Two Variable First-Order Logic over Ordered Domains

Martin Otto

revised, October 99

Abstract

The satisfiability problem for the two-variable fragment of first-order logic is investigated over finite and infinite linearly ordered, respectively wellordered domains, as well as over finite and infinite domains in which one or several designated binary predicates are interpreted as arbitrary wellfounded relations.

It is shown that FO^2 over ordered, respectively wellordered, domains or in the presence of one wellfounded relation, is decidable for satisfiability as well as for finite satisfiability. Actually the complexity of these decision problems is essentially the same as for plain unconstrained FO^2, namely non-deterministic exponential time.

In contrast FO^2 becomes undecidable for satisfiability and for finite satisfiability, if a sufficiently large number of predicates are required to be interpreted as orderings, wellorderings, or as arbitrary wellfounded relations. This undecidability result also entails the undecidability of the natural common extension of FO^2 and computation tree logic CTL.

1 Introduction

Two-variable first-order logic FO^2 stands out as one of the fragments of first-order logic whose satisfiability problem is decidable. In fact Mortimer [9] showed that FO^2 has the finite model property and hence is decidable for satisfiability. Recently the bound on model size has been improved by Grädel, Kolaitis and Vardi [4] to locate the complexity of the satisfiability problem for FO^2 in non-deterministic exponential time. The renewed interest in FO^2 is due to the fact that several process logics and terminological logics are closely related to FO^2 through their common core of propositional modal logic. While modal logic has long been extended – in various directions and answering various expressivity needs in applications – to more powerful yet still manageable systems, corresponding issues have only recently been studied for extensions of FO^2. On the one hand FO^2 is a natural manageable fragment of first-order logic to extend modal logic; on the other hand, the first-order closure properties of FO^2 would be a desirable semantic feature in some of the above-mentioned areas of applications. Several extensions of FO^2

have thus attracted attention, in particular extensions by mechanisms that were found useful in the modal context. The overall picture that emerged from these investigations is, however, that decidability of FO^2 is not nearly as robust as it is for modal logic [11, 6, 5].

As outlined in the survey [5], the satisfiability issue for several such extensions of FO^2 may be equivalently phrased as a satisfiability problem for FO^2 over particular classes of structures, or as a decidability problem concerning the FO^2-theory of such special classes. We here primarily investigate FO^2-satisfiability over

- finite or infinite ordered and wellordered domains,
- finite or infinite domains over which several designated binary predicates are required to be wellfounded.

We show that FO^2 is decidable for (finite) satisfiability over ordered, respectively wellordered domains. It should be noted that these results concern FO^2-satisfiability in the presence of other, and in particular in the presence of other binary, predicates besides the order relation.

Concerning wellordered domains, the constraint that some designated binary predicate is interpreted as a wellordering is actually equivalent w.r.t. FO^2-satisfiability with the constraint that some designated binary predicate be interpreted as a wellfounded relation. The investigation of wellfoundedness constraints involving several designated binary predicates is motivated by the equivalence of this variant of FO^2-satisfiability with the satisfiability problem for the two-variable logic CL^2 which arises as a natural common extension of FO^2 and computation tree logic CTL [5]. CTL is a decidable extension of propositional modal logic ML by certain mechanisms for quantification over paths. In expressive power CTL is intermediate between modal logic ML and the more expressive yet still decidable modal μ-calculus L_μ, all of which are modal two-variable logics in the sense of being preserved under bisimulation equivalence:

$$\mathrm{ML} \subsetneq \mathrm{CTL} \subsetneq \mathrm{L}_\mu.$$

Recent studies have focused on the possibility of lifting these extensions of ML to extensions of FO^2, as in

$$\mathrm{FO}^2 \subsetneq \mathrm{CL}^2 \subsetneq \mathrm{FP}^2.$$

The least common extension of FO^2 and L_μ, namely a weak two-variable least fixed-point logic FP^2, was shown to be undecidable in [6], and it remained open whether some natural common extension of FO^2 and CTL could maybe still realize the hope for a decidable lift at the level of CTL. Another natural candidate to extend CTL and FO^2, two-variable transitive closure logic TC^2, was also shown undecidable in [6] (also compare the survey [5] for these results, and [8] for the role of TC^2 in model checking with CTL). The main differences between the candidates TC^2 and CL^2 are the following:

(a) CTL is a sublogic of TC^2 only in restriction to finite structures; they are incomparable over infinite structures [6]. The inclusion $\mathrm{CTL} \subseteq \mathrm{CL}^2$ on the other hand is valid over all structures [5]. Actually $\mathrm{CL}^2 \subsetneq \mathrm{TC}^2$ over finite structures.

(b) Unlike TC^2, CL^2 is preserved under two-pebble game equivalence, and thus qualifies as a true two-variable logic. While TC^2 and FP^2 are incomparable [6], $\mathrm{CL}^2 \subsetneq \mathrm{FP}^2$ over arbitrary structures.

CL^2 is in fact the least extension of ML that simultaneously has the first-order closure properties of FO^2 *and* the power to express simple CTL-style path properties w.r.t. definable accessibilities in a multi-modal setting. It is shown in [5] that the following are equivalent via low-complexity logical reductions, both in their general versions and in restriction to finite structures:

- satisfiability for CL^2.
- satisfiability for FO^2 in models that interpret several designated predicates as wellfounded relations.
- satisfiability of FO^2 in models that interpret several designated predicates as wellorderings.

Our undecidability results for FO^2 over several wellorderings imply that CL^2 too is undecidable, thereby further corroborating the surprisingly negative overall picture concerning lifts of modally innocuous mechanisms to the first-order two-variable scenario. Note that the undecidability result for CL^2 is in a sense the strongest one obtained in this direction so far, in that CL^2 is strictly weaker than FP^2, and in that the undecidability proof given here also covers the case of TC^2 even though $\mathrm{CL}^2 \subseteq \mathrm{TC}^2$ only for finite structures, cf. Corollary 2.6. While similar in spirit to the above-mentioned undecidability proofs (via reductions from the domino tiling problem) the present undecidability proof is more involved, concerning a weaker logical framework and building on a weaker notion of interpreting grid structures for the reduction. It should be pointed out in this context that some of the technical results proved in this paper were announced without proof and discussed in a larger context in the survey paper [5] which precedes the present paper in publication.

1.1 Preliminaries

Formally FO^2, the two-variable fragment of first-order logic, comprises exactly those formulae of FO with equality, whose only variable symbols are x and y. We consider FO over *purely relational vocabularies*. Relational structures are denoted like $\mathfrak{A} = \big(A, R_1, R_2, \dots\big)$, A being the universe of $\mathfrak{A}$. Owing to the fact that FO^2-atoms in particular can only involve up to two distinct variables, one may without loss of generality further assume that the arities of the R_i are at most two, i.e. all our vocabularies consist of unary and binary predicates only. We also only deal with structures having at least two elements.

Writing $\varphi(x)$, $\varphi(y)$ for formulae, it is taken for granted that at most the displayed variables occur free in φ.

We use the standard notation $\mathit{sat}(\mathcal{L})$ and $\mathit{fin\text{-}sat}(\mathcal{L})$ to denote the subclasses of *satisfiable* and *finitely satisfiable* members in a class of sentences $\mathcal{L}$: $\varphi \in \mathcal{L}$ is in $\mathit{sat}(\mathcal{L})$ if there is a model $\mathfrak{A} \models \varphi$, and in $\mathit{fin\text{-}sat}(\mathcal{L})$ if there is a finite model $\mathfrak{A} \models \varphi$.

For a class $\mathcal{K}$ of structures, the *$\mathcal{L}$-theory* of $\mathcal{K}$ is the collection of those $\varphi \in \mathcal{L}$ that are satisfied in *all* $\mathfrak{A} \in \mathcal{K}$. We let $\mathcal{K}_{\mathrm{fin}}$ denote the subclass consisting of just the finite structures in $\mathcal{K}$, and correspondingly get an $\mathcal{L}$-theory of $\mathcal{K}_{\mathrm{fin}}$.

Of special interest will be the following classes of relational structures.

$$\begin{aligned}
\mathcal{O} &= \big\{\mathfrak{A} = (A, <, \dots) \bigm| < \text{linearly orders } A \big\},\\
\mathcal{WO} &= \big\{\mathfrak{A} = (A, <, \dots) \bigm| < \text{wellorders } A \big\},\\
\mathcal{WF} &= \big\{\mathfrak{A} = (A, E, \dots) \bigm| E \text{ is wellfounded} \big\},
\end{aligned}$$

and

$$\begin{aligned}
\mathcal{K}_1 &= \big\{\mathfrak{A} = \big(A, E_1, \dots, E_k, \dots\big) \bigm| E_i \text{ wellorders } A\ , 1 \leqslant i \leqslant k\big\},\\
\mathcal{K}_2 &= \big\{\mathfrak{A} = \big(A, E_1, \dots, E_k, \dots\big) \bigm| E_i \text{ wellfounded}\ , 1 \leqslant i \leqslant k\big\},\\
\mathcal{K}_3 &= \big\{\mathfrak{A} = \big(A, E_1, \dots, E_k, \dots\big) \bigm| E_i \text{ a partial ordering}\ , 1 \leqslant i \leqslant k\big\}.
\end{aligned}$$

In all these classes it is understood that there may be any number of other predicates, and in particular unconstrained binary predicates. Note that $\mathcal{O}_{\mathrm{fin}} = \mathcal{WO}_{\mathrm{fin}}$ and that $\mathcal{WF}_{\mathrm{fin}}$ consists of all finite $\mathfrak{A} = (A, E, \dots)$ with acyclic E.

Definition 1.1 The (finite) satisfiability problem for $\mathcal{L}$ over $\mathcal{K}$ is the problem to determine, for $\varphi \in \mathcal{L}$, whether φ has a model $\mathfrak{A}$ in $\mathcal{K}$, respectively in $\mathcal{K}_{\mathrm{fin}}$:

(a) $\mathit{sat}_{\mathcal{K}}(\mathcal{L}) = \big\{\varphi \in \mathcal{L} \bigm| \exists\, \mathfrak{A} \in \mathcal{K}, \mathfrak{A} \models \varphi\big\}$,

(b) $\mathit{fin\text{-}sat}_{\mathcal{K}}(\mathcal{L}) = \mathit{sat}_{\mathcal{K}_{\mathrm{fin}}}(\mathcal{L}) = \big\{\varphi \in \mathcal{L} \bigm| \exists\, \mathfrak{A} \in \mathcal{K}_{\mathrm{fin}}, \mathfrak{A} \models \varphi\big\}$.

