from the first/last element (if it exists) have S-predecessors/-successors.

An easy induction shows that, if S is the true successor relation for $<$ then stage X_α of the fixed-point generation consists of all elements of rank less than α in $(A, <)$. Indeed, $\xi(X_\alpha, x)$ describes just those elements, from which one cannot pass to an S-successor and from there go to smaller elements twice and still remain outside X_α. With S being the true successor and X_α an initial segment of $(A, <)$ this is true just for elements in X_α and for the least element above X_α. Therefore, in this case, $\mathrm{LFP}_{X,x}\xi(X, x)$ evaluates to the full set A.

Any S that satisfies $\varphi_2(<, S)$ can only fail to be the true successor because of the existence of S-edges that jump ahead in the sense that $(a_1, a_3) \in S$ and there is an a_2 such that $a_1 < a_2 < a_3$. It is easy to see that the least such a_1 will not enter the fixed-point, in fact $\mathrm{LFP}_{X,x}\xi(X, x)$ evaluates to the initial segment of $(A, <)$ below the least such a_1.

Putting these facts together, it is clear that the following characterizes well-orderings with successor:

$$\varphi_{\mathrm{wo,s}}(<, S) := \varphi_{\mathrm{wo}}(<) \wedge \varphi_2(<, S) \wedge \forall x \big[\mathrm{LFP}_{X,x}\xi(X, x) \big] x.$$

$\square$

The following technical corollary will serve to connect the above facts about definability in FP^2 to those applications of transitive closure operations that proved useful in the characterization of local grids in the last section.

Corollary 7.6 *The following admit projective characterizations in* FP^2*:*

(1) The class of matchings, i.e., of those directed graphs (V, E) whose edge predicate E is the graph of a bijective function with disjoint domain and range.

(2) The class of all structures (V, E_1, E_2, E_3) in which each of the E_i is a matching, $\mathrm{dom}(E_2) = \mathrm{range}(E_1)$*, and* $E_3 = E_1 \cdot E_2$*.*

Proof. Let $\psi_0 \in L^2$ say that the domain and range of E are disjoint. For (1) it suffices to assert in addition that $E \subseteq S$ for an expansion $(V, E, <, S)$ where $<$ is a well-ordering with successor relation S. Any matching may be extended to the successor relation of some suitably chosen well-ordering. Conversely, functionality and injectivity of a successor S carry over to $E \subseteq S$, and, together with ψ_0, force E to be a matching. The desired projective characterization is therefore given by the FP^2-sentence

$$\varphi_{\mathrm{match}}(E, <, S) := \varphi_{\mathrm{wo,s}}(<, S) \wedge \psi_0 \wedge \forall x \forall y \big(Exy \rightarrow Sxy\big).$$

For (2) use firstly a sentence $\psi_1 \in L^2$ to say that $\mathrm{dom}(E_i) \cap \mathrm{range}(E_i) = \emptyset$ for $i = 1, 2, 3$, $\mathrm{dom}(E_2) = \mathrm{range}(E_1)$, $\mathrm{dom}(E_3) = \mathrm{dom}(E_1)$, and $\mathrm{range}(E_3) = \mathrm{range}(E_2)$. Further introduce binary $<$, S, S' and unary U to say that

(i) S is the successor relation related to the well-ordering $<$,

(ii) $U = \mathrm{dom}(E_1) \cup \mathrm{range}(E_2)$,

(iii) $E_1 \cup E_2 \subseteq S$,

(iv) S' is the successor relation associated with the well-ordering $<\!\big| U$, and that

(v) $E_3 \subseteq S'$.

(i) is expressible with the sentence $\varphi_{\mathrm{wo,s}}(<, S)$ from Lemma 7.5; (ii), (iii) and (v) are plain L^2. For (iv) a relativization of $\varphi_{\mathrm{wo,s}}$ to U will do. This sentence $\big[\varphi_{\mathrm{wo,s}}(<, S')\big]^U$ is here directly obtained from $\varphi_{\mathrm{wo,s}}(<, S')$ by replacing all first-order quantifications of the form $\exists z \eta$ or $\forall z \eta$ by $\exists z(Uz \wedge \eta)$ and $\forall z(Uz \rightarrow \eta)$, and adding a conjunct Ux to the $\xi(X, x)$ that occurs as $\mathrm{LFP}_{X,x}\xi$ in $\varphi_{\mathrm{wo,s}}$.

We claim that the conjunction of ψ_1 with sentences to the effect of (i) — (v) provides a projective characterization as required for (2) of the corollary.

Let us first argue that any of the intended models (A, E_1, E_2, E_3) possesses an expansion satisfying the above conditions. The fact that the E_i are all matchings and that $E_3 = E_1 \cdot E_2$ implies that the connected components with respect to $E_1 \cup E_2 \cup E_3$ are triangles of the form depicted in (a) below.

Order each of these triangles so as to make the E_1- and E_2-edges consecutive S-successor edges. Then extend this local ordering to a global well-ordering $<$. With the approriate U and S' all conditions are met.

