Small Substructures and Decidability Issues for First-Order Logic with Two Variables

Emanuel Kieroński*
Institute of Computer Science
University of Wrocław
kiero@ii.uni.wroc.pl

Martin Otto†
Fachbereich Mathematik
Technische Universität Darmstadt
otto@mathematik.tu-darmstadt.de

Abstract

We study first-order logic with two variables FO^2 *and establish a small substructure property. Similar to the small model property for* FO^2 *we obtain an exponential size bound on embedded substructures, relative to a fixed surrounding structure that may be infinite. We apply this technique to analyse the satisfiability problem for* FO^2 *under constraints that require several binary relations to be interpreted as equivalence relations. With a single equivalence relation,* FO^2 *has the finite model property and is complete for non-deterministic exponential time, just as for plain* FO^2*. With two equivalence relations,* FO^2 *does not have the finite model property, but is shown to be decidable via a construction of regular models that admit finite descriptions even though they may necessarily be infinite. For three or more equivalence relations,* FO^2 *is undecidable.*

0. Introduction

The undecidability of the satisfiability problem for first-order logic has inspired the classification of various syntactic fragments of first-order logic with a view to delineating the boundary of decidability as well as to finding useful decidable fragments. The main programme of this kind, which has led to a complete classification, concerned the taxonomy of prenex normal form formulae w.r.t. quantifier prefixes [3]. Along an orthogonal direction, one may investigate fragments defined not with reference to prenex normal form, but in terms of other uniform structural constraints on the quantification patterns. Modal logics, with their characteristic relativisation of all quantifiers by binary edge predicates, provide a typical example of a benign fragment of this kind. A more recent extension led to the guarded fragment [1, 7]. More crudely, a mere restriction of the number of distinct variable symbols leads to the finite variable fragments FO^k of first-order logic. Interestingly, both the modal families of logics and the finite variable fragments can also be motivated in terms of model theoretic games – namely bisimulation games for modal logics, and k-pebble games for FO^k. Correspondingly, and unlike prefix classes, these fragments enjoy natural closure properties which support some characteristic model theory. In particular, the finite variable fragments play a prominent role in finite model theory. In terms of satisfiability, FO^2 is decidable while FO^3 is undecidable. FO^2 here stands for the fragment of first-order logic with equality with only two variable symbols x and y, in finite relational vocabularies (without constants or function symbols). Without loss of generality, we also only consider vocabularies of width 2, without relation symbols of arities greater than 2.

The first decidability proof for FO^2 was given by Scott [18], via a reduction to the so-called Gödel prefix class $\exists^*\forall\forall\exists^*$, which however is only decidable in the absence of equality [6]. Full decidability for FO^2 with equality is due to Mortimer [15]. Mortimer shows FO^2 has the finite model property, and in fact that every satisfiable FO^2 sentence has a model of size at most doubly exponential in the length of the sentence. This bound on the size of small models is improved to single exponential by Grädel, Kolaitis and Vardi in [8], which leads to their result that FO^2 is decidable in NEXPTIME, and in fact NEXPTIME-complete.

The study of FO^2 is also motivated by the fact that it embeds propositional modal logic K, via the standard translation. Numerous variants and extensions of modal logic find applications in various areas of computer science, including verification of software and hardware, distributed systems, knowledge representation and artificial intelligence. These applications are supported by the very good algorithmic and model-theoretic behaviour of modal logics, including their remarkably robust decidability which persists under various extensions towards greater expressiveness. Some of these extensions are equally well motivated in the context of two-variable logic. Description logics in particular naturally fit

* supported by EPSRC grant GR/R11896/01
† research partially supported by EPSRC grant GR/R11896/01

into the range between modal and two-variable logics, see [2] and in particular [4] for their connection with finite variable fragments. The extension of FO^2 by counting quantifiers [12], for instance, is decidable by [10, 17] although it does not have the finite model property. In analogy with graded modalities, it covers certain description logics with number constraints. But also in systematic terms the question naturally arises, to which extent FO^2 shares the good algorithmic behaviour of modal logics. The picture that emerged in [11] shows that, with the notable exception of the counting extension, most extensions of FO^2 – e.g. by mechanisms for fixed points or transitive closures, in analogy with the modal μ-calculus or computation tree logics – are undecidable, compare also [9]. In many cases, the results of these investigations can be phrased either for satisfiability of extensions of FO^2, or, alternatively, of FO^2 itself over restricted classes of structures. This interplay is fruitfully employed in [11, 16].

In connection with modal logics, or with applications of modal or two-variable logics in areas like knowledge representation or for description logics, a restriction of the underlying class of models is often very natural. Modal correspondence theory, for instance, associates transitivity of accessibility relations with the modal logic K4; equivalence relations with the modal logic S5. Multi-S5 systems with k equivalence relations among their accessibility relations can be used to model knowledge systems for k independent agents; linear orders as accessibility relations play an obvious role for linear temporal logics, etc. For FO^2 over such classes of structures, undecidability is established under several such constraints in [11, 9] and in particular in the presence of 4 equivalence relations. In the presence of a linear order, on the other hand, decidability is shown in [16]. In this paper we concentrate on a complete analysis of the important case of models with several equivalence relations, by clarifying the situation for up to three equivalence relations.

We look at finite relational vocabularies $\tau = \tau_0 \,\dot\cup\, \tau_{\mathrm{eq}}$ where τ_{eq} consists of a finite number of distinguished binary relations (typically E or $E_i, i = 1, 2, 3$). We let $\mathcal{EQ}[\tau_0; \tau_{\mathrm{eq}}]$ denote the class of all $\tau_0 \,\dot\cup\, \tau_{\mathrm{eq}}$ structures $\mathfrak{A} = (\mathfrak{A}_0, (E^{\mathfrak{A}})_{E \in \tau_{\mathrm{eq}}})$ (with at least two elements) that interpret the relations $E \in \tau_{\mathrm{eq}}$ as equivalence relations. We refer to such structures as *equivalence structures*.

$\mathrm{SAT}(\mathcal{L}, \mathcal{C})$, the satisfiability problem for $\mathcal{L}$ over the class $\mathcal{C}$, is the decision problem, for sentences $\varphi \in \mathcal{L}$, whether φ has a model in $\mathcal{C}$. We say that $\mathcal{L}$ has the finite model property (or a small model property) over $\mathcal{C}$, if every sentence $\varphi \in \mathcal{L}$ that has a model in $\mathcal{C}$ also has a finite (small) model in $\mathcal{C}$. In these terms our main results for FO^2 over equivalence structures are the following.

Theorem 1 (i) FO^2 *has an exponential model property over* $\mathcal{EQ}[\tau_0; E]$. $\mathrm{SAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E])$ *is* NEXPTIME-*complete.*

(ii) FO^2 *does not have the finite model property over* $\mathcal{EQ}[\tau_0; E_1, E_2]$. *However,* $\mathrm{SAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E_1, E_2])$ *is decidable in* 3NEXPTIME.

(iii) $\mathrm{SAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E_1, E_2, E_3])$ *is undecidable.*

Our decidability result for two equivalence relations should also be contrasted with another result by the first author, that FO^2 is undecidable in the presence of two transitive relations [14].

En route to the decidability result we establish a *small substructure property* for FO^2 which does not directly focus on entire models of FO^2-sentences but rather on small substructures that are parts of the actual models. This technique proves to be applicable also in the case of equivalence structures with two equivalence relations, where the finite model property for FO^2 fails. As we shall see, FO^2 sentences can force models in $\mathcal{EQ}[\tau_0; E_1, E_2]$ to have infinitely many classes as well as to have infinite classes. By our small substructure property the size of the equivalence classes of the common refinement $E = E_1 \cap E_2$ of E_1 and E_2 can be exponentially bounded. This serves as a crucial step towards the construction of regular infinite models that admit finite descriptions. The following is a truncated version of our small substructure property, compare Proposition 4 for a full statement. By $\mathfrak{A} \Rightarrow_{\forall\forall} \mathfrak{A}'$ we denote the transfer property that any prenex $\forall\forall$-formula satisfied in $\mathfrak{A}$ is also satisfied in $\mathfrak{A}'$; similarly for $\mathfrak{A} \Rightarrow_{\forall\exists} \mathfrak{A}'$.

Theorem 2 *Let* $\mathfrak{A}$ *be a* τ*-structure with universe* $A = B \,\dot\cup\, C$, $\mathfrak{B} := \mathfrak{A} \restriction B$. *Then there is a* τ*-structure* $\mathfrak{A}'$ *with universe* $A' = B' \,\dot\cup\, C$, *for some set* B' *of size exponential in* τ, *such that* $\mathfrak{A}' \restriction C = \mathfrak{A} \restriction C$, $\mathfrak{A} \Rightarrow_{\forall\forall} \mathfrak{A}'$ *and* $\mathfrak{A} \Rightarrow_{\forall\exists} \mathfrak{A}'$.

