An excursion to the border of decidability: between two- and three-variable logic*

Oskar Fiuk and Emanuel Kieroński

University of Wrocław, Poland

Abstract

With respect to the number of variables the border of decidability lies between 2 and 3: the two-variable fragment of first-order logic, FO^2, has an exponential model property and hence NExpTime-complete satisfiability problem, while for the three-variable fragment, FO^3, satisfiability is undecidable. In this paper we propose a rich subfragment of FO^3, containing full FO^2 (without equality), and show that it retains the finite model property and NExpTime complexity. Our fragment is obtained as an extension of the uniform one-dimensional variant of FO^3.

1 Introduction

In the realm of the fragments of first-order logic, FO, with decidable satisfiability problem the two-variable fragment, FO^2, is one of the most prominent. Its importance can be justified by at least two facts: (i) it is the maximal decidable fragment of FO with respect to the number of variables, and (ii) it embeds many modal and description logics (via the so-called standard translation), and hence it constitutes an elegant first-order framework for analysing those formalisms.

The first decidability proof for FO^2, working only in the absence of equality, was given by Scott [12] via a reduction to Gödel prefix class without equality, which in turn was shown decidable in the classical paper by Gödel [1]. The decidability of FO^2 with equality is due to Mortimer [11] who showed that FO^2 has the finite model property (fmp), more precisely, that every satisfiable FO^2 sentence has a model of size at most doubly exponential in its length. This bound on the size of small models was later improved to single exponential by Grädel, Kolaitis and Vardi in [3], which led to their result that FO^2 is decidable in NExpTime, and in fact NExpTime-complete, with the lower bound coming from the earlier work by Lewis [10].

The contrasting undecidability of FO^3 was first shown by Kahr, Moore, Wang [6]. Concretely, they showed that the prefix class $\forall\exists\forall$, that is the class of formulas $\forall x \exists y \forall z \phi(x, y, z)$, with quantifier-free ϕ, without equality, is undecidable. Worth mentioning here are two other sources of undecidability of FO^3. The first is that in FO^3 it is possible to specify that some binary relations are transitive (by the obvious prefix formulas starting with the pattern $\forall\forall\forall$). Indeed, already FO^2 (and even its guarded subfragment) with two transitive relations is undecidable (see Kazakov's Phd thesis [7] or [8]). The second is allowing for the use of equality in prefix formulas $\forall\forall\exists$, that is in a three-variable subfragment of the already mentioned Gödel class, but this time with equality, known to be undecidable from the paper by Goldfard [2].

One of the aims of our work is to better understand the reasons behind the undecidability of FO^3. This aim is obtained by proposing its rich decidable subfragment FO^3_-, fully containing FO^2 and not subsumed by any known decidable fragment. On the one hand, FO^3_- is a fragment showing what we can afford for without losing the decidability, on the other hand, we believe

*Supported by National Science Centre, Poland, grant 2021/41/B/ST6/00996.

that it may also serve as an interesting decidable specification language, orthogonal to the existing ones, allowing one to express many new natural properties. Roughly speaking, FO^3_- will be defined by limiting (but not forbidding!) the use of the two "dangerous" quantifier patterns $\forall\exists\forall$ and $\forall\forall\forall$ identified in the previous paragraph. We will concentrate here on the case without equality (only briefly discussing the case with equality in the concluding section), so the pattern $\forall\forall\exists$ does not lead directly to undecidability and hence we do not consider it as "dangerous".

A crucial role in our investigations will be played by the idea of the uniform one-dimensional fragment, UF_1, of Hella and Kuusisto [5]. In this fragment two restrictions are imposed on the formulas: (i) *one-dimensionality*: quantifiers are used in blocks, a single block is built out only of existential or only of universal quantifiers and leaves at most one variable free and (ii) *uniformity*: roughly speaking, Boolean combinations of atoms are allowed only if the atoms use precisely the same set of variables (more crudely: all the free variables of the subformula) or use just one variable. UF_1 was introduced as a "canonical" extension of FO^2 to contexts with relations of arity greater than two, like database theory. Indeed, UF_1 turned out to have the finite (exponential) model property and NExpTime-complete satisfiability problem, exactly as FO^2.

An interesting open question concerning UF_1 is if its extension AUF_1, in which the existential and universal quantifiers can be used in blocks in the alternating fashion, remains decidable. We start this paper by showing the decidability of its three-variable version AUF^3_1. AUF_1 and AUF^3_1 can express some interesting properties, which are typically not expressible in known decidable logics, for example, assuming that $PlayedIn$ is a ternary relation containing tuples (a, m, d), where a is the name of an actor, who played in a movie m directed by a director d, we can write that every director has an actor who played in its all famous movies: $\forall x\exists y\forall z(Director(x) \wedge FamousMovie(z) \rightarrow Actor(y) \wedge PlayedIn(y, z, x))$.

Our decidability proof for AUF^3_1 goes by a reduction of arbitrary formulas to their normal forms (conjunctions of formulas $QxQyQz\phi(x, y, z)$ with quantifier-free ϕ), and then by a two-level construction in which we build an exponentially bounded model for a given satisfiable formula. This implies NExpTime-completeness of the satisfiability and finite satisfiability problems. [1]

Next, we consider relaxing the uniformity condition in the normal form formulas. Namely, we leave the uniformity restriction only in those normal form conjuncts which start with the "dangerous" patterns $\forall\exists\forall$ or $\forall\forall\forall$. This will lead us to a definition of a rich subfragment FO^3_- of FO^3, in which in addition to what is expressible in AUF^3_1 we can specify some further properties, e.g., that every married couple has a child: $\forall x\forall y(Married(x, y) \rightarrow \exists z Child(x, z) \wedge Child(y, z))$. For a precise definition of this fragment see Section 2.2. Clearly FO^3_- is more expressive than FO^2 even over signatures built out only of unary and binary relation symbols, as demonstrated by the above formula (we remark that over such signatures AUF^3_1 syntactically collapses to FO^2). We will show the decidability and NExpTime-completeness of the satisfiability problem for FO^3_-. We also show that it possesses the finite, exponential model property. This time however, we will not construct models fully explicitly, as in the case of AUF^3_1, but, after preparing some building blocks, we will adapt the beautiful probabilistic approach used by Gurevich and Shelah in the case of Gödel class without equality [4] . This approach allows us to demonstrate fmp quite easily, while it seems to us that any explicit model construction, while possible, would need to be distractingly complicated.

[1] In the upcoming paper we show decidability of another fragment of AUF_1, orthogonal to AUF^3_1, in which the number of variables is not restricted, but the blocks of quantifiers (in negation normal form formulas) are required to be purely universal or to end with the existential quantifier.

The organization of the paper is as follows. In Section 2 we introduce our notation and terminology and define the relevant classes of formulas. In Section 3 we work with AUF^3_1 and in Section 4 with its extension FO^3_-. The latter has two main technical parts: in Subsection 4.1 we establish the complexity of FO^3_- by showing that for a given ϕ some kind of exponentially bounded structures may serve as *witnesses for satisfiability* of satisfiable formulas and can be unwound into their (infinite) models, and in Subsection 4.2 we prove the finite (exponential) model property for FO^3_- via the probabilistic method. We conclude the paper in Section 5.

With such an organization, some parts of the paper may seem redundant. Indeed, FO^3_- subsumes AUF^3_1, and hence our NExpTime-upper bound for FO^3_- implies the same bound for AUF^3_1, and similarly, the exponential model property for FO^3_- implies the exponential model property for AUF^3_1. We begin with analysing AUF^3_1 mostly to help the reader understand our ideas in a simpler settings (not requiring us to control too many factors at once) first. Moreover, this allows us to obtain some modularization of the paper, as the constructions for FO^3_- employ some crucial part of the construction for AUF^3_1 as a black box. Also, the probabilistic proof of the exponential model property could be used to infer a NExpTime-upper bound on the complexity of FO^3_-. We decided to keep Subsection 4.1 with an alternative approach for this bound, as the two approaches are quite different from each other, and together give a better insight into the reasons for decidability of FO^3_-.

2 Preliminaries

2.1 Notation, structures and pre-structures, types and pre-types

We assume that the reader is familiar with first-order logic, FO. We work with purely relational signatures with no constants and function symbols. We refer to structures using Fraktur capital letters, and to their domains using the corresponding Roman capitals. Given a structure $\mathfrak{A}$ and some $B \subseteq A$ we denote by $\mathfrak{A}{\restriction}B$ the structure being the restriction of $\mathfrak{A}$ to its subdomain B.

We usually use $a, b, \ldots$ to denote elements from the domains of structures, and x, y, $\ldots$ for variables, all of these possibly with some decorations. For a tuple of variables $\bar{x}$ we use $\psi(\bar{x})$ to denote that the free variables of ψ are in $\bar{x}$.

In the context of uniform logics it is convenient to speak about some partially defined (sub)structures which we will call *pre-(sub)structures*. A pre-structure over a signature σ consists of its domain A and a function specifying the truth-value of every fact $P(\bar{a})$, for $P \in \sigma$ and a tuple $\bar{a}$ of elements of A of length equal to the arity of P, such that $\bar{a}$ contains all elements of A or just one of them. The truth values of all the other facts remain unspecified. We will use Fraktur letters decorated with $*$ to denote pre-structures: a pre-structure with domain A will be denoted by $\mathfrak{A}^*$. If a structure $\mathfrak{A}$ is fully defined we denote $\mathfrak{A}^*$ its induced pre-structure. If $B \subseteq A$ is a subdomain of $\mathfrak{A}$ we denote $(\mathfrak{A}{\restriction}B)^*$ the induced pre-structure of the restriction $\mathfrak{A}{\restriction}B$.

For $k \geq 1$, an (atomic) *k-type* over a signature σ is a maximal consistent set of atomic or negated atomic formulas over σ using at most variables $x_1, \ldots, x_k$, containing additionally $x_i \neq x_j$, for all $i \neq j$. A *k-pre-type* over a signature σ is a maximal consistent set of those atomic or negated atomic formulas over σ that use precisely all variables $x_1, \ldots, x_k$ or just one of them, containing additionally $x_i \neq x_j$, for all $i \neq j$. We often identify a k-type or a k-pre-type with the conjunction of its elements. Given a k-type γ we denote $\gamma{\restriction}x_i$ the *restriction of γ to x_i*, that is the 1-type obtained from γ by first removing all literals containing a variable different from x_i and then renaming the remaining occurrences of x_i to x_1.

Let $\mathfrak{A}$ be a structure, and let $a_1, \ldots, a_k \in A$ be pairwise distinct elements. We denote by

$\text{tp}^{\mathfrak{A}}(a_1, \ldots, a_k)$ (resp. $\text{pretp}^{\mathfrak{A}}(a_1, \ldots, a_k)$, the unique atomic k-type (resp. k-pre-type) *realized* in $\mathfrak{A}$ by the tuple $a_1, \ldots, a_k$, *i.e.*, the k-type (resp. k-pre-type) $\alpha(x_1, \ldots, x_k)$ such that $\mathfrak{A} \models \alpha(a_1, \ldots, a_k)$.

