FO^2 with one transitive relation is decidable

Wiesław Szwast[1] and Lidia Tendera*[1]

1 **Institute of Mathematics and Informatics,**
Opole University, Oleska 48, 45-052 Opole, Poland
`[szwast,tendera]@math.uni.opole.pl`

Abstract

We show that the satisfiability problem for the two-variable first-order logic, FO^2, over transitive structures when only one relation is required to be transitive, is decidable. The result is optimal, as FO^2 over structures with two transitive relations, or with one transitive and one equivalence relation, are known to be undecidable, so in fact, our result completes the classification of FO^2-logics over transitive structures with respect to decidability. We show that the satisfiability problem is in 2-NExpTime. Decidability of the finite satisfiability problem remains open.

1998 ACM Subject Classification F.1.1 Models of Computation, F.4.1 Mathematical Logic, F.4.3 Formal Languages

Keywords and phrases classical decision problem, two-variable first-order logic, decidability, computational complexity

Digital Object Identifier 10.4230/LIPIcs.STACS.2013.317

1 Introduction

FO^2 is the restriction of the classical first-order logic over relational signatures to formulae with at most two distinct variables. It is well-known that FO^2 enjoys the finite model property [20], and its satisfiability (hence also finite satisfiability) problem is NExpTime-complete [5].

One particular drawback of FO^2 is that it can neither express transitivity of a binary relation nor say that a binary relation is a partial (or linear) order, or an equivalence relation. These natural properties are important for practical applications, thus research has started to investigate FO^2 over restricted classes of structures in which some distinguished binary symbols are required to be interpreted as transitive relations, orders, equivalences, etc. The idea comes from modal correspondence theory, where various conditions on the accessibility relations allow to restrict the class of Kripke structures considered, e.g. to transitive structures for the modal logic K4 or equivalence structures for the modal logic S5. Orderings, on the other hand, are very natural when considering temporal logics, where they model time flow, but they also are used in different scenarios, e.g. in databases or description logics, to compare objects with respect to some parameters.

Unfortunately, the remarkably robust decidability of modal logics and its various extensions towards greater expressibility does not transfer immediately to extensions of FO^2, and the picture for FO^2 is more complex and to some extent less understood. It appeared that both the satisfiability and the finite satisfiability problems for FO^2 are undecidable in the presence of several equivalence or several transitive relations [6, 7]. These results were later

* Lidia Tendera would like to acknowledge the support of Polish Ministry of Science and Higher Education grant N N206 37133.

30th Symposium on Theoretical Aspects of Computer Science (STACS'13).
Editors: Natacha Portier and Thomas Wilke; pp. 317–328
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

strengthened: FO^2 is undecidable in the presence of two transitive relations [11, 9], three equivalence relations [15], one transitive and one equivalence relation [17], or three linear orders [12].

On the positive side it is known that FO^2 with one or two equivalence relations is decidable [16, 17, 14]. The same holds for FO^2 with one linear order [22]. The intriguing questions left open by this research was the case of FO^2 with one transitive relation and FO^2 with two linear orders.

In this paper we answer the first question positively: we prove that the satisfiability problem for the extension of FO^2 where exactly one binary relation is required to be transitive, FO^2_T, is decidable in 2-NExpTime. The result completes the classification of variants of FO^2 over transitive structures with respect to decidability.

For the special case of two linear orders, ExpSpace-completeness of finite satisfiability is shown, subject to certain restrictions on signatures, in [24]. (The case of unrestricted signatures, and decidability of the general satisfiability problem are currently open.)

It is also worth to compare the above results with results concerning GF^2 i.e. the two-variable restriction of the *guarded fragment* GF [1] where quantifiers are guarded by atoms. GF+TG is the restriction of GF^2 with transitive relations where the transitive relation symbols are allowed to appear only in guards. As shown in [26] undecidability of FO^2 with transitivity transfers to GF^2 with transitivity; however, GF+TG is decidable irrespective of the number of transitive symbols. Moreover, as noted in [11], the decision procedure developed for GF^2+TG can be applied to GF^2 with one transitive relation that is allowed to appear also outside guards, giving 2-ExpTime-completeness of the latter fragment.

Also of note in this context is the interpretation of FO^2 over *data words* and *data trees* that appear e.g. in verification and XML processing. Decidability of FO^2 over data words with one additional equivalence relation was shown in [3]. For more results related to FO^2 over data words or data trees see e.g. [18, 24, 4, 21, 2].

It makes sense to also consider more expressive systems in which we may refer to the transitive closure of some relation. In fact, relatively few decidable fragments of first-order logic with transitive closure are known. One exception is the logic GF^2 with a transitive closure operator applied to binary symbols appearing only in guards [19]. This fragment captures the two-variable guarded fragment with transitive guards, GF^2+TG, preserving its complexity [27, 10]. Also decidable is the satisfiability problem for the logic $\exists\forall(\mathrm{DTC}^+[E])$, i.e. the prefix class $\exists\forall$ extended by the positive deterministic transitive closure operator of one binary relation, which is shown to enjoy the exponential model property [8]. Recently, it has been shown that the satisfiability problem for the two-variable universal fragment of first-order logic with constants remains decidable when extended by the transitive closure of a single binary relation [13]. Whether the same holds for full FO^2 is open.

Expressive power of FO^2_T. As has already been mentioned, FO^2 has the finite model property. Adding one transitive relation to GF^2 (even restricted only to guards) we can write infinity axioms, however models for this logic still enjoy the so called tree-like property, i.e. new elements required by $\forall\exists$-conjuncts can be added independently. Below we give an example of an infinity axiom in FO^2_T that enforces models where in some triples all elements depend on each other.

We use the transitive relation symbol T and two unary symbols P and Q. It is not difficult to formalize the following statements by an FO^2_T formula: (a) T is strictly antisymmetric. (b) Elements of P form one infinite chain. (c) Elements of Q are incomparable. (d) Every element of P has an incomparable element in Q. (e) Every element of Q is smaller than some element in P.

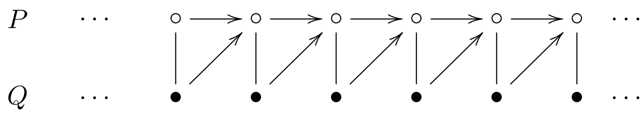

Figure 1 A model satisfying (a)-(e). Arrows depict elements related by T. Lines connect elements required by (d), not connected by T.

In any model satisfying (a)-(e) there is an infinite chain of elements in P that induces an infinite antichain of elements in Q (see Figure 1). Note also, that it suffices that the unique transitive relation is supposed to be a partial ordering.

Outline of the proof. Models for our logic, taking into account the interpretation of the transitive relation, can obviously be seen as partitioned into cliques. As usually for two-variable logics, we first establish a "Scott-type" normal form for FO^2_T: $\forall\forall \ \wedge \bigwedge \forall\exists$, allowing us to restrict the nesting of quantifiers to depth two, as well as to concentrate on the $\forall\exists$-conjuncts demanding "witnesses" for all elements in a model. The form of the $\forall\exists$-conjuncts enables to distinguish witnesses required inside cliques (i.e. realizing a 2-type containing both $T(x, y)$ and $T(y, x)$, see Section 2 for a precise definition) from witnesses outside cliques. We also establish a *small clique property* for FO^2_T (practically the same as in [15]), allowing us to restrict attention to models with cliques exponentially bounded in the size of the signature. Further constructions proceed on levels of cliques rather than individual elements. (An alternative approach would be to consider first the satisfiability problem over an antisymmetric relation T and then reduce the general problem to the aforementioned one taking into account the bound on the clique sizes.)