Small substructure properties of this kind may be interesting in their own right. In settings where some parts of a larger structure are controlled by specifications in a logic $\mathcal{L}$ in their relationship to that surrounding structure, it may be natural to consider variations of these parts while the surrounding structure stays unchanged. One may think of descriptions of distributed systems, in which only some "local" components are amenable to modifications while the overall or remote environment is regarded as fixed and not locally controllable. Then the question arises whether the local components can be replaced by smaller equivalent components, without violating the global specification and possibly preserving some logical features of the interface between the local components and their environment. Our small substructures property captures such a setting for FO^2 and yields exponential bounds on the size of crucial components, in a context where a global finite model is not available.

1. Preliminaries

Quantifier-free types and Scott normal form We use the term *type* to refer to quantifier-free types. Let $\mathfrak{A}$ be a τ-structure. For $a \in A$, the 1-type of a in $\mathfrak{A}$ is

$$\mathrm{tp}_{\mathfrak{A}}(a) = \{\varphi(x) \in \mathrm{FO}^2[\tau] \colon \varphi \text{ quantifier-free }, \mathfrak{A} \models \varphi[a]\}.$$

We let $\boldsymbol{\alpha}[\mathfrak{A}] = \{\mathrm{tp}_{\mathfrak{A}}(a) \colon a \in A\}$ be the set of all 1-types of $\mathfrak{A}$. Quantifier-free 2-types $\mathrm{tp}_{\mathfrak{A}}(a_1, a_2)$ of non-degenerate pairs, $a_1 \neq a_2$, are similarly defined. We let $\boldsymbol{\beta}[\mathfrak{A}] = \{\mathrm{tp}_{\mathfrak{A}}(a_1, a_2) \colon a_1 \neq a_2,\ a_1, a_2 \in A\}$ be the set of all 2-types of $\mathfrak{A}$.

We also write $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ for the sets of all 1-types and 2-types, across all $\mathfrak{A}$. We sometimes identify a type with a corresponding quantifier-free formula that determines it. Note that the size of the sets $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ is bounded by an exponential function in the size of the vocabulary.

We typically write $\alpha = \alpha(x)$ for a 1-type $\alpha \in \boldsymbol{\alpha}$, and $\beta = \beta(x, y)$ for a 2-type $\beta \in \boldsymbol{\beta}$. Then $\alpha(y)$ is the result of switching x for y in α. Also $\beta{\restriction}x$ is the 1-type consisting of all the $\varphi(x) \in \beta$; similarly for $\beta{\restriction}y$. However, we also write $\beta{\restriction}y = \alpha$ instead of the formally correct $\beta{\restriction}y = \alpha(y)$

Natural terminology with regard to types and their realisations applies. For instance, we say that an element $b \in \mathfrak{B}$ realises the type α if $\mathrm{tp}_{\mathfrak{B}}(b) = \alpha$ (or $\mathfrak{B} \models \alpha[b]$). For a type $\alpha \in \boldsymbol{\alpha}$, and a subset $S \subseteq A$ of a structure $\mathfrak{A}$ we denote as $\alpha[S] \subseteq A$ the set of all those $a \in S$ that have type α. Conversely $\boldsymbol{\alpha}[S]$, for a subset $S \subseteq \mathfrak{A}$, denotes the set of all 1-types realised in S, $\boldsymbol{\alpha}[S] = \{\mathrm{tp}_{\mathfrak{A}}(a) \colon a \in S\}$. Similarly $\boldsymbol{\beta}[S_1, S_2]$, for subsets $S_i \subseteq A$, is the set of all 2-types $\mathrm{tp}_{\mathfrak{A}}(a_1, a_2)$ with $a_i \in S_i$; $\boldsymbol{\beta}[a, S]$ for a subset $S \subseteq A$ is the set of all 2-types $\mathrm{tp}_{\mathfrak{A}}(a, a')$ with $a' \in S$, etc.

Proposition 3 (Scott normal form) *For $\varphi \in \mathrm{FO}^2[\tau]$ one can compute in polynomial time a* Scott normal form formula $\tilde{\varphi} \in \mathrm{FO}^2[\tilde{\tau}]$, *whose length is linear in the length of φ, of the form*

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

for quantifier-free formulae $\chi_i \in \mathrm{FO}^2[\tilde{\tau}]$, such that $\tilde{\varphi}$ is satisfiability equivalent with φ.

For normal form φ, whether or not $\mathfrak{A} \models \varphi$ is determined by the family of sets $\boldsymbol{\beta}[a, A] = \{\mathrm{tp}_{\mathfrak{A}}(a, b) \colon b \in A\}$ of types incident with a, for $a \in A$. Whether $\mathfrak{A}$ satisfies the $\forall\forall$ part of φ is determined by $\boldsymbol{\beta}[\mathfrak{A}] = \bigcup_{a \in A} \boldsymbol{\beta}[a, A]$; for the $\forall\exists$ parts, $\mathfrak{A}, a \models \exists \chi_i$ iff $\beta \models \chi_i$ for some $\beta \in \boldsymbol{\beta}[a, A]$.

2. A small substructure property for FO^2

Consider a structure $\mathfrak{A}$ and a fixed subset $B \subseteq A$ with induced substructure $\mathfrak{B} := \mathfrak{A}{\restriction}B$. We want to perform some surgery on $\mathfrak{A}$ which in effect replaces the substructure $\mathfrak{B}$ by some 'equivalent' $\mathfrak{B}'$ of bounded size.

Proposition 4 *Let $\mathfrak{A}$ be a τ-structure, $\mathfrak{B} = \mathfrak{A}{\restriction}B$ for some $B \subseteq A$, $C := A \setminus B$. Then there is a τ-structure $\mathfrak{A}'$ with universe $A' = B' \,\dot\cup\, C$ for some set B' of size polynomial in $|\boldsymbol{\beta}[\mathfrak{A}]|$ such that*

(i) $\mathfrak{A}'{\restriction}C = \mathfrak{A}{\restriction}C$.
(ii) $\boldsymbol{\alpha}[B'] = \boldsymbol{\alpha}[B]$, *whence* $\boldsymbol{\alpha}[\mathfrak{A}'] = \boldsymbol{\alpha}[\mathfrak{A}]$.
(iii) $\boldsymbol{\beta}[B'] = \boldsymbol{\beta}[B]$ *and* $\boldsymbol{\beta}[B', C] = \boldsymbol{\beta}[B, C]$, *whence* $\boldsymbol{\beta}[\mathfrak{A}'] = \boldsymbol{\beta}[\mathfrak{A}]$.
(iv) *for each $b' \in B'$ there is some $b \in B$ with $\boldsymbol{\beta}[b', A'] \supseteq \boldsymbol{\beta}[b, A]$.*
(v) *for each $a \in C$: $\boldsymbol{\beta}[a, B'] \supseteq \boldsymbol{\beta}[a, B]$.*

Note that (iii)–(v) imply that $\mathfrak{A} \models \varphi \ \Rightarrow \mathfrak{A}' \models \varphi$ for all normal form φ.

We point out that for plain FO^2 we reproduce the known small model property for FO^2 from [8]. Putting $B := A$ and observing that $\boldsymbol{\beta}[\mathfrak{A}] \subseteq \boldsymbol{\beta}$ is exponential in the vocabulary of (the normal form of) φ, we obtain an exponential size bound on $\mathfrak{A}' \models \varphi$. Also the proof of the proposition has crucial similarities with the proof by Grädel, Kolaitis, and Vardi for the small model property. However, our result is a proper extension, as will become apparent from its uses in settings where the surrounding structure cannot be made finite. In particular, our construction cannot be obtained by treating types over parameters in $C = A \setminus B$ as fixed.

Proof Let $m := |\boldsymbol{\beta}[\mathfrak{A}]|$ be the number of 2-types realised in $\mathfrak{A}$, and enumerate these as $\boldsymbol{\beta}[\mathfrak{A}] := \{\beta_1, \ldots, \beta_m\}$. Note that m is bounded by $|\boldsymbol{\beta}|$, which is exponential in $|\tau|$. W.l.o.g. assume that $B \subseteq \alpha[A]$ consists of elements of the same 1-type α. The general case may be reduced to this one, by repeated application to the sets $\alpha[B]$ for all relevant α. Also assume that $B = \alpha[B]$ contains at least two elements so that $\boldsymbol{\beta}[\mathfrak{B}] \neq \emptyset$. Otherwise put $B' := B$.