In this paper we will be interested in k-types and k-pre-types for $k = 1, 2, 3$. Note that the notions of k-types and k-pre-types coincide for $k = 1, 2$ and that to fully specify a structure over a signature containing relation symbols of arity at most 3 it suffices to define its domain, 1-types of all elements, 2-types of all pairs of distinct elements, and 3-pre-types of all tuples of pairwise distinct elements of length 3 (actually, information about the 1-types is stored also in 2-types and 3-pre-types, but we will usually start our model constructions with defining the 1-types of its domain elements).

In our proofs we will need a bound on the number of all possible k-types over a given signature. It is stated in the following Lemma, whose proof is obvious.

Lemma 1. *Let σ be a purely relational finite signature with no constants and function symbols and containing relation symbols of arity at most 3. Then the number of all possible k-types for $k = 1, 2, 3$ is an exponential function of $|\sigma|$.*

The pull operation. When defining logical structures we will use the following *pull operation*. Let $\mathfrak{C}$ be a partially defined structure over a domain C, let $\mathfrak{A}$ be a fully defined structure, $c_1, \ldots, c_k$ be a sequence of **pairwise distinct** elements from C, and $a_1, \ldots, a_k$ a sequence of **not necessarily pairwise distinct** elements of A. We will write $\mathfrak{C}\restriction\{c_1, \ldots, c_k\} \leftarrow pull(\mathfrak{A}, c_1 \leftarrow a_1, \ldots, c_k \leftarrow a_k)$, to specify that $\mathfrak{C}\restriction\{c_1, \ldots, c_k\}$ is copied from $\mathfrak{A}\restriction\{a_1, \ldots, a_k\}$. Formally, for every relation symbol R and a sequence of indices $i_1, \ldots, i_l$, where l is the arity of R we set $\mathfrak{C} \models R(c_{i_1}, \ldots, c_{i_l})$ iff $\mathfrak{A} \models R(a_{i_1}, \ldots, a_{i_l})$. We also introduce a version of the pull operation for specifying pre-substructures. Under the assumptions as above we write $(\mathfrak{C}\restriction\{c_1, \ldots, c_k\})^* \leftarrow pull^*(\mathfrak{A}, c_1 \leftarrow a_1, \ldots, c_k \leftarrow a_k)$ to mean that only the pre-substructure on $\{c_1, \ldots, c_k\}$ is defined. We will always use the pull operation in such contexts that no conflicts with previously defined parts of $\mathfrak{C}$ will arise; in particular the 1-types of the c_i will be always defined before a use of pull in such a way that for all i, $\text{tp}^{\mathfrak{C}}(c_i) = \text{tp}^{\mathfrak{A}}(a_i)$. The following fact easily follows from the definition.

Lemma 2. *Let $\mathfrak{A}$, $\mathfrak{C}$ be structures, $a_1, \ldots, a_k$ a sequence of elements of A and $c_1, \ldots, c_k$ a sequence of distinct elements of C. If the substructure $\mathfrak{C} \restriction \{c_1, \ldots, c_k\}$ was defined by $\mathfrak{C} \restriction \{c_1, \ldots, c_k\} \leftarrow pull(\mathfrak{A}, c_1 \leftarrow a_1, \ldots, c_k \leftarrow a_k)$ (resp. $(\mathfrak{C} \restriction \{c_1, \ldots, c_k\})^* \leftarrow pull^*(\mathfrak{A}, c_1 \leftarrow a_1, \ldots, c_k \leftarrow a_k)$) then for any first-order quantifier free formula without equality $\psi(x_1, \ldots, x_l)$ (resp. build out of atoms which use all the variables $x_1, \ldots, x_l$ or just one of them) and any sequence $i_1, \ldots, i_l$ of elements from $\{1, \ldots, k\}$ we have $\mathfrak{A} \models \psi(a_{i_1}, \ldots, a_{i_l})$ iff $\mathfrak{C} \models \psi(c_{i_1}, \ldots, c_{i_l})$.*

2.2 Fragments of logic

For a natural number $k > 0$, we denote by FO^k the fragment of FO consisting of those formulas that use at most k variables. As typical when working with bounded variable logics we assume that signatures for FO^k contain relation symbols of arity at most k. Given a class $\mathcal{L}$ of formulas in FO we denote $\mathcal{L}^k := \text{FO}^k \cap \mathcal{L}$. If not otherwise stated we assume that the considered fragments do not admit equality.

sUF_1. We now define the logic, sUF_1 (without equality), called in [9] the *strongly restricted uniform-one dimensional fragment*. Formally, for a relational signature σ, the set of σ-formulas of sUF_1 is the smallest set $\mathcal{F}$ such that:

(i) every σ-atom using at most one variable is in $\mathcal{F}$,

(ii) $\mathcal{F}$ is closed under Boolean connectives,

(iii) if $\phi(x_0, \ldots, x_k)$ is a Boolean combination of formulas in $\mathcal{F}$ with free variables in $\{x_0, \ldots, x_k\}$ and atoms built out of precisely all of the variables $x_0, \ldots, x_k$ (in an arbitrary order, possibly with repetitions) then $\exists x_0, \ldots, x_k \phi$ and $\exists x_1, \ldots, x_k \phi$ are in $\mathcal{F}$.

We will be mostly interested in the three-variable version of this logic, sUF_1^3. Example formulas in sUF_1^3 are:

$$\forall xyz(P(x) \wedge P(y) \wedge P(z) \rightarrow R(x,y,z) \vee \neg S(z,x,y)),$$

$$\forall x(P(x) \rightarrow \exists yz(\neg R(y,z,x) \wedge (\neg R(x,y,z) \vee P(y)))).$$

For interested readers we explain that the original *uniform one-dimensional fragment*, UF_1, is defined as above but in point (iii) of the definition the non-unary atoms must not necessarily use the whole set $\{x_0, \ldots, x_k\}$ of variables but rather all those atoms use the same subset of this set (see [9] for some discussion on the variations of UF_1).

$\mathbf{AUF_1}$. By AUF_1 we denote a new logic being the extension of sUF_1 with *alternation of quantifiers in blocks*. It is defined almost exactly as sUF_1. The only difference is that in point (iii) of the definition instead of $\exists x_0, \ldots, x_k \phi$ and $\exists x_1, \ldots, x_k \phi$ we write $\mathrm{Q}_1 x_1 \ldots \mathrm{Q}_k x_k \phi$, where the Q_i represent arbitrary quantifiers; this way we allow for arbitrary patterns of quantifiers in blocks. We will be mostly interested in the three-variable version of this logic, AUF_1^3. Example formulas are:

$$\forall x \exists y \forall z(P(x) \wedge \neg P(y) \rightarrow \neg R(x,y,z) \vee \neg P(z) \vee \neg R(y,z,x)),$$

$$\forall x(P(x) \rightarrow \exists y \forall z \neg R(x,y,z) \wedge \forall y S(y,x) \wedge \forall x \forall y R(z,y,x)).$$

$\mathbf{FO^3_-}$. Now we introduce the subfragment FO^3_- of FO^3, containing full FO^2 (without equality) and reaching into the area of FO^3 even further than AUF_1^3. Its definition may seem a bit unnatural, but our motivations should become clearer in Section 4. To avoid some distracting circumlocutions we allow to use negation only in literals. Formally, for a relational signature σ, the set of σ-formulas of FO^3_- is the smallest set $\mathcal{F}$ of formulas over variables x, y, z, such that:

(i) Every σ-literal using at most one variable is in $\mathcal{F}$.

(ii) $\mathcal{F}$ is closed under $\vee$ and $\wedge$.

(iii) If ϕ is a positive Boolean combination of formulas in $\mathcal{F}$ and σ-literals then $\exists v \phi$ is in $\mathcal{F}$, for any variable v.

(iv) Let v, v' be distinct variables. If ϕ is a positive Boolean combination of formulas in $\mathcal{F}$ with free variables in the set $\{v, v'\}$, and literals using at most variables v and v' then $\forall v \phi$ is in $\mathcal{F}$.

(v) If ϕ is a positive Boolean combination of formulas in $\mathcal{F}$ with at most one free variable and literals using precisely all of x, y, z then $\exists v \forall v' \phi$ and $\forall v \forall v' \phi$ are in $\mathcal{F}$, for any variables v, v'.

Note that in the above definition v, v' always represent some variables from $\{x, y, z\}$. The main restriction on FO^3 formulas we impose is limiting the use of quantifier patterns ending with $\forall$ followed by a three-variable subformula in (v). Existential quantification in (iii) and universal quantification for subformulas with two free variables in (iv) can be used quite freely. Unfortunately, this non-symmetric treatment of universal and existential quantifiers causes that

FO^3_- is not closed under negation. This is in contrast to, e.g., FO^2, UF_1, AUF^3_1 but similar to some other known decidable classes of first-order formulas, like the prefix classes or the unary negation fragment, UNFO, [13].

It is readily verified that (the class of negation normal form of formulas in) FO^2 and AUF^3_1 are indeed fragments of FO^3_-. For example, to show that $\forall x \forall y \neg R(x, y)$ is in FO^3_-, we first use rule (iv) (with $v = y$ and $v' = x$) to generate $\forall y \neg R(x, y)$ and then again use rule (iv) (with $v = x$ and $v' = y$). To show that $\forall x \exists y \forall z P(x, y, z)$ is in FO^3_- we first generate $\exists y \forall z P(x, y, z)$ using rule (v) (with $v = y$, $v' = z$) and then use rule (iv) (with $v = x$, and arbitrary v'). A more complicated example of a formula in FO^3_- is given below:

$$\forall x \forall y(\neg R(x, y) \vee \exists z(S(x, z) \wedge S(y, z) \wedge (\neg T(x, z, y) \vee \forall x S(x, z) \vee \exists x \forall y(T(z, y, x) \wedge P(y))))).$$

2.3 The satisfiability problems and the finite model property

Given a class of formulas (a fragment) of first-order logic $\mathcal{L}$, the (finite) satisfiability problem for $\mathcal{L}$ is defined as follows: given a sentence ϕ from $\mathcal{L}$ verify if ϕ has a (finite) model. We say that a fragment $\mathcal{L}$ has the *finite model property* (fmp) if every satisfiable sentence has a finite model. For a fragment with fmp the satisfiability and finite satisfiability problems coincide. Further, we say that $\mathcal{L}$ has the *exponential model property* if there is a fixed exponential function f, that is a function of the form $f(x) = 2^{p(x)}$ for some polynomial p, such that every satisfiable sentence φ in $\mathcal{L}$ has a finite model over a domain whose size if bounded by $f(|\varphi|)$ (where $|\varphi|$ is the length of, measured in any reasonable fashion). When speaking about the *size of a structure* we mean the size of its domain. We note that under our assumption that signatures contain relation symbols of arity at most 3, the size of a full description of a structure (a list of all tuples in all relations) is polynomial in the size of its domain.

3 Decidability of AUF^3_1

The aim of this section is to show the following result for AUF^3_1.

Theorem 3. AUF^3_1 *(without equality) has the exponentially model property. Hence the satisfiability problem (=finite satisfiability problem) for* AUF^3_1 *is NExpTime-complete.*

Clearly, the upper bound in the second part of this theorem above is implied by the first part: to verify that a given AUF^3_1 sentence is satisfiable we just guess a description of its appropriately bounded model and verify it (in a straightforward manner). We recall that the corresponding lower bound is retained from FO^2 without equality [10].