If $\mathcal{L}$ is closed under negation, then obviously φ has a (finite) model in $\mathcal{K}$ if and only if $\neg\varphi$ is not in the $\mathcal{L}$-theory of $\mathcal{K}$ (not in the $\mathcal{L}$-theory of $\mathcal{K}_{\mathrm{fin}}$). In other words, decidability of $\mathit{sat}_{\mathcal{K}}(\mathcal{L})$ is equivalent with the decidability of the $\mathcal{L}$-theory of $\mathcal{K}$, and decidability of $\mathit{fin\text{-}sat}_{\mathcal{K}}(\mathcal{L})$ is equivalent with the decidability of the $\mathcal{L}$-theory of $\mathcal{K}_{\mathrm{fin}}$.

The goals of this paper are the following theorems.

Theorem 1.2 *The satisfiability problem and the finite satisfiability problem for* FO^2 *over ordered and over wellordered domains are decidable.*

Indeed, $sat_{\mathcal{O}}(\mathrm{FO}^2)$*,* $sat_{\mathcal{WO}}(\mathrm{FO}^2)$*, and* $\text{fin-}sat_{\mathcal{O}}(\mathrm{FO}^2)$*, and similarly* $sat_{\mathcal{WF}}(\mathrm{FO}^2)$ *and* $\text{fin-}sat_{\mathcal{WF}}(\mathrm{FO}^2)$*, are in* NEXPTIME*; i.e. the* FO^2*-theories of* $\mathcal{O}$*,* $\mathcal{WO}$*,* $\mathcal{O}_{\mathit{fin}}$*,* $\mathcal{WF}$*, and* $\mathcal{WF}_{\mathit{fin}}$ *are in co-*NEXPTIME.

Here NEXPTIME is the union over $\mathrm{NTIME}(2^{p(n)})$ for all polynomials p. The following definition of a logic CL^2, designed to be a natural least common extension of FO^2 and CTL, is discussed and motivated in some detail in [5].

Definition 1.3 Let CL^2 be the extension of FO^2 which augments the FO^2-rules for the formation of formulae with the following modal closure operators: for any pair of formulae $\xi(x,y)$ and $\varphi(x)$ (with free variables as indicated), we have formulae $\psi_1(x) = \big(\langle\xi\rangle^\infty\varphi\big)(x)$ and $\psi_2(x) = \big([\xi]^\infty\varphi\big)(x)$. Their semantics is such that $(\mathfrak{A}, a) \models \psi_1$ (respectively $(\mathfrak{A}, a) \models \psi_2$) if there is a ξ-path from a to a vertex where φ is true (respectively if every infinite ξ-path from a eventually reaches a vertex where φ is true [1]).

Theorem 1.4 *For each of the above classes* $\mathcal{K} = \mathcal{K}_i$*,* $i = 1, 2, 3$ *and for* $k \geqslant 8$*, the* FO^2*-theories of* $\mathcal{K}$ *and of* $\mathcal{K}_{\mathit{fin}}$ *are undecidable. It follows in particular that* CL^2 *is undecidable for satisfiability and for finite satisfiability.*

In fact the claims of Theorem 1.4 can be strengthened to results about recursive inseparability, and it may also be shown that $sat(\mathrm{CL}^2)$ is even Σ^1_1-hard. Concerning the recurrent theme of trading wellorderings for arbitrary wellfounded relations in the context of FO^2-satisfiability, we observe that there are obvious logical reductions based on the following.

- $(A, <)$ is a wellordering if and only if $<$ is wellfounded and satisfies the FO^2-axiom of trichotomy: $\forall x \forall y \big(x = y \vee x < y \vee y < x\big)$.
- (A, E) is wellfounded if and only if there is a well-ordering $<$ on A such that $(A, E, <) \models \forall x \forall y (Exy \to x < y)$.

These imply in particular that e.g. $sat_{\mathcal{WO}}(\mathrm{FO}^2)$ and $sat_{\mathcal{WF}}(\mathrm{FO}^2)$ are LOGSPACE-reducible to each other, and so are e.g. the FO^2-theories of $\mathcal{K}_1$ and $\mathcal{K}_2$.

Theorem 1.4 and Theorem 1.2 are dealt with in Sections 2 and 3, respectively. These sections may be read independently.

[1] If instead we want to assert that every *maximal* ξ-path eventually reaches a vertex where φ is true, we may use $[\xi']^\infty\varphi$ where $\xi'(x,y) = \xi(x,y) \vee \neg\exists y Exy$.

2 The undecidability result

2.1 Domino systems

Recall that $\mathcal{K}_1 = \{\mathfrak{A} = (A, E_1, \dots, E_k, \dots) \mid E_i \text{ wellorders } A\ , 1 \leqslant i \leqslant k\}$. We are going to reduce tiling problems for domino systems to satisfiability of FO^2 over $\mathcal{K}_1$ in order to prove Theorem 1.4.

A *domino system* $\mathcal{D} = (D, R_H, R_V)$ consists of a finite set D of domino pieces together with two binary relations R_H and R_V over D, which specify horizontal and vertical adjacencies. A $\mathcal{D}$*-tiling* is a mapping t from $\mathbb{Z} \times \mathbb{Z}$ (or from $\mathbb{N} \times \mathbb{N}$) to D, $(p, q) \mapsto t(p, q)$, that respects horizontal and vertical adjacencies in the sense that always $(t(p, q), t(p+1, q)) \in R_H$ and $(t(p, q), t(p, q+1)) \in R_V$. It is a consequence of König's Lemma that, whenever there is a $\mathcal{D}$-tiling of $\mathbb{N} \times \mathbb{N}$ then there also is one of $\mathbb{Z} \times \mathbb{Z}$. As we are only interested in the existence or nonexistence of tilings, we need not distinguish between the two notions.

A tiling t is *periodic* if there are natural numbers $r, s \geqslant 1$ such that $t(n+r, m) = t(n, m+s) = t(n, m)$ for all $n, m \in \mathbb{Z}$. In that case the tiling may be regarded as a tiling of the finite torus $(\mathbb{Z} \bmod r) \times (\mathbb{Z} \bmod s)$. By taking the least common multiple of the two periods r and s one can always regard a periodic tiling as one whose fundamental domain is a square, or as a tiling of some $(\mathbb{Z} \bmod m) \times (\mathbb{Z} \bmod m)$, $m \geqslant 1$.

Theorem 2.1 (Berger; Gurevich/Koryakov) *The following domino problems are undecidable:*

(i) (tiling)*: given $\mathcal{D}$, determine whether there is a $\mathcal{D}$-tiling.*

(ii) (periodic tiling)*: given $\mathcal{D}$, determine whether there is a periodic $\mathcal{D}$-tiling.*

Indeed, those $\mathcal{D}$ that admit no tiling are recursively inseparable *from those that admit a periodic tiling.*

Inseparability, due to Gurevich and Koryakov [7], follows from Trakhtenbrot's Theorem, via a reduction

$$\varphi \in \mathrm{FO} \longmapsto \mathcal{D}_\varphi$$

such that $\varphi \in \mathit{sat}(\mathrm{FO})$ (respectively $\mathit{fin\text{-}sat}(\mathrm{FO})$) iff $\mathcal{D}_\varphi$ admits a tiling (respectively a periodic tiling). Trakhtenbrot's Theorem on the inseparability of the complement of $\mathit{sat}(\mathrm{FO})$ versus $\mathit{fin\text{-}sat}(\mathrm{FO})$ thus translates into inseparability of "no tiling" versus "periodic tiling". Many undecidability proofs concerning *sat*- and *fin-sat*-problems for weak logical systems $\mathcal{L}$ on the other hand are carried out through a reduction

$$\mathcal{D} \longmapsto \psi_{\mathcal{D}} \in \mathcal{L}$$

such that $\mathcal{D}$ admits a (periodic) tiling if and only if $\psi_{\mathcal{D}} \in sat(\mathcal{L})$ ($\mathit{fin}\text{-}sat(\mathcal{L})$). The existence of such a reduction not only implies that $sat(\mathcal{L})$ and $\mathit{fin}\text{-}sat(\mathcal{L})$ are undecidable, but also that the complement of $sat(\mathcal{L})$ is recursively inseparable from $\mathit{fin}\text{-}sat(\mathcal{L})$, as well as the existence of a reduction from FO to $\mathcal{L}$ which simultaneously preserves satisfiability and satisfiability in finite models (thus showing that $\mathcal{L}$ is a *conservative reduction class*, cf. [2] for this and for general background on the relevant classical notions).

The straightforward way to achieve the desired reduction $\mathcal{D} \mapsto \psi_{\mathcal{D}}$ consists in providing descriptions of correctly $\mathcal{D}$-tiled grid structures in the logic at hand. Think of a grid structure as a structure of type $\mathfrak{G} = (G, H, V)$ with binary predicates H and V for the horizontal and vertical successor relations in the grid. The idea is to formalize a $\mathcal{D}$-tiling of $\mathfrak{G}$ by an expansion with monadic predicates P_d, one for each tile type $d \in D$. This approach is particularly suited to logics that extend FO^2, since the correctness of a tiling is easily expressible in first-order with just two variables by

$$\begin{aligned} \varphi_{\mathcal{D}} \;=\; & \forall x \big(\textstyle\bigvee_d P_d x \wedge \bigwedge_{d \neq d'} \neg (P_d x \wedge P_{d'} x) \big) \\ & \wedge \forall x \forall y \big(Hxy \rightarrow \textstyle\bigvee_{(d,d') \in R_H} (P_d x \wedge P_{d'} y) \big) \\ & \wedge \forall x \forall y \big(Vxy \rightarrow \textstyle\bigvee_{(d,d') \in R_V} (P_d x \wedge P_{d'} y) \big). \end{aligned} \tag{1}$$

It therefore essentially remains to axiomatize in the logic at hand classes of structures $\mathfrak{G}$ that are sufficiently close to the canonical grid structures on $\mathbb{Z} \times \mathbb{Z}$ and $(\mathbb{Z} \bmod m) \times (\mathbb{Z} \bmod m)$. This is carried out for a number of interesting extensions of FO^2 in [6, 5].