We want to find a suitable set B' of new realisations of α, and links between B' and C and within B' in accordance with (iii)–(v). Technically $\mathfrak{A}'$ is specified by a consistent choice of 2-types for any pair involving at least one new element $b' \in B'$. For $\alpha' \in \boldsymbol{\alpha}[\mathfrak{A}]$ let

$$\begin{aligned} \boldsymbol{\beta}(\alpha') &:= \{\mathrm{tp}_{\mathfrak{A}}(b, a) \colon b \in B, a \in \alpha'[C]\} \subseteq \boldsymbol{\beta}[\mathfrak{A}], \\ N(\alpha') &:= |\boldsymbol{\beta}(\alpha')|, \\ n(\alpha') &:= \min\big(N(\alpha'), |\alpha'[C]|\big). \end{aligned}$$

So $\boldsymbol{\beta}[\alpha']$ is the set of 2-types linking elements of B with elements of type α' outside. Note that $n(\alpha') \leqslant N(\alpha') \leqslant m$. For $\alpha' \in \boldsymbol{\alpha}[C]$ choose subsets $M(\alpha') \subseteq \alpha'[C]$ of size $n(\alpha')$. Let M be the union of the (disjoint) sets $M(\alpha')$. We replace B by

$$B' = \{0, 1, 2\} \times \{1, \ldots, m\} \times M,$$

consisting of $3m|M| \leqslant 3m^3$ elements that realise type α. It remains to allocate 2-types over $B' \times B'$ and $B' \times C$ in a consistent fashion and such that requirements (iii)–(v) are met. This is done in stages.

(1) We allocate 2-types for some pairs $(b', a) \in B' \times M$ to settle (v) for all $a \in M$. Fix the injection $f\colon (k,a) \mapsto (0,k,a)$ of $\{1,\dots,m\} \times M$ into B'. For each $a \in M(\alpha') \subseteq M$ and if $\beta_k \in \boldsymbol{\beta}(\alpha')$ is such that $\beta_k \restriction y = \alpha'$, put $\mathrm{tp}(f(k,a),a) := \beta_k$.

Note that for each $b' \in B'$, $\mathrm{tp}(b', a')$ has been set for at most one element a' (necessarily $a' \in M$), and that in this case $\mathrm{tp}(b', a') = \mathrm{tp}_{\mathfrak{A}}(b, a)$ for some $(b, a) \in B \times C$.

(2) For each $b' \in B'$, we allocate 2-types to some further pairs (b', a) with $a \in M$, in such a way as to guarantee (iv) at b'. We treat one b' at a time.

– If no 2-type involving b' has been determined in stage (1), pick some element $b \in B$ and for each 2-type $\beta \in \boldsymbol{\beta}[b, C]$ and $\alpha' = \beta{\restriction}y$, select a fresh element $a = a(b', \beta) \in M(\alpha')$ and put $\mathrm{tp}(b', a) := \beta$. There are sufficiently many elements in $M(\alpha')$ by the definition of $n(\alpha')$.

– If one 2-type involving b' (and an element of M) has been determined in stage (1), $\mathrm{tp}_{\mathfrak{A}'}(b', a') = \mathrm{tp}_{\mathfrak{A}}(b, a)$ for some reference elements $b \in B$ and $a \in C$. We realise further 2-types $\beta \in \boldsymbol{\beta}[b, C] \subseteq \bigcup_{\alpha'} \boldsymbol{\beta}(\alpha')$ at b' with partner elements $a \in M(\alpha')$ for the appropriate α'. By the choice of $M(\alpha')$ there are sufficiently many distinct target elements available in each $M(\alpha')$ (just as there were in $\alpha'[C]$ for the reference element b). This makes sure that $\boldsymbol{\beta}[b', C] \supseteq \boldsymbol{\beta}[b, C]$.

(3) Allocation of all remaining 2-types for pairs $(b', a) \in B' \times M$. Choose $b_0 \in B$. For each $b' \in B'$ and $a \in M$, whose 2-type has not been attributed, put $\mathrm{tp}(b', a) := \mathrm{tp}_{\mathfrak{A}}(b_0, a)$.

(4) Allocation of 2-types to pairs $(b', a) \in B' \times (C \setminus M)$. For each $a \in C \setminus M$ of type α', pick $a_0 \in M(\alpha')$ and set $\mathrm{tp}(b', a) := \mathrm{tp}(b', a_0)$ for all $b' \in B$. Together with (1) this settles (v) for all $a \in C$.

(5) Allocation of 2-types to pairs in $B' \times B'$. For $b_1' = (i, j, a) \in B'$ and $\beta_k \in \boldsymbol{\beta}[\mathfrak{B}]$ put $\mathrm{tp}((i,j,a),(i',k,a)) := \beta_k$ where $i' = (i+1) \bmod 3$. For any two distinct elements of B' whose type has not been allocated, put arbitrary $\beta \in \boldsymbol{\beta}[\mathfrak{B}]$. This settles (iv). □

3. Equivalence structures

Types and Scott normal form for equivalence structures When we are interested in models in $\mathcal{EQ}[\tau_0, \tau_{\text{eq}}]$, we only admit types that are realisable in these. In the presence of one or more equivalence relations we want to distinguish 2-types according to equivalences/non-equivalences between x and y. For $\tau_{\text{eq}} = \{E\}$, we distinguish $\boldsymbol{\beta}^+$ and $\boldsymbol{\beta}^-$ such that $\boldsymbol{\beta} = \boldsymbol{\beta}^+ \dot\cup \boldsymbol{\beta}^-$, where all $\beta \in \boldsymbol{\beta}^+$ contain the formula Exy, while those in $\boldsymbol{\beta}^-$ contain its negation. In the case of $\tau_{\text{eq}} = \{E_1, E_2\}$ we correspondingly distinguish four sets: $\boldsymbol{\beta}^{++}, \boldsymbol{\beta}^{+-}, \boldsymbol{\beta}^{-+}, \boldsymbol{\beta}^{--}$, such that for instance $\beta \in \boldsymbol{\beta}^{-+}$ iff $(\neg E_1 xy \wedge E_2 xy) \in \beta$.

We also use superscripts $+$ and $-$ to indicate for an individual quantifier-free formula which equivalences/non-equivalences it stipulates. For instance, with $\tau_{\text{eq}} = \{E_1, E_2\}$ and for a quantifier-free formula $\chi = \chi(x,y) \in \mathrm{FO}^2[\tau_0]$ we let

$$\chi^{+-}(x,y) := \chi(x,y) \wedge (\neg x = y \rightarrow (E_1 xy \wedge \neg E_2 xy)).$$

χ^+ and χ^- over $\tau_{\text{eq}} = \{E\}$ are similarly defined. This decomposition of formulae leads to the following variant of the Scott normal form adapted to structures in $\mathcal{EQ}[\tau_0; \tau_{\text{eq}}]$.

A $\mathcal{EQ}[\tau_0, E_1, E_2]$ Scott normal form sentence is of the form

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

for quantifier-free $\chi \in \mathrm{FO}[\tau_0; E_1, E_2]$ and $\chi_i \in \mathrm{FO}[\tau_0]$ and $s_i \in \{+,-\} \times \{+,-\}$. For $\tau_{\text{eq}} = \{E\}$, similarly, $\mathcal{EQ}[\tau_0, E]$ Scott normal form is $\forall x \forall y\ \chi_0 \wedge \bigwedge_{i=1}^{m} \forall x \exists y\ \chi_i^{s_i}$ for quantifier-free $\chi \in \mathrm{FO}[\tau_0; E]$ and $\chi_i \in \mathrm{FO}[\tau_0]$ and $s_i \in \{+,-\}$.

Proposition 5 (Scott normal form) *For $\varphi \in \mathrm{FO}^2[\tau_0 \cup \tau_{eq}]$ one computes in polynomial time a $\mathcal{EQ}[\tau; \tau_{eq}]$ Scott normal form sentence $\tilde\varphi \in \mathrm{FO}^2[\tilde\tau_0 \cup \tau_{eq}]$, whose length is linear in the length of φ such that φ and $\tilde\varphi$ are satisfiability equivalent over $\mathcal{EQ}[\tau; \tau_{eq}]$.*

4. One equivalence relation

We consider models in $\mathcal{EQ}[\tau_0; E]$, with $\tau_{\text{eq}} = \{E\}$. Structures in $\mathcal{EQ}[\tau_0; E]$ are of the form $\mathfrak{A} = (A, E^{\mathfrak{A}}, \ldots)$, $E^{\mathfrak{A}}$ an equivalence relation over A. If $B \subseteq A$ is an equivalence class w.r.t. $E^{\mathfrak{A}}$, we refer also to the substructure $\mathfrak{B} = \mathfrak{A}{\restriction}B$ as an equivalence class of $\mathfrak{A}$.