We begin with introducing a normal form, bounding the quantifier-depth of formulas. We say that an AUF^3_1 sentence ϕ is in normal form if $\phi = \phi_1 \wedge \phi_2$ for

$$\phi_1 = \bigwedge_{i \in \mathcal{I}_1} \forall x \forall y \phi_i(x, y) \wedge \bigwedge_{i \in \mathcal{I}_2} \forall x \forall y \forall z \phi_i(x, y, z) \wedge \bigwedge_{i \in \mathcal{I}_3} \forall x \exists y \forall z \phi_i(x, y, z), \text{ and}$$

$$\phi_2 = \bigwedge_{i \in \mathcal{I}_4} \forall x \exists y \phi_i(x, y) \wedge \bigwedge_{i \in \mathcal{I}_5} \forall x \forall y \exists z \phi_i(x, y, z) \wedge \bigwedge_{i \in \mathcal{I}_6} \forall x \exists y \exists z \phi_i(x, y, z),$$

where the $\mathcal{I}_i$ are pairwise disjoint sets of indices, and all the ϕ_i are *uniform* Boolean combinations of atoms over the appropriate sets of variables ($\{x, y, z\}$ or $\{x, y\}$), that is of atoms containing either all of the free variables of ϕ_i or just one of them. Conjuncts indexed by $\mathcal{I}_i$, $1 \leq i \leq 6$, will be sometimes called the $\mathcal{I}_i$-conjuncts. The following lemma shows that for our purposes we can restrict attention to normal form sentences.

Lemma 4. *(i) The satisfiability problem for* AUF_1^3 *can be reduced in nondeterministic polynomial time to the satisfiability problem for normal form* AUF_1^3 *sentences. (ii) If the class of all normal form* AUF_1^3 *sentences has the finite (exponential) model property then also the whole* AUF_1^3 *has the finite (exponential) model property.*

Proof. (Sketch). Take an AUF_1^3 sentence ψ, assuming wlog. that it is in negation normal form. Take its innermost subformula ψ_0 starting with a maximal block of quantifiers. If it has a free variable, that is, is of the form, e.g., $\mathrm{Q}_1 x \mathrm{Q}_2 y \eta(x, y, z)$ replace it by $P(z)$, for a fresh unary symbol P, and add the following normal form conjunct ϕ_{ψ_0} (partially) axiomatising P: $\forall z \mathrm{Q}_1 x \mathrm{Q}_2 y (P(y) \to \eta(x, y, z))$. In other words, ϕ is replaced by $\phi(P(z)/\psi_0) \wedge \phi_{\psi_0}$. We proceed similarly with other possible cases of ψ_0 with a free variable–when its quantifier block has length 1 or quantifies different variables. If ψ_0 is a proper subsentence, that is, it is of the form, e.g., $\mathrm{Q}_1 x \mathrm{Q}_2 y \mathrm{Q}_3 z \eta(x, y, z)$, then we replace it by E, for a fresh 0-ary symbol E and add the conjunct $E \to \mathrm{Q}_1 x \mathrm{Q}_2 y \mathrm{Q}_3 z \eta(x, y, z)$.

Repeat this process as long as possible, obtaining eventually a formula of the shape $\eta_0 \wedge \bigwedge_i \bar{Q}\eta_i \wedge \bigwedge_i E_i \to \bar{Q}\eta_i$, where η_0 is a (positive) Boolean combination of 0-are predicates. The described process is similar to the one used, e.g., in the case of FO^2 [3]. It should be clear that the obtained formula is satisfiable over the same domains as the initial formula ψ.

We further nondeterministically eliminate the 0-are predicates by guessing their truth-values so that η_0 is true, obtaining a formula of the shape $\bigwedge_i \bar{Q}\eta_i$. Finally, the conjuncts starting with the existential quantifier, like $\exists x \mathrm{Q} y \mathrm{Q} z \eta_i$ are replaced by $\forall x \exists y P(y) \wedge \forall x \mathrm{Q} y \mathrm{Q} z (P(x) \to \eta_i)$, for a fresh symbol P. □

For the rest of this section we fix an AUF_1^3 normal form sentence $\phi = \phi_1 \wedge \phi_2$ without equality and its model $\mathfrak{A}$. We show how to construct a small model of ϕ. Compared to small model constructions for other fragments of first-order logic, one non-typical problem we need to deal with here is respecting $\mathcal{I}_3$-conjuncts, that is conjuncts starting with the quantifier patter $\forall\exists\forall$. To help the reader understand what happens, we split our construction into two parts. We first build (Subsection 3.1) a small model $\mathfrak{B} \models \phi_1$; here we indeed need to concentrate mostly on the $\mathcal{I}_3$-conjuncts. Then (Subsection 3.2) we will use $\mathfrak{B}$ as a building block in the construction of a finite, exponentially bounded, model for the whole ϕ. By Lemma 4, this suffices to prove Theorem 3.

3.1 Dealing with $\forall\exists\forall$-conjuncts.

Here, we give a construction of a model $\mathfrak{B}$ satisfying all the $(\mathcal{I}_1 \cup \mathcal{I}_2 \cup \mathcal{I}_3)$-conjuncts (i.e. $\mathfrak{B} \models \phi_1$). To properly handle the $\mathcal{I}_3$-conjuncts, we pick for every $j \in \mathcal{I}_3$ a witness function $\mathfrak{f}_j : A \to A$ such that for every $a \in A$ we have $\mathfrak{A} \models \forall z \phi_j(a, \mathfrak{f}_j(a), z)$. An element $\mathfrak{f}_j(a)$ will be sometimes called a *witness* for a and ϕ_j (or for a and the appropriate $\mathcal{I}_3$-conjunct). Similar terminology will be also used later for conjuncts from ϕ_2. Here and in the rest of this paper we will use the convention that when referring to the indexes of $\mathcal{I}_3$-conjuncts, we will use the letter j, while for the other groups of conjuncts (or in situations where it is not relevant which group of conjuncts is considered) we will usually use the letter i.

Let us introduce the following definition of a *generalized type* of an element $a \in A$. Formally, denoting $\boldsymbol{\alpha}$ the set of 1-types realized in $\mathfrak{A}$, a generalized type is the following tuple:

$$\tau = (\alpha(\tau), \beta(\tau), j(\tau), s(\tau)) \in \boldsymbol{\alpha} \times (\boldsymbol{\alpha} \cup \{\bot\}) \times (\{1, \dots |\mathcal{I}_3|\} \cup \{\bot\}) \times \boldsymbol{\alpha}^{|\mathcal{I}_3|}.$$

We say that an element $a \in A$ realizes a generalized type τ if:

- $\alpha(\tau)$ is the 1-type of a;
- if $\beta(\tau) \neq \bot$ then there exists an element $b \in A$ such that $\beta(\tau)$ is the 1-type of b and a is the witness for b and the $\mathcal{I}_3$-conjunct with index $j(\tau)$, i.e. $f_{j(\tau)}(b) = a$; and if $\beta(\tau) = \bot$ then we do not impose any requirements, except that $j(\tau) = \bot$;
- $s(\tau) = (s(\tau)_1, \ldots, s(\tau)_{|\mathcal{I}_3|})$ is the sequence of 1-types of witnesses for a for all the $\mathcal{I}_3$-conjuncts, i.e., $s(a)_j = \text{tp}^{\mathfrak{A}}(f_j(a))$ for all $1 \leq j \leq |\mathcal{I}_3|$.

Intuitively, if a realizes τ, then τ captures the information about the 1-types of a itself and all the witnesses of a, from the perspective of an element b witnessed by a (if such an element b exists). We want to stress that one element might realize potentially *many* generalized types, since the choices of $\beta(\tau)$ and $j(\tau)$ might not be unique. The set of generalized types realized by a in $\mathfrak{A}$ is denoted $\text{gtp}(a)$.

Let $\text{gtp}(\mathfrak{A}) := \bigcup_{a \in A} \text{gtp}(a)$ be the set of all generalized types realized in $\mathfrak{A}$. For each $\tau \in \text{gtp}(\mathfrak{A})$ pick a single representative $\mathfrak{pat}_{\text{gtp}}(\tau) = a_\tau$ such that $\tau \in \text{gtp}(a)$ (we do not require $\mathfrak{pat}_{\text{gtp}}$ to be 1-1).

Stage 0: The domain. Let the domain of $\mathfrak{B}$ be the set

$$B = \text{gtp}(\mathfrak{A}) \times \{0, 1, 2, 3\}.$$

We will sometimes call the sets $B_m = \{(\tau, m) : \tau \in \text{gtp}(\mathfrak{A})\}$ *layers*. We can naturally lift the function $\mathfrak{pat}_{\text{gtp}}$ to the function $\mathfrak{pat} : B \to A$: $\mathfrak{pat}((\tau, m)) = \mathfrak{pat}_{\text{gtp}}(\tau)$.

Before proceeding to the construction, we need one more definition, namely the definition of *witnessing graph* which will contain a *declaration* about who is a witness for whom for $\mathcal{I}_3$-conjuncts in the final structure $\mathfrak{B}$. More precisely, a *pre-witnessing graph* is a directed graph $\mathcal{G}'_w = (B, E')$ and an edge-labelling function $\ell : E' \to \{1, \ldots, |\mathcal{I}_3|\}$ such that $b_1 = (\tau_1, m_1)$ is connected to $b_2 = (\tau_2, m_2)$ by an arc e with label $\ell(e) = j(\tau_2)$ if $m_1 + 1 = m_2 (\text{mod } 4)$, $\beta(\tau_2) = \alpha(\tau_1)$, and $s(\tau_1)_{j(\tau_2)} = \alpha(\tau_2)$. As the reader may notice the pre-witnessing graph connects only elements which are in consecutive layers (modulo 4), for every vertex there is at least one outgoing arc in every color, and there are no multiple arcs between the same pair of elements. Finally, we "functionalize" the graph by restricting the arc-set E' to contain exactly one outgoing arc in each color for every vertex. Let this new arc-set be E. Thus, we obtain a *witnessing graph* $\mathcal{G}_w = (B, E)$ and write $b_1 \rightsquigarrow_j b_2$ if b_1 is connected to b_2 with an arc colored by j.

The construction itself will be given by the following iterative algorithm. The role of the generalized types and the witnessing graph is to guide the algorithm in selecting patterns from $\mathfrak{A}$ to be pulled into $\mathfrak{B}$, avoiding potential conflicts between different conjuncts (especially) of type $\mathcal{I}_3$. For the clarity of the presentation, all the conditions present in the algorithm need to be taken modulo permutation of the considered elements in each step, and by $\mathfrak{f}_j^{-1}(a)$ we mean to take any $b \in A$ satisfying $\mathfrak{f}_j(b) = a$. The algorithm:

Stage 1: 1-types. We set the 1-types of elements b in $\mathfrak{B}$, as given by $\alpha(b)$.