Crucial to our argument is this property: any infinitely satisfiable sentence has an infinite *narrow* model, i.e. a model whose universe can be partitioned into segments (i.e. sets of cliques) $S_0, S_1, \ldots$, each of doubly exponential size, such that every element in $\bigcup_{i=0}^{j-1} S_i$ requiring a witness outside its clique has the witness either in S_0 or in S_j (so, in every S_k, $k \geq j$, Def. 16). This immediately implies that, when needed, every single segment S_j ($j > 0$) can be removed from the structure, to yield a model with new properties.

To prove existence of narrow models, we first make some useful observations. In particular, we show that a single clique can be duplicated, provided its type called *splice* appears at least twice in a model (Claim 7). The property is used to show the main technical result (Claim 10 and Corollary 11). Next, the idea is generalized in Lemma 13 to show that for any finite subset F of elements, the model can be extended by a fixed number of cliques (depending only on the signature, and not depending on the cardinality of F) providing all required witnesses for elements from F.

As the main result of the paper, we show that from any narrow model we can build a *canonical model* where every two segments of the infinite partition (except the first) are isomorphic and they are connected using at most two distinct similarity types (Def. 19). In fact, these constructions can be seen as an application of the infinite Ramsey theorem [23], where segments of the models are considered to be nodes in a colored graph, and similarity types of pairs of segments are colors of edges.

The above properties suffice to obtain the 2-NExpTime decision procedure for the satisfiability problem for FO^2_T given in Theorem 21 and Corollary 22. We note however that the best lower bound coming from GF^2+TG is 2-ExpTime, thus our result leaves a gap in complexity. We also note that our decision procedure cannot be straightforwardly generalized to solve the finite satisfiability problem for FO^2_T and to the best of our knowledge, the latter problem remains open (see Outlook for some discussion).

2 Preliminaries

We denote by FO^2 the two-variable fragment of first-order logic (with equality) over relational signatures. By FO^2_T we understand the set of FO^2_T-formulas over any signature $\sigma = \sigma_0 \cup \{T\}$, where T is a distinguished binary predicate. The semantics for FO^2_T is as for FO^2, subject to the restriction that T is always interpreted as a *transitive* relation.

In this paper, σ-structures are denoted by Gothic capital letters and their universes by corresponding Latin capitals. Where a structure is clear from context, we frequently equivocate between predicates and their realizations, thus writing, for example, R in place of the technically correct $R^{\mathfrak{A}}$. If $\mathfrak{A}$ is a σ-structure and $B \subseteq A$, then $\mathfrak{A} \restriction B$ denotes the substructure of $\mathfrak{A}$ with the universe B.

An (atomic and proper) *k-type* (over a given signature) is a maximal consistent set of atoms or negated atoms over k distinct variables not containing equality atoms $x_i = x_j$ with $i \neq j$. If $\beta(x, y)$ is a 2-type over variables x and y, then $\beta \restriction x$ (respectively, $\beta \restriction y$) denotes the unique 1-type that is obtained from β by removing atoms with the variable y (respectively, the variable x). We denote by $\boldsymbol{\alpha}$ the set of all 1-types and by $\boldsymbol{\beta}$ the set of all 2-types (over a given signature). Note that $|\boldsymbol{\alpha}|$ and $|\boldsymbol{\beta}|$ are bounded exponentially in the size of the signature. We often identify a type with the conjunction of all its elements.

For a given σ-structure $\mathfrak{A}$ and $a \in A$ we say that *a realizes* a 1-type α if α is the unique 1-type such that $\mathfrak{A} \models \alpha[a]$. We denote by $tp^{\mathfrak{A}}(a)$ the 1-type realized by a. Similarly, for distinct $a, b \in A$, we denote by $tp^{\mathfrak{A}}(a, b)$ the unique 2-type *realized* by the pair a, b, i.e. the 2-type β such that $\mathfrak{A} \models \beta[a, b]$. In general, for finite $B, C \subseteq A$, $B \cap C = \emptyset$, by $tp^{\mathfrak{A}}(B, C)$ we denote the similarity type of the substructure $\mathfrak{A} \restriction (B \cup C)$ (or, in other words, its $card(B \cup C)$-type).

Assume $\mathfrak{A}$ is a σ-structure and $B, C \subseteq A$. We denote by $\boldsymbol{\alpha}^{\mathfrak{A}}$ (respectively, $\boldsymbol{\alpha}^{\mathfrak{A}}[B]$) the set of all 1-types realized in $\mathfrak{A}$ (respectively, realized in $\mathfrak{A} \restriction B$), and by $\boldsymbol{\beta}^{\mathfrak{A}}$ (respectively, $\boldsymbol{\beta}^{\mathfrak{A}}[B]$) the set of all 2-types realized in $\mathfrak{A}$ (respectively, realized in $\mathfrak{A} \restriction B$). We denote by $\boldsymbol{\beta}^{\mathfrak{A}}[a, B]$ the set of all 2-types $tp^{\mathfrak{A}}(a, b)$ with $b \in B$, and by $\boldsymbol{\beta}^{\mathfrak{A}}[B, C]$ the set of all 2-types $tp^{\mathfrak{A}}(b, c)$ with $b \in B, c \in C$.

Let γ be a σ-sentence of the form $\forall x \exists y\, \psi(x, y)$ and $a \in A$. We say that an element $b \in A$ is a *γ-witness* for a in the structure $\mathfrak{A}$ if $\mathfrak{A} \models \psi(a, b)$; b is a *proper γ-witness*, if b is a γ-witness and $a \neq b$.

Scott normal form. As with FO^2, so too with FO^2_T, analysis is facilitated by the availability of normal forms.

▶ **Definition 1.** An FO^2-sentence Ψ is in *Scott normal form* if it is of the following form: $\forall x \forall y\, \psi_0(x, y) \wedge \bigwedge_{i=1}^{M} \forall x \exists y\, \psi_i(x, y)$, where every ψ_i is quantifier-free and includes unary and binary predicate letters only.

Without loss of generality we suppose that for $i \geq 1$, $\psi_i(x, y)$ entails $x \neq y$ (replacing $\psi_i(x, y)$ with $(\psi_i(x, y) \vee \psi_i(x, x)) \wedge x \neq y$, which is sound over all structures with at least two elements).

Two formulas are said to be *strongly equisatisfiable* if they are satisfiable over the same universe. The following Lemma is typical for two-variable logics.

▶ **Lemma 2** ([25, 5]). *For every formula $\varphi \in \mathrm{FO}^2$ one can compute in polynomial time a strongly equisatisfiable normal form formula $\psi \in \mathrm{FO}^2$ over a new signature whose length is linear in the length of φ.*

Suppose the signature σ consists of predicates of arity at most 2. To define a σ-structure $\mathfrak{A}$, it suffices to specify the 1-types and 2-types realized by elements and pairs of elements from the universe A. In the presence of a transitive relation, we classify 2-types according to the transitive connection between x and y. And so, we distinguish $\boldsymbol{\beta}^{\rightarrow}$, $\boldsymbol{\beta}^{\leftarrow}$, $\boldsymbol{\beta}^{\leftrightarrow}$ and $\boldsymbol{\beta}^{-}$ such that $\boldsymbol{\beta} = \boldsymbol{\beta}^{\rightarrow} \dot\cup \boldsymbol{\beta}^{\leftarrow} \dot\cup \boldsymbol{\beta}^{\leftrightarrow} \dot\cup \boldsymbol{\beta}^{-}$ and for instance: $\beta \in \boldsymbol{\beta}^{\rightarrow}$ iff $(T(x,y) \wedge \neg T(y,x)) \in \beta$, $\beta \in \boldsymbol{\beta}^{\leftrightarrow}$ iff $(T(x,y) \wedge T(y,x)) \in \beta$, etc.