Conversely consider any expanded structure $(A, E_1, E_2, E_3, <, S, S', U)$ satisfying ψ_1 and (i) — (iv). The E_i must be bijections since they can be embedded into successor relations according to (iii) and (v). By ψ_1 then they are matchings. ψ_1 and (ii) then force local configurations with respect to E_1, E_2 and U as in (b) of the sketch below. As $E_1 \cup E_2 \subseteq S$ by (iii), condition (iv) implies that the diagonal $E_1 \cdot E_2$-edge is part of

S'. There must also be an E_3-edge going out from the lower left node, since $\mathrm{dom}(E_1) = \mathrm{dom}(E_3)$. If $E_3 \subseteq S'$ according to (v), then the diagonal $E_1 \cdot E_2$-edge and this E_3-edge must coincide. This implies that $E_3 = E_1 \cdot E_2$, globally. $\square$

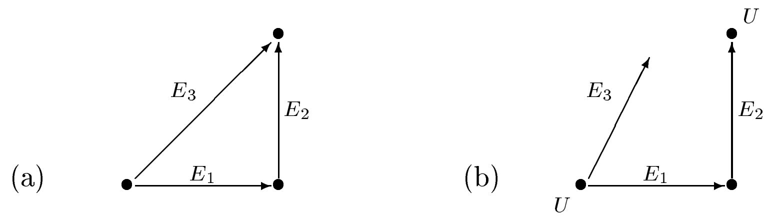

7.3 Local grids in FP^2

Proposition 7.7 *There is a projective characterization in* FP^2 *of the class of even local grids without boundary.*

Sketch of proof. We aim at an axiomatization of those expansions of even local grids without boundary that locally look like $\mathcal{G}_{\mathbb{Z}}{}^*$ (compare the sketch in the proof of Proposition 5.5).

In addition to the obvious L^2-definable conditions (cf. the sentence φ_0 in the proof of Proposition 5.5), use projective characterizations according to Corollary 7.6 (i) to enforce that the V_i, H_i and D_{ij} are matchings. It remains to enforce commutativity of h and v. Recall the formulae $\eta_{ij}(x)$ that describe the four types of vertices with respect to incoming and outgoing H_i- and V_j-edges in expansions of even local grids that locally look like $\mathcal{G}_{\mathbb{Z}}{}^*$. Once more, we look at derived edge relations that are L^2-definable from the H_i and V_j:

$$\begin{aligned} H_{ij}xy \quad &:\Longleftrightarrow \quad H_ixy \wedge \eta_{ij}(x), \\ V_{ij}xy \quad &:\Longleftrightarrow \quad V_jxy \wedge \eta_{ij}(x). \end{aligned}$$

Consider a vertex of type $(0,0)$, i.e. one that satisfies η_{00}. Using Corollary 7.6 (ii), we may assert projectively the following (which obviously hold true in the intended models):

- H_{00}, V_{10} and D_{00} are matchings for which $D_{00} = H_{00} \cdot V_{10}$

- V_{00}, H_{01} and D_{00} are matchings for which $D_{00} = V_{00} \cdot H_{01}$.

$$\begin{tikzcd}
(0,1) \arrow[r, "H_{01}"] & (1,1) \\
(0,0) \arrow[u, "V_{00}"] \arrow[r, "H_{00}"'] \arrow[ur, "D_{00}"] & (1,0) \arrow[u, "V_{10}"']
\end{tikzcd}$$

It is also clear that these assertions force H and V to commute in vertices of type $(0,0)$, since D_{00} is a matching (and thus in particular functional). $\square$

7.4 Σ_1^1-hardness

Proposition 7.8 FP^2 *admits a projective characterization of $\mathcal{G}_{\mathbb{Z}}$ (up to isomorphism). By Theorem 4.9, therefore, sat(FP^2) is Σ_1^1-hard.*

We do not know whether this can be strengthened to a completeness result.

Sketch of proof. The proof is similar to that given for DTC^2 above, see Proposition 6.6. Introducing monadic B, C, and O, and binary $<_B$ and $<_C$, we may stipulate the following in FP^2 (recall that $R^{\leftrightarrow}$ denotes the reflexive symmetric closure of R):

(i) $O = B \cap C$ consists of a single vertex.

(ii) B is closed with respect to $H^{\leftrightarrow}$, C with respect to $V^{\leftrightarrow}$.

(iii) $<_B$ is a well-ordering on those elements $b \in B$ such that there is an H-path from an element in O to b, with least element in O and with successor H and without limit points; similarly for $<_C$ and V.

(iv) Every vertex is connected to O (with respect to $H^{\leftrightarrow} \cup V^{\leftrightarrow}$).

The last condition simply is $\forall x \big[\mathrm{LFP}_{X,x}\big(Ox \vee \exists y \big(Xy \wedge (H^{\leftrightarrow}xy \vee V^{\leftrightarrow}xy)\big)\big)\big]x$. A local grid without boundary that satisfies these conditions is isomorphic to $\mathcal{G}_{\mathbb{Z}}$. $\square$

8 Excursion on related process logics: CTL and L_μ

This aside contains some background material that serves to relate the two main extensions of L^2 discussed so far, TC^2 and FP^2, to two well known decidable extensions of modal logic that play an important role for the specification and verification of concurrent processes: computation tree logic CTL and the propositional μ-calculus L_μ. While

L_μ extends ML by a monadic least fixed-point construct, CTL allows one to make reachability assertions essentially through path quantification. Decidability of CTL is due to Emerson and Halpern [8], decidability of L_μ to Kozen and Parikh [30]. Both CTL and L_μ have the finite model property (as well as a tree model property), see [29] for L_μ.

The relative expressive strengths of these logics are represented in the following diagram, where arrows indicate strict inclusion in expressive power. The information in the diagram is complete in the sense that relations that are not settled by sequences of arrows may be read as incomparability statements even over finite structures (see in particular Fact 5.1, the discussion following Fact 7.2, and Proposition 8.1 below). It is clear that FP^2 qualifies as *the* natural common extension of L^2 and L_μ. Similarly, over finite structures, TC^2 is a canonical common extension of L^2 and CTL [27].