Stage 2: 2-types. For all pairs $\{b_1 = (\tau_1, m_1), b_2 = (\tau_2, m_2)\}$ of elements from B, $b_1 \neq b_2$ do:

- if $\text{tp}^{\mathfrak{B}}(b_1, b_2)$ is already defined then continue.
- if $b_1 \rightsquigarrow_j b_2$ for some j then set $(\mathfrak{B}\restriction\{b_1, b_2\}) \leftarrow \textit{pull}(\mathfrak{A}, b_1 \leftarrow a_{\tau_1}, b_2 \leftarrow \mathfrak{f}_j(a_{\tau_1}))$.
- otherwise, set $(\mathfrak{B}\restriction\{b_1, b_2\}) \leftarrow \textit{pull}(\mathfrak{A}, b_1 \leftarrow a_{\tau_1}, b_2 \leftarrow a_{\tau_2})$.

Stage 3: 3-pre-types. For all triples $\{b_1 = (\tau_1, m_1), b_2 = (\tau_2, m_2), b_3 = (\tau_3, m_3)\}$ of pairwise distinct elements from B do:

- if $\text{pretp}^{\mathfrak{B}}(b_1, b_2, b_3)$ is already defined then continue.
- if $b_1 \rightsquigarrow_j b_2$ and $b_2 \rightsquigarrow_k b_3$, for some j, k then set

$$(\mathfrak{B}\restriction\{b_1, b_2, b_3\})^* \leftarrow \mathit{pull}^*(\mathfrak{A}, b_1 \leftarrow \mathfrak{f}_j^{-1}(a_{\tau_2}), b_2 \leftarrow a_{\tau_2}, b_3 \leftarrow \mathfrak{f}_k(a_{\tau_2})).$$

- if $b_1 \rightsquigarrow_j b_2$ and $b_1 \rightsquigarrow_k b_3$ for some j, k then set

$$(\mathfrak{B}\restriction\{b_1, b_2, b_3\})^* \leftarrow \mathit{pull}^*(\mathfrak{A}, b_1 \leftarrow a_{\tau_1}, b_2 \leftarrow \mathfrak{f}_j(a_{\tau_1}), b_3 \leftarrow \mathfrak{f}_k(a_{\tau_1})).$$

- if $b_2 \rightsquigarrow_j b_1$ and $b_3 \rightsquigarrow_k b_1$ then set

$$(\mathfrak{B}\restriction\{b_1, b_2, b_3\})^* \leftarrow \mathit{pull}^*(\mathfrak{A}, b_1 \leftarrow a_{\tau_1}, b_2 \leftarrow \mathfrak{f}_j^{-1}(a_{\tau_1}), b_3 \leftarrow \mathfrak{f}_k^{-1}(a_{\tau_1})).$$

- if $b_1 \rightsquigarrow_j b_2$ for some j and b_3 is connected to neither b_1 nor b_2 in the witnessing graph then set

$$(\mathfrak{B}\restriction\{b_1, b_2, b_3\})^* \leftarrow \mathit{pull}^*(\mathfrak{A}, b_1 \leftarrow a_{\tau_1}, b_2 \leftarrow \mathfrak{f}_j(a_{\tau_1}), b_3 \leftarrow a_{\tau_3}).$$

- otherwise, set $(\mathfrak{B}\restriction\{b_1, b_2, b_3\})^* \leftarrow \mathit{pull}^*(\mathfrak{A}, b_1 \leftarrow a_{\tau_1}, b_2 \leftarrow a_{\tau_2}, b_3 \leftarrow a_{\tau_3})$.

Claim 5. $\mathfrak{B} \models \phi_1$.

Proof. Most of the proof follows from the same case analysis as in the algorithm itself. Hence, we give just a sketch and leave some details for the reader.

First notice that in the final structure $\mathfrak{B}$ all the $\mathcal{I}_1$- and $\mathcal{I}_2$-conjuncts are satisfied, since we have defined all the 2-types and 3-pre-types in $\mathfrak{B}$ by copying (pulling) them from $\mathfrak{A}$. More precisely, let $\forall x \forall y \forall z \phi_i(x, y, z)$ be one of the $\mathcal{I}_2$-conjuncts (for $\mathcal{I}_1$-conjucts the situation is analogous). Thanks to the uniformity of our logic, all the atoms in ϕ_i either use all the variables x, y, z or just one of them. Hence, if we evaluate $\phi_i(a, b, c)$ on any three (not necessary distinct) elements $a, b, c \in B$, the value of $\phi_i(a, b, c)$ depends only on: the 1-types of a, b, c, and the connections joining all the three elements a, b, c simultaneously. Since, the 1-types are fixed beforehand, and the connections joining all the elements a, b, c are pulled as a pre-structure from a valid model $\mathfrak{A}$, then $\mathfrak{B} \models \phi_i(a, b, c)$ must hold by Lemma 2. One can also notice, that no conflicts were introduced when copying different pre-structures. Indeed, in the algorithm we defined pre-structure on every at most three-element subset of B exactly once, and the pre-structures cannot speak about any proper subset of elements of its domain, except the 1-types which where fixed already in *Step 1*.

Let us now explain that the $\mathcal{I}_3$-conjucts are satisfied. Let $\forall x \exists y \forall z \phi_j(x, y, z)$ be any $\mathcal{I}_3$-conjuct. Fix an element an assume that $b \in B_m$. We need to see that in the next layer, i.e. $B_{(m+1)\text{mod } 4}$, there exists an element that can serve as a witness for b. It is equivalent to having an outgoing edge from b colored by j. Indeed, if there exists such an edge then the witnessing graph forces us to pick an appropriate pattern from a model $\mathfrak{A}$. Recall that each layer B_m consists of pairs (τ, m), where τ ranges over all generalized types realized in $\mathfrak{A}$. Therefore, there must be a representative a realizing τ in $\mathfrak{A}$. Hence, we have an element $\mathfrak{f}_j(a)$, such that $\mathfrak{A} \models \forall z \phi_j(a, \mathfrak{f}_j(a), z)$, which itself realizes some other generalized type τ'. Thus, in $B_{(m+1)\text{mod } 4}$ we have an element $b' = (\tau', (m+1)\text{mod } 4)$ satisfying $b \rightsquigarrow_j b'$. The fact that $\mathfrak{B} \models \forall z \phi_j(b, b', z)$ comes from the same line of arguments as in the previous paragraph. □

In the resulting structure $\mathfrak{B}$, similarly as for the structure $\mathfrak{A}$, we define witness functions $\mathfrak{f}_j : B \to B$ for each $\mathcal{I}_3$-conjuct. More precisely, we can set $\mathfrak{f}_j(b) = b'$ if $b \rightsquigarrow_j b'$; from the construction it follows that $\mathfrak{B} \models \forall z \phi_j(b, \mathfrak{f}_i(b), z)$. Note that the new functions $\mathfrak{f}_j$ are not necessarily 1-1, but they are without fixpoints and have disjoint images.

Finally, we give a bound on the size of the produced structure $\mathfrak{B}$. The number of generalized types can be trivially bounded by $t(t+1)(|\mathcal{I}_3|+1)t^{|\mathcal{I}_3|}$, where $t = |\boldsymbol{\alpha}|$. As mentioned in the Preliminaries (Lemma 1) t can be bounded by a function exponential in $|\phi|$, and $|\mathcal{I}_3|$ is clearly polynomial in $|\phi|$. Hence:

Claim 6. *The domain of* $\mathfrak{B}$ *is bounded exponentially in* $|\phi|$.

3.2 FMP for AUF_1^3

Now we build a model $\mathfrak{C} \models \phi_1 \wedge \phi_2$. It will be composed out of some number of copies of the structure $\mathfrak{B} \models \phi_1$ from the previous subsection, carefully connected to provide witnesses for conjuncts from ϕ_2.

Stage 0: The domain. Denote $\mathcal{I} = \mathcal{I}_4 \cup \mathcal{I}_5 \cup \mathcal{I}_6$. Let us form $\mathfrak{C}$ over the following finite domain.

$$C = B \times \mathcal{I} \times \{0,1\} \times \{0,\ldots,4\}.$$

For every $i \in \mathcal{I}, u = 0, 1$, and $0 \leq m \leq 4$ make the substructure $\mathfrak{C}\restriction(B \times \{i\} \times \{u\} \times \{m\})$ isomorphic to $\mathfrak{B}$, via the natural isomorphism sending each (b,i,u,m) to b. We also naturally transfer the functions $\mathfrak{f}_j$ and $\mathfrak{pat}$ from $\mathfrak{B}$ to $\mathfrak{C}$, setting $\mathfrak{f}_j((b,i,u,m)) = (\mathfrak{f}_j(b),i,u,m)$ and $\mathfrak{pat}((b,i,u,m)) = \mathfrak{pat}(b)$, for all b,i,u,m. For convenience we split C into five subsets $C_m = B \times \mathcal{I} \times \{0,1\} \times \{m\}$, for $m = 0,\ldots,4$.

In the next stages we will successively take care of the all types of conjuncts of ϕ. We will provide witnesses for $\mathcal{I}_4$-, $\mathcal{I}_5-$, and $\mathcal{I}_6$-conjuncts, using a strategy, which guarantees no conflicts, e.g., if a 3-pre-type is defined on a tuple c_1, c_2, c_3 to make c_3 a witness for c_1, c_2 for an $\mathcal{I}_5$-conjunct, then this tuple will not be used any more for a similar task (e.g., to make c_1 a witness for c_2, c_3). To design such a strategy we fix an auxiliary $1-1$ function $\mathfrak{chc}$ taking a subset of $\{0,\ldots,4\}$ of size two and extending it by 1 element from $\{0,\ldots,4\}$.[2] Having provided the witnesses we will carefully complete the structure ensuring that for every $c \in C$, $\mathfrak{f}_j(c)$ remains a witness for c and the $\mathcal{I}_3$-conjunct indexed by j.

Stage 1: Witnesses for $\mathcal{I}_4$-conjuncts. For every $c \in C$ and every $i \in \mathcal{I}_4$ repeat the following. Let m be the index such that $c \in C_m$. Let $a = \mathfrak{pat}(c)$. Let $a' \in A$ be such that $\mathfrak{A} \models \phi_i(a,a')$. It may happen that $a' = a$. Take as c' any element from $B \times \{i\} \times \{0\} \times \{m+1 (\mathrm{mod}\ 5)\}$ with the same 1-type as a'. Set $(\mathfrak{C}\restriction\{c,c'\}) \leftarrow \mathit{pull}(\mathfrak{A}, c \leftarrow a, c' \leftarrow a')$.

Stage 2: Witnesses for $\mathcal{I}_5$-conjuncts. For every unordered pair of elements $c_1, c_2 \in C$, $c_1 \neq c_2$ and every $i \in \mathcal{I}_5$ repeat the following. Let m_1, m_2 be the indices such that $c_1 \in C_{m_1}$ and $c_2 \in C_{m_2}$. We choose two pattern elements $a_1, a_2 \in A$ with the same 1-types as c_1 and, resp., c_2, and an index m' as follows.

(a) If $c_2 = \mathfrak{f}_j(c_1)$ for some j (in this case $m_1 = m_2$), then let $a_1 = \mathfrak{pat}(c_1)$ and let $a_2 = \mathfrak{f}_j(a_1)$; set $m' = m_1 + 1 (\mathrm{mod}\ 5)$.

(b) If $c_1 = \mathfrak{f}_j(c_2)$ for some j ($m_1 = m_2$), then let $a_2 = \mathfrak{pat}(c_2)$ and let $a_1 = \mathfrak{f}_j(a_2)$; set $m' = m_1 + 1 (\mathrm{mod}\ 5)$.