In [5] we also delineate a method to derive an additional result along the way, through the introduction of some simple class of structures $\mathcal{K}$ which mediates the connection between $\mathcal{L}$ and grids as follows. Suppose that

- $\mathcal{K}$ is (projectively[2]) axiomatizable in $\mathcal{L}$.
- a sufficiently rich class of grids is (projectively[2]) FO^2-axiomatizable over $\mathcal{K}$.

In this case we infer not only undecidability results and recursive inseparability results concerning $sat(\mathcal{L})$ and $\mathit{fin}\text{-}sat(\mathcal{L})$, but moreover an inseparability result concerning the FO^2-theories of $\mathcal{K}$ and $\mathcal{K}_{\mathrm{fin}}$. Namely, the following are recursively inseparable: the complement of the FO^2-theory of $\mathcal{K}_{\mathrm{fin}}$ and the FO^2-theory of $\mathcal{K}$. In other words: the FO^2-theory of $\mathcal{K}$ is *strongly undecidable*. This approach is exemplified in [5] for the class $\mathcal{K}$ of those structures that interpret several basic binary relations as equivalence relations.

While the reductions presented in [5] actually achieve an FO^2-axiomatization, relative to the appropriate $\mathcal{K}$, of a sufficiently rich class of structures $\mathfrak{G} = (G, H, V)$ in

[2]Recall that a class $\mathcal{K}$ of τ-structures is *projectively* axiomatized by some sentence ψ in an extended vocabulary $\tau' \supseteq \tau$ if $\mathcal{K} = \{ \mathfrak{A} \restriction \tau \mid \mathfrak{A} \models \psi \}$.

which H and V are the graphs of two *commuting, bijective functions*, we here make do with a weaker form of grid-like behaviour. In particular, functionality does not seem to be even projectively characterizable in CL^2, nor in FO^2 relative to several wellorderings.

2.2 A reduction criterion

Let $\mathfrak{G}_{\mathbb{Z}}$ be the canonical grid structure on $\mathbb{Z} \times \mathbb{Z}$:

$$\mathfrak{G}_{\mathbb{Z}} = \big(\mathbb{Z}^2, H, V\big), \qquad \begin{array}{l} H = \big\{\big((p,q),(p+1,q)\big) \bigm| p,q \in \mathbb{Z}\big\}, \\ V = \big\{\big((p,q),(p,q+1)\big) \bigm| p,q \in \mathbb{Z}\big\}, \end{array}$$

and let $\mathfrak{G}_{\mathbb{N}}$ denote the standard grid on $\mathbb{N} \times \mathbb{N}$, which is just the restriction of $\mathfrak{G}_{\mathbb{Z}}$ to $\mathbb{N} \times \mathbb{N}$. Let $\mathfrak{G}_m$ denote the standard grid on a finite $m \times m$ torus:

$$\mathfrak{G}_m = \big((\mathbb{Z} \bmod m) \times (\mathbb{Z} \bmod m), H, V\big), \qquad \begin{array}{l} H = \big\{\big((p,q),(p',q)\big) \bigm| p' - p \equiv 1 \bmod m\big\}, \\ V = \big\{\big((p,q),(p,q')\big) \bigm| q' - q \equiv 1 \bmod m\big\}. \end{array}$$

Let $\mathfrak{G}_i = (G_i, H_i, V_i)$, $i = 1, 2$. $\mathfrak{G}_1$ is *homomorphically embeddable* into $\mathfrak{G}_2$ if there is a homomorphism $\pi\colon \mathfrak{G}_1 \to \mathfrak{G}_2$, i.e. a mapping π such that for all $v, v' \in G_1$: $(v,v') \in H_1 \Rightarrow (\pi(v),\pi(v')) \in H_2$ and $(v,v') \in V_1 \Rightarrow (\pi(v),\pi(v')) \in V_2$. Note in particular that $\mathfrak{G}_{\mathbb{N}}$ is homomorphically embeddable into $\mathfrak{G}_{\mathbb{Z}}$, and that both are homomorphically embeddable into every $\mathfrak{G}_m$.

Definition 2.2 (a) An infinite structure $\mathfrak{G} = (G, H, V)$ is called *grid-like* if $\mathfrak{G}_{\mathbb{N}}$ is homomorphically embeddable into $\mathfrak{G}$; a finite $\mathfrak{G}$ is grid-like if some $\mathfrak{G}_m$ is homomorphically embeddable into $\mathfrak{G}$.

(b) A class $\mathcal{G}$ of grid-like structures is called *rich* if at least one of $\mathfrak{G}_{\mathbb{N}}$ or $\mathfrak{G}_{\mathbb{Z}} \in \mathcal{G}$, and if for all $n \geqslant 1$ there is some k such that $\mathfrak{G}_{k \cdot n} \in \mathcal{G}$.

Lemma 2.3 *Let $\mathcal{L}$ be an extension of* FO^2 *and assume that $\mathcal{L}$ projectively axiomatizes a rich class $\mathcal{G}$ of grid-like structures. Then* fin-sat$(\mathcal{L})$ *is recursively inseparable from the complement of* sat$(\mathcal{L})$ *(i.e. $\mathcal{L}$ is a conservative reduction class).*

If $\mathcal{K}$ is projectively $\mathcal{L}$-axiomatizable, and if a rich class $\mathfrak{G}$ of grid-like structures is projectively FO^2*-axiomatizable over $\mathcal{K}$, then moreover the* FO^2*-theory of $\mathcal{K}$ is recursively inseparable from the complement of the* FO^2*-theory of $\mathcal{K}_{\mathit{fin}}$ (i.e. the* FO^2*-theory of $\mathcal{K}$is strongly undecidable).*

Proof. It suffices to show the following, for every rich class $\mathcal{G}$ and domino system $\mathcal{D}$, and with $\varphi_{\mathcal{D}}$ as given in (1) above:

(i) $\mathcal{D}$ admits a tiling if and only if some $\mathfrak{G} \in \mathcal{G}$ can be expanded with unary predicates P_d to a model $(G, H, V, (P_d)_{d \in D})$ of $\varphi_{\mathcal{D}}$.

(ii) $\mathcal{D}$ admits a periodic tiling if and only if some finite $\mathfrak{G} \in \mathcal{G}$ can be expanded to a model of $\varphi_{\mathcal{D}}$.

The implications from the existence of $\mathcal{D}$-tilings to the existence of (finite) $\mathfrak{G} \in \mathcal{G}$ with expansions to models of $\varphi_{\mathcal{D}}$ follows directly from the richness conditions in Definition 2.2. Considering converse assume that $\big(\mathfrak{G}, (P_d)_{d \in D}\big) \models \varphi_{\mathcal{D}}$ for some $\mathfrak{G} \in \mathcal{G}$. By Definition 2.2 there is a homomorphism $\pi \colon \mathfrak{G}_{\mathbb{N}} \to \mathfrak{G}$. But then the following is a correct tiling:

$$\begin{array}{rcll} t \colon \mathbb{N} \times \mathbb{N} & \longrightarrow & D & \\ (n, m) & \longmapsto & d & \text{if } \pi(n, m) \in P_d. \end{array}$$

If $\mathfrak{G}$ is finite, we even have a homomorphism $\pi \colon \mathfrak{G}_m \to \mathfrak{G}$ to obtain an induced tiling of some $\mathfrak{G}_m$. □

The following provides a simple and sufficient local criterion for grid-likeness in the sense of Definition 2.2. We say that H is *complete over* V in $\mathfrak{G} = (G, H, V)$ if $\mathfrak{G}$ satisfies

$$\forall x \forall y \forall x' \forall y' \big((Hxy \wedge Vxx' \wedge Vyy') \longrightarrow Hx'y' \big).$$

Lemma 2.4 *Assume that* $\mathfrak{G} = (G, H, V)$ *satisfies the* FO^2*-axiom* $\forall x (\exists y Hxy \wedge \exists y Vxy)$. *If* H *is complete over* V, *then* $\mathfrak{G}$ *is grid-like.*

Proof. We first construct a homomorphism $\pi \colon \mathfrak{G}_{\mathbb{N}} \to \mathfrak{G}$. Using the fact that $\forall x \exists y Hxy$, one may homomorphically embed $(\mathbb{N} \times \{0\}, H)$ into $\mathfrak{G}$. This embedding is then inductively extended to $(\mathbb{N} \times \mathbb{N}, H, V)$ as follows. Assuming that π is as desired on $\mathbb{N} \times [0, m]$, it is extended by choosing $\pi(n, m+1)$ in $\{w \mid (\pi(n, m), w) \in V\}$ (this uses $\forall x \exists y Vxy$). It follows from completeness of H over V that π is a homomorphism w.r.t. H also over $\mathbb{N} \times [0, m+1]$.

If $\mathfrak{G}$ is finite, choose a homomorphic embedding of $(\mathbb{N} \times \{0\}, H)$ into $\mathfrak{G}$ as above. It follows that there are $r' < r$ for which $\pi(r, 0) = \pi(r', 0)$. Without loss of generality assume that $r' = 0$. We may regard $\pi \restriction [0, r) \times \{0\}$ as a homomorphic embedding of $(\mathbb{Z} \bmod r, H)$. Inductively we extend π so that $\pi \restriction [0, r) \times \mathbb{N}$ becomes a homomorphic embedding of $\big((\mathbb{Z} \bmod r) \times \mathbb{N}, H, V\big)$ into $\mathfrak{G}$: just as above it suffices to choose inductively $\pi(n, m+1) \in \{w \mid (\pi(n, m), w) \in V\}$.

