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

Emanuel Kieronski* Martin Otto†

March 4, 2009

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, 8]. 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

*Institute of Computer Science, University of Wrocław, kiero@ii.uni.wroc.pl, supported by EPSRC grant GR/R11896/01 and Polish Ministry of Science and Higher Education research project nr N206 022 31/3660 2006/2009

†Department of Mathematics, Technische Universität Darmstadt, otto@mathematik.tu-darmstadt.de, research partially supported by EPSRC grant GR/R11896/01 and DFG project OT147/5-1

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 [23], via a reduction to the so-called Gödel prefix class $\exists^*\forall\forall\exists^*$, which however is only decidable in the absence of equality [7]. Full decidability for FO^2 with equality is due to Mortimer [20]. 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 [9], 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 into the range between modal and two-variable logics, see [2] and in particular [5] for their connection with finite variable fragments. The extension of FO^2 by counting quantifiers [14], for instance, is decidable by [11, 22] 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 [12] 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 [10]. 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 [12, 21].

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 [12, 10] 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 [21]. 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 $\tau = \tau_0 \dot{\cup} \tau_{\text{eq}}$ where τ_{eq} consists of a finite number of distinguished binary relations (typically E or $E_i, i = 1, 2, 3$). We let $\mathcal{EQ}[\tau_0; \tau_{\text{eq}}]$ denote the class of all $\tau_0 \dot{\cup} \tau_{\text{eq}}$ structures $\mathfrak{A} = (\mathfrak{A}_0, (E^{\mathfrak{A}})_{E \in \tau_{\text{eq}}})$ (with at least two elements) that interpret the relations $E \in \tau_{\text{eq}}$ 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 $\mathcal{C}$. $\mathrm{FINSAT}(\mathcal{L}, \mathcal{C})$ stands for the satisfiability problem for $\mathcal{L}$ in restriction to finite structures from $\mathcal{C}$, so that the finite model property for $\mathcal{L}$ over $\mathcal{C}$ is equivalent to $\mathrm{FINSAT}(\mathcal{L}, \mathcal{C}) = \mathrm{SAT}(\mathcal{L}, \mathcal{C})$. 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* $\mathcal{EQ}[\tau_0; E]$. *Hence* $\mathrm{SAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E])$ *and* $\mathrm{FINSAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E])$ *are* 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* 3-NEXPTIME.

(iii) $\mathrm{SAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E_1, E_2, E_3])$ *and* $\mathrm{FINSAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E_1, E_2, E_3])$ *are undecidable; in fact* FO^2 *over* $\mathcal{EQ}[\tau_0; E_1, E_2, E_3]$ *forms a conservative reduction class.*

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 [17]. This fact was independently proved in [15]. A further improvement was obtained lately in [19] where undecidability of FO^2 in the presence of one equivalence and one transitive relation is shown. In [19] it is also shown that $\mathrm{FINSAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E_1, E_2])$ is decidable.

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 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 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

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

We let $\boldsymbol{\alpha}[\mathfrak{A}] = \big\{\mathrm{tp}_{\mathfrak{A}}(a) \colon a \in A\big\}$ 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 $\boldsymbol{\beta}[\mathfrak{A}] = \big\{\mathrm{tp}_{\mathfrak{A}}(a_1, a_2) \colon a_1 \neq a_2,\, a_1, a_2 \in A\big\}$ 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, $\boldsymbol{\alpha}[S] = \big\{\mathrm{tp}_{\mathfrak{A}}(a) \colon a \in S\big\}$. 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*

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

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 [9]. 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

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

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

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

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). $\square$

3 Equivalence structures

We now restrict attention to finite relational vocabularies $\tau = \tau_0 \dot\cup \tau_{\text{eq}}$ where τ_{eq} consists of a finite number of distinguished binary relations to be interpreted as equivalence relations (we use E or $E_i, i = 1, 2, 3$ for these).