(c) If none of the above holds then let $a_1 = \mathfrak{pat}(c_1)$, $a_2 = \mathfrak{pat}(c_2)$ (it may happen that $a_1 = a_2$); if $m_1 = m_2$ then set $m' = m_1 + 1 (\mathrm{mod}\ 5)$, otherwise set m' to be the only element of $\mathfrak{chc}(\{m_1, m_2\}) \setminus \{m_1, m_2\}$.

[2] Example of $\mathfrak{chc}$: 01→012, 02→023, 03→034, 04→024, 12→123, 13→013, 14→014, 23→234, 24→124, 34→134.

Let $a' \in A$ be such that $\mathfrak{A} \models \phi_i(a_1, a_2, a')$ and let a'' be such that $\mathfrak{A} \models \phi_i(a_2, a_1, a'')$. Take as c' any element from $B \times \{i\} \times \{0\} \times \{m'\}$ with the same 1-type as a', and as c'' any element from $B \times \{i\} \times \{1\} \times \{m'\}$ with the same 1-type as a''. Set $(\mathfrak{C}\restriction\{c_1, c_2, c'\})^* \leftarrow pull^*(\mathfrak{A}, c_1 \leftarrow a_1, c_2 \leftarrow a_2, c' \leftarrow a'))$ and set $(\mathfrak{C}\restriction\{c_2, c_1, c''\})^* \leftarrow pull^*(\mathfrak{A}, c_2 \leftarrow a_2, c_1 \leftarrow a_1, c'' \leftarrow a'')$.

Additionally, for any $c \in C$ and any $i \in \mathcal{I}_5$ repeat the following. Let m be the index such that $c \in C_m$. Let $a = \mathfrak{pat}(c)$. Let $a' \in A$ be such that $\mathfrak{A} \models \phi_i(a, a, a')$. Take as c' any element from $B \times \{i\} \times \{0\} \times \{m + 1 (\text{mod } 5)\}$ with the same 1-type as a'. Set $(\mathfrak{C}\restriction\{c, c'\}) \leftarrow pull(\mathfrak{A}, c \leftarrow a, c' \leftarrow a')$.

Stage 3: Witnesses for $\mathcal{I}_6$-conjuncts. For every $c \in C$, and every $i \in \mathcal{I}_6$ repeat the following. Let m be the index such that $c \in C_m$. Let $a = \mathfrak{pat}(c)$. Let $a', a'' \in A$ be such that $\mathfrak{A} \models \phi_i(a, a', a'')$. Take as c' any element in C_m, in a different copy of $\mathfrak{B}$ than c, with the same 1-type as a'. Take as c'' any element in $B \times \{i\} \times \{0\} \times \{m + 1 (\text{mod } 5)\}$ with the same 1-type as a'' and set $(\mathfrak{C}\restriction\{c, c', c''\})^* \leftarrow pull^*(\mathfrak{A}, c \leftarrow a, c' \leftarrow a', c'' \leftarrow a'')$.

Stage 4: Taking care of the $\mathcal{I}_3$-conjuncts in $\mathfrak{C}$. For every $c \in C$, for every $c' \in C$ and every $j \in \mathcal{I}_3$ such that the pre-substructure on $\{c, \mathfrak{f}_j(c), c'\}$ has not yet been defined (note that the elements $c, \mathfrak{f}_j(c), c'$ must be pairwise distinct in this case) find any $a' \in A$ with the same 1-type as c', take $a = \mathfrak{pat}(c)$ and set $(\mathfrak{C}\restriction\{c, \mathfrak{f}_j(c), c'\})^* \leftarrow pull(\mathfrak{A}, c \leftarrow a, \mathfrak{f}_j(c) \leftarrow \mathfrak{f}_j(a), c' \leftarrow a')$.

Stage 5: Completing $\mathfrak{C}$. For any tuple of pairwise distinct elements $c_1, c_2, c_3 \in C$ for which the pre-substructure of $\mathfrak{C}$ on has not yet been defined set $(\mathfrak{C} \restriction \{c_1, c_2, c_3\})^* \leftarrow pull(\mathfrak{A}, c_1 \leftarrow \mathfrak{pat}(c_1), c_2 \leftarrow \mathfrak{pat}(c_2), c_3 \leftarrow \mathfrak{pat}(c_3))$. Similarly, for any pair of distinct elements $c_1, c_2 \in C$ for which the substructure of $\mathfrak{C}$ has not yet been defined, set $\mathfrak{C}\restriction\{c_1, c_2\} \leftarrow pull(\mathfrak{A}, c_1 \leftarrow \mathfrak{pat}(c_1), c_2 \leftarrow \mathfrak{pat}(c_2))$.

Claim 7. $\mathfrak{C} \models \phi$.

Proof. The purely universal conjuncts of ϕ ($\mathcal{I}_1-$ and $\mathcal{I}_2$-conjuncts) are satisfied since every 1-type, 2-type and 3-pre-type defined in $\mathfrak{C}$ is copied (pulled) from $\mathfrak{A}$ (either in this subsection or in the previous one where $\mathfrak{B}$ was constructed), which is a model of ϕ. All elements have witnesses for $\mathcal{I}_4$- $\mathcal{I}_5$-, and $\mathcal{I}_6$-conjuncts, as this is guaranteed in Stages 1, 2, and 3, respectively. Finally, for every c and every $j \in \mathcal{I}_3$, $\mathfrak{f}_j(c)$ is a witness for c and $\forall x \exists y \forall z \phi_j(x, y, z)$ in its copy of $\mathfrak{B}$. In Stage 4 we carefully define 3-pre-types on tuples consisting of $c, \mathfrak{f}_j(c)$ and elements from the other copies of $\mathfrak{B}$, so that $\mathfrak{f}_j(c)$ remains a witness for c in the whole structure $\mathfrak{C}$.

The only point which remains to be explained is that there are no conflicts in assigning types to tuples of elements. First note that there are no conflicts when assigning 2-types. They are defined in Stage 1 and in the last paragraph of Stage 2 (and completed in Stage 5–but this Stage may not be a source of conflicts as it sets only the types which have not been set before). The strategy is that the elements from C_m look for witnesses in $C_{m+1 \text{ mod } 5}$ and if an element c looks for witnesses for conjuncts indexed by different i_1, i_2 then it always looks for them in different copies of $\mathfrak{B}$. (A reader familiar with [3] may note that the way we deal with 2-types is similar to the one used in that paper, where three sets of elements are used and the witnessing scheme is modulo 3.) Considering 3-pre-types, they are set in Stages 2, and 3 (plus the completion in Stage 4 and Stage 5). No pre-type definition from Stage 2 can conflict with a pre-type definition from Stage 3: in Stage 2 we define pre-types on tuples consisting of elements either belonging two three different C_m, or two of them belongs to the same C_m and the third one to C_{m+1}, to a copy of $\mathfrak{B}$ indexed by $\mathcal{I}_5$; in Stage 3 we set pre-types on tuples consisting of two elements from the same C_m and an element from C_{m+1}, but from a copy of $\mathfrak{B}$ indexed by $\mathcal{I}_6$. Definitions from the same stage, involving two different C_m do not conflict with each other due to our circular witnessing scheme requiring elements from C_m to look for witnesses in C_{m+1}. The most interesting case is setting 3-pre-types for elements from 3-different C_m-s (Stage 2). This is done without conflict since the function $\mathfrak{chc}$ is $1-1$. □

Since $|C| = 8|B||\mathcal{I}|$, $|\mathcal{I}|$ is linear in $|\phi|$ and $|B|$ is exponential in ϕ we see that $|C|$ is bounded exponentially in $|\phi|$, which completes the proof of Thm. 3.

4 Reaching beyond AUF_1^3

The main technical advance in this Section will be showing that the decidability is retained if we relax the uniformity conditions in normal form for AUF_1^3, by allowing the formulas ϕ_i in $\mathcal{I}_5$- and $\mathcal{I}_6$-conjuncts to be arbitrary (not necessarily uniform) Boolean combinations of atoms. We believe that this result is interesting in itself, but additionally it allows us to define a pretty rich decidable class of FO^3 formulas with nested quantification, which we call FO^3_{-}. Recall that FO^3_{-} is defined in Preliminaries, where we also observed that it contains both FO^2 and AUF_1^3.

Let us see that we are indeed able to reduce FO^3_{-} sentences to normal form resembling that of AUF_1^3. We say that an FO^3_{-} sentence is in normal form if $\phi = \phi_1 \wedge \phi_2$ for

$$\phi_1 = \bigwedge_{i\in\mathcal{I}_2} \forall x\forall y\forall z\phi_i(x,y,z) \wedge \bigwedge_{i\in\mathcal{I}_3} \forall x\exists y\forall z\phi_i(x,y,z), \text{ and } \phi_2 = \bigwedge_{i\in\mathcal{I}_5} \forall x\forall y\exists z(x \neq y \to \phi_i(x,y,z)),$$

where the ϕ_i for $i \in \mathcal{I}_2 \cup \mathcal{I}_3$ are positive Boolean combinations of literals using either one variable or all variables from $\{x,y,z\}$ (uniform combinations), and the ϕ_i for $i \in \mathcal{I}_5$ are of the form $\phi'_i(x,y,z) \wedge \phi'_i(x,x,z)$, for some positive Boolean combination of arbitrary literals ϕ'_i (without equality). We note that even though we generally do not admit equality we allow ourselves to use $x \neq y$ as a premise in $\mathcal{I}_5$-conjuncts; as we will see this use of equality simplifies model constructions. Those conjuncts are actually equivalent to some equality-free conjuncts $\forall x\forall y\exists z\phi'_i(x,y,z)$, and hence Lemma 2, concerning the pull operation, holds for normal form FO^3_{-} formulas.