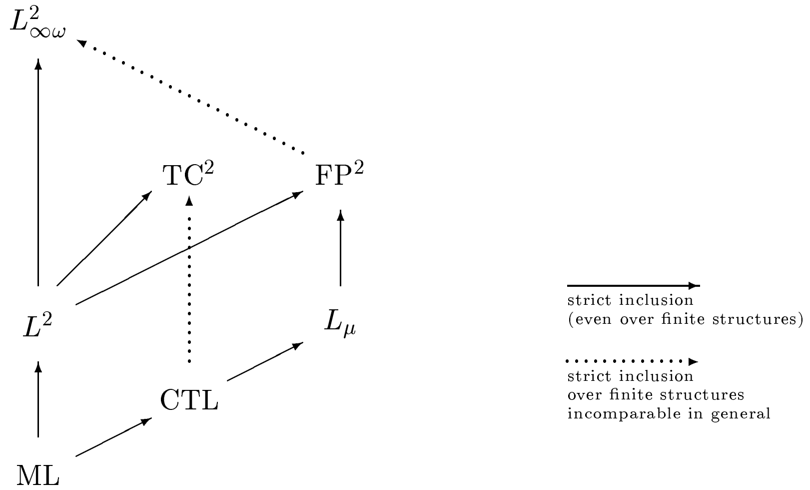

Kripke structures. A *Kripke structure* is a structure $(V, E_1, \dots, E_m, P_1, \dots, P_n)$ with binary E_k and unary P_l. The elements of V are usually thought of as *worlds* or *states*. The P_l correspond to basic propositions that may or may not hold of a world or state. The E_k describe accessibility relations between worlds or states.

Modal logic. It is well known that propositional modal logic ML corresponds to that fragment of L^2 that is generated from the atomic formulae $P_i x$ and $P_i y$ through closure under Boolean connectives and modal quantification w.r.t. the accessibility relations E_k, according to

$$\Box_k \varphi(x) \equiv \forall y \big(E_k xy \to \varphi(y)\big) \qquad \text{and} \qquad \Diamond_k \varphi(x) \equiv \exists y \big(E_k xy \land \varphi(y)\big).$$

The propositional μ-calculus. Viewing ML as a fragment of L^2, we may think of L_μ as the monadic least fixed-point extension of the modal fragment of L^2. This enriches $\mathrm{ML} \subseteq L^2$ with the usual rule for the formation of least fixed points: if $\varphi(X, x)$ is positive in the (propositional, i.e. monadic predicate) variable X, then $[\mu_{X,x}\varphi](x)$ is a formula whose semantics is that of $[\mathrm{LFP}_{X,x}\varphi](x)$. It is obvious that $L_\mu \subseteq \mathrm{FP}^2$, just as $\mathrm{ML} \subseteq L^2$.

Sometimes L_μ is presented as an extension of process logics like propositional dynamic logic PDL, which allow one to define new *transition formulae*, corresponding to formulae $\alpha(x, y)$ in two free variables, which are then inductively used like the basic accessibility relations E_k in modal constructs $\Box_\alpha$ and $\Diamond_\alpha$ with the obvious semantics. The typical constructors considered are relational composition, union, and transitive closure. It is therefore clear that for the corresponding transition formulae, inclusion in FP^2 and even in $L^2_{\infty\omega}$ fails. As far as *state formulae* are concerned, however, the above inclusion in FP^2 obtains. In other words, when restricted to applications in modal operators, then even those transition formulae which are not in FP^2 — like transitive closures of the basic accessibilities E_k — can only give rise to state formulae that are in FP^2. Consider $\psi = \Diamond_\alpha\varphi$ where $\alpha = \mathrm{TC}(E_k)$. It is directly checked that ψ is equivalent with $\mu_{X,x}\big(\varphi \vee \Diamond_k Xx\big)$. The use of transitive closures in modal operators is thus seen to reduce to a well-foundedness assertion.

Computation tree logic. It is customary here to consider only one basic accessibility relation E. The syntax and semantics of CTL involves *state formulae* and *path formulae*. The latter are assertions about paths in Kripke structures, where a path is an infinite E-path. For this semantic intuition, the Kripke structures under consideration should not have dead ends, in the sense that they satisfy the L^2-axiom $\forall x \exists y Exy$. The translation into a first-order framework is less straightforward for CTL, since paths are not first-order objects. For the semantics of state formulae, however, path constructs can be eliminated through μ-constructs, and — in the case of finite structures — through the use of transitive closure operators. Let path formulae be denoted $\pi(p)$ where p is treated as a formal variable ranging over paths $p = v_0v_1v_2\ldots$, where $v_i \in V$ and $(v_i, v_{i+1}) \in E$.

Beyond the modal rules, which apply to the formation of state formulae, CTL has the following rules for the formation of path formule from state formulae, and rules for the formation of state formule from path formulae:

(i) If φ_1 and φ_2 are state formulae, then $\big(\mathrm{next}\,\varphi\big)(p)$ and $\big(\varphi_1 \text{ until } \varphi_2\big)(p)$ are path formulae.

(ii) If $\pi(p)$ is a path formula then $\forall p\pi(p)$ and $\exists p\pi(p)$ are state formulae.

The semantics of these is the following: $\big(\mathrm{next}\,\varphi\big)$ holds of $p = v_0v_1v_2\ldots$ if φ holds at v_1. $\big(\varphi_1 \text{ until } \varphi_2\big)$ holds of $p = v_0v_1v_2\ldots$ if for some j, φ_2 holds at v_j and φ_1 holds at v_i for all $i < j$. $\forall p\pi(p)$ is true at v if all paths of the form $p = vv_1v_2\dots$ satisfy π. Similarly, $\exists p\pi(p)$ asserts the existence of a path of the form $p = vv_1v_2\dots$ satisfying π.

It is well known that CTL $\subseteq L_\mu$ (w.r.t. state formulae), and Immerman and Vardi have recently applied a translation of CTL into TC^2 over finite structures to model checking considerations [27].