Note that if $\mathfrak{A} \models \varphi$, where φ is in normal form according to Proposition 3, then each equivalence class $\mathfrak{B} = \mathfrak{A} \restriction B$ is a model of the $\forall\forall$ constituent as well as of all the $\forall\exists$ constituents of type χ^+.

We shall find small models, whose size is exponential in terms of $|\tilde\tau_0|$ or in the length of φ, in two stages:

(*small classes*) replacing each individual equivalence class $\mathfrak{A}{\restriction}B$ in $\mathfrak{A}$ by a small (exponential size) structure, while retaining the remainder of $\mathfrak{A}$ unchanged; for this we apply the small substructure property, which preserves any normal form φ.

(*few classes*) building a new structure from an exponential number of isomorphic copies of classes from $\mathfrak{A}$, such that again any normal form φ is preserved.

As there are only exponentially many distinct 1-types, one can afford to realise exactly the same 1-types in the target structure as in the given model. At the level of equivalence classes, however, the given model may have doubly exponentially many types, distinguished by their composition in terms of 1-types. Out of these one needs to select a set of just exponentially many types to be used in the new small model.

Small classes Let $\mathfrak{A} \in \mathcal{EQ}[\tau_0; E]$. We replace a single equivalence class $\mathfrak{B}$ in $\mathfrak{A}$ by a new substructure $\mathfrak{B}'$ that also consists of a single E-class, of size exponential in $|\tau_0|$, and such that any normal form formula satisfied in $\mathfrak{A}$ is also satisfied in the resulting structure.

For a fixed class $\mathfrak{B}$, apply Proposition 4 to $\mathfrak{A}$ and $\mathfrak{B} \subseteq \mathfrak{A}$ to obtain $\mathfrak{A}'$ in which $\mathfrak{B}$ has been replaced by an exponential size $\mathfrak{B}'$ in such a way that in particular (cf. (iii) in the proposition)

$$\boldsymbol{\beta}[B'] = \boldsymbol{\beta}[B] \quad \text{and} \quad \boldsymbol{\beta}[B', A \setminus B'] = \boldsymbol{\beta}[B, A \setminus B].$$

$\boldsymbol{\beta}[B'] = \boldsymbol{\beta}[B] \subseteq \boldsymbol{\beta}^+$ and $\boldsymbol{\beta}[B', A \setminus B'] = \boldsymbol{\beta}[B, A \setminus B] \subseteq \boldsymbol{\beta}^-$ imply that $\mathfrak{B}' \subseteq \mathfrak{A}'$ also forms an equivalence class, and that $\mathfrak{A}' \in \mathcal{EQ}[\tau_0, E]$.

Repeated application of the process to all equivalence classes of a countable $\mathfrak{A} \models \varphi$ in $\mathcal{EQ}[\tau_0, E]$ for normal form φ yields a new model for φ in $\mathcal{EQ}[\tau_0, E]$ whose classes are all bounded by an exponential function in $|\tau_0|$ and $|\varphi|$.

Few classes Fix a model $\mathfrak{A} \in \mathcal{EQ}[\tau_0; E]$, whose equivalence classes are of exponential size in $|\tau_0|$. There may still be doubly exponentially many non-isomorphic classes (distinguished even by their composition in terms of 1-types $\boldsymbol{\alpha}[B]$).

Call an equivalence class B of $\mathfrak{A}$ *singular* if for some $\alpha, \alpha' \in \boldsymbol{\alpha}[\mathfrak{B}]$ there is no $\beta \in \boldsymbol{\beta}^-[\mathfrak{A}]$ for which $\beta \restriction x = \alpha$ and $\beta \restriction y = \alpha'$. Clearly any equivalence structure $\mathfrak{A}'$ with $\boldsymbol{\beta}[\mathfrak{A}'] = \boldsymbol{\beta}[\mathfrak{A}]$ can realise at most one class with this $\boldsymbol{\beta}[\mathfrak{B}]$.

Lemma 6 *The number of singular classes in any given $\mathfrak{A} \in \mathcal{EQ}[\tau_0; E]$ is bounded by $|\boldsymbol{\alpha}[\mathfrak{A}]| \leqslant |\boldsymbol{\alpha}|$, exponential in the size of the vocabulary.*

Proof For each singular class B of $\mathfrak{A}$ there is at least one $\alpha \in \boldsymbol{\alpha}[B]$ that is not realised in any other class B' of $\mathfrak{A}$. Picking one such $\alpha_B \in \boldsymbol{\alpha}[B]$ for each class B, we obtain an injection $B \mapsto \alpha_B$ of $A/E^{\mathfrak{A}}$ into $\boldsymbol{\alpha}[\mathfrak{A}]$. □

Let $\mathfrak{A}_s \subseteq \mathfrak{A}$ be the substructure formed by the union of the singular classes. Note that A_s is exponentially bounded. Let C be a union of equivalence classes of $\mathfrak{A}$ such that $A_s \subseteq C$ and for all $a \in A_s$ and $\beta \in \boldsymbol{\beta}^-[a, A]$ we have $\beta \in \boldsymbol{\beta}^-[a, C]$. As A_s and $\boldsymbol{\beta}^-[\mathfrak{A}]$ as well as all classes of $\mathfrak{A}$ are exponential, C can be chosen of exponential size. Let $\mathfrak{C} \subseteq \mathfrak{A}$ be the corresponding substructure of $\mathfrak{A}$, a union of full classes and containing in particular all singular classes.

Call $\alpha \in \boldsymbol{\alpha}[\mathfrak{A}]$ non-singular if $\alpha \in \boldsymbol{\alpha}[B]$ for some non-singular class B of $\mathfrak{A}$. Note that for any two non-singular α, α' (not necessarily distinct) there is some 2-type $\beta \in \boldsymbol{\beta}^-[\mathfrak{A}]$ compatible with $\alpha(x)$ and $\alpha'(y)$.

Let $\mathfrak{B} \subseteq \mathfrak{A}$ be a union of equivalence classes such that every non-singular α is realised by some b_α in $\mathfrak{B}$. Again $\mathfrak{B}$ can be chosen of exponential size. Observe that $\boldsymbol{\alpha}[\mathfrak{A}] = \boldsymbol{\alpha}[\mathfrak{C}] \cup \boldsymbol{\alpha}[\mathfrak{B}]$.

We construct $\mathfrak{A}'$ from the disjoint union of $\mathfrak{C}$ and an exponential number of isomorphic copies of $\mathfrak{B}$ in such a way that

(i) $\boldsymbol{\beta}[\mathfrak{A}'] \subseteq \boldsymbol{\beta}[\mathfrak{A}]$.
(ii) for each $c \in \mathfrak{C} \subseteq \mathfrak{A}'$, $\boldsymbol{\beta}[c, \mathfrak{A}'] \supseteq \boldsymbol{\beta}[c, \mathfrak{A}]$.
(iii) for each isomorphic copy $\mathfrak{B}' \subseteq \mathfrak{A}'$ of $\mathfrak{B}$, and for each $b \in \mathfrak{B}$, $\boldsymbol{\beta}[b', \mathfrak{A}'] \supseteq \boldsymbol{\beta}[b, \mathfrak{A}]$.

It follows that $\mathfrak{A}'$ satisfies any normal form sentence satisfied in $\mathfrak{A}$.

Let $\boldsymbol{\beta}^-[\mathfrak{A}] = \{\beta_1, \ldots, \beta_m\}$. We build $\mathfrak{A}'$ from

$$\mathfrak{C} \cup \{0, 1, 2\} \times \{1, \ldots, m\} \times \mathfrak{B}$$

by allocating 2-types $\beta \in \boldsymbol{\beta}^-[\mathfrak{A}]$ between any two elements form different parts. Again we proceed in several stages.

(1) For any $b \in \mathfrak{B}$ and $\beta \in \boldsymbol{\beta}^-[b, A]$ such that $\alpha' = \beta \restriction y$ is singular, put $\mathrm{tp}_{\mathfrak{A}'}((i, k, b), c) := \beta$ for all $c \in A_s$ with $\mathrm{tp}_{\mathfrak{A}}(b, c) = \beta$ (there are such, as α' is only realised in A_s in $\mathfrak{A}$).

(2) For $c \in C \setminus A_s$ and $\beta_k \in \boldsymbol{\beta}^-[c, A]$, if $\beta_k \notin \boldsymbol{\beta}^-[c, C]$ then $\alpha = \beta \restriction y$ must be non-singular and we put $\mathrm{tp}_{\mathfrak{A}'}(c, (0, k, b_\alpha)) := \beta_k$. This settles (ii).