But again, since $\mathfrak{G}$ is finite, there are $s' < s$ such that $\pi \restriction [0, r) \times \{s'\} = \pi \restriction [0, r) \times \{s\}$. Assuming again w.l.o.g. that $s' = 0$, we may regard $\pi \restriction [0, r) \times [0, s)$ as a homomorphic embedding of $\big((\mathbb{Z} \bmod r) \times (\mathbb{Z} \bmod s), H, V\big)$ into $\mathfrak{G}$. Extending the domain of this embedding to the least common multiple m of r and s, we obtain a homomorphic embedding of $\mathfrak{G}_m$ into $\mathfrak{G}$, as desired. □

As above, let
$$\begin{array}{l} \mathcal{K}_1 = \big\{ \mathfrak{A} = (A, E_1, \ldots, E_k, \ldots) \;\big|\; E_i \text{ wellorders } A, 1 \leqslant i \leqslant k \big\}, \\ \mathcal{K}_2 = \big\{ \mathfrak{A} = (A, E_1, \ldots, E_k, \ldots) \;\big|\; E_i \text{ wellfounded}, 1 \leqslant i \leqslant k \big\}, \end{array}$$

where k remains to be specified, and the vocabulary of the $\mathcal{K}_i$ may contain a number of other binary as well as unary predicates, apart from the designated E_i. Recall from Section 1.1 that $\mathcal{K}_1$ is FO^2-axiomatizable over $\mathcal{K}_2$ by the FO^2-assertion of trichotomy. Both classes are clearly axiomatizable in CL^2, indeed $\mathcal{K}_2$ is axiomatized by the universally quantified, generalized CTL-formulae

$$\forall x \big[\forall (E_i^{-1}\text{-paths}\, p) \big(\top \text{ until} \neg \exists y E_i y x \big) \big].$$

The interior part of this formula is equivalent to the CTL formula $\mathsf{AF}\,\mathsf{AX}\,\bot$ (all paths eventually hit a dead end), if interpreted w.r.t. the converse of the transition relation E_i (thus referring to backward E_i-paths) and waiving the requirement that paths have to be infinite (which is not essential to CTL semantics). The following proposition demonstrates once more the power of transitivity assertions towards a characterization of grid-like structures – here with emphasis on the case that transitivity is only available in the very special shape of a wellordering.

Proposition 2.5 *There is a rich class $\mathcal{G}$ of grid-like structures that is* FO^2*-axiomatizable over $\mathcal{K}_1$ (for $k = 8$). In fact there is an* FO^2*-sentence η in an extended vocabulary including eight designated binary relations, such that*

- *(a) $\mathfrak{G}_{\mathbb{Z}}$ and all $\mathfrak{G}_{2m}$ for $m \geqslant 1$ can be expanded by eight wellorderings (and other extra relations) to models of η.*
- *(b) Every model of η that interprets the eight designated relations as transitive (!) is grid-like.*

As $\mathcal{K}_1$ is FO^2-axiomatizable over $\mathcal{K}_2$, a rich class of grid-like structures is similarly FO^2-axiomatizable over $\mathcal{K}_2$. But the asymmetry in the stronger statements (a) and (b) actually implies that instead of $\mathcal{K}_1$ we may also use for instance the following

$$\begin{aligned}
\mathcal{K}_3 &= \big\{ \mathfrak{A} = \big(A, E_1, \dots, E_k, \dots\big) \;\big|\; E_i \text{ a partial ordering }, 1 \leqslant i \leqslant k \big\}, \\
\mathcal{K}_4 &= \big\{ \mathfrak{A} = \big(A, E_1, \dots, E_k, \dots\big) \;\big|\; E_i \text{ transitive }, 1 \leqslant i \leqslant k \big\},
\end{aligned}$$

as well as other 'intermediate' classes between wellorderings and mere transitive relations.

Corollary 2.6 *The* FO^2*-theories of $\mathcal{K}_1$, $\mathcal{K}_2$, $\mathcal{K}_3$, and $\mathcal{K}_4$ are strongly undecidable (for $k = 8$), i.e. for these $\mathcal{K}$ the complement of the* FO^2*-theory of $\mathcal{K}_{\mathit{fin}}$ is recursively inseparable from the* FO^2*-theory of $\mathcal{K}$.*

This immediately implies our main undecidability result Theorem 1.4, and in particular the undecidability of *fin-sat*(CL^2) and *sat*(CL^2). Using the methods of [6] it is actually not hard to extend the undecidability result for CL^2 to show that $sat(\mathrm{CL}^2)$ is Σ^1_1-hard, i.e. hard for the first level of the analytical hierarchy.

Proof of Proposition 2.5. The proof proceeds as follows: we first expand the standard grids $\mathfrak{G}_{\mathbb{Z}}$ and $\mathfrak{G}_{2m}$ by several relations including eight wellorderings. In the course of doing so, we collect a number of sentences from the FO^2-theory of these expansions into some $\eta \in \mathrm{FO}^2$. Finally it is shown that η, in models which interpret the 8 designated relations as wellorderings (or in fact even as just transitive relations), implies that H is complete over V. Putting $\mathcal{G}$ to be the class of reducts of models of η in which the designated relations are wellorderings, this proves the proposition: $\mathfrak{G}_{\mathbb{Z}}$ and all the $\mathfrak{G}_{2m}$ are in $\mathcal{G}$ by construction, and Lemma 2.4 ensures that all structures in $\mathcal{G}$ are indeed grid-like.

Consider first the expansion of the standard grid $\mathfrak{G}_{\mathbb{Z}}$ by the following predicates H_{ij}, V_{ij} (binary), and A_{ij} (unary), for $i, j \in \{0,1\}$ (compare the sketch below):

$$\begin{array}{rcl} A_{ij} & = & \big\{(n,m) \mid n \equiv i \bmod 2, m \equiv j \bmod 2\big\}, \\ H_{ij} & = & H \cap (A_{ij} \times \mathbb{Z}), \\ V_{ij} & = & V \cap (A_{ij} \times \mathbb{Z}). \end{array}$$

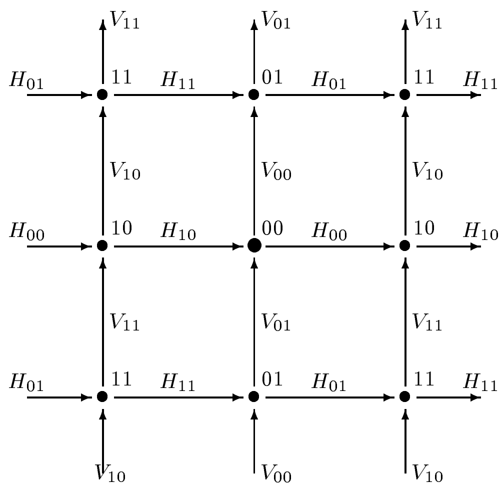

Clearly the expansion satisfies the FO^2-sentence η_0 which is the conjunction of the following (for all pairs i, j):

$$\begin{array}{l} \forall x \bigvee_{ij} A_{ij}x \wedge \forall x \bigwedge_{(i,j)\neq(i',j')} \neg(A_{ij}x \wedge A_{i'j'}x) \\ \forall x \exists y Hxy \wedge \forall x \exists y Vxy \\ \forall x \forall y \big(Hxy \leftrightarrow \bigvee_{ij} H_{ij}xy\big) \wedge \forall x \forall y \big(Vxy \leftrightarrow \bigvee_{ij} V_{ij}xy\big) \\ \forall x \forall y \big(H_{ij}xy \rightarrow (A_{ij}x \wedge A_{1-i\,j}y)\big) \\ \forall x \forall y \big(V_{ij}xy \rightarrow (A_{ij}x \wedge A_{i\,1-j}y)\big). \end{array}$$

Moreover, it is not hard to see that all finite standard grids $\mathfrak{G}_{2m}$, whose fundamental domain has even dimensions, can also be expanded to a model of η_0 which locally looks just like the given expansion of $\mathfrak{G}_{\mathbb{Z}}$. In what follows we treat the case of $\mathfrak{G}_{\mathbb{Z}}$ and that of the $\mathfrak{G}_{2m}$ in parallel without explicit distinction.

Consider for every pair (i, j) the reduct of the above expansions to the vocabulary consisting of the edge relations H_{ij}, V_{ij}, $H_{i\,1-j}$, $V_{1-i\,j}$. Note that each of these collections of edge relations forms a graph that is the disjoint union of quadrangles; call these quadrangles the quadrangles of type (i, j).

Let $\sim_{ij}$ be the symmetric transitive closure of $H_{ij} \cup V_{ij} \cup H_{i\,1-j} \cup V_{1-i\,j}$. Note that $\sim_{ij}$ is an equivalence relation whose classes are precisely the disjoint quadrangles of type (i, j). Let $\prec_{ij}$ be a wellordering of the quotient w.r.t. $\sim_{ij}$. We introduce two wellorderings, $<_{ij}$ and $<'_{ij}$, for each index pair (i, j) in such a way that between different $\sim_{ij}$-classes both orderings coincide with $\prec_{ij}$, and such that within each $\sim_{ij}$-class $<'_{ij}$ is the converse of $<_{ij}$. Specifically, let $<_{ij}$ in restriction to individual $\sim_{ij}$-classes be the ordering that puts the element of colour $A_{i\,1-j}$ first, that of colour A_{ij} second, that of colour $A_{1-i\,j}$ third, and that of colour $A_{1-i\,1-j}$ last. Thinking of a $\sim_{ij}$-class as a quadrangle whose lower left corner is coloured A_{ij}, this ordering corresponds to the counterclockwise succession starting in the upper left corner.

For the constructed expansions of $\mathfrak{G}_{\mathbb{Z}}$ and $\mathfrak{G}_{2m}$ by the A_{ij}, H_{ij}, V_{ij}, $<_{ij}$ and $<'_{ij}$, we have the following:

(i) the $<_{ij}$ and $<'_{ij}$ are wellorderings.

(ii) $x \sim_{ij} y$ if and only if $<_{ij}$ and $<'_{ij}$ disagree on x and y; in particular, the following two FO^2-sentences are satisfied:

$$\eta_1 \;=\; \bigwedge_{i,j} \forall x \forall y \left[\begin{array}{l} \big(Hxy \wedge (A_{ij}x \vee A_{i\,1-j}x)\big) \rightarrow (x <_{ij} y \wedge y <'_{ij} x) \\ \wedge\big((Vxy \wedge A_{ij}x) \rightarrow (y <_{ij} x \wedge x <'_{ij} y)\big) \\ \wedge\big((Vxy \wedge A_{1-i\,j}x) \rightarrow (x <_{ij} y \wedge y <'_{ij} x)\big) \end{array} \right],$$

$$\eta_2 \;=\; \bigwedge_{i,j} \forall x \forall y \Big[\big(x <_{ij} y \wedge y <'_{ij} x \wedge A_{i\,1-j}x \wedge A_{1-i\,1-j}y\big) \rightarrow H_{i\,1-j}xy\Big].$$

Conversely, we claim that in every model of the FO^2-sentence $\eta = \bigwedge_{i=0}^{2} \eta_i$ in which the $<_{ij}$ and $<'_{ij}$ are interpreted as *transitive* relations, H is complete over V.