Lemma 8. *(i) The satisfiability problem for* FO^3_{-} *can be reduced in nondeterministic polynomial time to the satisfiability problem for normal form* FO^3_{-} *sentences. (ii) If the class of all normal form* FO^3_{-} *sentences has the finite (exponential) model property then also the whole* FO^3_{-} *has the finite (exponential) model property.*

Proof. (Sketch) We proceed as in the case of AUF_1^3, introducing new relation symbols to represent subformulas starting with a block of quantifiers. This time we need symbols of arity 0, 1 and 2. In particular, subformulas of the form $\exists v\forall v'\eta$ ($\forall v\forall v'\eta$) with a free variable v'' are replaced by $P(v'')$ for a fresh symbol P and axiomatized by $\forall v''\exists v\forall v'(P(v'') \to \eta)$ ($\forall v''\forall v\forall v'(P(v'') \to \eta)$), and subformulas $\exists v\eta(v,v',v'')$ are replaced by $R(v',v'')$ which is axiomatized by $\forall v\forall v'\exists v''(R(v',v'') \to \eta)$). This way, possibly by renaming the variables, we get a formula of the shape as in AUF_1-normal form but allowing the formulas ϕ_i for $i \in \mathcal{I}_5 \cup \mathcal{I}_6$ to be arbitrary (not necessarily uniform) Boolean combinations of atoms. We can now eliminate $\mathcal{I}_1$-, $\mathcal{I}_4$- and $\mathcal{I}_6$-conjuncts. We first proceed with $\mathcal{I}_6$-conjuncts $\forall x\exists y\exists z\phi_i(x,y,z)$, replacing them with the conjunctions of $\forall x\exists yR(x,y)$ and $\forall x\forall y\exists z(R(x,y) \to \phi_i(x,y,z))$, for a fresh binary symbol R. Then for $\mathcal{I}_1$-conjuncts $\forall x\forall y\phi_i(x,y)$, we just equip them with a dummy existential quantifier: $\forall x\forall y\exists z\phi_i(x,y)$, and for $\mathcal{I}_4$-conjuncts $\forall x\exists y\phi_i(x,y)$ we add a dummy universal quantifier and rename variables: $\forall x\forall y\exists z\phi_i(x,z)$. We end up with only $\mathcal{I}_5$-conjuncts in the ϕ_2 part of ϕ. Finally, any conjunct of the form $\forall x\forall y\exists z\phi_i(x,y,z)$ is replaced by $\forall x\forall y\exists z(x \neq y \to \phi_i(x,y,z) \wedge \phi_i(x,x,z))$. It should be clear that our transformations are sound over the domains consisting of at least two elements. (Singleton models can be enumerated and checked separately.) □

4.1 The complexity of FO^3_-

In this Section we show that FO^3_- is decidable, moreover its satisfiability problem is still in NExpTime. Concerning the obtained results, this subsection is a bit redundant, as in the next subsection, applying an alternative, probabilistic approach, we demonstrate the exponential model property for FO^3_-, which immediately gives the same NExpTime-upper bound on the complexity of the satisfiability problem, and additionally proves the same bound on the complexity of the finite satisfiability problem.

We decided to keep this subsection (although presenting some constructions and proofs in a slightly more sketchy way), since it gives some additional insight in the nature of the considered problems. We proceed here in a rather classical way: we show that a normal form FO^3_- sentence ϕ is satisfiable iff it has a certain *witness for satisfiability*, of exponentially bounded size, as stated in the following lemma. Such a witness of satisfiability is a structure which is "almost" a model of ϕ. It will turn out that one can quite easily (i) extract a witness from a given model of ϕ and (ii) unwind a witness into a proper infinite model of ϕ. In both directions we use a chase-like procedure.

Lemma 9. *A normal form FO^3_- sentence $\phi = \phi_1 \wedge \phi_2$ is satisfiable iff there exists a structure $\mathfrak{D}$ of size bounded by $h(|\phi|)$ for some fixed exponential function h, and functions without fixed points $\mathfrak{f}_j : D \to D$ for $j \in \mathcal{I}_3$, such that:*

- *(a)* $\mathfrak{D} \models \phi_1$
- *(b) for every $d, d' \in D$ and every $j \in \mathcal{I}_3$ we have $\mathfrak{D} \models \phi_j(d, \mathfrak{f}_j(d), d')$.*
- *(c) for every $i \in \mathcal{I}_5$, for every 2-type β realized in $\mathfrak{D}$, there is a realization (d, d') of β and an element d'' such that $d'' \neq d$, $d'' \neq d'$ and $\mathfrak{D} \models \phi_i(d, d', d'')$.*
- *(c') for every $i \in \mathcal{I}_5$, for every 2-type β realized in $\mathfrak{D}$ by a pair $(d_0, \mathfrak{f}_j(d_0))$ for some d_0 and j, there is $d \in D$ such that for $d' = \mathfrak{f}_j(d)$ the pair (d, d') has type β and there is an element d'' such that $d'' \neq d$, $d'' \neq d'$ and $\mathfrak{D} \models \phi_i(d, d', d'')$.*

A structure $\mathfrak{D}$, together with the functions $\mathfrak{f}_j$ meeting the conditions of the above lemma will be called *a witness of satisfiability for* ϕ. As we see $\mathfrak{D}$ satisfies the $\mathcal{I}_2$- and $\mathcal{I}_3$-conjuncts of ϕ (Conditions (a), (b)), but need not to satisfy the $\mathcal{I}_5$-conjuncts. However, if a pair d_0, d_0' does not have a witness for an $\mathcal{I}_5$-conjunct, then there is an appropriate "similar" pair d, d' having such a witness (Conditions (c), (c')).

Let us first prove the left-to-right direction of Lemma 9.

Left-to-right: extracting a witness of satisfiability from a model.

Stage 1: The building block $\mathfrak{B}$. Let $\mathfrak{A}$ be a model of a normal form FO^3_- formula $\phi = \phi_1 \wedge \phi_2$. Let $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ be the sets of all 1-types and all 2-types realized in $\mathfrak{A}$, respectively. Since our logic does not support equality, wlog. we assume that for each 1-type $\alpha \in \boldsymbol{\alpha}$, there exists a pair $a, b \in \mathfrak{A}$ realizing the 2-type $\beta \in \boldsymbol{\beta}$ which "unwinds" α, which means that: $\mathrm{tp}^{\mathfrak{A}}(a) = \mathrm{tp}^{\mathfrak{A}}(b) = \alpha$, and for each relation symbol R in the signature it holds: $\mathfrak{A} \models R(a, b)$ iff $\mathfrak{A} \models R(a, a)$, and $\mathfrak{A} \models R(b, a)$ iff $\mathfrak{A} \models R(a, a)$.

Let $\mathfrak{B}$ be the finite structure constructed for ϕ_1 and $\mathfrak{A}$ as in Section 3.1. It will serve as a building block in the construction of a witness of satisfiability $\mathfrak{D}$ for ϕ. Inspecting the construction of $\mathfrak{B}$ we note that the set of 1-types realized in $\mathfrak{B}$ is precisely $\boldsymbol{\alpha}$ and the set of

2-types is contained in $\boldsymbol{\beta}$; and regarding the size of $\mathfrak{B}$, it is exponentially bounded in $|\phi|$. Let $\mathfrak{f}_j : B \to B$ for $j \in \mathcal{I}_3$ be the witness functions defined as in Section 3.1.

Stage 2: The domain of $\mathfrak{D}_s$*.* The witness of satisfiability $\mathfrak{D}$ will consists of some number of copies of the structure $\mathfrak{B}$, say k (will be clarified later). More precisely, we will construct a sequence of structures $\mathfrak{D}_0, \mathfrak{D}_1, \ldots, \mathfrak{D}_{k-1}$, where $\mathfrak{D}_0$ is isomorphic to $\mathfrak{B}$ and $\mathfrak{D}_{k-1} = \mathfrak{D}$. Each $\mathfrak{D}_{s+1}$ will be created from $\mathfrak{D}_s$ by attaching a fresh disjoint copy of $\mathfrak{B}$, providing an $\mathcal{I}_5$-witness for some pair of elements of $(\mathfrak{D}_{s+1} \restriction D_s)$, and then completing the structure.

Formally, the domain of $\mathfrak{D}_s$ is the set

$$D_s = B \times \{0, \ldots, s\}.$$

We make the structure $\mathfrak{D}_s \restriction (B \times \{s\})$ isomorphic to $\mathfrak{B}$ and for $s > 0$ we make the structure $\mathfrak{D}_s \restriction (B \times \{0, \ldots, s-1\})$ isomorphic to $\mathfrak{D}_{s-1}$. We should think that $B \times \{0\}$ is the "main" copy of $\mathfrak{B}$ and $B \times \{m\}$ for $m > 0$ are auxiliary copies that will help us to satisfy the conditions (c) and (c'). We can adopt the functions $\mathfrak{f}_j : B \to B$ to the functions $\mathfrak{f}_j : D_s \to D_s$ by letting $\mathfrak{f}_j((b, m)) = (\mathfrak{f}_j(b), m)$.

We introduce a notion of an *extended 2-type* $(\beta, j) \in \boldsymbol{\beta} \times (\mathcal{I}_3 \cup \{\bot\})$ which except the information about the 2-type contains an information about whether one of the element (of a pair realizing this 2-type) is an $\mathcal{I}_3$-witness for the other. We say that $a_1, a_2 \in \mathfrak{A}$ ($a_1, a_2 \in \mathfrak{D}_s$) realize an extended 2-type (β, j) in $\mathfrak{A}$ (resp. $\mathfrak{D}_s$) if (1) a_1, a_2 realize 2-type β, and (2) either $i = \bot$, or $\mathfrak{f}_j(a_1) = a_2$, or $\mathfrak{f}_j(a_2) = a_1$. Denote by $\boldsymbol{\beta}^+$ the set of all extended 2-types realized in $\mathfrak{A}$. For each $i \in \mathcal{I}_5$ select a pattern function $\mathfrak{pat}_i^{\mathfrak{A}} : \boldsymbol{\beta}^+ \to A^3$ such that $\mathfrak{pat}_i^{\mathfrak{A}}(\beta, j) = (a_1, a_2, a_3)$ if $\{a_1, a_2\}$ realizes (β, j) in $\mathfrak{A}$, and $\mathfrak{A} \models \phi_i(a_1, a_2, a_3)$.

Stage 3: Constructing $\mathfrak{D}_{s+1}$ *from* $\mathfrak{D}_s$*.* Let $b_1, b_2 \in \mathfrak{D}_s$ be a pair of distinct elements realizing an extended 2-type (β, j) for which (c) or (c') fails for some $i \in \mathcal{I}_5$. Let $(a_1, a_2, a_3) = \mathfrak{pat}_i^{\mathfrak{A}}(\beta, j)$. Set $(\mathfrak{D}_{s+1} \restriction \{b_1, b_2, b_3\}) \leftarrow \mathit{pull}(\mathfrak{A}, b_1 \leftarrow a_1, b_2 \leftarrow a_2, b_3 \leftarrow a_3)$, where $b_3 \in B \times \{s+1\}$ is any element of the freshly added copy of $\mathfrak{B}$ with the same 1-type as a_3. After providing a witness for b_1, b_2 this way, we need to complete $\mathfrak{D}$, i.e., we need to specify the relations between all pairs and triples of elements for which we have not done it yet. Take any subset $S \subseteq D_{s+1}$ of size 2 or 3, and search for an appropriate pattern in $\mathfrak{A}$, then pull it into $\mathfrak{D}_{s+1} \restriction S$. We will omit the details here. Consult the Section 3.1 and the algorithm defined there for the very similar line of the reasoning. Therefore, $\mathfrak{D}_{s+1} \models \phi_1$ (Condition (a)), and $\mathfrak{D}_{s+1} \models \phi_i(d, \mathfrak{f}_j(d), d')$ for each $j \in \mathcal{I}_3$ (Condition (b)).

Stage 4: Finalization. Conditions (a) and (b) are satisfied after every step, and after at most $|\boldsymbol{\beta}^+|$ many steps (the number k), which is polynomial in $|\boldsymbol{\beta}|$ and $|\phi|$, all the 2-types will satisfy (c) and (c'). Thus, all the conditions are met for $\mathfrak{D}_k$. The size of the domain of $\mathfrak{D}_k$ is at most exponential in $|\phi|$, since $\mathfrak{D}_k$ consists in fact of $|\boldsymbol{\beta}^+|$ copies of $\mathfrak{B}$, and both $\mathfrak{B}$ and $|\boldsymbol{\beta}^+|$ are exponentially bounded in ϕ.