(3) For $b \in B$ and $\beta_k \in \boldsymbol{\beta}^-[b, A]$, if $\beta_k \notin \boldsymbol{\beta}^-[b, A_s]$ then $\alpha = \beta \restriction y$ must be non-singular and we put $\mathrm{tp}_{\mathfrak{A}'}((i, j, b), (i', k, b_\alpha)) := \beta_k$ for $i' = (i + 1) \bmod 3$. This settles (iii).

(4) For all remaining pairs of undeclared 2-type find $\beta \in \boldsymbol{\beta}[\mathfrak{A}]$ compatible with the given 1-types. We check that this is possible: pairs within $\mathfrak{C}$ are settled; pairs of one element from $\mathfrak{A}_s$ and one from an isomorphic copy of $\mathfrak{B}$ may be given the 2-type of a corresponding pair in $\mathfrak{A}$; all remaining pairs involve two non-singular 1-types and always admit a matching 2-type from $\boldsymbol{\beta}^-[\mathfrak{A}]$.

Corollary 7 *There is an exponential function f such that for any normal form sentence $\varphi \in \mathrm{FO}^2[\tau_0; E]$:*

If φ is satisfiable in $\mathcal{EQ}[\tau_0; E]$ then it has a model in $\mathcal{EQ}[\tau_0; E]$ of size bounded by $f(|\tau_0|)$.

It follows that $\mathrm{SAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E])$ *is in* NEXPTIME, *hence* NEXPTIME*-complete.*

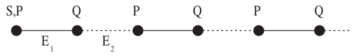

Figure 1. A model of φ.

5. Two equivalence relations

We consider models in $\mathcal{EQ}[\tau_0; E_1, E_2]$, with $\tau_{\text{eq}} = \{E_1, E_2\}$ consisting of two binary relations which have to be interpreted as equivalence relations. We let E stand for the equivalence relation $E = E_1 \cap E_2$ which is the common refinement of E_1 and E_2. The equivalence classes of E are called *intersections*. If B is an equivalence class of E_1 or E_2 in $\mathfrak{A}$, and I, H intersections in $\mathfrak{A}$, we always let $\mathfrak{B}$, $\mathfrak{I}$, $\mathfrak{H}$ denote the induced substructures of $\mathfrak{A}$.

5.1. Failure of the finite model property

We show that over $\mathcal{EQ}[\tau_0; E_1, E_2]$, FO^2 does not have the finite model property.

We firstly exhibit an infinity axiom that requires an infinite number of E_1-classes and of E_2-classes. Let τ_0 consist of three unary predicates P, Q, S, and let $\varphi \in \mathrm{FO}^2[\tau_0; E_1, E_2]$ say that

(1) P and Q are disjoint and each E_2-class contains at most one element from P and one from Q; similarly for E_1-classes. The E_2-class of any element of S is trivial (a singleton).

(2) every element of P is E_1 equivalent to one in Q; every element of Q is E_2 equivalent to one in P.

(3) $S \cap P \neq \emptyset$.

It is easy to formalise (1),(2),(3) in FO^2 over $\mathcal{EQ}[P, Q, S; E_1, E_2]$, and in fact even in the guarded fragment with two variables. Figure 1 shows a model of this sentence φ.

Conversely, any model of φ embeds this infinite chain. Starting from an element of $S \cap P$, one finds new elements in Q and P along links in E_1 and E_2, respectively, in an alternating fashion, by appeal to condition (2); these will indeed always have to be new elements, i.e. distinct from previous position in the chain because of (1).

A simple modification φ' of the previous sentence φ even enforces the existence of an infinite E_1-class. For a new unary predicate R, we can say that

(4) each E_2-class contains an element from R and all of R is contained in a single E_1-class.

One checks that this, together with φ, forces the chain to have links to an infinite sequence of distinct elements in R, all contained in an infinite E_1-class.

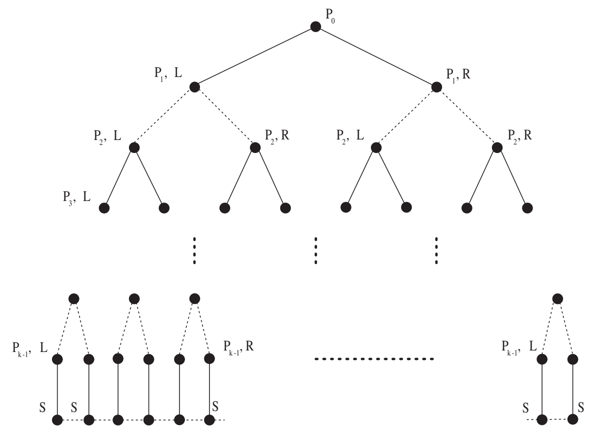

Figure 2. A large finite class.

5.2. Large finite classes

The tree-like structure in Figure 2 contains an E_2-class S with exactly 2^{k-1} elements of the same type along its leaves. L, R and S are unary predicates of τ_0. The predicates P_i serve to number consecutive levels of the tree-like structure; they will be definable in terms of other basic predicates such that the number of levels k can exceed the size of $|\tau_0|$.

In the first instance, we describe how to axiomatise the given structure for $k = 2^n$ by a formula φ whose length is linear in n. For this we code the values i for the levels P_i in binary by means of n extra unary predicates $Q_0, \dots, Q_{n-1}$. Let φ say that

(1) there is exactly one element in P_0.

(2) (top-down requirements) every $a \in P_{2i}$, $2i < k-1$, has E_1-links to two elements of P_{2i+1}, one in L and one in R; similarly for $a \in P_{2i+1}$ and E_2-links to two successors in P_{2i+2}. Moreover, we want each element to link to two fresh elements on the next level, which may be enforced as in the sentence discussed in section 5.1.

(3) (bottom-up requirements) that, dual to (2), every $a \in P_{i+1}$ is appropriately linked to some element on level P_i.

(4) the E_1-class of every $a \in P_{k-1}$ contains just a and one element in S distinct from a, and every element of S is linked by E_1 to some element on level P_k.

(5) all of S is contained in a single E_2-class.

The construction may be extended to cover $k = 2^{2^n}$. For this consider the black circles from Fig. 2 as intersections, each intersection consisting of 2^n elements numbered by means of another set of unary predicates $U_1, \dots, U_{n-1}$. Values of P_i are encoded (in binary) by another unary predicate V. An element, whose U-value equals ℓ, is put into V if, and only if, the ℓ-th bit of i is 1. It is then possible to express, in a linear size FO^2 formula, the condition that a pair of intersections within the same E_i-class encodes a pair of consecutive values in $\{0, \ldots, 2^{2^n} - 1\}$.

It is worth mentioning that a similar construction may be used to encode computations of an exponential space alternating Turing machine. This yields a 2EXPTIME-lower bound for the satisfiability problem for FO^2 over $\mathcal{EQ}[\tau_0; E_1, E_2]$, and in fact even for the two-variable guarded fragment without equality. This result will be presented in another paper [14].

5.3. Small intersections

Let $\mathfrak{A} \in \mathcal{EQ}[\tau_0; E_1, E_2]$ be an equivalence structure. From the previous subsections we now that, in contrast to the case of $\mathcal{EQ}[\tau_0; E]$, if we want to preserve (normal form) FO^2 sentences we cannot replace $\mathfrak{A}$ by a structure with finite classes. However, we can transform $\mathfrak{A}$ into a structure with small intersections.

We replace a single intersection $\mathfrak{I}$ in $\mathfrak{A}$ by a new substructure $\mathfrak{I}'$ of size exponential in $|\tau_0|$. This is done in such a way that the E_i-links to the full E_i-classes of $\mathfrak{I}$ are reproduced by $\mathfrak{I}'$, which will then again be the intersection of these classes. We also preserve any normal form formula in this modification.

For a fixed intersection $\mathfrak{I}$ extend τ_0 by new unary symbols U_1, U_2 for the E_1- and E_2-classes of $\mathfrak{I}$, and expand $\mathfrak{A}$ accordingly. In other words put $U_i^{\mathfrak{A}} := B_i$ where B_i is the E_i-class of $\mathfrak{I}$ in $\mathfrak{A}$. We now apply Proposition 4 to this expansion $\hat{\mathfrak{A}} = (\mathfrak{A}, U_1^{\mathfrak{A}}, U_2^{\mathfrak{A}})$ and to the induced substructure $\hat{\mathfrak{I}} = \hat{\mathfrak{A}} \restriction I$, which is the trivial expansion of $\mathfrak{I}$ with $U_i^{\mathfrak{I}} = I$. By Proposition 4 we obtain a new structure $\hat{\mathfrak{A}}' = (\mathfrak{A}', U_1^{\mathfrak{A}'}, U_2^{\mathfrak{A}'})$ by replacing $\hat{\mathfrak{I}}$ in $\hat{\mathfrak{A}}$ by an exponential size $\hat{\mathfrak{I}}'$ in such a way that in particular (cf. (iii) in the proposition) $\boldsymbol{\beta}[I'] = \boldsymbol{\beta}[I]$ and $\boldsymbol{\beta}[I', A \setminus I'] = \boldsymbol{\beta}[I, A \setminus I]$.