To see this, consider a configuration of four vertices a, b, a' and b', such that $(a,b) \in H$ and $(a,a'),(b,b') \in V$. Assuming for instance that $a \in A_{00}$, we know from η_0 that $(a,b) \in H_{00}$, $(a,a') \in V_{00}$, $(b,b') \in V_{10}$, $b \in A_{10}$, $a' \in A_{01}$ and $b' \in A_{11}$. According to η_1, we thus find that $a' <_{00} a <_{00} b <_{00} b'$, and $b' <'_{00} b <'_{00} a <'_{00} a'$. It follows from transitivity of $<_{00}$ and $<'_{00}$ that $a' <_{00} b'$ and $b' <'_{00} a'$, whence η_2 implies $(a',b') \in H_{01}$ and therefore in H. $\Box$

This also completes the proof of Theorem 1.4. While one might suspect that one could do with fewer than eight wellorderings, we have at least not managed to reduce this number to two. That two is indeed a lower bound on the number of wellorderings that would render FO^2 undecidable will follow from the results of the following section.

3 The decidability result

3.1 Scott's normal form

An important tool in the analysis of FO^2-satisfiability is a Skolemization procedure reducing arbitrary FO^2-sentences to a normal form consisting of conjunctions of prenex FO^2-sentences of quantifier prefixes $\forall\forall$ and $\forall\exists$. This satisfiability preserving normal form is due to Scott [10] and reduces FO^2-satisfiability to the Gödel case of the classical decision problem (which at the time was believed to be decidable even in the presence of equality).

Theorem 3.1 (Scott) *There is a* PTIME *reduction* $\mathrm{NF}\colon \mathrm{FO}^2 \to \mathrm{FO}^2$ *mapping every sentence* $\varphi \in \mathrm{FO}^2$ *to a sentence* $\mathrm{NF}(\varphi)$ *in an extended vocabulary and of the form*

$$\forall x \forall y \chi_0 \wedge \bigwedge_{i=1}^{m} \forall x \exists y \chi_i$$

with quantifier-free χ_i*, and such that* φ *has a (finite) model if and only if* $\mathrm{NF}(\varphi)$ *has a (finite) model. More precisely,* $\mathrm{NF}(\varphi) \models \varphi$ *and any model of* φ *can be expanded to a model of* $\mathrm{NF}(\varphi)$*. The length of* $\mathrm{NF}(\varphi)$ *is linearly bounded in the length of* φ*.*

When using this normal form, we shall further assume without loss of generality that the χ_i for $i \geqslant 1$ satisfy $\forall x \forall y (\chi_i(x,y) \to \neg x = y)$. To achieve this one would simply have to replace χ_i with $(\chi_i(x,y) \vee \chi_i(x,x)) \wedge \neg x = y$. This replacement is sound over all structures with at least two elements (and we may safely disregard satisfiability over one-element structures).

Let φ be in normal form over vocabulary τ. Models of φ are to be analyzed in terms of the basic types they realize. We keep in mind that τ obviously depends on φ, but choose to suppress this in our notation.

A *basic* 1-*type* of vocabulary τ is a maximally consistent finite set of atomic and negated atomic formulae in the single variable x. We write $\alpha = \mathrm{tp}_{\mathfrak{A}}(a)$ for the basic 1-type realized by a in $\mathfrak{A}$, and $\boldsymbol{\alpha}$ for the finite set of all basic 1-types (in the fixed finite relational vocabulary τ), and $\boldsymbol{\alpha}_{\mathfrak{A}}$ for the set of all basic 1-types realized in $\mathfrak{A}$.

A *basic* 2-*type* of vocabulary τ similarly is a maximally consistent finite set of atomic and negated atomic formulae in variables x and y, containing $\neg x = y$ rather than $x = y$ (for non-degeneracy of the 2-type). Standard notation for the basic 2-types of a non-degenerate pair (a, b) of elements in some τ-structure $\mathfrak{A}$ is $\beta = \mathrm{tp}_{\mathfrak{A}}(a, b)$. We let $\boldsymbol{\beta}$ denote the finite set of all basic 2-types (in the fixed finite relational vocabulary τ), and $\boldsymbol{\beta}_{\mathfrak{A}}$ for the set of all basic 2-types realized in $\mathfrak{A}$.

Since basic types are finite collections of quantifier-free FO^2-formulae we may identify them with the conjunctions over these sets and regard them as quantifier-free FO^2-formulae. Corresponding notation is self-explanatory. E.g. we write $(\mathfrak{A}, a) \models \alpha(x)$ if $\mathrm{tp}_{\mathfrak{A}}(a) = \alpha$, or $\beta(x, y) \models \chi(x, y)$ if any realization of β necessarily satisfies χ.

Obviously, $\mathfrak{A}$ satisfies a normal form FO^2-sentence $\varphi = \forall x \forall y \chi_0 \wedge \bigwedge_{i=1}^{m} \forall x \exists y \chi_i$, if and only if

(i) all $\beta \in \boldsymbol{\beta}_{\mathfrak{A}}$ satisfy $\beta \models \chi_0$.

(ii) for every i, $1 \leqslant i \leqslant m$, and every every a in $\mathfrak{A}$ there is some $b \neq a$ in $\mathfrak{A}$ such that $\mathrm{tp}_{\mathfrak{A}}(a, b) \models \chi_i$.

One may think of the task of constructing a model for φ over some given domain A as the task of allocating basic 2-types to all non-degenerate pairs (a, b) in A in a consistent manner (i.e. with agreement on the basic 1-types in shared elements) and such that (i) and (ii) are satisfied.

Let, for any 2-type β, $\beta|_x$ and $\beta|_y$ denote the unique 1-types of its x- and y-component. Clearly $\boldsymbol{\alpha}_{\mathfrak{A}} = \{\beta|_x \,|\, \beta \in \boldsymbol{\beta}_{\mathfrak{A}}\} = \{\beta|_y \,|\, \beta \in \boldsymbol{\beta}_{\mathfrak{A}}\}$ (assuming, as we always will, that $|A| \geqslant 2$). Elements whose basic 1-type is realized exactly once in $\mathfrak{A}$ play a special role. It has become customary to address them as *kings*, and their basic 1-types as *royal*. Obviously $\alpha \in \boldsymbol{\alpha}_{\mathfrak{A}}$ is royal if and only if there is *no* $\beta \in \boldsymbol{\beta}_{\mathfrak{A}}$ such that $\beta|_x = \beta|_y = \alpha$.

The following theorem from [4] improves on Mortimer's result in giving a better and essentially tight bound on the size of small models for FO^2-sentences. We choose a technically slightly more elaborate variant of the standard formulation, which is however implicit in the proof of this theorem as given in [4] (see also the outline in [5]).

Theorem 3.2 (Grädel,Kolaitis,Vardi) *Any satisfiable* FO^2*-sentence in normal form* $\varphi = \forall x \forall y \chi_0 \wedge \bigwedge_{i=1}^{m} \forall x \exists y \chi_i$ *has a model whose size is at most* $(4m+1)|\boldsymbol{\alpha}|$. *If the set of* 1*-types that can be realized in models of* φ *may a priori be restricted to a subset of* $\boldsymbol{\alpha}$ *of cardinality* ν, *then this bound accordingly improves to* $(4m+1)\nu$.

3.2 An analysis of (well-)ordered models

It turns out to be convenient always to have several consecutive realizations of non-royal 1-types. This can be achieved through the introduction of a series of indistinguishable copies of individual realizations. The observation that a corresponding extension of a structure preserves its normal form FO^2-theory is already used by Mortimer in [9].

Lemma 3.3 *Let c be a positive constant. For any $\mathfrak{A} \in \mathcal{O}$ (respectively $\mathfrak{A} \in \mathcal{WO}$) there is some $\mathfrak{B} \in \mathcal{O}$ (respectively $\mathfrak{B} \in \mathcal{WO}$) satisfying exactly the same normal form* FO^2-*sentences as $\mathfrak{A}$, and such that for every non-royal $b \in B$ there is an interval $I(b)$ of at least c consecutive elements around b, such that for all $b', b'' \in I(b)$ and all kings k, $\mathrm{tp}_{\mathfrak{B}}(k, b') = \mathrm{tp}_{\mathfrak{B}}(k, b'')$.*[3] *$\mathfrak{B}$ may be chosen as an extension of $\mathfrak{A}$ of size $|B| \leqslant c|A|$.*

Proof. The desired model $\mathfrak{B}$ is obtained through repeated application of the following construction. Let $a \in A$ be non-royal, which implies that there is some $a' \in A$, such that $a' \neq a$ but $\mathrm{tp}_{\mathfrak{A}}(a') = \mathrm{tp}_{\mathfrak{A}}(a)$. Let $\mathfrak{A}' \supseteq \mathfrak{A}$ have universe $A \dot\cup \{a''\}$ with a new element a'', and put $\mathrm{tp}_{\mathfrak{A}'}(a'', b) = \mathrm{tp}_{\mathfrak{A}}(a, b)$ for all $b \in A \setminus \{a\}$. Finally put $\mathrm{tp}_{\mathfrak{A}'}(a, a'') = \mathrm{tp}_{\mathfrak{A}}(a, a')$. It is checked that $<$ is an ordering of $\mathfrak{A}'$, and that a and a'' are direct neighbours w.r.t. $<$. Basic 2-types of elements from A are preserved in $\mathfrak{A}'$, and $\mathfrak{A}$ and $\mathfrak{A}'$ satisfy exactly the same normal form FO^2-sentences.

It is straightforward to verify that a $(c-1)$-fold extension of every non-royal element of $\mathfrak{A}$ according to the above pattern yields $\mathfrak{B}$ as desired. $\square$

Let $\boldsymbol{\alpha}_{\mathfrak{A}}$ be the set of basic 1-types of $\mathfrak{A}$, $\boldsymbol{\alpha}_{\mathfrak{A}}^K \subseteq \boldsymbol{\alpha}_{\mathfrak{A}}$ the set of royal 1-types, $K \subseteq A$ the set of kings in $\mathfrak{A}$, and $\boldsymbol{\beta}_{\mathfrak{A}}$ the set of basic 2-types of $\mathfrak{A}$.