Proposition 8.1 *While* CTL $\subseteq \mathrm{TC}^2$ *over finite structures, CTL and* TC^2 *are incomparable in general, and even over just countably infinite structures.*

Sketch of proof. For the inclusion, the only steps in the generation of state formulae that are of interest are the quantifications over paths $\exists p\pi(p)$ and $\forall p\pi(p)$. We indicate how to reduce these with respect to the build-up of the path formula π. Clearly, a path quantifier applied to $\pi(p) = \mathrm{next}\,\varphi$, correspondes exactly to a modal quantifier applied to φ. Consider then $\pi(p) = \varphi_1$ until φ_2. $\exists p\pi(p)$ is obviously equivalent (over arbitrary structures) to $\varphi_2(x) \vee \exists y\big(\mathrm{TC}(Exy \wedge \varphi_1(x)) \wedge \varphi_2(y)\big)$. $\forall p(\varphi_1$ until $\varphi_2)$, on the other hand, exactly defines the set of states in the well-founded support of the relation defined by

$$\xi(x,y) = \neg\varphi_2(y) \wedge \Big(\big(\varphi_1(y) \wedge Eyx \wedge \neg\varphi_2(x)\big) \vee \big(\neg\varphi_1(y) \wedge x = y\big)\Big),$$

which is essentially the restriction of E^{-1} to $\neg\varphi_2$. Consider first the states that have no predecessors under ξ (those in the first level of an inductive generation of $\mathrm{WF}(\xi)$): these are those in φ_2 plus those in φ_1 that have E-edges only to states in φ_2. In the further inductive generation of $\mathrm{WF}(\xi)$, one need only consider states in $\neg\varphi_2$. These can only be in $\mathrm{WF}(\xi)$ if they satisfy φ_1 (otherwise they have ξ-loops). But for states in $\neg\varphi_2 \wedge \varphi_1$, $\xi(x,y)$ becomes equivalent with Eyx.

Over finite structures, therefore, CTL $\subseteq \mathrm{TC}^2$: the formation of the well-founded support reduces to a TC^2-application over finite structures by Remark 7.4. An example to the effect that CTL $\not\subseteq \mathrm{TC}^2$ over countable structures is directly constructed using Remark 7.4: $\forall p(Px\,\mathrm{until}\,\neg Px)$ precisely expresses well-foundedness of $E^{-1}|P$. $\square$

9 Two-variable logic with cardinality comparison

It is well-known that one of the major limitations of first-order logic (with respect to expressiveness) is its inability to count. For instance, there is no first-order sentence that defines parity, in the sense that its finite models are precisely the structures of even cardinality. In fact even much stronger languages like fixed-point logic or the infinitary logic with bounded number of variables $L^{\omega}_{\infty\omega}$ cannot express parity; this can be shown by a simple argument using Ehrenfeucht-Fraïssé games (see [6]). This issue is particularly relevant in finite model theory and in databases, where one of the major goals is the design of logics or query languages that capture complexity classes, such as polynomial time or logarithmic space. Since counting is a computationally simple task it is natural to investigate logics with counting constructs (such as counting quantifiers, counting terms or generalized quantifiers). We refer to [14, 35] for a detailed discussion of logics with counting.

Let C^2 be the extension of L^2 by counting quantifiers of the form $\exists^{\geqslant n}$ and $\exists^{\leqslant n}$, for arbitrary $n \in \omega$. The semantics of these quantifiers is the obvious one. It is not difficult to see that C^2 contains infinity axioms. For instance, there is a C^2-sentence with one binary predicate F, expressing that F is the graph of an injective but not surjective unary function: $\forall x \exists^{=1} y Fxy \wedge \forall y \exists^{\leqslant 1} x Fxy \wedge \exists y \forall x \neg Fxy$.

There are other forms of adding counting to a logic, for instance via *counting terms* or via the *cardinality comparison quantifier*, also called *Härtig quantifier*.

Counting terms. In logics with counting terms we have the possibility to build for every formula φ and every variable x a term $\#_x[\varphi]$, taking cardinals as values. The free variables of φ different from x remain free in $\#_x[\varphi]$.

The semantics of $\#_x[\varphi]$ is defined as follows. As usual, the notation $\varphi(x, \bar{z})$ indicates that the free variables of φ are among $x, \bar{z}$. Given a formula $\varphi(x, \bar{z})$, a structure $\mathfrak{A}$ and a valuation $\bar{c}$ for $\bar{z}$, then the meaning of $\#_x[\varphi]$ for $\mathfrak{A}$ and $\bar{c}$ is

$$\#_x[\varphi]^{\mathfrak{A}, \bar{c}} := |\{a : \mathfrak{A} \models \varphi[a, \bar{c}]\}|.$$

Fixing a logic L and any collection of basic relations Q on cardinals we can build a counting extension of L containing formulae of the form $Qt_1 \cdots t_m$ where $t_1, \ldots, t_m$ are counting terms. The simplest case is obtained by just allowing equality of counting terms.

The Härtig quantifier. A slightly different presentation of first-order logic (or other logics) with equality of counting terms involves a particular generalized quantifier, namely the cardinality comparison quantifier I, introduced by Härtig [21]. For a survey on this quantifier, we refer to [24]; more background on generalized quantifiers can be found in [5].

The extension $L[I]$ of a logic L by the Härtig quantifier is defined by adding to L the following rule for building formulae: given two formulae φ, ψ and two (not necessarily distinct) variables x, y, the expression

$$(Ix, y\ \varphi, \psi)$$

is a formula, saying that the set of elements x satisfying φ has the same cardinality as the set of elements y satisfying ψ.