Using the U_i, we see that this implies $\boldsymbol{\beta}[I'] \subseteq \boldsymbol{\beta}^{++}$, $\boldsymbol{\beta}[I', B_1 \setminus I'] \subseteq \boldsymbol{\beta}^{+-}$, $\boldsymbol{\beta}[I', B_2 \setminus I'] \subseteq \boldsymbol{\beta}^{-+}$, and $\boldsymbol{\beta}[I', A \setminus (B_1 \cup B_2)] \subseteq \boldsymbol{\beta}^{--}$. This guarantees that $\mathfrak{A}' \in \mathcal{EQ}[\tau_0; E_1, E_2]$ and that $\mathfrak{I}'$ is an intersection of $\mathfrak{A}'$.

We apply the above process to all intersections of a countable $\mathfrak{A} \models \varphi$ in $\mathcal{EQ}[\tau_0, E_1, E_2]$. For normal form φ this yields a new model for φ in $\mathcal{EQ}[\tau_0; E_1, E_2]$ whose intersections are all exponential in $|\tau_0|$.

5.4. Regular models

Fix $\mathfrak{A} \in \mathcal{EQ}[\tau_0, E_1, E_2]$ whose intersections are of exponential size in $|\tau_0|$. We build a new regular model $\mathfrak{A}'$ such that for every normal form sentence φ: $\mathfrak{A} \models \varphi \Rightarrow \mathfrak{A}' \models \varphi$. For the rest of this section we let $m = |\boldsymbol{\beta}[\mathfrak{A}]| \leqslant \boldsymbol{\beta}$, exponential in $|\tau_0|$.

The construction of $\mathfrak{A}'$ works at the level of intersections, rather than elements. In particular when we put a new element a' to $\mathfrak{A}'$ we always add an entire intersection, which will always be an isomorphic copy of an intersection from $\mathfrak{A}$. Thus appropriate $\boldsymbol{\beta}^{++}$ types will automatically be imported.

We use *isomorphism types* of intersections from $\mathfrak{A}$. Let Δ be the set of all isomorphism types of intersections in $\mathfrak{A}$. Since $\mathfrak{A}$ has exponential size intersections, $|\Delta|$ is at most doubly exponential in $|\varphi|$.

Two intersection $\mathfrak{I}_1$ and $\mathfrak{I}_2$ within an $\mathcal{EQ}$ structure are in a *free relation* if $\boldsymbol{\beta}[I_1, I_2] \subseteq \boldsymbol{\beta}^{--}$, i.e., if they are not part of the same E_i-class for either $i = 1$ or $i = 2$.

Plan of the construction As the first step we distinguish some finite set of *special classes* of $\mathfrak{A}$ which contain all the intersections whose types may be realised possibly only in a finite number of classes in $\mathfrak{A}'$. This includes the equivalence classes of those intersections whose types are realised exactly once in $\mathfrak{A}$, as well as equivalence classes containing all realizations of some $\delta \in \Delta$.

In the transition from $\mathfrak{A}$ to $\mathfrak{A}'$ we replace equivalence classes by regular versions which – owing to a regular pattern of connections between the constituent intersections – admit finite descriptions. Each such regular class will itself consist of a finite initial part whose isomorphism type is explicitly described, and of a possibly infinite part of a repetitive (regular) nature.[1]

The overall construction of the new model in organised in a countable sequence of levels. We start with level L_0 which comprises in particular a finite substructure of $\mathfrak{A}$ which covers (at least) the finite initial parts of all special classes.[2] Further levels $L_1, L_2, \ldots$ are governed by a finite supply of *patterns* for the extension of intersections created at level L_i to complete their E_1- and E_2-classes as well as their $\boldsymbol{\beta}^{--}$ requirements within level L_{i+1} by the creation of new intersections. Some such new intersections may, however, have to be joined to existing special E_1- or E_2-classes, as such classes cannot be replicated on demand. This constraint is responsible for the complicated nature of the overall construction, as well as being the source of infinite classes in the new model.

Special classes We distinguish a set S of equivalence classes and a set $\Delta_0 \subseteq \Delta$ of isomorphism types of intersections in $\mathfrak{A}$.

If all realisations of δ are contained in a single class B, then put $B \in S$ and $\delta \in \Delta_0$. We saturate S and Δ_0 by the following iteration, until no further extension occurs:

[1] Our construction reflects some key elements of the model construction in [8], at several levels: here the parallel is with the royal part (kings and court) versus the regular remainder, inside an individual equivalence class – and we have to make provisions for classes to be infinite, albeit in a regular fashion.

[2] Corresponding to the royal part of the overall model, if we draw a parallel with [8].

(1) If for $\delta \in \Delta_0$ there exists I of type δ such that I belongs to an E_1-class (resp. E_2-class) $B_1 \in S$ and I is the only intersection of type δ in B_1 such that its E_2-class (resp. E_1-class) B_2 does not belong to S, then let $S := S \cup \{B_2\}$.

(2) If for some $\delta \notin \Delta_0$ all realisations of δ are contained in classes in S, then let $\Delta_0 := \Delta_0 \cup \{\delta\}$.

The members of S are called *special classes*, members of Δ_0 are *special types* and realizations of special types are *special intersections*.

Lemma 8 *The size of S is at most exponential in $|\Delta|$ (triply exponential in $|\tau_0|$).*

Extended types Because some intersections have to belong to special classes we use 'extended' types to identify those classes. The *extended type* of an intersection I in $\mathfrak{A}$ is $\bar{\delta}(I) := (\delta(I), s(I))$ where $\delta(I) \in \Delta$ is the isomorphism type of I and $s(I) \subseteq S$ is a set of special classes in which I is contained: $s(I) = \{B \in S : I \subseteq B\}$, a set of at most two elements. $\bar{\delta} = (\delta, s)$ is an extension of δ, $\bar{\Delta}$ denotes the set of all extended types realised in $\mathfrak{A}$.

We continue to refer to isomorphism types of intersections as *types*; and explicitly refer to *extended types* where these become relevant.

Regular classes The equivalence classes of $\mathfrak{A}'$ are regular versions of classes in $\mathfrak{A}$. We describe the process of constructing a regular counterpart B' of a class B of $\mathfrak{A}$, relative to some tuple $(I_1, \ldots, I_k)$ of distinguished intersections contained in B which are treated as parameters. We explicitly consider the case of an E_1-class B. Let $I_{k+1}, \ldots, I_l$ be all the intersections in $B \setminus (I_1 \cup \ldots \cup I_k)$ whose types are realised exactly once in $\mathfrak{B}$. For all the extended types realised at least twice in B, say $\bar{\delta}_1, \ldots, \bar{\delta}_r$, choose intersections $H_1, \ldots H_r$ in B of the respective extended type.

Some initial part of the regular version $\mathfrak{B}'$ of $\mathfrak{B}$ is chosen as a finite substructure $\mathfrak{C} \subseteq \mathfrak{B}$ as follows. Let $C \subseteq B$ consist of $I_1 \cup \ldots \cup I_l$ together with a (minimal) set of intersections such that for every $a \in I_i$, $1 \leq i \leq k$, and $\beta \in \boldsymbol{\beta}^{+-}[a, B \setminus (I_1 \cup \ldots \cup I_l)]$ there exists $b \in I \subseteq C$ such that $\mathrm{tp}_{\mathfrak{B}}(a, b) = \beta$.

$$B' := C \cup \bigcup_{p=1}^{r} \{0, 1, 2\} \times \{1, \ldots, m, m+1\} \times H_p,$$

where $m = |\boldsymbol{\beta}[\mathfrak{A}]|$. We let the the isomorphism type of $\mathfrak{B}' \restriction (\{(i, j)\} \times H_p$ be δ_p. Let

$$D_i = \bigcup_{p=1}^{r} \{i\} \times \{1, \ldots, m, m+1\} \times H_p \subseteq B'.$$

We allocate the remaining 2-types inside B', in cyclic fashion following the idea in [8], as follows. For $a \in C \setminus (I_1 \cup \ldots \cup I_l)$ and $\beta \in \boldsymbol{\beta}^{+-}[a, B \setminus (I_1 \cup \ldots \cup I_l)]$ select some b in D_0 and set $\mathrm{tp}_{\mathfrak{B}'}(a, b) := \beta$.