For every $\alpha \in \boldsymbol{\alpha}$ the following formulae $\alpha^{\perp}(x)$ and $\alpha^{\top}(x)$ define the minimal, respectively maximal, realization of α. Let $\boldsymbol{\alpha}_{\mathfrak{A}}^{\perp}, \boldsymbol{\alpha}_{\mathfrak{A}}^{\top} \subseteq \boldsymbol{\alpha}_{\mathfrak{A}}$ be the sets of those α that possess a minimal, respectively maximal, realization in $\mathfrak{A}$. Note that $\boldsymbol{\alpha}_{\mathfrak{A}}^{\perp} = \boldsymbol{\alpha}_{\mathfrak{A}}^{\top} = \boldsymbol{\alpha}_{\mathfrak{A}}$ for finite $\mathfrak{A}$, and $\boldsymbol{\alpha}_{\mathfrak{A}}^{\perp} = \boldsymbol{\alpha}_{\mathfrak{A}}$ for $\mathfrak{A} \in \mathcal{WO}$.

$$\begin{aligned}
\alpha^{\perp}(x) &= \alpha(x) \wedge \forall y\big((\alpha(y) \wedge x \neq y) \rightarrow x < y\big), \\
\alpha^{\top}(x) &= \alpha(x) \wedge \forall y\big((\alpha(y) \wedge x \neq y) \rightarrow y < x\big).
\end{aligned}$$

Suppose that $\mathfrak{A} \in \mathcal{O}$ satisfies the normal form sentence $\varphi = \forall x \forall y \chi_0 \wedge \bigwedge_{i=1}^{m} \forall x \exists y \chi_i$. W.l.o.g. we assume that $\mathfrak{A}$ realizes every non-royal 1-type in at least three consecutive elements, all of which moreover are indistinguishable w.r.t. the 2-types they realize with the kings $k \in K \subseteq A$ (cf. Lemma 3.3, to be applied with $c = 3$).

We fix some finite subset $C \subseteq A$ composed of witnesses for the $\forall x \exists y \chi_i$-requirements at kings $k \in K$ as follows. For $k \in K$ and $1 \leqslant i \leqslant m$ pick some $a \in A$ such that

[3] This condition in particular implies $\mathrm{tp}_{\mathfrak{B}}(b') = \mathrm{tp}_{\mathfrak{B}}(b'')$.

$\mathfrak{A} \models \chi_i[k, a]$. If a happens to be a king itself, it is collected into C. If a is a non-royal element of $\mathfrak{A}$, we choose for c the central element of a group of three consecutive elements of the same 1-type and such that $\mathfrak{A} \models \chi_i[k, c]$. This choice implies in particular that a non-royal c satisfies none of the $\alpha^\perp$ or $\alpha^\top$.

Let C consist of K together with the elements c obtained in this way. Note that $|C| \leqslant (m+1)|K| \leqslant (m+1)|\boldsymbol{\alpha}|$. Let $\mathfrak{A}_C$ denote $\mathfrak{A}$ with such choice for C, and let $\mathfrak{C} \subseteq \mathfrak{A}$ be the induced substructure whose universe is C. In the suggestive terminology of [4], $\mathfrak{C}$ forms a *court* for the kings of $\mathfrak{A}$.

Let $(I_s)_{0 \leqslant s \leqslant N}$ be the smallest family of consecutive intervals [4], enumerated in the natural order, that partitions A in such a way that

– every element $c \in C$ gives rise to a singleton interval $I_{\sigma(c)} = \{c\}$ in the family,
– for every $\alpha \in \boldsymbol{\alpha}_{\mathfrak{A}}$ the smallest interval containing all realizations of α is a union of intervals from this family.

It is readily checked that this family is well-defined, and that its length is linearly bounded in $|\boldsymbol{\alpha}|$: any division between consecutive intervals is located directly below or above some $c \in C$, or is one of the Dedekind cuts induced by the realizations of some α, $D = \bigcup_{a \models \alpha}\{b \in A \mid b \leqslant a\}$ or $D = \bigcup_{a \models \alpha}\{b \in A \mid b \geqslant a\}$; it follows that there are at most $2(|\boldsymbol{\alpha}| + |C|)$ such loci of separation, whence for C as above $N \leqslant (2m+4)|\boldsymbol{\alpha}|$. Let for $0 \leqslant s \leqslant N$

$$\boldsymbol{\alpha}_s = \big\{\mathrm{tp}_{\mathfrak{A}}(a) \;\big|\; a \in I_s\big\}$$

be the non-empty subset of $\boldsymbol{\alpha}_{\mathfrak{A}}$ consisting of the 1-types realized in I_s, and let

$$\boldsymbol{\alpha}_s^- = \boldsymbol{\alpha}_s \setminus \bigcup_{t<s} \boldsymbol{\alpha}_t \qquad \text{and} \qquad \boldsymbol{\alpha}_s^+ = \boldsymbol{\alpha}_s \setminus \bigcup_{t>s} \boldsymbol{\alpha}_t$$

be the sets of those 1-types in $\boldsymbol{\alpha}_s$ that are not realized below or above I_s. The following facts about $\mathfrak{C} \subseteq \mathfrak{A}$, the $\boldsymbol{\alpha}_s$, $\boldsymbol{\alpha}^K := \boldsymbol{\alpha}^K_{\mathfrak{A}}$, $\boldsymbol{\alpha}^\perp := \boldsymbol{\alpha}^\perp_{\mathfrak{A}}$, $\boldsymbol{\alpha}^\top := \boldsymbol{\alpha}^\top_{\mathfrak{A}}$, $\sigma \colon C \to \{0, \dots, N\}$, and $\boldsymbol{\beta}_0 := \boldsymbol{\beta}_{\mathfrak{A}}$ will be useful later.

(i) $\boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top \subseteq \bigcup \boldsymbol{\alpha}_s \subseteq \{\alpha \in \boldsymbol{\alpha} \mid \alpha \models \neg x < x\}$;
if $\boldsymbol{\alpha}_s \cap \boldsymbol{\alpha}^K \neq \emptyset$, then $s = \sigma(c)$ for some $c \in C$.

(ii) For $c \in C$, $\boldsymbol{\alpha}_{\sigma(c)} = \{\mathrm{tp}_{\mathfrak{C}}(c)\}$, and either $\mathrm{tp}_{\mathfrak{C}}(c) \in \boldsymbol{\alpha}^K$, or $\boldsymbol{\alpha}^-_{\sigma(c)} = \boldsymbol{\alpha}^+_{\sigma(c)} = \emptyset$.

(iii) For $s < t$ and $\alpha \in \boldsymbol{\alpha}_s$, $\alpha' \in \boldsymbol{\alpha}_t$, there is some $\beta \in \boldsymbol{\beta}_0$ containing $x < y$ and such that $\beta|_x = \alpha$ and $\beta|_y = \alpha'$.

(iv) If $\boldsymbol{\alpha}_s$ does not consist of a single royal 1-type and if $\alpha, \alpha' \in \boldsymbol{\alpha}_s$, then there are $\beta_1, \beta_2 \in \boldsymbol{\beta}_0$ both containing $x < y$ and such that $\beta_1|_x = \alpha$, $\beta_1|_y = \alpha'$, and $\beta_2|_x = \alpha'$, $\beta_2|_y = \alpha$.

[4] An *interval* is a subset I s.t. $(x < z < y \wedge x \in I \wedge y \in I) \to z \in I$.

If $\mathfrak{A} \in \mathcal{WO}$, moreover:

(v) $|\boldsymbol{\alpha}_s^-| \leqslant 1$.

(vi) $\boldsymbol{\alpha}^\perp = \bigcup \boldsymbol{\alpha}_s$.

And if $\mathfrak{A}$ is finite, moreover:

(vii) $|\boldsymbol{\alpha}_s^+| \leqslant 1$.

(viii) $\boldsymbol{\alpha}^\top = \bigcup \boldsymbol{\alpha}_s$.

The reasons for these are given in the following:

(i) is obvious.

(ii) is a direct consequence of the choice of $C \subseteq A$, and in particular of the choice of non-royal elements in C: if $c \in I_s$ has non-royal 1-type α, then c has an immediate predecessor and an immediate successor, both of type α; by construction the predecessor is a member of I_{s-1} and the successor belongs to I_{s+1}.

(iii) is obvious.

(iv) is obvious for $\alpha = \alpha'$. Assume therefore that $\alpha \neq \alpha'$ and that, contrary to what is claimed in (iv), for instance all realizations of α' precede all those of α. The defining conditions for the sequence of the I_s imply that in this case α and α' cannot be realized in the same interval.

For (v), in the case of wellordered $\mathfrak{A}$, note that if $\alpha \in \boldsymbol{\alpha}_s^-$, then the $<$-least realization of α must be a member of I_s. The defining conditions for the I_s imply that no I_s can contain the minimal realizations of more than one type.

(vi) is obvious for wellordered $\mathfrak{A}$.

For (vii) assume that $\mathfrak{A}$ is finite, and let $\alpha \in \boldsymbol{\alpha}_s^+$. Then, for reasons as in (v), the $<$-maximal realization of α has to be a member of I_s, and this can happen for at at most one α.

(viii) is obvious for finite $\mathfrak{A}$.

Definition 3.4 Consider a tuple Γ consisting of a finite, linearly ordered structure $\mathfrak{C}$ and a family of non-empty subsets $\boldsymbol{\alpha}_s \subseteq \boldsymbol{\alpha}$ (indexed by $0 \leqslant s \leqslant N$ for some N) together with subsets $\boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top \subseteq \boldsymbol{\alpha}$, an injection $\sigma \colon C \to \{0, \dots, N\}$, and a subset $\boldsymbol{\beta}_0 \subseteq \boldsymbol{\beta}$. $\Gamma = \langle \mathfrak{C}, (\boldsymbol{\alpha}_s), \boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top, \sigma, \boldsymbol{\beta}_0 \rangle$ is

(a) *admissible for* $\mathcal{O}$, if conditions (i)–(iv) above are satisfied.