The set of free variables of this formula is

$$\text{free}((Ix, y\ \varphi, \psi)) = (\text{free}(\varphi) - \{x\}) \cup (\text{free}(\psi) - \{y\}).$$

Formally, the semantics of such formulae can be defined by the equivalence

$$(Ix, y\ \varphi, \psi) \equiv (\#_x[\varphi] = \#_y[\psi]).$$

Thus, given a formula $(Ix, y\ \varphi(x, y, \bar{z}), \psi(x, y, \bar{z}))(x, y, \bar{z})$, a structure $\mathfrak{A}$ (of appropriate vocabulary) and valuations $a, b, \bar{c}$ for $x, y, \bar{z}$, respectively, then

$$\mathfrak{A} \models (Ix, y\ \varphi, \psi)[a, b, \bar{c}] \quad \text{iff} \quad \#_x[\varphi]^{\mathfrak{A}, b, \bar{c}} = \#_y[\psi]^{\mathfrak{A}, a, \bar{c}}.$$

We prove that the satisfiability problem for $L^2[I]$ is undecidable. Note that this does not carry over to C^2. The two counting logics $L^2[I]$ and C^2 are incomparable with respect to expressive power, and in fact we prove in a companion paper to the present one that the satisfiability problem and the finite satisfiability problem for C^2 are decidable [15].

The power of the Härtig quantifier in the context of L^2 results from the fact that – unlike the counting quantifiers $\exists^{\geqslant n}$ or $\exists^{\leqslant n}$ – application of I does not necessarily reduce the number of free variables. Indeed if x, y are free in $\varphi(x, y)$ and $\psi(x, y)$, then both x and y are free also in $(Ix, y\ \varphi, \psi)(x, y)$.

For instance, we can axiomatize in $L^2[I]$ the class of *regular graphs* $G = (V, E)$ by the formula

$$\psi_{\mathrm{reg}} := \forall x \forall y (\neg Exx \wedge (Exy \rightarrow Eyx) \wedge (Ix, y\ Eyx, Exy)(x, y)).$$

Indeed $G \models \psi_{\mathrm{reg}}$ if and only if E is irreflexive and symmetric and for all pairs of nodes u, v the number $\#_x[Evx]$ of neighbours of v is the same as the number $\#_y[Euy]$ of neighbours of u. (In fact, due to symmetry of E, we could just as well use the subformula $(Ix, y\ Exy, Exy)$ rather than $(Ix, y\ Eyx, Exy)$.)

Theorem 9.1 *$L^2[I]$ is a conservative reduction class. Thus the satisfiability problem and the finite satisfiability problem for $L^2[I]$ are undecidable.*

Proof. By Theorem 4.9 it suffices to show that $L^2[I]$ admits a projective characterization of a rich class of local grids.

Let E be a binary relation symbol. We write $N_E(x)$ for $\#_y[Exy]$ and similarly $N_E(y)$ for $\#_x[Eyx]$ to specify the number of outgoing E-edges at points x, y. Clearly the statements $N_E(y) = 0$ and $N_E(y) = N_E(x)$ are expressible by a formula in $L^2[I]$. We observe that the Härtig quantifier allows us to express not just equality of counting terms, but also a kind of successor relation when applied to irreflexive relations.

Lemma 9.2 *Let $\mathcal{C}$ be a class of structures such that $\mathfrak{A} \models \forall x \neg Exx$ for all $\mathfrak{A} \in \mathcal{C}$. Then there exists a formula $\psi(x, y) \in L^2[I]$ expressing on $\mathcal{C}$ that $N_E(y) = N_E(x) + 1$.*

Proof. Take $\psi(x,y) := [Ix,y\ Eyx, Exy \vee (x=y)](x,y)$. $\square$

We first present a projective characterization of the infinite grid $\mathcal{G}_{\mathbb{N}}$ by a sentence of $L^2[I]$ containing the binary predicates H, V, E, F. Here H and V are the horizontal and vertical adjacency relations and from each point (m,n) we have precisely m outgoing E-edges and n outgoing F-edges.

The desired $L^2[I]$-sentence is the conjunction of the following.

$$\begin{aligned}\alpha \ := \ & \forall x \exists y Hxy \wedge \forall x \exists y Vxy \\ & \wedge \forall x(\neg Hxx \wedge \neg Vxx \wedge \neg Exx \wedge \neg Fxx) \\ & \wedge \exists x \forall y(\neg Exy \wedge \neg Fxy) \\ & \wedge \forall x \forall y\big((N_E(x) = N_E(y) \wedge N_F(x) = N_F(y)) \rightarrow x = y\big).\end{aligned}$$

$$\begin{aligned}\beta \ := \ & \forall x \forall y\big(Hxy \rightarrow (N_E(y) = N_E(x)+1 \wedge N_F(y) = N_F(x))\big) \\ & \wedge \forall x \forall y\big(Vxy \rightarrow (N_E(y) = N_E(x) \wedge N_F(y) = N_F(x)+1)\big).\end{aligned}$$

Note that α enforces irreflexivity of E and F. Thus the statements $N_E(y) = N_E(x)+1$ and $N_F(y) = N_F(x)+1$ are indeed expressible in $L^2[I]$ by Lemma 9.2. By expanding $\mathcal{G}_{\mathbb{N}}$ with any pair of relations E, F such that (m,n) has precisely m outgoing E-edges and n outgoing F-edges we obtain a model of $\alpha \wedge \beta$. Conversely, if $\mathfrak{A} = (A, H, V, E, F) \models \alpha \wedge \beta$ then $(A, H, V) \cong \mathcal{G}_{\mathbb{N}}$, since α enforces that for all $m, n \in \mathbb{N}$ there is precisely one point in A with m outgoing E-edges and n outgoing F-edges, and β axiomatizes the horizontal and vertical adjacency relations in the correct way.

Thus, $(A, H, V) \cong \mathcal{G}_{\mathbb{N}}$ if and only if there exist relations $E, F \subseteq A \times A$ such that $\mathfrak{A} = (A, H, V, E, F) \models \alpha \wedge \beta$. This already suffices to establish the undecidability of $sat(L^2[I])$.