For a quantifier-free formula $\varphi(x,y)$ we use superscripts $^{\rightarrow}$, $^{\leftarrow}$, $^{\leftrightarrow}$ and $^{-}$ to define new formulas that explicitly specify the transitive connection between x and y. For instance, for a quantifier-free formula $\varphi(x,y) \in \mathrm{FO}^2_T$ we let $\varphi^{\rightarrow}(x,y) := \varphi(x,y) \wedge T(x,y) \wedge \neg T(y,x)$.

This conversion of FO^2_T-formulae leads to the the following variant of the Scott normal form:

$$\forall x \forall y\, \psi_0 \wedge \bigwedge_{i=1}^{m} \gamma_i \wedge \bigwedge_{i=1}^{\overline{m}} \delta_i \tag{1}$$

where $\gamma_i = \forall x \exists y\, \psi_i^{d_i}(x,y)$ with $d_i \in \{^{\rightarrow}, ^{\leftarrow}, ^{-}\}$, and $\delta_i = \forall x \exists y\, \psi_i^{\leftrightarrow}(x,y)$.

For a fixed sentence Ψ in normal form (1) we often write $\gamma_i \in \Psi$ to indicate that γ_i is a conjunct of Ψ of the form $\forall x \exists y\, \psi_i^{d_i}(x,y)$.

▶ **Lemma 3.** *Let φ be an FO^2_T-formula over a signature τ. We can compute, in polynomial time, a strongly equisatisfiable FO^2_T-formula in normal form, over a signature σ consisting of τ together with a number of additional unary and binary predicates.*

Sketch. We employ the standard technique of renaming subformulas familiar from [25] and [5], noting that any formula $\exists y \psi$ is logically equivalent to $\exists y \psi^{\rightarrow} \vee \exists y \psi^{\leftarrow} \vee \exists y \psi^{\leftrightarrow} \vee \exists y \psi^{-}$. ◀

The following trivial observation will be very useful in the paper.

▶ **Proposition 4.** *Assume $\mathfrak{A}$ is a σ-structure and Ψ is a FO^2_T-sentence over σ in normal form (1). Then $\mathfrak{A} \models \Psi$ if and only if*

(a) for every $a \in A$, for every γ_i $(1 \leq i \leq m)$ there is a γ_i-witness for a in $\mathfrak{A}$,
(b) for every $a \in A$, for every δ_i $(1 \leq i \leq \overline{m})$ there is a δ_i-witness for a in $\mathfrak{A}$,
(c) for every $a, b \in A$, $tp^{\mathfrak{A}}(a,b) \models \psi_0$,
(d) $T^{\mathfrak{A}}$ is transitive in $\mathfrak{A}$.

A small clique property for FO^2_T. Let $\mathfrak{A}$ be a σ-structure. A subset B of A is called *T-connected* if $\boldsymbol{\beta}[B] \subseteq \boldsymbol{\beta}^{\leftrightarrow}[\mathfrak{A}]$. Maximal T-connected subsets of A are called *cliques*. Note that if $\boldsymbol{\beta}[a,\mathfrak{A}] \cap \boldsymbol{\beta}^{\leftrightarrow}[A] = \emptyset$, for some $a \in A$, then $\{a\}$ is a clique. We prove the following *small clique property*.

▶ **Lemma 5.** *Let Ψ be a satisfiable FO^2_T-sentence in normal form, over a signature σ. Then there exists a model of Ψ in which the size of each clique is bounded exponentially in $|\sigma|$.*

We first show how to replace a single clique in models of normal-form FO^2_T-sentences by an equivalent small one. The idea is not new, it was used in [27] to show that T-cliques in models of GF^2+TG can be replaced by appropriate small structures called T-petals (Lemma 17). Later, in [16] it was proved that for any structure $\mathfrak{A}$ and its substructure $\mathfrak{B}$, one may replace $\mathfrak{B}$ by an alternative structure $\mathfrak{B}'$ of a bounded size in such a way that the obtained structure $\mathfrak{A}'$ and the original structure $\mathfrak{A}$ satisfy exactly the same normal form FO^2 formulas. Due to space limitations, a precise statement of the latter lemma and the proof of the small clique model property will appear in the full version of the paper.

3 Splices and duplicability

In the remainder of the paper we fix a relational signature σ and assume Ψ is an FO^2_T-sentence in normal form (1). By Lemma 5, we may already assume (and we do so) that models of Ψ have the small clique property. In the next two sections we assume that Ψ is satisfiable and, if not stated otherwise, $\mathfrak{A} \models \Psi$.

In this section we analyze properties of models of Ψ on the level of cliques rather than individual elements. We give here the key technical argument of the paper (Corollary 11). It says, roughly speaking, that if $\mathfrak{A} \models \Psi$ and elements of a finite subset F of the universe A have their γ_i-witnesses in several "similar" cliques then $\mathfrak{A}$ can be extended by one new clique, where all the elements of F have their γ_i-witnesses.

First, we need to introduce some new notions and notation. For $a \in A$ denote by $Cl^{\mathfrak{A}}(a)$ the unique clique $C \subseteq A$ with $a \in C$. When $F \subseteq A$, denote $Cl^{\mathfrak{A}}(F) = \{Cl^{\mathfrak{A}}(a) : a \in F\}$ and finally, $Cl^{\mathfrak{A}} = Cl^{\mathfrak{A}}(A)$. Note that whenever $B \in Cl^{\mathfrak{A}}$, and $a \in A$ is an element outside the clique B, then the 2-types between a and any element $b \in B$ belong to the same subset of $\boldsymbol{\beta}$, i.e. either to $\boldsymbol{\beta}^{\rightarrow}$, $\boldsymbol{\beta}^{\leftarrow}$ or $\boldsymbol{\beta}^{-}$. So, we might speak about elements of $\mathfrak{A}$ connected with the clique B using "directed" edges. Similarly, we can identify cliques connected with B using "incoming" and "outgoing" edges.

▶ **Definition 6.** Let $\mathfrak{A}$ be a σ-structure and let $B \in Cl^{\mathfrak{A}}$. Define:
- $In^{\mathfrak{A}}(B) \stackrel{def}{=} \{tp^{\mathfrak{A}}(C) : C \in Cl^{\mathfrak{A}} \text{ and } \boldsymbol{\beta}[C,B] \subseteq \boldsymbol{\beta}^{\rightarrow}\}$,
- $Out^{\mathfrak{A}}(B) \stackrel{def}{=} \{tp^{\mathfrak{A}}(C) : C \in Cl^{\mathfrak{A}} \text{ and } \boldsymbol{\beta}[C,B] \subseteq \boldsymbol{\beta}^{\leftarrow}\}$,
- $sp^{\mathfrak{A}}(B) = \langle tp^{\mathfrak{A}}(B), In^{\mathfrak{A}}(B), Out^{\mathfrak{A}}(B)\rangle$,
- $Sp^{\mathfrak{A}} = \{sp^{\mathfrak{A}}(B) : B \in Cl^{\mathfrak{A}}\}$. Elements of $Sp^{\mathfrak{A}}$ are called $\mathfrak{A}$-*splices*.

Splices define cliques reachable from a given clique via T.