(b) *admissible for* $\mathcal{WO}$, if (i)–(vi) are satisfied.

(c) *admissible for* $\mathcal{O}_{\mathit{fin}}$, if (i)–(viii) are satisfied.

Definition 3.5 With $\mathfrak{A}_C$ associate, as characteristic data, the following

$$\mathrm{char}(\mathfrak{A}_C) = \langle \mathfrak{C}, (\boldsymbol{\alpha}_s)_{1 \leqslant s \leqslant N}, \boldsymbol{\alpha}_{\mathfrak{A}}^K, \boldsymbol{\alpha}_{\mathfrak{A}}^\perp, \boldsymbol{\alpha}_{\mathfrak{A}}^\top, \sigma, \boldsymbol{\beta}_{\mathfrak{A}} \rangle.$$

Clearly $\mathrm{char}(\mathfrak{A}_C)$ — prepared according to Lemma 3.3 and with an appropriate choice for the court $\mathfrak{C}$ — is admissible for $\mathcal{K}$ if $\mathfrak{A} \in \mathcal{K}$, where $\mathcal{K}$ is any of $\mathcal{O}$, $\mathcal{WO}$ or $\mathcal{O}_{\mathrm{fin}}$.

3.3 Reconstruction of an ordered model

Towards a proof of Theorem 1.2 we shall reconstruct an appropriately ordered model for a given normal-form FO^2-sentence

$$\varphi = \forall x \forall y \chi_0 \wedge \bigwedge_{i=1}^{m} \forall x \exists y \chi_i$$

from characteristic data as above, and from an arbitrary model of a suitable auxiliary FO^2-sentence associated with φ. Towards finding this auxiliary sentence, we expand our prepared model $\mathfrak{A}_C$ of φ with unary predicates U_s, for $0 \leqslant s \leqslant N$ for the partition of A into the intervals I_s. In what follows, basic types $\alpha \in \boldsymbol{\alpha}$ or $\beta \in \boldsymbol{\beta}$ still refer to the original vocabulary, while we keep in mind that these are not complete basic types in terms of the expanded structure.

Let $\Gamma = \big\langle \mathfrak{C}, (\boldsymbol{\alpha}_s)_{0 \leqslant s \leqslant N}, \boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top, \sigma, \boldsymbol{\beta}_0 \big\rangle$ and let $\varphi^*[\Gamma]$ be a normal form FO^2-sentence in the expanded vocabulary expressing the following (which are all satisfied in the proposed expansion of $\mathfrak{A}_C$ and for $\Gamma = \mathrm{char}(\mathfrak{A})$):

(1) φ
(2) the U_s partition the universe,
(3) $\boldsymbol{\alpha}_s$ is precisely the set of basic 1-types realized in U_s,
(4) each $U_{\sigma(c)}$ is a singleton set containing an element that realizes $\mathrm{tp}_{\mathfrak{C}}(c)$,
(5) each $\alpha \in \boldsymbol{\alpha}^K$ is realized exactly once,
(6) for $\alpha \in \boldsymbol{\alpha}^K$: $\bigwedge_{i=1}^m \forall x \exists y \big(\alpha(x) \rightarrow (\bigvee_c U_{\sigma(c)} y \wedge \chi_i)\big)$,
(7) for all $c \neq c'$ in $\mathfrak{C}$ with $\mathrm{tp}_{\mathfrak{C}}(c, c') = \beta$: $\forall x \forall y (U_{\sigma(c)} x \wedge U_{\sigma(c')} y \rightarrow \beta(x, y))$
(8) for $\alpha \in \boldsymbol{\alpha}_s^- \cap \boldsymbol{\alpha}^\perp$: $\exists x \big(\alpha^\perp(x) \wedge U_s x\big)$,
(9) for $\alpha \in \boldsymbol{\alpha}_s^+ \cap \boldsymbol{\alpha}^\top$: $\exists x \big(\alpha^\top(x) \wedge U_s x\big)$,
(10) $\forall x \forall y (x < y \vee y < x \vee x = y) \wedge \forall x \forall y \neg (x < y \wedge y < x)$,
(11) for all $s < t$: $\forall x \forall y \big((U_s x \wedge U_t y) \rightarrow x < y\big)$.

Note in particular that a model of $\varphi^*[\Gamma]$ contains an isomorphic copy of $\mathfrak{C}$, by (7).

As satisfiability of φ over $\mathcal{O}$, $\mathcal{WO}$, or $\mathcal{O}_{\mathrm{fin}}$ is going to be reduced to satisfiability of suitable $\varphi^*[\Gamma]$, it is important to bound the crucial parameters that govern the model checking complexity for $\varphi^*[\Gamma] = \varphi^*[\mathfrak{C}, (\boldsymbol{\alpha}_s), \boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top, \sigma]$ over small candidate models. These parameters are

- the length of $\varphi^*[\Gamma]$, which is polynomially bounded in $N + |\boldsymbol{\alpha}| + |\boldsymbol{\beta}|$.
- the size of small models of $\varphi^*[\Gamma]$. Here Theorem 3.2 provides a bound that is linear in both the number of 1-types that may be realized, and in the number of $\forall\exists$-requirements in $\varphi^*[\Gamma]$. Both are polynomially bounded in $N + |\boldsymbol{\alpha}|$.

As we have seen above, N can be linearly bounded in $|\boldsymbol{\alpha}|$, and as $|\boldsymbol{\alpha}|$ and $|\boldsymbol{\beta}|$ are exponentially bounded in $|\varphi|$ (whether in normal form or not), we find that both, model size and formula length, can be exponentially bounded. The relevant model checking tasks

$$\mathfrak{B} \models \varphi^*[\mathfrak{C}, (\boldsymbol{\alpha}_s), \boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top, \sigma]$$

are thus in EXPTIME w.r.t. $|\varphi|$, as model checking for any variable-bounded fragment of first-order is in PTIME, cf. e.g. [5]. The admissibility conditions in Definition 3.4 for $\Gamma = \big\langle \mathfrak{C}, (\boldsymbol{\alpha}_s), \boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top, \sigma, \boldsymbol{\beta}_0 \big\rangle$ are checkable in time polynomial in $N + |\boldsymbol{\alpha}| + |\boldsymbol{\beta}|$. Thus there is a NEXPTIME algorithm for checking whether there is an admissible Γ of the appropriate type such that $\varphi^*[\Gamma]$ is satisfiable.

It remains to argue that satisfiability of some such $\varphi^*[\Gamma]$ is a valid criterion for satisfiability of φ over $\mathcal{O}$, $\mathcal{WO}$, and $\mathcal{O}_{\text{fin}}$, respectively. This is brought out in the following proposition.

Proposition 3.6 *Let $\varphi = \forall x \forall y \chi_0 \wedge \bigwedge_{i=1}^m \forall x \exists y \chi_i$ be a normal form FO^2-sentence, and let $\mathcal{K}$ be either $\mathcal{O}$ or $\mathcal{WO}$ or $\mathcal{O}_{fin}$. Then the following are equivalent:*

(a) $\varphi \in sat_{\mathcal{K}}(\mathrm{FO}^2)$.

(b) $\varphi^*[\Gamma] \in sat(\mathrm{FO}^2)$ *for some* $\Gamma = \big\langle \mathfrak{C}, (\boldsymbol{\alpha}_s), \boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top, \sigma, \boldsymbol{\beta}_0 \big\rangle$ *that is admissible for $\mathcal{K}$ with $\boldsymbol{\beta}_0 \subseteq \big\{ \beta \mid \beta \models \chi_0 \big\}$.*

Corollary 3.7 *$sat_{\mathcal{O}}(\mathrm{FO}^2)$, $sat_{\mathcal{WO}}(\mathrm{FO}^2)$ and fin-$sat_{\mathcal{O}}(\mathrm{FO}^2)$ are in* NEXPTIME*; in fact* NEXPTIME*-complete.*

Proof of the proposition. The implications (a) $\Rightarrow$ (b) are already settled by the considerations that led us to $\varphi^*[\mathrm{char}(\mathfrak{A}_C)]$. For the converse implications let $\Gamma = \big\langle \mathfrak{C}, (\boldsymbol{\alpha}_s), \boldsymbol{\alpha}^K, \boldsymbol{\alpha}^\perp, \boldsymbol{\alpha}^\top, \sigma, \boldsymbol{\beta}_0 \big\rangle$ be admissible for the appropriate class $\mathcal{K}$, $\boldsymbol{\beta}_0 \subseteq \big\{ \beta \mid \beta \models \chi_0 \big\}$, and let $\mathfrak{B} \models \varphi^*[\Gamma]$. We construct an ordered model $\mathfrak{A} \models \varphi$, which in the case that Γ is admissible for $\mathcal{WO}$ is actually wellordered, and actually finite if Γ is even admissible for $\mathcal{O}_{\text{fin}}$.

The ordered universe of $\mathfrak{A}$. Let $(A, <)$ be the ordered sum of ordered intervals $(I_s, <)$, $(A, <) = \sum_s (I_s, <)$. Along with the construction of the I_s we designate the basic 1-types, which the elements of I_s are to realize.

- If $\boldsymbol{\alpha}_s$ contains a single royal type, then I_s consists of a single element of that type.

For the other I_s, which contain no royal types (by admissibility (i)), we distinguish three cases, according to the degree of admissibility. In all cases we use as building blocks the following finite linear orderings $(J, <)$, $(J^-, <)$, and $(J^+, <)$: $(J, <)$ consists of consecutive $3m$-blocks of elements, one block for each $\alpha \in \boldsymbol{\alpha}_s$. Similarly $(J^-, <)$ and $(J^+, <)$ consist of $3m$-blocks for all $\alpha \in \boldsymbol{\alpha}_s^-$, respectively $\boldsymbol{\alpha}_s^+$. The ordering among the blocks is fixed in an arbitrary way. As $\boldsymbol{\alpha}_s^-$ and $\boldsymbol{\alpha}_s^+$ may be empty, so may J^- and J^+.