A projective characterization of $\{\mathcal{G}_{\mathbb{N}}\} \cup \{\mathcal{G}_m : m \in \mathbb{N}\}$ (a rich class of local grids) can be obtained by adding two unary predicates U, R and by replacing β with a formula γ that is the conjunction of the following:

$$\begin{aligned}& \exists x Ux \leftrightarrow \exists x Rx, \\ & \forall x \forall y\big(Hxy \rightarrow (Ux \leftrightarrow Uy)\big), \\ & \forall x \forall y\big(Vxy \rightarrow (Rx \leftrightarrow Ry)\big), \\ & \forall x \forall y\Big(\big(Hxy \wedge \neg Rx\big) \rightarrow \big(N_F(y) = N_F(x) \wedge N_E(y) = N_E(x)+1\big)\Big), \\ & \forall x \forall y\Big(\big(Hxy \wedge Rx\big) \rightarrow \big(N_F(y) = N_F(x) \wedge N_E(y) = 0\big)\Big), \\ & \forall x \forall y\Big(\big(Vxy \wedge \neg Ux\big) \rightarrow \big(N_E(y) = N_E(x) \wedge N_F(y) = N_F(x)+1\big)\Big), \\ & \forall x \forall y\Big(\big(Vxy \wedge Ux\big) \rightarrow \big(N_E(y) = N_E(x) \wedge N_F(y) = 0\big)\Big), \\ & \forall x\big((Ux \wedge Rx) \rightarrow N_F(x) = N_E(x)\big).\end{aligned}$$

Note that U and R play, respectively, the role of the uppermost row and the rightmost column of the space to be tiled.

Indeed, for any structure (A, H, V, R, U) there exist relations E, F such that

$$(A, H, V, R, U, E, F) \models \alpha \wedge \gamma$$

if and only if either $(A, H, V) \cong \mathcal{G}_{\mathbb{N}}$ and $R = U = \varnothing$, or $(A, H, V) \cong \mathcal{G}_m$ (for some $m \in \mathbb{N}$) and there is an isomorphism $\pi : (A, H, V) \to \mathcal{G}_m$ with $\pi(R) = \{(m-1, j) : j = 0, \ldots, m-1\}$ (the rightmost column) and $\pi(U) = \{(i, m-1) : i = 0, \ldots, m-1\}$ (the uppermost row).

Thus $\alpha \wedge \gamma$ axiomatizes a rich class of grids. □

Further, we can also eliminate equality. Let $L^2[I]^-$ be the set of formulae in $L^2[I]$ that do not contain the equality sign.

Theorem 9.3 *$L^2[I]^-$ is a conservative reduction class.*

Proof. In the proof of Theorem 9.1, equality was used at two places. To express that $N_E(y) = N_E(x) + 1$ and to ensure that points with the same number of E- and F-neighbours coincide (see the last clause in the definition of α).

We modify the construction as follows.

Let A, B be unary predicates. In $L^2[I]^-$ we can express that A and B are nonempty, disjoint sets of the same cardinality such that $|A| \neq |A \cup B|$. This forces A and B to be finite.

We then add the clause $\forall x \forall y (Exy \to \neg Ay)$ and replace the condition $N_E(y) = N_E(x) + 1$ by $N_E(y) = N_E(x) + |A|$ which, given that no E-arc has its endpoint in A, is expressible by

$$(Ix, y\ Eyx, Exy \vee Ay).$$

Now the grids can be axiomatized in a similar way as above. In any model of φ, the points $(m, n) \in \mathbb{N} \times \mathbb{N}$ are represented by elements with $m|A|$ E-neighbours and $n|A|$ F-neighbours.

The clause saying that points with the same number of E- and F-neighbours coincide is not really needed. It can be replaced by the weaker condition that such points are indistinguishable by atomic formulae (e.g., that they are tiled with same domino). □

A variant of a cardinality comparison quantifier is the *Rescher quantifier J*. With the Rescher quantifier we build formulae

$$(Jx, y\ \varphi, \psi)$$

expressing that $\#_x[\varphi] < \#_y[\psi]$. Clearly the Härtig quantifier is expressible with the Rescher quantifier since

$$(Ix, y\ \varphi, \psi) \equiv \neg(Jx, y\ \varphi, \psi) \wedge \neg(Jy, x\ \psi, \varphi).$$

Corollary 9.4 *$L^2[J]$ is a conservative reduction class.*

Given that the satisfiability problem for $L^2[I]$ is undecidable, we may ask about $L^1[I]$, first-order logic with just one variable, extended with cardinality comparison. Of course we can restrict attention to monadic predicates, and in fact, also cardinality comparison can be reduced to statements $|F| = |G|$ for atomic predicates F, G. Indeed, every subformula $(Ix, x\ \varphi, \psi)$ has no free variable and may be replaced by

$$\forall x(F_\varphi x \leftrightarrow \varphi) \wedge \forall x(F_\psi x \leftrightarrow \psi) \wedge |F_\varphi| = |F_\psi|$$

where F_φ and F_ψ are new predicates. In this way any formula of $L^1[I]$ is transformed into a formula with cardinality comparison restricted to atomic predicates, which is satisfiable if and only if the original formula is.

Note that $L^1[I]$ does not have the finite model property. Indeed the sentence

$$\forall x(Fx \rightarrow Gx) \wedge \exists x(Gx \wedge \neg Fx) \wedge |F| = |G|$$

is an infinity axiom. Since we can say that two sets F, G are infinite and of different cardinality, it does not have the Löwenheim-Skolem property either, i.e., there are satisfiable sentences without countable models. Nevertheless, since a sentence in $L^1[I]$ can express only inclusion relations and cardinality comparisons of finitely many sets it is clear that the satisfiability problem for $L^1[I]$ is decidable.