We say that two cliques $B, B' \in Cl^{\mathfrak{A}}$ *realize the same splice*, written $B \equiv^{\mathfrak{A}} B'$, if $sp^{\mathfrak{A}}(B) = sp^{\mathfrak{A}}(B')$. When $\mathfrak{A}$ is understood we often omit the superscript in $\equiv^{\mathfrak{A}}$ and write $\equiv$. Note that $\equiv^{\mathfrak{A}}$ is an equivalence relation on $Cl^{\mathfrak{A}}$. Moreover, if we have an a priori upper bound on the size of cliques in $Cl^{\mathfrak{A}}$, then $Cl^{\mathfrak{A}}/_{\equiv}$ is finite (and of bounded cardinality).

Additionally, we distinguish the set $\mathbb{K}(\mathfrak{A})$ of *unique cliques* in $\mathfrak{A}$: $\mathbb{K}(\mathfrak{A}) = \{B \in Cl^{\mathfrak{A}} : card([B]_{\equiv}) = 1\}$ and the corresponding subset $K(\mathfrak{A})$ of the universe of $\mathfrak{A}$, that consists of the elements of the unique cliques: $K(\mathfrak{A}) = \bigcup_{B \in \mathbb{K}(\mathfrak{A})} B$.

Our task is now to show that any model of Ψ containing a non-unique clique B can be extended into a new model of Ψ by adding a copy of B. The copy of B is added in such a way that it also provides, for all conjuncts of the form γ_i, all the witnesses for elements outside the two cliques that have been provided by B. This property will be explored later, when new models will be constructed by removing *segments* (i.e. sets of cliques) from given ones.

For every $\gamma_i \in \Psi$ (recall $\gamma_i = \forall x \exists y\, \psi_i^{d_i}(x,y)$ with $d_i \in \{^{\rightarrow}, ^{\leftarrow}, ^{-}\}$) and for every $a \in A$ we define $W_i^{\mathfrak{A}}(a)$ as the set of all proper γ_i-witnesses for a in $\mathfrak{A}$:

$$W_i^{\mathfrak{A}}(a) \stackrel{def}{=} \{b \in A : \mathfrak{A} \models \psi_i(a,b),\, b \neq a\}.$$

Similarly, for every $\gamma_i \in \Psi$ and for every $F \subseteq A$ we define $W_i^{\mathfrak{A}}(F) \stackrel{def}{=} \bigcup_{a \in F} W_i^{\mathfrak{A}}(a)$.

The following claim states that every non-unique clique in a given model can be properly duplicated, as informally described above.

▶ **Claim 7** (Duplicability)**.** *Assume $\mathfrak{A} \models \Psi$, $B_1 \in Cl^{\mathfrak{A}}$ and $B_1 \notin \mathbb{K}(\mathfrak{A})$. There is an extension $\mathfrak{A}'$ of $\mathfrak{A}$ by one new clique D such that*

1. $\mathfrak{A}' \models \Psi$,

2. *for every conjunct γ_i of Ψ, for every $a \in A$ we have:*
$$B_1 \cap W_i^{\mathfrak{A}}(a) \neq \emptyset \quad \textit{iff} \quad D \cap W_i^{\mathfrak{A}'}(a) \neq \emptyset, \textit{ and}$$
3. $sp^{\mathfrak{A}'}(D) = sp^{\mathfrak{A}'}(B_1) = sp^{\mathfrak{A}}(B_1)$.

Proof. Let $\mathfrak{A} \models \Psi$, $B_1 \in Cl^{\mathfrak{A}}$. Since $B_1 \notin \mathbb{K}(\mathfrak{A})$, there exists $B_2 \in Cl^{\mathfrak{A}}$, $B_2 \neq B_1$ such that $sp^{\mathfrak{A}}(B_2) = sp^{\mathfrak{A}}(B_1)$. Assume $\mathfrak{D}$ is a duplicate of $\mathfrak{A} \restriction B_1$, $D \cap A = \emptyset$, $f_1 : D \mapsto B_1$ and $f_2 : D \mapsto B_2$ are appropriate isomorphism functions. Let $\mathfrak{A}'$ be an extension of $\mathfrak{A}$ with the universe $A \dot\cup D$ such that:

- $tp^{\mathfrak{A}'}(d,b) \stackrel{def}{=} tp^{\mathfrak{A}}(b, f_2(d))$, for every $b \in B_1$, $d \in D$ (note that $\boldsymbol{\beta}[d, B_1] = \boldsymbol{\beta}[B_1, f_2(d)]$ and so $\boldsymbol{\beta}[D, B_1] = \boldsymbol{\beta}[B_1, B_2]$),
- $tp^{\mathfrak{A}'}(d,a) \stackrel{def}{=} tp^{\mathfrak{A}}(f_1(d), a)$, for every $a \in A \setminus (B_1 \cup D)$, $d \in D$ (note that $\boldsymbol{\beta}[d, A \setminus B_1] = \boldsymbol{\beta}[f_1(d), A \setminus B_1]$ and so $\boldsymbol{\beta}[D, A \setminus B_1] = \boldsymbol{\beta}[B_1, A \setminus B_1]$).

To see that $\mathfrak{A}' \models \Psi$ one can show that conditions (a)–(d) of Proposition 4 hold for $\mathfrak{A}'$. ◀

Using the above claim we may build *saturated models* in the following sense.

▶ **Definition 8.** Assume $\mathfrak{A} \models \Psi$. We say that $\mathfrak{A}$ is *witness-saturated*, if $\mathfrak{A}$ has the small clique property and for every $a \in A$, for every $\gamma_i \in \Psi$ $(1 \leq i \leq m)$
$$W_i(a) \subseteq K(\mathfrak{A}) \quad \text{or} \quad W_i(a) \text{ is infinite.}$$
Note that if a witness-saturated model $\mathfrak{A}$ is finite then $A = K(\mathfrak{A})$. By Lemma 5 and by iterative application of Claim 7 we get the following.

▶ **Lemma 9** (Saturated model)**.** *Every satisfiable normal form sentence Ψ has a countable witness-saturated model. Additionally, if $\mathfrak{A}$ is witness-saturated, then the extension $\mathfrak{A}'$ given by Claim 7 is also witness-saturated.*

The above Lemma is essential for the proof of the key technical tool for the paper, Corollary 11, given below. It says that when several elements $a_1, a_2, \ldots, a_n$ of a model $\mathfrak{A}$ have γ_i-witnesses in several distinguished cliques that realize the same splice, one can extend the model $\mathfrak{A}$ by a single clique D (realizing the same splice) in which $a_1, a_2, \ldots, a_n$ have their γ_i-witnesses. The proof is based on a more subtle (than in Claim 7) analysis of models of a normal form sentence Ψ given in Claim 10. In the Claim note that whenever $\boldsymbol{\beta}(C_1, C_2) \in \boldsymbol{\beta}^{\leftarrow}$ then $\boldsymbol{\beta}(C_2, C_1) \notin \boldsymbol{\beta}^{\leftarrow}$.

▶ **Claim 10.** *Assume $\mathfrak{A}$ is countable witness-saturated and $\gamma_i \in \Psi$. Let $C_1, C_2, B_1, B_2 \in Cl^{\mathfrak{A}}$, $\boldsymbol{\beta}(C_1, C_2) \notin \boldsymbol{\beta}^{\leftarrow}$, $B_1 \neq B_2$, $B_1 \cap W_i(C_1) \neq \emptyset$, $B_2 \cap W_i(C_2) \neq \emptyset$, $C_1 \notin \mathbb{K}(\mathfrak{A})$, $C_1 \equiv C_2$ and $B_1 \equiv B_2$. Then, there exists an extension $\mathfrak{A}_1$ of $\mathfrak{A}$ by at least one clique D such that*
(i) $\mathfrak{A}_1 \models \Psi$ *and* $\mathfrak{A}_1$ *is witness-saturated,*
(ii) *for every* $a \in A$: *if* $B_1 \cap W_i^{\mathfrak{A}}(a) \neq \emptyset$ *then* $D \cap W_i^{\mathfrak{A}_1}(a) \neq \emptyset$,
(iii) *for every* $a \in C_2$: *if* $B_2 \cap W_i^{\mathfrak{A}}(a) \neq \emptyset$ *then* $D \cap W_i^{\mathfrak{A}_1}(a) \neq \emptyset$,
(iv) $sp^{\mathfrak{A}_1}(D) = sp^{\mathfrak{A}_1}(B_1) = sp^{\mathfrak{A}}(B_1)$.