- If Γ is just admissible for $\mathcal{O}$, let $(I_s, <)$ be the ordering in which $(J, <)$ is preceded by countably many copies of $(J^-, <)$ which are ordered like ω^*, i.e. of the converse order type of the natural numbers, and succeeded by countably many copies of $(J^+, <)$ which are ordered like ω, i.e. in the order type of the naturals:

$$(I_s, <) = (J^-, <) \cdot (\omega, <)^* + (J, <) + (J^+, <) \cdot (\omega, <).$$

- If Γ is admissible for $\mathcal{WO}$, there can be at most one $\alpha \in \boldsymbol{\alpha}_s^-$ (admissibility (v)). Let

$$(I_s, <) = (J^-, <) + (J, <) + (J^+, <) \cdot (\omega, <).$$

- If Γ finally is even admissible for $\mathcal{O}_{\text{fin}}$, then both $\boldsymbol{\alpha}_s^-$ and $\boldsymbol{\alpha}_s^+$ can have at most one element (admissibility (v) and (vii)). We put

$$(I_s, <) = (J^-, <) + (J, <) + (J^+, <).$$

An embedding of $\mathfrak{C}$. If c realizes a royal 1-type, identify c with the corresponding king in $(A, <)$. Otherwise, i.e. if $\mathrm{tp}_{\mathfrak{C}}(c) \notin \boldsymbol{\alpha}^K$, I_s consists of a single $3m$-block of that type (admissibility (ii)), which is also the type realized by the element marked $U_{\sigma(c)}$ in $\mathfrak{B}$ (compare (4)). We identify c with the first element in I_s. Note that this identification automatically respects the ordering of $\mathfrak{C}$ (by (11)). We isomorphically embed $\mathfrak{C}$ into $\mathfrak{A}$ by putting $\mathrm{tp}_{\mathfrak{A}}(c, c') = \mathrm{tp}_{\mathfrak{B}}(c, c')$.

The $\forall\exists$-requirements. Having specified $(A, <)$, declared basic 1-types $\mathrm{tp}_{\mathfrak{A}}(a)$ for all $a \in A$, and isomorphically embedded $\mathfrak{C}$, we proceed to assign 2-types β to some pairs, with a view to satisfying the requirements $\exists y \chi_i(x, y)$ at every $a \in A$. The ensuing partial interpretation will be completed in the final step.

- At kings, all requirements have already been taken care of, as they are even satisfied within $\mathfrak{C}$ (by (6)).
- Now for non-kings: let $a \in I_s$ realize some $\alpha \in \boldsymbol{\alpha}_s$, $\alpha \notin \boldsymbol{\alpha}^K$.
  - Suppose first that $\alpha \notin \boldsymbol{\alpha}_s^{+/-}$, which (by admissibility (ii)) is in particular the case for non-royal $a = c \in C$. If $a = c \in C$, let b be the corresponding element of $U_{\sigma(c)}$ in $\mathfrak{B}$, otherwise choose for b any realization of α in U_s. Further choose elements b_i such that $\mathfrak{B} \models \chi_i[b, b_i]$. ($\mathfrak{B} \models \varphi$ by (1).) We shall put $\mathrm{tp}_{\mathfrak{A}}(a, a_i) = \mathrm{tp}_{\mathfrak{B}}(b, b_i)$ for suitable a_i of matching position and type in A, which are located as follows:

    If b_i is royal, choose a_i as the corresponding king in A. Putting $\mathrm{tp}_{\mathfrak{A}}(a, a_i) = \mathrm{tp}_{\mathfrak{B}}(b, b_i)$ cannot clash with previous stipulations even if $a \in C$, as $\mathfrak{B}$ and $\mathfrak{A}$ both respect $\mathfrak{C}$ (by (7)).

If b_i is not royal, first locate an appropriate $3m$-block of realizations of $\mathrm{tp}_{\mathfrak{B}}(b_i)$ as follows. If $b_i \in U_t$ where $t \neq s$, select such a $3m$-block in I_t. If $b_i \in U_s$, and if $\mathfrak{B} \models b_i < b$, then either $\mathrm{tp}_{\mathfrak{B}}(b_i) \in \boldsymbol{\alpha}_s^-$ and we may choose a $3m$-block of type $\mathrm{tp}_{\mathfrak{B}}(b_i)$ that precedes the block of a within I_s, or $\mathrm{tp}_{\mathfrak{B}}(b_i) \notin \boldsymbol{\alpha}_s^-$ and we may choose a $3m$-block of that type in some I_t where $t < s$. If $b_i \in U_s$, and if $\mathfrak{B} \models b < b_i$, then either $\mathrm{tp}_{\mathfrak{B}}(b_i) \in \boldsymbol{\alpha}_s^+$ and we may choose a $3m$-block of that type above the block of a within I_s, or $\mathrm{tp}_{\mathfrak{B}}(b_i) \notin \boldsymbol{\alpha}_s^+$ and we find a $3m$-block of that type in some I_t where $t > s$.

Having thus located an appropriate $3m$-block, choose a_i within that $3m$-block as the i-th element from the second/third/first group of m elements, if a is from the first/second/third group of m elements within its $3m$-block. This specification makes sure that we never encounter clashes between different tasks (including previous stipulations concerning the elements of C), since different tasks involve disjoint sets of pairs.

We now let $\mathrm{tp}_{\mathfrak{A}}(a, a_i) = \mathrm{tp}_{\mathfrak{B}}(b, b_i)$, so that $\mathfrak{A} \models \chi_i[a, a_i]$.

- If $\mathrm{tp}_{\mathfrak{A}}(a) = \alpha \in \boldsymbol{\alpha}_s^-$, we may go through the same steps after choosing b as an arbitrary realization of α within U_s in the case of admissibility for just $\mathcal{O}$, and as a realization of $\alpha^{\perp}(x)$ in U_s in the case of admissibility for $\mathcal{WO}$ (by (8) and admissibility (vi)). This choice makes sure that $3m$-blocks of appropriate type and position matching the non-royal b_i in $\mathfrak{B}$ are available in A. In the $\mathcal{WO}$ case it now follows that if $b_i < b$, then $\mathrm{tp}_{\mathfrak{B}}(b_i) \notin \boldsymbol{\alpha}_s^-$, so there is an appropriate $3m$-block in some I_t for $t < s$; for $b_i > b$ we go through the same distinction of cases as above.
- If $\mathrm{tp}_{\mathfrak{A}}(a) = \alpha \in \boldsymbol{\alpha}_s^+$, we choose b as an arbitrary realization of α in U_s in the case of admissibility for $\mathcal{O}$, and as a realization of $\alpha^{\top}(x)$ in U_s (by (9) and admissibility (viii)) in the case of admissibility for $\mathcal{O}_{\mathrm{fin}}$. In both cases we find $3m$-blocks of appropriate type and position matching the non-royal b_i, by an argument entirely analogous to the above.

Completion of the interpretation. It remains to settle all remaining $\mathrm{tp}_{\mathfrak{A}}(a, a')$ in accordance with χ_0. The availability of appropriate β is guaranteed by the admissibility conditions (iii) and (iv) pertaining to $\boldsymbol{\beta}_0$. □

Remark It is rather obvious how to modify the above proof in order to treat some other natural classes of linear orderings apart from wellorderings, finite orderings and the full class of all linear orderings. One interesting such class of models is the class $\mathcal{K}_\omega$ of structures that are linearly ordered of order type ω. A treatment of the satisfiability problem for FO^2 over $\mathcal{K}_\omega$, which is otherwise strictly analogous to the above, can be based on the following notion of admissibility. Call $\Gamma = \langle \mathfrak{C}, (\boldsymbol{\alpha}_s), \boldsymbol{\alpha}^K, \boldsymbol{\alpha}^{\perp}, \boldsymbol{\alpha}^{\top}, \sigma, \boldsymbol{\beta}_0 \rangle$ admissible for $\mathcal{K}_\omega$ if it is admissible for $\mathcal{WO}$ and if in addition

(i) only the last member $\boldsymbol{\alpha}_N$ of the family of the $\boldsymbol{\alpha}_s$ may have $|\boldsymbol{\alpha}_N^+| > 1$,

(ii) for $0 \leqslant s < N$, $\boldsymbol{\alpha}_s^+ \subseteq \boldsymbol{\alpha}^\top$, and $\boldsymbol{\alpha}_N \not\subseteq \boldsymbol{\alpha}^K$.

The analogue of Proposition 3.6 then shows that also FO^2-satisfiability over $\mathcal{K}_\omega$ is in NEXPTIME. This actually implies a result of Etessami, Vardi, and Wilke that FO^2-satisfiability over structures $(\omega, <, P_1, \dots, P_l) \in \mathcal{K}_\omega$, where $<$ is the only binary predicate in the vocabulary, is in NEXPTIME [3].

References

1. R. BERGER, *The undecidability of the domino problem*, Mem. AMS, 66, 1966.

2. E. BÖRGER, E. GRÄDEL, AND Y. GUREVICH, *The Classical Decision Problem*, Springer 1997.

3. K. ETESSAMI, M. VARDI, AND T. WILKE, *First-order logic with two variables and unary temporal logic.* In Proc. 12th IEEE Symp. Logic in Compter Science, 1997, pp. 228–235.

4. E. GRÄDEL, P. KOLAITIS, AND M. VARDI, *On the complexity of the decision problem for two-variable first-order logic.* Bulletin of Symbolic Logic, 3, 1997, pp. 53–69.

5. E. GRÄDEL AND M. OTTO, *On logics with two variables.* Theoretical Computer Science, 224, 1999, pp. 73–113.

6. E. GRÄDEL, M. OTTO, AND E. ROSEN, *Undecidability results on two-variable logics*, Archive of Mathematical Logic, 38, 1999, pp. 313–354. A preliminary version appeared in Proceedings of 14th Symposium on Theoretical Aspects of Computer Science 1997, LNCS, vol. 1200, Springer 1997, pp. 249–260.

7. Y. GUREVICH AND I. KORYAKOV, *Remarks on Berger's paper on the domino problem*, Siberian Math. Journal, 13 (1972), pp. 319–321.

8. N. IMMERMAN AND M. VARDI, *Model checking and transitive closure logic*, Proc. 9th Conf. Computer-Aided Verification, 1997, pp. 291–302.

9. M. MORTIMER, *On languages with two variables*, Zeitschr. f. math. Logik u. Grundlagen d. Math., 21, 1975, pp. 135–140.

10. D. SCOTT, *A decision method for validity of sentences in two variables*, Journal of Symbolic Logic, 27, 1962, p. 377.

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