Finally, note that the axiom $\alpha \wedge \beta$ in the proof of Theorem 9.1 gives a projective characterization of $\mathcal{G}_{\mathbb{N}}$. By Theorem 4.9 we conclude:

Corollary 9.5 *$sat(L^2[I])$ is Σ^1_1-hard.*

10 Two-variable logic with the choice operator

First-order logic with Hilbert's choice operator (the so called ε-operator) extends first-order logic with equality by an additional term building rule. If ψ is a formula and x a variable, then $\varepsilon x\psi$ is a term, read: one x such that ψ. The interpretation of ε over a given domain A is defined by a choice function $F : \mathcal{P}(A) \rightarrow A$. (Recall that a choice function satifies the condition that $F(X) \in X$ for all nonempty X.) Thus, $\varepsilon x\psi$ is interpreted by $F(\{x : \psi(x)\})$.

This logic is also known as ε-logic [31] and is closely related to Hermes' term logic [23]. Note that the syntax of ε-logic is rather unusual since it allows one to build terms from formulae.

The ε-operators plays a role similar to that of a quantifier. In fact, existential and universal quantifiers can be easily expressed in terms of the ε-operator:

$$\begin{array}{rcl}\exists x\psi(x) & \equiv & \psi(\varepsilon x\psi) \\ \forall x\psi(x) & \equiv & \psi(\varepsilon x\neg\psi)\end{array}$$

A sentence ψ of ε-logic is *satisfiable* if there exists a first-order structure $\mathfrak{A}$ and a choice function F on the domain of $\mathfrak{A}$ such that $(\mathfrak{A}, F) \models \psi$. Heidler [22] proved that the satisfiability problem for ε-logic is undecidable even without any function or predicate symbol besides equality. We prove here a stronger result imposing the additional restriction that the formulae contain only two variables.

Definition 10.1 Let ε-L^2 be the fragment of ε-logic consisting of the formulae that contain only the two variables x, y.

The following theorem shows that the choice operator and equality alone, without any other function, predicate or quantifier, suffice for yielding an unsolvable decision problem for two-variable logic.

Theorem 10.2 *The class of sentences in ε-L^2 built from the variables x, y using only equality, Boolean connectives and the ε-operator is a conservative reduction class.*

Proof. In ε-L^2 we can build the following terms:

$$\begin{array}{rcl} c & := & \varepsilon y(y = y) \\ ft & := & \varepsilon y(y \neq t) \\ gt & := & \varepsilon y(y \neq t \land y \neq c)\end{array}$$

where t is any previously defined term (without free occurrences of y).

Note that the functions f, g defined in this way can be composed arbitrarily, e.g., fgx is an abbreviation for the term

$$\varepsilon y(y \neq \varepsilon y(y \neq x \land y \neq \varepsilon y(y = y))).$$

Over any domain A with choice function F, these terms define an element $c = F(A) \in A$ and functions $f, g : A \to A$ such that

$$f(a) = F(A - \{a\}), \qquad g(a) = F(A - \{a, c\})$$

for all $a \in A$. If A has at least three elements, then c, f, g always satisfy the following condition, for all $a \in A$:

$$(*) \qquad f(a) \neq a,\ g(a) \neq a,\ g(a) \neq c,\ f(c) = g(c).$$

Conversely, for every domain A of cardinality greater than two, with an element c and functions $f, g : A \to A$ satisfying $(*)$ there exists a choice function F defining c, f, g as described.

The rest is routine. One could define more functions in a similar way as f and g and encode the domino problem along the lines described in the previous sections.

A simpler way to complete the proof is by reduction from the conservative reduction class formed by the sentences of form $\forall x \varphi$ where φ is quantifier-free and contains two unary functions h, h' and equality. It was proved by Gurevich [16] that this class is a conservative reduction class. In fact, the class remains conservative when restricted to structures $\mathfrak{A} = (A, h, h')$ where the two functions h, h' have no fixed points (see [3, Sect. 3.1]).

Now translate any sentence $\psi = \forall x \varphi$ of this form into

$$\psi^* := \exists x(x \neq c) \land \forall x(x \neq c \to (fx \neq c \land gx \neq c \land \varphi[h/f, h'/g]))$$

of ε-L^2, where $\varphi[h/f, h'/g]$ is obtained from φ by replacing the functions h and h' by f and g, respectively.

The quantifiers can be defined in terms of the ε-operator as described above. To complete the proof it remains to show that the reduction from ψ to ψ^* preserves satisfiability and finite satisfiability.

Suppose that $\mathfrak{A} = (A, h, h') \models \psi$ where h and h' have no fixed points (and therefore $|A| \geqslant 2$). Let $B := A \cup \{c\}$ where c is a new element. Then there exists a choice function $F : \mathcal{P}(B) \to B$ such that the functions defined by f, g on (B, F) coincide on A with h, h'. Thus $(B, F) \models \psi^*$.

Conversely, suppose that $(B, F) \models \psi^*$. Let $A = B - \{c\}$ and let h, h' be the restrictions of the functions f, g to A (note that A is nonempty and closed under f, g since $f(a) \neq c$ and $g(a) \neq c$ for all $a \in A$). Obviously, $(A, h, h') \models \psi$. □

References

[1] F. Baader, H. Bürckert, B. Hollunder, A. Laux, and W. Nutt, *Terminologische Logiken*, KI, 3/92 (1992), pp. 23–33.

[2] R. Berger, *The undecidability of the domino problem*, Mem. AMS, 66 (1966).

[3] E. Börger, E. Grädel, and Y. Gurevich, *The Classical Decision Problem*, Springer, 1997.

[4] A. Borgida, *On the relative expressiveness of description logics and predicate logics*, Artificial Intelligence, 82, 1996, pp. 353–367.

[5] H.-D. Ebbinghaus, *Extended logics: The general framework*, in Model-Theoretic Logics, J. Barwise and S. Feferman, eds., Springer, 1985, pp. 25–76.