Proof. We have several cases. In each case we add a duplicate D of the clique B_1 where both C_1 and C_2 will get their γ_i-witnesses. We sketch only one of the interesting cases.
Case 1. $\boldsymbol{\beta}[C_1, B_1] \subseteq \boldsymbol{\beta}^{-}$ and $\boldsymbol{\beta}[C_1, C_2] \subseteq \boldsymbol{\beta}^{\rightarrow} \cup \boldsymbol{\beta}^{-}$. In this case $C_1 \neq C_2$ and $\gamma_i = \forall x \exists y\, \psi_i^{-}(x,y)$. So, $\boldsymbol{\beta}[C_1, B_1] \subseteq \boldsymbol{\beta}^{-}$ and $\boldsymbol{\beta}[C_2, B_2] \subseteq \boldsymbol{\beta}^{-}$.
Subcase 1.a. $\boldsymbol{\beta}[C_1, C_2] \in \boldsymbol{\beta}^{\rightarrow}$, $\boldsymbol{\beta}[B_1, B_2] \subseteq \boldsymbol{\beta}^{-}$, $\boldsymbol{\beta}[C_1, B_2] \subseteq \boldsymbol{\beta}^{\rightarrow}$ and $\boldsymbol{\beta}[B_1, C_2] \in \boldsymbol{\beta}^{\rightarrow}$.

The construction proceeds in four steps.

Step 1 (copying of B_1). Let $\mathfrak{A}'$ be the witness saturated extension of $\mathfrak{A}$ by one new clique D – a duplicate of the clique B_1 given by Claim 7 and Lemma 9. Observe that all conditions *(i)–(iv)* hold in $\mathfrak{A}'$ except *(iii)* since by construction of $\mathfrak{A}'$ $\boldsymbol{\beta}^{\mathfrak{A}'}[C_2, D] = \boldsymbol{\beta}^{\mathfrak{A}}[C_2, B_1] \subseteq \boldsymbol{\beta}^{\leftarrow}$.

Step 2 (modification of $tp^{\mathfrak{A}'}(C_2, D)$). To ensure that *(iii)* holds, a new structure $\mathfrak{A}_2$ is built by defining $tp^{\mathfrak{A}_2}(C_2, D) \stackrel{def}{=} tp^{\mathfrak{A}'}(C_2, B_2)$.
Step 3 (transitivity correction). To ensure that T is transitive we construct a structure $\mathfrak{A}_3$: for every $X \in Cl^{\mathfrak{A}_2}$, if $\boldsymbol{\beta}[D, X] \subseteq \boldsymbol{\beta}^{\rightarrow}$ and $\boldsymbol{\beta}[X, C_2] \subseteq \boldsymbol{\beta}^{\rightarrow}$ then replace $tp^{\mathfrak{A}_2}(D, X)$ by $tp^{\mathfrak{A}_2}(B_2, X)$. One can observe that $\boldsymbol{\beta}[X, B_2] \subseteq \boldsymbol{\beta}^{-}$, so T is transitive in $\mathfrak{A}_3$.
Step 4 (γ_j-witness in X correction). Let $X \in Cl^{\mathfrak{A}_2}$ be such that the type $tp^{\mathfrak{A}_2}(D, X)$ is changed in Step 3. We show that then there is a clique $X_2 \in Cl^{\mathfrak{A}_3}$ such that $tp^{\mathfrak{A}_3}(X_2) = tp^{\mathfrak{A}_3}(X)$ and $\boldsymbol{\beta}^{\mathfrak{A}_3}[D, X_2] \subseteq \boldsymbol{\beta}^{\rightarrow}$. For, observe that $tp^{\mathfrak{A}_3}(X) \in Out^{\mathfrak{A}_3}(B_1)$ and so $tp^{\mathfrak{A}_3}(X) \in Out^{\mathfrak{A}_3}(B_2)$ since $B_1 \equiv B_2$. Let $X_1 \in Cl^{\mathfrak{A}_3}$, $tp^{\mathfrak{A}_3}(X_1) = tp^{\mathfrak{A}_3}(X)$ and $\boldsymbol{\beta}[B_2, X_1] \subseteq \boldsymbol{\beta}^{\rightarrow}$. Since $\boldsymbol{\beta}[C_1, B_2] \subseteq \boldsymbol{\beta}^{\rightarrow}$ we have $\boldsymbol{\beta}[C_1, X_1] \subseteq \boldsymbol{\beta}^{\rightarrow}$ and so $tp^{\mathfrak{A}_3}(X_1) \in Out^{\mathfrak{A}_3}(C_1)$. Hence, since $C_1 \equiv C_2$ there exist $X_2 \in Cl^{\mathfrak{A}_3}$ such that $tp^{\mathfrak{A}_3}(X_2) = tp^{\mathfrak{A}_3}(X_1)$ and $\boldsymbol{\beta}[C_2, X_2] \subseteq \boldsymbol{\beta}^{\rightarrow}$. Since $\boldsymbol{\beta}[B_1, C_2] \subseteq \boldsymbol{\beta}^{\rightarrow}$, so $\boldsymbol{\beta}[B_1, X_2] \subseteq \boldsymbol{\beta}^{\rightarrow}$, and hence, by construction of $\mathfrak{A}'$, $\boldsymbol{\beta}[D, X_2] \subseteq \boldsymbol{\beta}^{\rightarrow}$. To obtain the required model $\mathfrak{A}_1$ replace $tp^{\mathfrak{A}_3}(D, X_2)$ by $tp^{\mathfrak{A}_2}(D, X)$ $(= tp^{\mathfrak{A}}(B_1, X))$.

Correctness proof of the above construction, as well as other cases, will appear in the full version of the paper. ◀

▶ **Corollary 11.** *Assume $\mathfrak{A}$ is countable witness-saturated, $\gamma_i \in \Psi$ and $X \in Sp^{\mathfrak{A}}$. Let $F \subseteq A \setminus K(\mathfrak{A})$ be a finite set such that for every $a \in F$ there is $b \in W_i^{\mathfrak{A}}(a)$ such that $b \notin K(\mathfrak{A})$ and $sp(Cl^{\mathfrak{A}}(b)) = X$. Then, there exists an extension $\mathfrak{A}'$ of $\mathfrak{A}$ by at least one clique D such that*

(i) $\mathfrak{A}' \models \Psi$ *and* $\mathfrak{A}'$ *is witness-saturated,*
(ii) $D \cap W_i^{\mathfrak{A}'}(a) \neq \emptyset$*, for every* $a \in F$*,*
(iii) $sp^{\mathfrak{A}'}(D) = X$.