Definition 5 $\mathcal{EQ}[\tau_0; \tau_{\text{eq}}]$ denotes the class of all *equivalence structures* in the vocabulary $\tau = \tau_0 \dot\cup \tau_{\text{eq}}$, i.e., the class of all $\tau_0 \dot\cup \tau_{\text{eq}}$ structures $\mathfrak{A} = (\mathfrak{A}_0, (E^{\mathfrak{A}})_{E \in \tau_{\text{eq}}})$ with at least two elements that interpret the relations $E \in \tau_{\text{eq}}$ as equivalence relations.

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 $(\neg E_1xy \wedge E_2xy) \in \beta$.

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 $\chi = \chi(x, y) \in \text{FO}^2[\tau_0]$ we let

$$\chi^{+-}(x, y) := \chi(x, y) \wedge (\neg x = y \to (E_1xy \wedge \neg E_2xy)).$$

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 $\mathcal{EQ}[\tau_0; \tau_{\text{eq}}]$.

An $\mathcal{EQ}[\tau_0, E_1, E_2]$ 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 $\chi \in \text{FO}[\tau_0; E_1, E_2]$ and $\chi_i \in \text{FO}[\tau_0]$ 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 $\chi \in \text{FO}[\tau_0; E]$ and $\chi_i \in \text{FO}[\tau_0]$ and $s_i \in \{+, -\}$.

Proposition 6 ($\mathcal{EQ}$ Scott normal form) *For $\varphi \in \text{FO}^2[\tau_0 \cup \tau_{eq}]$ one computes in polynomial time an $\mathcal{EQ}[\tau; \tau_{eq}]$ Scott normal form sentence $\tilde\varphi \in \text{FO}^2[\tilde\tau_0 \cup \tau_{eq}]$, whose length is linear in the length of φ such that φ and $\tilde\varphi$ are satisfiability equivalent over $\mathcal{EQ}[\tau; \tau_{eq}]$.*

In terms of this normal form, φ can be thought of as stipulating constraints on the sets of atomic 1-types and 2-types that may occur (the $\forall\forall$ requirements in χ_0) and requirements for certain *witnesses* to be provided at every element (the $\forall\exists$ requirements in the $\chi_i^{s_i}$ for $1 \geqslant 1$). Among the $\forall\exists$ requirements we distinguish, for instance in the case of two equivalence relations E_1 and E_2, by the superscripts $s_i \in \{+, -\}^2$ requirements that need to be met by witnesses

– within the same intersection of E_1- and E_2-classes $(++)$,

– within the same E_1-class but in a distinct E_2-class $(+-)$,
– within the same E_2-class but in a distinct E_1-class $(-+)$, or
– in a distinct E_1-class and in a distinct E_2-class $(--)$.

Correspondingly, such witnesses are linked to the base point by 2-types in $\boldsymbol{\beta}^{++}$, $\boldsymbol{\beta}^{+-}$, $\boldsymbol{\beta}^{-+}$, or $\boldsymbol{\beta}^{--}$, respectively. For instance, a witness at $a \in \mathfrak{A}$ for an $\forall\exists$ requirement $\forall x \exists y \chi_i^{+-}(x,y)$ in φ is an element a' such that $\mathfrak{A} \models \chi_i^{+-}[a,a']$, which means that $\mathrm{tp}_{\mathfrak{A}}(a,a') \in \boldsymbol{\beta}^{+-}$ contains $\chi_i^{+-}(x,y)$. If there is such a witness, we also say that that the requirement is *fulfilled* at a.

4 One equivalence relation

We consider models in $\mathcal{EQ}[\tau_0; E]$, with $\tau_{\mathrm{eq}} = \{E\}$. 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^+.

4.1 An exponential model property

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.

Together these two steps allow us transform any given $\mathcal{EQ}[\tau_0; E]$ structure into an exponentially size bounded $\mathcal{EQ}[\tau_0; E]$ structure while preserving the truth of all normal form sentences. We thus get the following. The proof according to the two steps outlined above is provided in the next sections.