Let us now turn to the right-to-left direction of Lemma 9.

Right-to-left: building a model from a witness of satisfiability

Let $\mathfrak{D}$ be a witnesses of satisfiability. If ϕ does not have any $\mathcal{I}_5$-conjucts then $\mathfrak{D}$ is already a model of ϕ. In the other case, we will show how to construct a new witness of satisfiability $\mathfrak{D}'$, but bigger. Iterating this construction indefinitely, and taking the limit structure, we will obtain an infinite model of $\mathfrak{A}$.

The construction of a new witnesses of satisfiability is very similar to the process of extracting a witnesses of satisfiability from a model. Hence, we will just sketch the construction, and stress the differences. The domain of $\mathfrak{D}'$ is the set[3]

$$D' = D \times \{(\bot, \emptyset)\} \cup D \times \mathcal{I}_5 \times \binom{D}{2},$$

and let the structure on $\mathfrak{D}' \restriction D \times \{s\}$ be isomorphic to $\mathfrak{D}$ for each $s \in \{(\bot, \emptyset)\} \cup (\mathcal{I}_5 \times \binom{D}{2})$. Then we take care of witnesses for all $\mathcal{I}_5$-conjuncts for all the pairs of elements in $\mathfrak{D}_{\bot,\emptyset}$. The witness for a conjunct ϕ_i and a pair $\{d_1, d_2\}$ is provided in $D_{i,\{d_1,d_2\}}$. The major difference is that this time we pull pre-substructures not from an existing model of ϕ, but rather from $\mathfrak{D}$ itself; this is always possible due to Conditions (c) and (c'). Then we complete the structure, similarly as in our previous constructions, by setting the relations between all the pairs and the triples of elements for which we have not done it yet. Again, in this completion process the pre-substructures are pulled from $\mathfrak{D}$, and not from an external structure. So after that all the $\mathcal{I}_2$- and $\mathcal{I}_3$-conjuncts are satisfied, and all the pairs of elements from D have witnesses for all $\mathcal{I}_5$-conjuncts. We claim that the newly obtained structure is again a witness of satisfiability, with all elements from D having the required witnesses, but with a strictly larger domain. This allows us to continue the process ad infinitum, with the natural limit structure being a model of ϕ.

Having shown Lemma 9, and noting that checking if a given structure $\mathfrak{D}$ is a witness of satisfiability for ϕ can be easily done in time polynomial in $|D|$ and $|\phi|$, we get:

Theorem 10. *The satisfiability problem for* FO^3_- *is* NExpTime*-complete.*

4.2 Finite model property via probabilistic method

The probabilistic model generation we present in this subsection is inspired by Gurevich and Shelah's fmp proof for Gödel class without equality [4]. The main differences are that we will use finite building blocks $\mathfrak{B} \models \phi_1$ constructed deterministically and generate whole 3-pre-substructures at once rather than sample individual 3-ary atoms. We will also present an analysis of the size of a minimal model guaranteed by our reasoning.

Let $\phi = \phi_1 \wedge \phi_2$ be a normal form FO^3_--sentence and $\mathfrak{A}$ its model. Wlog. we assume that any 1-type realized in $\mathfrak{A}$ is realized by at least three distinct elements. If it is not the case, we replace $\mathfrak{A}$ by $3\mathfrak{A}$–the structure with universe $A \times \{0,1,2\}$ such that (i) $\mathrm{tp}^{3\mathfrak{A}}((a,m)) = \mathrm{tp}^{\mathfrak{A}}(a)$ (ii) $3\mathfrak{A} \restriction \{(a_1,m_1),(a_2,m_2)\}$ is obtained as $\mathit{pull}(\mathfrak{A}, (a_1,m_1) \leftarrow a_1, (a_2,m_2) \leftarrow a_2)$ and $(3\mathfrak{A} \restriction \{(a_1,m_1),(a_2,m_2),(a_3,m_3)\})^*$ is obtained by $\mathit{pull}^*(\mathfrak{A}, (a_1,m_1) \leftarrow a_1, (a_2,m_2) \leftarrow a_2, (a_3,m_3) \leftarrow a_3)$. As previously, we pick for every $j \in \mathcal{I}_3$ a witness function $\mathfrak{f}_j : A \to A$ such that for every $a \in A$ we have $\mathfrak{A} \models \forall z \phi_j(a, \mathfrak{f}_j(a), z)$. Let $\boldsymbol{\alpha}, \boldsymbol{\beta}$ and $\boldsymbol{\gamma}^*$ be the set of 1-types realized in $\mathfrak{A}$, the set of 2-types realized in $\mathfrak{A}$, and, resp., the set of 3-pre-types realized in $\mathfrak{A}$.

Let us construct a model $\mathfrak{B} \models \phi_1$, together with functions $\mathfrak{pat} : B \to A$ and $\mathfrak{f}_j : B \to B$, for $j = 1, \ldots, |\mathcal{I}_3|$, the latter without fixed points, as in Section 3.1. We may assume that the set of 1-types realized in $\mathfrak{B}$ is $\boldsymbol{\alpha}$ (cf. Section 4.1, Left-to-right, Stage 1).

We now define a sequence of (partly) random structures $\mathfrak{C}_n$, for $n = 1, 2, \ldots$.

Stage 0: The domain. $C_n = B \times \{1, \ldots, n\}$. We make the substructures $\mathfrak{C}_n \restriction B \times \{m\}$ isomorphic to $\mathfrak{B}$, for all $m = 1, 2, \ldots, n$, and set no other connections in this stage. We naturally lift functions $\mathfrak{pat}$ and the $\mathfrak{f}_j$ to $\mathfrak{C}_n$.

[3] By $\binom{A}{2}$, where A is a set, we mean the set of all 2-element subsets of A.

Stage 1: 2-types. For every pair $c_1, c_2 \in C_n$ of distinct elements for which $\mathrm{tp}^{\mathfrak{C}_n}(c_1, c_2)$ has not yet been defined (note that c_1 and c_2 belong then to different copies of $\mathfrak{B}$ in this case) proceed as follows. Let $\alpha_1 = \mathrm{tp}^{\mathfrak{C}_n}(c_1)$, $\alpha_2 = \mathrm{tp}^{\mathfrak{C}_n}(c_2)$ and let $\boldsymbol{\beta}_0$ be the set of 2-types β realized in $\mathfrak{A}$ such that $\beta \restriction x_1 = \alpha_1$ and $\beta \restriction x_2 = \alpha_2$. Clearly $\boldsymbol{\beta}_0$ is non-empty, since even if $\alpha_1 = \alpha_2$ we will have two distinct realizations of this 1-type by our assumption about $\mathfrak{A}$. Choose $\mathrm{tp}^{\mathfrak{C}_n}(c_1, c_2)$ at random from $\boldsymbol{\beta}_0$, with uniform probability.

Stage 2a: 3-pre-types. For every tuple $c_1, c_2, c_3 \in C_n$ such that $c_2 = \mathfrak{f}_j(c_1)$ for some j, and $\mathrm{pretp}^{\mathfrak{C}_n}(c_1, c_2, c_3)$ has not yet been defined proceed as follows. Recall that by the construction of $\mathfrak{B}$, c_2 may not be in the image of any $\mathfrak{f}_k$ for $j \neq k$ in this case. Let $a_1 = \mathfrak{pat}(c_1)$ and $a_2 = \mathfrak{f}_j(a_1)$, and $\alpha_3 = \mathrm{tp}^{\mathfrak{C}_n}(c_3)$. Let $\boldsymbol{\gamma}_0^*$ be the set of 3-pre-types realized in $\mathfrak{A}$ by tuples a_1, a_2, a_3, for all $a_3 \in A$ having 1-type α_3. Note that $\boldsymbol{\gamma}_0^*$ is non-empty–even if a_1 and a_2 have type α_3 then at least one appropriate a_3, different from a_1, a_2 exists since we assumed that every 1-type realized in $\mathfrak{A}$ is realized at least three times. Choose $\mathrm{pretp}^{\mathfrak{C}_n}(c_1, c_2, c_3)$ at random from $\boldsymbol{\gamma}_0^*$, with uniform probability.

Stage 2b: 3-pre-types contd. For every tuple $c_1, c_2, c_3 \in C_n$ for which $\mathrm{pretp}^{\mathfrak{C}_n}(c_1, c_2, c_3)$ has not yet been defined proceed as follows. Let $\alpha_1 = \mathrm{tp}^{\mathfrak{C}_n}(a_1)$, $\alpha_2 = \mathrm{tp}^{\mathfrak{C}_n}(a_2)$, $\alpha_3 = \mathrm{tp}^{\mathfrak{C}_n}(a_3)$ and let $\boldsymbol{\gamma}_0^*$ be the set of 3-pre-types γ^* realized in $\mathfrak{A}$ such that $\gamma^* \restriction x_1 = \alpha_1$ and $\gamma^* \restriction x_2 = \alpha_2$ and $\gamma^* \restriction x_3 = \alpha_3$. Again, using the assumption about $\mathfrak{A}$ we note that $\boldsymbol{\gamma}_0^*$ is non-empty. Choose $\mathrm{pretp}^{\mathfrak{C}_n}(c_1, c_2, c_3)$ at random from $\boldsymbol{\gamma}_0^*$, with uniform probability.

Let us now estimate the probability that $\mathfrak{C}_n \not\models \phi$. It is not difficult to observe that, regardless of our random choices, $\mathfrak{C}_n \models \phi_1$. Indeed, the conjuncts indexed by $\mathcal{I}_3$ are respected since they are respected in $\mathfrak{B}$ and we ensure that $\mathfrak{f}_j(c)$ remains an appropriate witness for c in $\mathfrak{C}_n$, for every c in Stage 2a; the conjuncts indexed by $\mathcal{I}_2$ are respected since they are respected in $\mathfrak{B}$ and all pre-types defined in $\mathfrak{C}$ are copied from $\mathfrak{A}$ (Stages 2a, 2b). Hence, we need to take into consideration only the satisfaction of ϕ_2. Let $F^i_{c_1,c_2}$, for a pair of distinct elements $c_1, c_2 \in C_n$ and $i \in \mathcal{I}_5$, denotes the event "$\mathfrak{C}_n \not\models \exists z \phi_i(c_1, c_2, z)$". We first estimate its probability.