Proof. Let $F = \{a_1, a_2, \ldots, a_p\}$, where a_1 denotes an element of F such that for every $a \in F \setminus \{a_1\}$, $tp^{\mathfrak{A}}(a_1, a) \notin \boldsymbol{\beta}^{\leftarrow}$. (Note that a_1 can always be found since F is finite and T is transitive in $\mathfrak{A}$.) We iteratively apply Claim 10. Denote $\mathfrak{A}^1 = \mathfrak{A}$ and for $k = 2, 3, \ldots, p$ let $\mathfrak{A}^k$ and D^k be the extension of $\mathfrak{A}_1^{k-1}$ by at least one clique D^k given by Claim 10 for a_1 and a_k. Obviously, for every k $(2 \leq k \leq p)$ we have $D^k \cap W_i^{\mathfrak{A}'}(a) \neq \emptyset$, for every $a \in \{a_1, a_2, \ldots, a_k\}$. ◀

4 Canonical models

In this section we analyze properties of models of Ψ on the level of segments which consist of several cliques, and constitute a partition $S_0, S_1, \ldots$ of the universe of a model. Every segment S_j has a fixed (doubly exponential) size and is meant to contain all γ_i-witnesses for elements from earlier segments $S_0, S_1, \ldots, S_{j-1}$. On this level of abstraction cliques and splices of a model become less important.

▶ **Definition 12.** A finite subset $S \subset A$ is a *segment* in $\mathfrak{A}$ if $Cl^{\mathfrak{A}}(a) \subseteq S$ for every $a \in S$.

In the following we reserve the letter S (possibly decorated) to denote segments. Define $s = |Sp(\sigma)|$ and denote by h the bound of the size of each clique in a small-clique σ-structure, given by Lemma 5. Note that s is doubly exponential and h is exponential in $|\sigma|$.

We first prove a generalization of Corollary 11. It says, roughly speaking, that if $\mathfrak{A} \models \Psi$ and F is a finite subset of A, then it is possible to extend $\mathfrak{A}$ by a segment of fixed cardinality in which all elements of F have their γ_i-witnesses, for every i $(1 \leq i \leq m)$.

▶ **Lemma 13** (Witnesses compression)**.** *Assume $\mathfrak{A}$ is a countable witness-saturated model of Ψ and $F \subseteq A \setminus K(\mathfrak{A})$ is finite. There is a witness-saturated extension $\mathfrak{A}'$ of $\mathfrak{A}$ by a segment S such that*

1. $\mathfrak{A}' \models \Psi$,
2. $|S| \leq m \cdot s \cdot h$,
3. *for every* $\gamma_i \in \Psi$, *for every* $a \in F$, *if* $W_i^{\mathfrak{A}}(a) \setminus K(\mathfrak{A}) \neq \emptyset$, *then* $W_i^{\mathfrak{A}'}(a) \cap S \neq \emptyset$.

Proof. First, for given i $(1 \leq i \leq m)$ and $X \in Sp^{\mathfrak{A}}$ let $F_i^X \subseteq S$ be a maximal subset of F such that for every $a \in F_i^X$ there is $b \in W_i^{\mathfrak{A}}(a)$ such that $b \notin K(\mathfrak{A})$ and $sp(Cl^{\mathfrak{A}}(b)) = X$. Now, for every i $(1 \leq i \leq m)$ and for every $X \in Sp^{\mathfrak{A}}$ iteratively apply Corollary 11 for F_i^X and denote each new clique added in the process by D_i^X. Let S be the segment consisting of elements of the new cliques: $S \stackrel{def}{=} \bigcup_{1 \leq i \leq m} \bigcup_{X \in Sp^{\mathfrak{A}}} D_i^X$.

Condition (i) of Lemma 11 implies that $\mathfrak{A}' \models \Psi$. Obviously, $|S| \leq m \cdot s \cdot h$. To show that condition 3 of our lemma holds, assume $\gamma_i \in \Psi$, $a \in F$ and there exists $b \in W_i^{\mathfrak{A}}(a)$ such that $b \notin K(\mathfrak{A})$. So $a \in F_i^X$, where $X = sp(Cl^{\mathfrak{A}}(b))$. By condition (ii) of Lemma 11 we obtain $D_i^X \cap W_i^{\mathfrak{A}'}(a) \neq \emptyset$, and so, $W_i^{\mathfrak{A}'}(a) \cap S \neq \emptyset$. ◀

▶ **Definition 14.** A segment $S \subsetneq A$ is *redundant in* $\mathfrak{A}$, if for every $a \in A \setminus S$ and for every $\gamma_i \in \Psi$ we have: $W_i^{\mathfrak{A}}(a) \cap S \neq \emptyset$ implies there exists $c \in A \setminus S$ such that $c \in W_i^{\mathfrak{A}}(a)$.

▶ **Claim 15.** *If* $\mathfrak{A} \models \Psi$ *and* $S \subsetneq A$ *is redundant in* $\mathfrak{A}$, *then* $\mathfrak{A} \restriction (A \setminus S) \models \Psi$.

Proof. Every subgraph of a transitive graph is also transitive. Conditions (a)–(c) of Proposition 4 obviously hold for $\mathfrak{A} \restriction (A \setminus S)$. ◀

▶ **Definition 16.** A model $\mathfrak{A}$ of Ψ is *narrow* if there is an infinite partition $P_A = \{S_0, S_1, \ldots\}$ of the universe A such that:
1. $K(\mathfrak{A}) \subset S_0$, $|S_0| \leq (m+1) \cdot s \cdot h$,
2. $|S_j| \leq m \cdot s \cdot h$, for every $j \geq 1$,
3. for every $j \geq 0$, for every $e \in \bigcup_{k=0}^{j} S_k$ and for every $\gamma_i \in \Psi$,
   if $W_i^{\mathfrak{A}}(e) \cap S_0 = \emptyset$, then $W_i^{\mathfrak{A}}(e) \cap S_{j+1} \neq \emptyset$.

▶ **Lemma 17.** *Every infinitely satisfiable sentence* Ψ *has a narrow model.*

Proof. Assume $\mathfrak{A}$ is an infinite witness-saturated model of Ψ that exists by Lemma 13. For $\gamma_i \in \Psi$ and $a \in A$ denote by $\bar{\gamma}_i(a)$ an arbitrarily chosen element $b \in W_i^{\mathfrak{A}}(a)$. Define $\mathfrak{A}_0 = \mathfrak{A}$ and $S_0 = K(\mathfrak{A}) \cup \bigcup_{a \in K(\mathfrak{A})} Cl^{\mathfrak{A}}(\bar{\gamma}_i(a))$.
Now, for $j = 0, 1, 2, \ldots$ define:
- $\mathfrak{A}_{j+1} = \mathfrak{A}_j'$, where $\mathfrak{A}_j'$ is the extension of $\mathfrak{A}_j$ given by Lemma 13 for $F = \bigcup_{k=0}^{j} S_k$,
- $S_{j+1} = B$, where B is the finite set given by Lemma 13; B extends $\mathfrak{A}_j$ to $\mathfrak{A}_{j+1}$ in such a way, that:
  - $\mathfrak{A}_{j+1} \models \Psi$,
  - $|B| \leq m \cdot s \cdot h$,
  - for every $\gamma_i \in \Psi$, for every $a \in F$, if $W_i^{\mathfrak{A}}(a) \setminus K(\mathfrak{A}_j) \neq \emptyset$, then $W_i^{\mathfrak{A}_{j+1}}(a) \cap B \neq \emptyset$.