For $a' \in D_i$ select an element $a \in B$ of the same 1-type. For $\beta \in \boldsymbol{\beta}^{+-}[a, B]$ find an element b' of the appropriate 1-type in $I_1, \ldots, I_k$ or in $D_{i'}$ for $i' = (i+1) \bmod 3$, to put $\mathrm{tp}_{\mathfrak{B}'}(a', b') := \beta$.

The above procedure may be performed without conflicts: every isomorphism type of an intersection and thus also every 1-type of elements in D_i is realised at least $m+1$ times and $|\boldsymbol{\beta}^{+-}[\mathfrak{B}]| \leq |\boldsymbol{\beta}[\mathfrak{A}]| = m$. The remaining 2-types within B' may be set in any way consistent with $\boldsymbol{\beta}^{+-}[\mathfrak{B}]$.

The replacement $\mathfrak{B}'$ for $\mathfrak{B}$ will be used exactly as just constructed in the case of non-special classes $\mathfrak{B}$. In the case of special $\mathfrak{B}$, however, this $\mathfrak{B}'$ will only form the finite initial part produced in level L_0 of the full E_1-class, to which further special intersections may be added as required in the generation of higher levels. Such further extensions are unproblematic as they only occur when an intersection I of extended type $\bar{\delta} = (\delta(I), s(I))$ is looking for its E_1-class, in cases where $\bar{\delta} = \bar{\delta}_i$ for some $1 \leq i \leq r$ and $B \in s(I)$. Further copies of an intersection of this extended type can always be added to $\mathfrak{B}'$ without conflict.

Level L_0 We distinguish a finite substructure $\mathfrak{C}$ of $\mathfrak{A}$ that will be copied identically into $\mathfrak{A}'$. Let K be the set of those intersections of $\mathfrak{A}$ whose E_1- and E_2-classes are both special. The universe C of $\mathfrak{C}$ consists of the union of K with some (minimal) choice of intersections such that for every $a \in K$ and $\beta \in \boldsymbol{\beta}^{--}[a, \mathfrak{A}]$, there is some $b \in C$ with $\mathrm{tp}_{\mathfrak{A}}(a, b) = \beta$. We let $\mathfrak{A}' \restriction C := \mathfrak{A} \restriction C$.

For special classes B, we also put in L_0 its finite (initial) regular version B'. For this we apply the general construction method for regular classes, where all the intersections in $B \cap C$ are treated as distinguished parameters. (Remember that some special intersections may be added to B' later and that B' may become infinite.)

Let L_0 be the union of C and all the initial regular versions of special classes. L_0 is finite, of exponential size in $|\Delta|$.

Patterns For every $\bar{\delta} \in \bar{\Delta}$ we define sets of substructures $\mathbb{P}^{+-}(\bar{\delta})$, $\mathbb{P}^{-+}(\bar{\delta})$, $\mathbb{P}^{--}(\bar{\delta})$. Their purpose is to say how an intersection of extended type $\bar{\delta}$ is to be extended to its E_1-class, E_2-class, and how to realise its $\boldsymbol{\beta}^{--}$ requirements. $\mathbb{P}^{+-}(\bar{\delta})$ may be empty if every E_1-class containing a realisation of $\bar{\delta}$ is special; otherwise $\mathbb{P}^{+-}(\bar{\delta})$ consists of the regular version $\mathfrak{B}'$ of one non-special E_1-class $\mathfrak{B}$ in which a realisation $\mathfrak{I} \subseteq \mathfrak{B}$ of the extended type $\bar{\delta}$ is distinguished. Similarly for $\mathbb{P}^{-+}(\bar{\delta})$ w.r.t. E_2-classes. $\mathbb{P}^{--}(\bar{\delta})$ is a finite set of substructures of the form $\mathfrak{A} \restriction (I \cup I')$ for $\mathfrak{I}$ of extended type $\bar{\delta}$ and $\mathfrak{I}'$ in free relation to $\mathfrak{I}$.

We find these pattern sets as follows, for $\bar{\delta} \in \bar{\Delta}$.

$\mathbb{P}^{+-}(\bar{\delta})$: if there is a realization I of $\bar{\delta}$ in $\mathfrak{A}$ such that the E_1-class B of I is non-special, choose one such I and let $\mathbb{P}^{+-}(\bar{\delta})$ consist of the regular version $\mathfrak{B}'$ of $\mathfrak{B}$ with distinguished intersection I. Retain in $\mathfrak{B}'$ the extended types of special intersections and of I; for non-special intersections other than I, put $s := \emptyset$.

$\mathbb{P}^{-+}(\bar{\delta})$ is analogously obtained.

$\mathbb{P}^{--}(\bar{\delta})$: choose any realisation I of $\bar{\delta}$ and select (a minimal number of) intersections $I_1, \dots, I_k$ such that for every $a \in I$ and $\beta \in \boldsymbol{\beta}^{--}[a, \mathfrak{A}]$ there is some $b \in I_1 \cup \dots \cup I_k$ with $\mathrm{tp}_{\mathfrak{A}}(a, b) = \beta$. Let $\mathbb{P}^{--}(\bar{\delta}) := \{\mathfrak{A}\restriction(I \cup I_i) \colon 1 \le i \le k\}$, each substructure with distinguished I. Retain the extended types for I and for special intersections I_i; for non-special I_i put $s := \emptyset$.

Successor levels The construction of levels L_{k+1} is recursive. As the procedure for members of $C \subseteq L_0$ is slightly different, we treat them separately.

Let $I \subseteq C \subseteq L_0$. If I belongs to K then the initial regular parts of both its classes have been defined in L_0 and all $\boldsymbol{\beta}^{--}$ requirements are met in L_0.

Consider $I \subseteq C \setminus K$.

(1) $\boldsymbol{\beta}^{--}$ requirements: let $I_1, \dots I_k$ be some (minimal) choice of intersections of $\mathfrak{A}$ such that for every $a \in I$ and $\beta \in \boldsymbol{\beta}^{--}[a, \mathfrak{A}]$ there is some $b \in I_1 \cup \dots \cup I_k$ with $\mathrm{tp}_{\mathfrak{A}}(a, b) = \beta$.

If I_i of extended type $\bar{\delta} = (\delta(I_i), s(I_i))$ is not a member of C but is special, then $B \in s(I_i)$ for some special class B. Let I_i' be one of the intersections of type δ in the set D_0 in the regular version of the special class B.

If I_i of type $\bar{\delta}$ is not a special intersection then put a new intersection I_i' to level L_1 and define its type as δ and extended type as $(\delta, \emptyset)$. Make $\mathfrak{A}'\restriction(I \cup I_i') \simeq \mathfrak{A}\restriction(I \cup I_i)$.

(2) Extending I to its classes ($\boldsymbol{\beta}^{+-}$ and $\boldsymbol{\beta}^{-+}$ requirements): if the E_1-class of I has not been completed, then let B be the E_1-class of I in $\mathfrak{A}$ and $I_1, \dots, I_k$ the intersections in $C \cap B$. Let $\mathfrak{B}'$ be the regular version of $\mathfrak{B}$ with distinguished intersections $I, I_1, \dots, I_k$. We retain extended types of special intersections in $\mathfrak{B}'$, but put $s := \emptyset$ for non-special intersections. All new intersections from B' are put into L_1. If for one of those, say I', its extended type has $B_2 \in s$ for a special E_2-class B_2 then join I' to $\mathfrak{B}_2'$ as discussed above.

If the E_2-class of I has not been completed, proceed in the same way.

We turn to the generic successor step from L_k to L_{k+1}, and consider I of extended type $\bar{\delta}$ in L_k, not in C. To fulfil the $\boldsymbol{\beta}^{--}$ requirements of I we use pattern $\mathbb{P}^{--}(\bar{\delta})$, for its extensions to its E_1- or E_2-classes patterns $\mathbb{P}^{+-}(\bar{\delta})$ or $\mathbb{P}^{-+}(\bar{\delta})$, respectively.

(1) $\boldsymbol{\beta}^{--}$ requirements. For every $\mathfrak{A}\restriction(I \cup I') \in \mathbb{P}^{--}(\bar{\delta})$, we want to locate some intersection I'' of the same extended type as I' and connect it to I isomorphic to the given pattern.