[6] H.-D. Ebbinghaus and J. Flum, *Finite Model Theory*, Springer, 1995.

[7] A. Ehrenfeucht, *Decidability of the theory of one function*, Notices of the AMS, 6 (1959), p. 268.

[8] E. Emerson and J. Halpern, *Decision procedures and expressiveness in the temporal logic of branching time*, Journal of Computer and System Sciences, 30 (1985), pp. 1–24.

[9] M. Fürer, *The computational complexity of the unconstrained limited domino problem (with implications for logical decision problems)*, in Logic and Machines: Desision Problems and Complexity, Lecture Notes in Computer Science No. 171, Springer, 1984, pp. 312–319.

[10] E. Grädel, *Dominoes and the complexity of subclasses of logical theories*, Annals of Pure and Applied Logic, 43 (1989), pp. 1–30.

[11] ——, *On transitive closure logic*, in Proceedings of 5th Workshop on Computer Science Logic CSL '91, Bern 1991, Lecture Notes in Computer Science No. 626, Springer, 1991, pp. 149–163.

[12] E. Grädel, P. Kolaitis, and M. Vardi, *On the decision problem for two-variable first-order logic*, Bulletin of Symbolic Logic 3 (1997), 53–69.

[13] E. Grädel and G. McColm, *On the power of deterministic transitive closures*, Information and Computation, 119 (1995), pp. 129–135.

[14] E. Grädel and M. Otto, *Inductive definability with counting on finite structures*, in Proceedings of 6th Workshop on Computer Science Logic CSL '92, San Miniato 1992, Selected Papers, Lecture Notes in Computer Science No. 702, Springer, 1993, pp. 231–247.

[15] E. Grädel, M. Otto, and E. Rosen, *Two-variable logic with counting is decidable*, Proceedings of 12th IEEE Symposium on Logic in Computer Science LICS '97, Warsaw (1997).

[16] Y. Gurevich, *The decision problem for standard classes*, Journal of Symbolic Logic, 41 (1976), pp. 460–464.

[17] Y. Gurevich and I. Koryakov, *Remarks on Berger's paper on the domino problem*, Siberian Math. Journal, 13 (1972), pp. 319–321.

[18] Y. Gurevich and S. Shelah, *Fixed-point extensions of first-order logic*, Annals of Pure and Applied Logic, 32 (1986), pp. 265–280.

[19] D. Harel, *Recurring dominoes: Making the highly undecidable highly understandable*, Annals of Discrete Mathematics, 24 (1985), pp. 51–72.

[20] ——, *Effective transformations on infinite trees, with applications to high undecidability, dominoes and fairness*, Journal of the ACM, 33 (1986), pp. 224–248.

[21] K. Härtig, *Über einen Quantifikator mit zwei Wirkungsbereichen*, in Colloquium on the foundations of mathematics, mathematical machines and their applications, L. Kalmár, ed., Akadémiai Kiadó, Budapest, 1962, pp. 31–36.

[22] K. Heidler, *Untersuchungen zur Reduktionstheorie des Entscheidungsproblems in der Prädikaten- und Termlogik*, PhD thesis, Universität Freiburg, 1973.

[23] H. Hermes, *Term logic with choice operator*, Springer Lecture Notes in Mathematics No. 6 (2nd ed.), 1970.

[24] H. Herre, M. Krynicki, A. Pinus, and J. Väänänen, *The Härtig quantifier: a survey*, Journal of Symbolic Logic, 56 (1991), pp. 1153–1183.

[25] N. Immerman, *Relational queries computable in polynomial time*, Information and Control, 68 (1986), pp. 86–104.

[26] ——, *Languages that capture complexity classes*, SIAM Journal of Computing, 16 (1987), pp. 760–778.

[27] N. Immerman and M. Vardi, *Model Checking and Transitive Closure Logic*, Proc. 9th Conf. on Computer-Aided Verification CAV '97, Lecture Notes in Computer Science No. 1254, Springer, 1997, pp.258–272.

[28] A. Kahr, E. Moore, and H. Wang, *Entscheidungsproblem reduced to the* $\forall\exists\forall$ *case*, Proc. Nat. Acad. Sci. U.S.A., 48 (1962), pp. 365–377.

[29] D. Kozen, *A finite model theorem for the propositional* μ*-calculus*, Studia Logica, 47 (1983), pp. 233–241.

[30] D. Kozen and R. Parikh, *A decision procedure for the propositional* μ*-calculus*, in Logics of Programs, E. Clarke and D. Kozen, eds., Lecture Notes in Computer Science No. 164, Springer, 1984, pp. 313–325.

[31] A. Leisenring, *Mathematical logic and Hilbert's* ε*-symbol*, Gordon & Breach, New York, 1969.

[32] H. Lewis, *Complexity Results for Classes of Quantificational Formulas*, Journal of Computer and System Sciences21 (1980), pp. 317–353.

[33] E. López-Escobar, *On defining well-orderings*, Fund. Math., 59, 1966, pp. 13–21.

[34] M. Mortimer, *On languages with two variables*, Zeitschr. f. math. Logik u. Grundlagen d. Math., 21 (1975), pp. 135–140.

[35] M. Otto, *Bounded variable logics and counting — a study in finite models*, Lecture Notes in Logic No. 9, Springer, 1997.

[36] M. Rabin, *Decidability of second-order theories and automata on infinite trees*, Trans. Amer. Math. Soc., 141 (1969), pp. 1–35.

[37] B. Trakhtenbrot, *On recursive separability*, Dokl. Akad. Nauk SSSR, 88 (1953), pp. 953–955. In Russian.

[38] M. Vardi, *The complexity of relational query languages*, in Proc. 14th ACM Symp. on Theory of Computing, 1982, pp. 137–146.

[39] M. Vardi, *Why is modal logic so robustly decidable*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science 31, American Math. Society, 1997, 149–184.