Now, define $\mathfrak{A}' = (\bigcup_{k=0}^{\infty} \mathfrak{A}_k) \restriction \bigcup_{k=0}^{\infty} S_k$. By Claim 15 and Lemma 13, it is easy to see, that $\mathfrak{A}'$ is a narrow model of Ψ with partition $P_{A'} = \{S_0, S_1, \ldots\}$. ◀

▶ **Definition 18.** Assume $\mathfrak{A}$ is a σ-structure, $x, y \in \mathbb{N}^+$ and B, B', C, C' are finite subsets of A with fixed orderings: $B = \{b_1 < \ldots < b_x\}$, $B' = \{b_1' < \ldots < b_y'\}$, $C = \{c_1 < \ldots < c_x\}$, $C' = \{c_1' < \ldots, c_y'\}$ such that $B \cap B' = \emptyset$, $C \cap C' = \emptyset$.

A *connection type* of B and B' in $\mathfrak{A}$ is the structure $\langle B, B' \rangle \stackrel{def}{=} \mathfrak{A} \restriction (B \cup B')$. Two connection types $\langle B, B' \rangle$ and $\langle C, C' \rangle$ are *the same connection types in* $\mathfrak{A}$, denoted $\langle B, B' \rangle \cong^{\mathfrak{A}} \langle C, C' \rangle$, if the function $f : B \cup B' \mapsto C \cup C'$ defined by $f(b_j) = c_j$ and $f(b_j') = c_j'$ is an isomorphism of $\langle B, B' \rangle$ and $\langle C, C' \rangle$.

▶ **Definition 19.** Assume $\mathfrak{A}$ is a narrow model of Ψ and $P_A = \{S_0, S_1, \ldots\}$ is any partition satisfying conditions 1-3 of Definition 16. We say that $\mathfrak{A}$ is *canonical* if for every $j, k \in \mathbb{N}^+$, $0 < j < k$, we have $\langle S_j, S_0 \rangle \cong^{\mathfrak{A}} \langle S_1, S_0 \rangle$, and $\langle S_k, S_j \rangle \cong^{\mathfrak{A}} \langle S_2, S_1 \rangle$.

▶ **Lemma 20.** *Every infinitely satisfiable sentence Ψ has a canonical model.*

Proof. Let $\mathfrak{A}$ be a narrow model of Ψ with partition $P_A = \{S_0, S_1, \ldots\}$ given by Definition 16. Additionally, assume that in every segment S_j, $j \geq 0$, there is a fixed linear ordering.

Observe first that for every $p > 0$, S_p is redundant in $\mathfrak{A}$. For, assume (cf. Definition 14) $b \in S_p$, $a \in A \setminus S_p$ and $b \in W_i^{\mathfrak{A}}(a)$. Assume $a \in S_q$ and take $j \in \mathbb{N}^+$ such that $j > \max\{p, q\}$. By Definition 16, we have that if $W_i^{\mathfrak{A}}(a) \cap S_0 = \emptyset$, then $W_i^{\mathfrak{A}}(a) \cap S_{j+1} \neq \emptyset$. So, there is $c \in S_0 \cup S_{j+1}$ such that $c \in W_i^{\mathfrak{A}}(a)$. Similarly, for every infinite $V \subset \mathbb{N}^+$, the segment $\bigcup_{j \in \mathbb{N}^+ \setminus V} S_j$ is redundant in $\mathfrak{A}$ and, by Claim 15, $\mathfrak{A} \restriction \bigcup_{j \in V \cup \{0\}} S_j \models \Psi$.

To construct the canonical model we first find an infinite set $V \subset \mathbb{N}^+$ such that for every $j,\ l \in V$, $j \neq l$, $\langle S_l, S_0 \rangle \cong^{\mathfrak{A}} \langle S_j, S_0 \rangle$. Observe that the set V does exist (by the infinite Ramsey theorem, there are infinitely many segments in P_A and only a finite number of similarity types).

Secondly, we find an infinite set $W \subseteq V$ such that for every $i, j, k, l \in W$ with $i < j$ and $k < l$ we have $\langle S_j, S_i \rangle \cong^{\mathfrak{A}} \langle S_l, S_k \rangle$. (Again, W exists by the infinite Ramsey theorem).

Finally, we define $\mathfrak{A}' = \mathfrak{A} \restriction \bigcup_{j \in W \cup \{0\}} S_j$. By Claim 15, $\mathfrak{A}' \models \Psi$. Obviously, $\mathfrak{A}'$ is canonical with partition $P_{A'} = \{S_j : j \in \mathbb{N}^+, j = \#p\}$, where $\#p$ is the position number of $p \in W \cup \{0\}$. ◀

5 Decidability and complexity

From Lemma 20 we get immediately the following theorem.

▶ **Theorem 21.** *An FO^2_T-sentence Ψ is satisfiable if and only if there exist a σ-structure $\mathfrak{A}$ and $S_0, S_1, S_2, S_3 \subseteq A$, such that:*
1. $|A| \leq (4m + 1) \cdot s \cdot h$,
2. *either $S_1 = S_2 = S_3 = \emptyset$, or $\{S_0, S_1, S_2, S_3\}$ is a partition of A and then*
   a. $\langle S_1, S_0 \rangle \cong^{\mathfrak{A}} \langle S_2, S_0 \rangle \cong^{\mathfrak{A}} \langle S_3, S_0 \rangle$,
   b. $\langle S_2, S_1 \rangle \cong^{\mathfrak{A}} \langle S_3, S_2 \rangle \cong^{\mathfrak{A}} \langle S_3, S_1 \rangle$,
3. *for every $a, b \in A$, $tp^{\mathfrak{A}}(a, b) \models \psi_0$,*
4. *$T^{\mathfrak{A}}$ is transitive in $\mathfrak{A}$,*
5. *for every $j = 0, 1, 2$, for every $e \in S_j$ and for every $\gamma_i \in \Psi$,*
   if $W_i^{\mathfrak{A}}(e) \cap S_0 = \emptyset$, then $W_i^{\mathfrak{A}}(e) \cap S_{j+1} \neq \emptyset$.

Proof. ($\Rightarrow$) There are two cases. Either Ψ has only finite models and then, by Lemma 9, Ψ has a witness-saturated model $\mathfrak{A}$ with $A = K(\mathfrak{A})$. In this case, we put $S_0 = A$ and $S_1 = S_2 = S_3 = \emptyset$. Or, Ψ has infinite models, and then, by Lemma 20, Ψ has a canonical model $\mathfrak{A}'$ with partition $P_{A'} = \{S_0, S_1, \ldots\}$. In this case, we define $\mathfrak{A} \stackrel{def}{=} \mathfrak{A}' \restriction (S_0 \dot\cup S_1 \dot\cup S_2 \dot\cup S_3)$. Note that in either case $|S_0| \leq s \cdot h + m \cdot s \cdot h = (m + 1) \cdot s \cdot h$.

($\Leftarrow$) Define a structure $\mathfrak{A}'$ such that $A' \stackrel{def}{=} S_0 \ \dot\cup\ S_1 \ \dot\cup\ S_2 \ \dot\cup\ S_3 \ \dot\cup\ \dot\bigcup_{j=4}^{\infty} S_j$ and, for every $j, k \in \mathbb{N}^+$ $(0 < j < k) : \langle S_j, S_0 \rangle \cong^{\mathfrak{A}'} \langle S_1, S_0 \rangle$ and $\langle S_k, S_j \rangle \cong^{\mathfrak{A}'} \langle S_2, S_1 \rangle$. It is obvious that $\mathfrak{A}'$ satisfies conditions (a)–(c) of Proposition 4. To show that T is transitive in $\mathfrak{A}'$ it suffices to prove that for every j, k, l $(0 \leq j \leq k \leq l)$, $T^{\mathfrak{A}' \restriction (S_j \cup S_k \cup S_l)}$ is transitive in $\mathfrak{A}' \restriction (S_j \cup S_k \cup S_l)$. The latter condition can be easily verified; hence, $\mathfrak{A}' \models \Psi$. ◀

▶ **Corollary 22.** *SAT(FO^2_T) $\in$ 2-NExpTime.*

Proof. To check whether a given FO^2_Tsentence Ψ is satisfiable we follow Theorem 21 and we obtain a nondeterministic double exponential time procedure, as described below.