If $s(I') = \emptyset$, which means that I' is non-special, we just let $I'' \in L_{k+1}$ be a new intersection of the same type as I'. If $s(I') = \{B_1, B_2\}$, then $I' \subseteq K$, and we let $I'' = I'$. If $s(I') = \{B\}$, depending on whether I itself belongs to the D_i-part of some special class or not, either find an intersection I'' of the extended type of I' in the $D_{i'}$-part, for $i' = i + 1 \bmod 3$, or in the D_0-part of the special class B, such that I'' is in free relation with I. In each case $\mathfrak{A}'\restriction(I \cup I'')$ can be made isomorphic to the given pattern. (Observe that this can be done without conflicts; for instance, even in the case where I belongs to an E_2-class which is the extension of an intersection I_1 in the special E_1-class B belonging to the set D_0 of B – so that I_1 cannot be used – there are m more copies in D_0 that may be used, all of them in a free relationship with I.)

(2) Extending I to its classes ($\boldsymbol{\beta}^{+-}$ and $\boldsymbol{\beta}^{-+}$ requirements). If the E_1-class of I has not been completed, then there is some non-special E_1-class B in $\mathbb{P}^{+-}(\bar{\delta})$ with a distinguished intersection of the extended type of I. We extend I to an isomorphic copy of B by identifying I with this distinguished intersection of B and putting all other intersections of this new class into L_{k+1}. For intersections I' in B of non-trivial extended type, i.e., with an E_2-class $B_2 \in s(I')$, join I' to the special E_2-class $\mathfrak{B}_2'$.

The extension to a complete E_2-class is handled in the same fashion, using $\mathbb{P}^{-+}(\bar{\delta})$.

Nothing bad happens By a careful analysis of cases the following can be proved.

Lemma 9 *In our regular model $\mathfrak{A}'$, if intersection I_2 is made to satisfy a $\boldsymbol{\beta}^{--}$ requirement of intersection I_1, then:*

(i) *I_1 and I_2 are in a free relation in $\mathfrak{A}'$.*

(ii) *only for I_1 and I_2 both in C could simultaneously I_1 be necessary to satisfy a $\boldsymbol{\beta}^{--}$ requirement of I_2.*

Remaining 2-types We define a set $\mathbb{T}$ of patterns for free relations between pairs of intersections. For every pair $\delta_1, \delta_2 \in \Delta$ if they are realised in $\mathfrak{A}$ in a free relation choose such a pair of realizations I_1, I_2 and add the substructure $\mathfrak{A}\restriction(I_1 \cup I_2)$ to $\mathbb{T}$. We set all non-specified links between intersections of $\mathfrak{A}'$ according to $\mathbb{T}$. The following lemma shows that this is always possible.

Lemma 10 *If δ_1, $\delta_2 \in \Delta$ are realised in $\mathfrak{A}'$ by intersections I_1', I_2' in free relation, then there are I_1, I_2 realising δ_1, δ_2 in $\mathfrak{A}$ also in free relation.*

Finite descriptions We introduce the notion of a *certificate* for satisfiability. A *valid* certificate for a normal form sentence φ is a finite description of a regular $\mathcal{EQ}[\tau_0, E_1, E_2]$-model of φ. The above extraction of a regular model from any given model of φ shows that a satisfiable φ has a valid certificate of triply exponential size. Conversely, from a valid certificate a model for φ can be reconstructed.

A certificate has the following components:

(1) finite sets X_1, X_2 of (names for) special E_1-classes, and E_2-classes, respectively;

(2) a set $\bar{\Delta} = \{(\delta_1, s_1), \dots (\delta_l, s_l)\}$ of extended types of intersections w.r.t. $X_1 \cup X_2$;

(3) a finite $\mathcal{EQ}[\tau_0 \cup X_1 \cup X_2, E_1, E_2]$ structure $\mathfrak{D}$;

(4) for every $\bar{\delta} = (\delta, s) \in \bar{\Delta}$ for which $|s| \leq 1$ a pattern set $\mathbb{P}^{--}(\bar{\delta})$ of structures, and appropriate pattern sets: $\mathbb{P}^{+-}(\bar{\delta})$ containing a finite E_1-class, or $\mathbb{P}^{-+}(\bar{\delta})$ containing a finite E_2-class, or both $\mathbb{P}^{+-}(\bar{\delta})$ and $\mathbb{P}^{-+}(\bar{\delta})$;

(5) a set $\mathbb{T}$ of structures of the form $\mathfrak{I} \cup \mathfrak{I}'$.

Intersections within structures appearing in (3) and (4) are coloured with $X_1 \cup X_2$ saying to which special classes they should belong.

From a regular model $\mathfrak{A}'$ as constructed above, one obtains a valid certificate as follows: X_i is the set of special E_i-classes of $\mathfrak{A}'$; $\bar{\Delta}$ the set of extended types realized in $\mathfrak{A}'$; $\mathfrak{D}$ the substructure of $\mathfrak{A}'$ comprising $\mathfrak{C}$, full non-special E_i-classes of all intersections in $\mathfrak{C}$, intersections fulfilling $\boldsymbol{\beta}^{--}$ requirements of elements from C, and the finite, initial parts of all special classes; $\mathbb{P}^{+-}(\bar{\delta})$, $\mathbb{P}^{-+}(\bar{\delta})$, $\mathbb{P}^{--}(\bar{\delta})$ and $\mathbb{T}$ are from the construction of $\mathfrak{A}'$.

Conversely, a certificate is a blueprint for the construction of a regular equivalence structure, in analogy with the construction of $\mathfrak{A}'$ outlined above.

A certificate is *valid* for the normal form sentence φ if the structure thus obtained is a model of φ – a condition which can be checked in terms of the 2-types stipulated in the certificate. We thus obtain the following.

Corollary 11 *The satisfiability problem for* FO^2 *over* $\mathcal{EQ}[\tau_0, E_1, E_2]$ *is in* 3NEXPTIME.

6. Three equivalence relations

It is shown in [9] that FO^2 is undecidable over equivalence structures with four equivalence relations. We sharpen this result by reducing the number of equivalence relations to three, thus completing the picture. In fact we obtain the stronger result below. Its proof uses a special adaptation of a reduction of the domino tiling problem [3] particularly suited for FO^2, as presented in [16] based on work in [11].

Proposition 12 *The two-variable guarded fragment without equality over* $\mathcal{EQ}[\tau_0; E_1, E_2, E_3]$ *forms a conservative reduction class, even for* τ_0 *consisting of just unary predicates.*

For more on the guarded fragment in restricted classes of models see [5, 19, 13, 14].

References

[1] H. Andréka, J. van Benthem, I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic* 27:217-274, 1998.

[2] F. Baader et al. (eds). *The Description Logic Handbook*. Cambridge University Press, 2003.

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

[4] A. Borgida. On the relative expressive power of description logics and predicate logic. *Artificial Intelligence*, 82:353-367, 1996.

[5] H. Ganzinger, C. Meyer, M. Veanes. The Two-Variable Guarded Fragment with Transitive Relations, *LICS*, pages 24-34, 1999.

[6] W. Goldfarb. The unsolvability of the Gödel's class with identity. *Journal of Symbolic Logic*, 49:1237-1252, 1984.

[7] E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64:1719-1742, 1999.

[8] E. Grädel, P. Kolaitis, M. Vardi. On the decision problem for two-variable first order logic. *Bulletin of Symbolic Logic*, 3:53-96, 1997.

[9] E. Grädel, M. Otto. On logics with two variables, *Theoretical Computer Science*, 224:77-113, 1999.

[10] E. Grädel, M. Otto, E. Rosen. Two-variable logic with counting is decidable. *LICS*, pages 306-317, 1997.

[11] E. Grädel, M. Otto, E. Rosen. Undecidability results on two-variable first-order logic. *Archive of Mathematical Logic*, 38:313-354, 1999.

[12] N. Immerman, E. Lander. Describing graphs: a first-order approach to graph canonization. in: A. Selman (ed), *Complexity Theory Retrospective*, Springer, 59-81, 1990.

[13] E. Kieroński. The two-variable guarded fragment with transitive guards is 2EXPTIME-hard. *FOSSACS*, pages 299-312, 2003.

[14] E. Kieroński. Results on the guarded fragment with transitive or equivalence relations. Preprint, 2005.

[15] M. Mortimer. On languages with two variables. *Zeitschr. f. Logik und Grundlagen d. Mathematik*, 21:135-140, 1975.

[16] M. Otto. Two variable first-order logic over ordered domains. *Journal of Symbolic Logic*, 66:685-702, 2001.

[17] L. Pacholski, W. Szwast, L. Tendera. Complexity of two-variable logic with counting. *LICS*, pages 318-327, 1997.

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

[19] W. Szwast, L. Tendera. The Guarded Fragment with Transitive Guards. *Annals of Pure and Applied Logic*, 128:227-276, 2004.