Take a pair of distinct elements $c_1, c_2 \in C_n$ and $i \in \mathcal{I}_5$. There are three cases: (a) $c_2 = f_j(c_1)$ for some j, (b) $c_1 = f_j(c_2)$ for some j, and (c) none of the above holds. Let us go in details through case (a). Let $a_1 = \mathfrak{pat}(c_1)$, $a_2 = \mathfrak{f}_j(a_1)$. In this case c_1 and c_2 belong to the same copy of $\mathfrak{B}$, say to $\mathfrak{B} \times \{k\}$, and $\mathrm{tp}^{\mathfrak{C}_n}(c_1, c_2)$ was set in Stage 0 to be equal to $\mathrm{tp}^{\mathfrak{A}}(a_1, a_2)$. If $\mathfrak{A} \models \phi_i(a_1, a_2, a_1)$ then $\mathfrak{C}_n \models \phi_i(c_1, c_2, c_1)$, and if $\mathfrak{A} \models \phi_i(a_1, a_2, a_2)$ then $\mathfrak{C}_n \models \phi_i(c_1, c_2, c_2)$, so the pair a_1, a_2 has a witness for the ϕ_i. Otherwise, there is an element a_3, $a_3 \notin \{a_1, a_2\}$ such that $\mathfrak{A} \models \phi_i(a_1, a_2, a_3)$. Let $\alpha_3 = \mathrm{tp}^{\mathfrak{A}}(a_3)$. Consider now any element c_3 of type α_3 in $\mathfrak{B} \times (\{1, \ldots k-1, k+1, \ldots n)$. During the construction of $\mathfrak{C}_n$ we randomly chose $\mathrm{tp}^{\mathfrak{C}_n}(c_1, c_3)$ and $\mathrm{tp}^{\mathfrak{C}_n}(c_2, c_3)$ (Stage 1) and $\mathrm{pretp}^{\mathfrak{C}_n}(c_1, c_2, c_3)$ (Stage 2a). We had at most $|\boldsymbol{\beta}|$ choices for $\mathrm{tp}^{\mathfrak{C}_n}(c_1, c_3)$ and at least one of them was $\mathrm{tp}^{\mathfrak{A}}(a_1, a_3)$. Analogously, for $\mathrm{tp}^{\mathfrak{C}_n}(c_2, c_3)$. We had at most $|\boldsymbol{\gamma}^*|$ choices for $\mathrm{pretp}^{\mathfrak{C}_n}(c_1, c_2, c_3)$ and at least one of them was $\mathrm{pretp}^{\mathfrak{A}}(a_1, a_2, a_3)$. Hence the probability that $\{c_1 \to a_1, c_2 \to a_2, c_3 \to a_3\}$ is a partial isomorphism and hence $\mathfrak{C}_n \models \phi_i(c_1, c_2, c_3)$ is at least $\epsilon = 1/(|\boldsymbol{\beta}|^2 \cdot |\boldsymbol{\gamma}^*|)$. As there are at least $n-1$ elements of type α_3 in $\mathfrak{B} \times (\{1, \ldots k-1, k+1, \ldots n)$ (at least one in each copy of $\mathfrak{B}$), and the structures on c_1, c_2, c_3 and c_1, c_2, c_3' for different such elements c_3, c_3' are chosen independently, it follows that $\mathrm{Prob}(F^i_{c_1,c_2}) \leq (1-\epsilon)^{n-1}$.

Note that the event "$\mathfrak{C}_n \not\models \phi$" is equal to $\bigcup_{c_1 \neq c_2} \bigcup_{i \in \mathcal{I}_5} F^i_{c_1,c_2}$. Hence

$$\mathrm{Prob}(\mathfrak{C}_n \not\models \phi) = \mathrm{Prob}(\bigcup_{c_1 \neq c_2} \bigcup_{i \in \mathcal{I}_5} F^i_{c_1,c_2}) \leq \sum_{c_1 \neq c_2} \sum_{i \in \mathcal{I}_5} \mathrm{Prob}(F^i_{c_1,c_2}) = n|B|(n|B|-1)|\mathcal{I}_5|(1-\epsilon)^{n-1}.$$

The limit of the above estimation when n approaches infinity is 0 (note that ϵ, $|B|$ and $|\mathcal{I}_5|$ do not depend on n). In particular for some $n \in \mathbb{N}$ we have that $\mathrm{Prob}(\mathfrak{C}_n \not\models \phi) < 1$ and thus one of the possible (finite number) of choices of $\mathfrak{C}_n$ is indeed a finite model of ϕ.

We can refine this reasoning and achieve an exponential bound on the size of a minimal finite model by applying the following Lemma:

Lemma 11. *There exists an exponential function f such that Prob$(\mathfrak{C}_n \not\models \phi) < 1$ for $n \geq f(|\phi|)$.*

Proof. Denote by $|\phi|$ the size of ϕ in any reasonable encoding. In Section 3.1 we estimated the size of the domain of $\mathfrak{B}$ to be exponential in $|\phi|$. Hence, there exist a polynomial $p_1(x)$ such that $|B| \leq 2^{p_1(|\phi|)}$. Similarly, from Lemma 1 we get that there exists polynomials $q_1(x)$ and $q_2(x)$ such that $|\boldsymbol{\beta}| \leq 2^{q_1(|\phi|)}$ and $|\boldsymbol{\gamma}^*| \leq 2^{q_2(|\phi|)}$. Wlog. we can assume that $p_1(x), q_1(x), q_2(x) \geq 10x$.

Let us state the bound on the probability of the event "$\mathfrak{C}_{n+1} \not\models \phi$" in terms of $p(x) := 2p_1(x)$, and $q(x) := 2q_1(x) + q_2(x)$:

$$\mathrm{Prob}(\mathfrak{C}_{n+1} \not\models \phi) = (n+1)|B|((n+1)|B|-1)|\mathcal{I}_5|(1-\epsilon)^n \leq 10n^2 2^{p(|\phi|)}|\phi|(1-\varepsilon(|\phi|))^n, \quad (1)$$

where $\varepsilon(x)$ is defined as $\varepsilon(x) := 1/2^{q(x)}$ and therefore $\varepsilon(|\phi|) \leq 1/(|\boldsymbol{\beta}|^2 \cdot |\boldsymbol{\gamma}^*|) = \epsilon$.

We need to show that taking $n = f(|\phi|)$ for some exponential function f makes the right-hand side of (1) strictly less than one.[4] Let $f(x) = 10\ln(2) \cdot 2^{2q(x)}p(x)$. Then taking log on both sides of:

$$10f(x)^2 2^{p(x)} x(1-\varepsilon(x))^{f(x)} < 1$$

gives us:

$$\log(10) + 2\log(f(x)) + p(x) + \log(x) + f(x)\log(1-\varepsilon(x)) < 0.$$

Recall that the Mercator Series states that $\ln(1-x) = -\sum_{i=1}^{\infty} x^i/i$ for all $x \in (-1,1)$. Hence we want to see that:

$$\log(10) + 2\log(f(x)) + p(x) + \log(x) < \frac{f(x)}{\ln(2)} \sum_{i=1}^{\infty} \frac{\varepsilon(x)^i}{i}.$$

Since $f(x)$ and $\varepsilon(x)$ are always positive, the above inequality is implied by:

$$\log(10) + 2\log(f(x)) + p(x) + \log(x) < f(x)\varepsilon(x)/\ln(2).$$

Expanding the definition of $\varepsilon(x)$ and $f(x)$:

$$\log(10) + 2\log(10\ln(2)) + 4q(x) + 2\log(p(x)) + p(x) + \log(x) < 10\ln(2) \cdot 2^{2q(x)}p(x) \cdot 2^{-q(x)}/\ln(2).$$

Finally we get:

$$\log(10) + 2\log(10\ln(2)) + 4q(x) + 2\log(p(x)) + p(x) + \log(x) < 10 \cdot 2^{q(x)}p(x),$$

and it is easy to see that each term on the left-hand side is bounded by $2^{q(x)}p(x)$ (possibly for a large enough x), so the inequality holds. □

Recalling Lemma 8 we conclude.

Theorem 12. FO^3_{-} *has the finite (exponential) model property. The finite satisfiability problem (=satisfiability problem) for* FO^3_{-} *is* NExpTime*-complete.*

[4] We did not tried to optimize this bound for the sake of simplicity.

5 Concluding remarks

5.1 Notes on equality

It is clear that allowing equality in FO^3_- spoils the decidability, as FO^3 contains formulas of the form $\forall x \forall y \exists z \phi$, for quantifier-free ϕ, which form an undecidable subclass of Gödel class with equality [2].

Turning to AUF^3_1, we have two options of adding equality: first, we can treat the equality symbol as the relation symbols from the signature and allow only its uniform use; second: we can allow for free use of equality, as e.g. in the formula $\forall xyz(R(x,y,z) \rightarrow x \neq y)$. As for the former, note that the syntactic restrictions allow the equality to be used only in $\mathcal{I}_1$-conjuncts and (less importantly) $\mathcal{I}_4$-conjuncts. $\mathcal{I}_1$-conjuncts allow one to say that some 1-types are realized at most once. We believe that treating the realizations of such types with appropriate care and respect (extending the approach with *kings* and *court* from [3]) one can show the finite model property for the obtained variant and its decidability in NExpTime. We leave the details to be worked out. Regarding the free use of equality, we observe that the finite model property is lost. Consider the following formula:

$$\exists x S(x) \wedge \forall x \exists y \forall z (\neg S(y) \wedge R(x,y,z) \wedge (x = z \vee \neg R(z,y,x))$$

It is satisfied in the model whose universe is the set of natural numbers, S is true only at 0 and $R(x,y,z)$ is true iff $y = x+1$. It is readily verified that there are no finite models (every x needs to take a fresh y as a witness since otherwise $R(z,y,x)$ would not hold for one of the earlier z, $z \neq x$). The above example shows that satisfiability and finite satisfiability are different problems in this case. We do not know if they are decidable, however. This issue is left for further investigations.

5.2 Future work

Except the potential addition of equality, the main question which remains open is if full AUF_1 has decidable satisfiability and the finite model property. We partially answer this question in the already mentioned upcoming paper, in which we show decidability for a subvariant of AUF_1 in which every block of quantifiers is purely universal or ends with the existential quantifier. The general case is left open.

References

[1] K. Gödel. Zum entscheidungsproblem des logischen funktionenkalkuils. *Monatshefte fur Mathematik und Physik*, 40:433–443, 1933.

[2] W. D. Goldfarb. The unsolvability of the Gödel class with identity. *J. Symb. Logic*, 49:1237–1252, 1984.

[3] E. Grädel, P. Kolaitis, and M. Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.

[4] Y. Gurevich and S. Shelah. Random models and the Gödel case of the decision problem. *J. Symbolic Logic*, 48(4):1120–1124, 1983.

[5] L. Hella and A. Kuusisto. One-dimensional fragment of first-order logic. In *Proceedings of Advances in Modal Logic, 2014*, pages 274–293, 2014.

[6] A.S. Kahr, E.F. Moore, and H. Wang. Entscheidungsproblem reduced to the ∀∃∀ case. *Proc. Nat. Acad. Sci. U.S.A.*, 48:365–377, 1962.

[7] Y. Kazakov. *Saturation-based decision procedures for extensions of the guarded fragment.* PhD thesis, Universität des Saarlandes, Saarbrücken, Germany, 2006.

[8] E. Kieroński. Results on the guarded fragment with equivalence or transitive relations. In *CSL*, volume 3634 of *LNCS*, pages 309–324. Springer, 2005.

[9] E. Kieronski and A. Kuusisto. Uniform one-dimensional fragments with one equivalence relation. In *CSL*, volume 41 of *LIPIcs*, pages 597–615, 2015.

[10] H. R. Lewis. Complexity results for classes of quantificational formulas. *Journal of Computer and System Sciences*, 21(3):317 – 353, 1980.

[11] M. Mortimer. On languages with two variables. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:135–140, 1975.

[12] D. Scott. A decision method for validity of sentences in two variables. *Journal Symbolic Logic*, 27:477, 1962.

[13] B. ten Cate and L. Segoufin. Unary negation. *Logical Methods in Comp. Sc.*, 9(3), 2013.