1. `Guess` a σ-structure $\mathfrak{A}$ of cardinality $|A| \leq (4m+1) \cdot s \cdot h$ and partition $P_A = \{S_0, S_1, S_2, S_3\}$;
2. `Guess` enumerations of every segment S_0, S_1, S_2, S_3;
3. `If not`:
   a. $\langle S_1, S_0\rangle \cong^{\mathfrak{A}} \langle S_2, S_0\rangle \cong^{\mathfrak{A}} \langle S_3, S_0\rangle$ and
   b. $\langle S_2, S_1\rangle \cong^{\mathfrak{A}} \langle S_3, S_2\rangle \cong^{\mathfrak{A}} \langle S_3, S_1\rangle$

   `then` **reject**;
4. `For every` $a, b \in A$, `if` $tp^{\mathfrak{A}}(a,b) \not\models \psi_0$ `then` **reject**;
5. `For every` $a, b, c \in A$, `if not` $(T^{\mathfrak{A}}(a,b) \wedge T^{\mathfrak{A}}(b,c) \Rightarrow T^{\mathfrak{A}}(a,c))$ `then` **reject**;
6. `For every` $j = 0, 1, 2$, `for every` $e \in S_j$, `for every` $\gamma_i \in \Psi$ such that $W_i^{\mathfrak{A}}(e) \cap S_0 = \emptyset$ `if` $W_i^{\mathfrak{A}}(e) \cap S_{j+1} = \emptyset$ `then` **reject**;
7. **Accept**; ◀

6 Outlook

Since the finite model property fails for FO^2_T, an interesting question is whether the finite satisfiability problem is also decidable. Immediately from Lemma 17 we have the following observation.

▶ **Corollary 23.** *An FO^2_T-sentence Ψ is satisfiable if and only if Ψ has a model of cardinality $\leq s \cdot h$, or Ψ has an infinite model.*

The $s \cdot h$ bound on the size of the finite model of Ψ depends on the number of different σ-splices and the size of cliques in a structure with the small clique property. Unfortunately, this observation does not suffice to answer the finite satisfiability problem, as in general, one can imagine that a finite model contains several realizations of the same splice. So, to the best of our knowledge, the finite satisfiability problem for FO^2_T remains open. We believe that the latter problem is decidable; however, we suppose that an essential extension of the above approach is required to get the proof.

We also note that the 2-NExpTime bound for the satisfiability problem leaves a gap in complexity, as the best lower bound coming from the two-variable guarded logic with transitive guards is 2-ExpTime [10]. We believe that the upper bound proved in our paper can be improved; however, as FO^2_T does not enjoy the tree-like property, standard techniques using alternating machines cannot be applied directly.

Acknowledgements The authors acknowledge insightful comments from the anonymous referees that helped to improve the presentation and simplify some proofs in Section 4.

References

1 H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *J. Philos. Logic*, 27:217–274, 1998.

2 M. Bojańczyk, C. David, A. Muscholl, T. Schwentick, and L. Segoufin. Two-variable logic on data trees and xml reasoning. In *Proc. of PODS-06*, pages 10–19, New York, NY, USA, 2006. ACM.

3 M. Bojańczyk, C. David, A. Muscholl, T. Schwentick, and L. Segoufin. Two-variable logic on words with data. In *LICS-06*, pages 7–16, 2006.

4 C. David, L. Libkin, and Tony Tan. On the satisfiability of two-variable logic over data words. In Christian G. Fermüller and Andrei Voronkov, editors, *LPAR (Yogyakarta)*, volume 6397 of *Lecture Notes in Computer Science*, pages 248–262. Springer, 2010.

5 E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bull. of Symb. Logic*, 3(1):53–69, 1997.
6 E. Grädel and M. Otto. On Logics with Two Variables. *Theoretical Computer Science*, 224:73–113, 1999.
7 E. Grädel, M. Otto, and E. Rosen. Undecidability results on two-variable logics. *Arch. Math. Log.*, 38(4-5):313–354, 1999.
8 N. Immerman, A. M. Rabinovich, T. W. Reps, S. Sagiv, and G. Yorsh. The boundary between decidability and undecidability for transitive-closure logics. In *CSL*, pages 160–174, 2004.
9 Y. Kazakov. *Saturation-based decision procedures for extensions of the guarded fragment.* PhD thesis, Universität des Saarlandes, Saarbrücken, Germany, 2006.
10 E. Kieroński. The two-variable guarded fragment with transitive guards is 2EXPTIME-Hard. In *Proc. of FOSSACS*, volume LNCS 2620, pages 299–312, 2003.
11 E. Kieroński. Results on the guarded fragment with equivalence or transitive relations. In *Computer Science Logic*, volume 3634, pages 309–324. Springer Verlag, 2005.
12 E. Kieroński. Decidability issues for two-variable logics with several linear orders. In *CSL-11*, pages 337–351, 2011.
13 E. Kieroński and J. Michaliszyn. Two-variable universal logic with transitive closure. In *CSL-12*, LIPIcs, pages 337–351, 2012.
14 E. Kieroński, J. Michaliszyn, I. Pratt-Hartmann, and L. Tendera. Two-variable first-order logic with equivalence closure. In *Proc. of LICS2012*, pages 431–440. IEEE, 2012.
15 E. Kieroński and M. Otto. Small substructures and decidability issues for first-order logic with two variables. In *Proc. of LICS2005*, pages 448–457, 2005.
16 E. Kieroński and M. Otto. Small substructures and decidability issues for first-order logic with two variables. *Journal of Symbolic Logic*, 77 (3):729–765, 2012.
17 E. Kieroński and L. Tendera. On finite satisfiability of two-variable first-order logic with equivalence relations. In *Proc. of LICS2009*, pages 123–132, 2009.
18 A. Manuel. Two variables and two successors. In Petr Hlinený and Antonín Kucera, editors, *MFCS*, volume 6281 of *Lecture Notes in Computer Science*, pages 513–524. Springer, 2010.
19 J. Michaliszyn. Decidability of the guarded fragment with the transitive closure. In *Proc. of ICALP2009*, pages 261–272, 2009.
20 M. Mortimer. On languages with two variables. *Zeitschr. f. Logik und Grundlagen d. Math.*, 21:135–140, 1975.
21 M. Niewerth and T. Schwentick. Two-variable logic and key constraints on data words. In Tova Milo, editor, *ICDT*, pages 138–149. ACM, 2011.
22 M. Otto. Two-variable first-order logic over ordered domains. *Journal of Symbolic Logic*, 66:685–702, 2001.
23 F. Ramsey. On a problem of formal logic. *Proc. London Math. Soc. series 2*, 30:264–286, 1930.
24 T. Schwentick and T. Zeume. Two-variable logic with two order relations - (extended abstract). In *CSL*, pages 499–513, 2010.
25 D. Scott. A decision method for validity of sentences in two variables. *J. Symb. Logic*, 27:477, 1962.
26 W. Szwast and L. Tendera. On the decision problem for the guarded fragment with transitivity. In *Proc. of LICS2001*, pages 147–156, 2001.
27 W. Szwast and L. Tendera. The guarded fragment with transitive guards. *Annals of Pure and Applied Logic*, 128:227–276, 2004.