Proposition 7 *There is an exponential function f such that any normal form sentence $\varphi \in \mathrm{FO}^2[\tau_0; E]$ that is satisfiable in $\mathcal{EQ}[\tau_0; E]$ also 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.*

4.1.1 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)

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

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

Let $\mathfrak{A}$ be a countable $\mathcal{EQ}[\tau_0, E]$-structure, with infinitely many equivalence classes enumerated as $(B_i)_{i\in\mathbb{N}}$ without repetition. We may apply the above process inductively to E-classes B_i, one at a time, to obtain a sequence of structures $\mathfrak{A}_n \in \mathcal{EQ}[\tau_0, E]$, starting form $\mathfrak{A}_0 = \mathfrak{A}$, such that for all $n \in \mathbb{N}$:

(i) the universe of $\mathfrak{A}_n$ consists of the disjoint union of old and new equivalence classes of the form $\bigcup_{i<n} B'_i \cup \bigcup_{i\geqslant n} B_i$, where the new B'_i are exponentially size bounded in $|\tau_0|$;

(ii) $\mathfrak{A}_n \restriction \bigcup_{i\geqslant n} B_i \simeq \mathfrak{A} \restriction \bigcup_{i\geqslant n} B_i$ and $\mathfrak{A}_n \restriction \bigcup_{i<n} B'_i = \mathfrak{A}_{n+1} \restriction \bigcup_{i<n} B'_i$;

(iii) $\boldsymbol{\alpha}[\mathfrak{A}_{n+1}] = \boldsymbol{\alpha}[\mathfrak{A}_n]$ and $\boldsymbol{\beta}[\mathfrak{A}_{n+1}] = \boldsymbol{\beta}[\mathfrak{A}_n]$;

(iv) for every $b' \in B'_n$ there is some $b \in B_n$ such that $\boldsymbol{\beta}[b', \mathfrak{A}_{n+1}] \supseteq \boldsymbol{\beta}[b, \mathfrak{A}_n]$;

(v) $\boldsymbol{\beta}[b, B'_n] \supseteq \boldsymbol{\beta}[b, B_n]$ for all $b \in \bigcup_{i<n} B'_i$.[1]

Let $\mathfrak{A}'$ be the natural limit structure induced by this sequence on $A' := \bigcup_{n\in\mathbb{N}} B'_i$. It is easy to see that $\mathfrak{A}'$ is an $\mathcal{EQ}[\tau_0, E]$-structure with equivalence classes B'_i, which are exponentially size bounded in $|\tau_0|$. Moreover, due to construction,

$$\mathfrak{A} \models \varphi \implies \mathfrak{A}_n \models \varphi \text{ for all } n \in \mathbb{N} \implies \mathfrak{A}' \models \varphi$$

for all normal form φ.

The first implication follows inductively, essentially from (iii), (iv) and (v), with the same argument as before. For the second implication, concerning the limit, note that $\boldsymbol{\alpha}[\mathfrak{A}'] = \boldsymbol{\alpha}[\mathfrak{A}]$ and $\boldsymbol{\beta}[\mathfrak{A}'] = \boldsymbol{\beta}[\mathfrak{A}]$ directly follow from the corresponding equalities at all intermediate $\mathfrak{A}_n$, because all 1- and 2-types are eventually static by (ii). For $\forall\exists$ requirements in normal form φ, (v) guarantees that witnesses for such requirements provided for some $b \in B_i$ in B_j for some $j > i$ will be reproduced in B'_j and then persist in the limit by (ii).

We have thus proved the following, in the more interesting case in which we start with a model that has infinitly many equivalence classes; the argument for the case of finitely many equivalence classes is analogous but simpler as we do not need to go to a limit.

[1] Note that the first set of 2-types may (by (ii)) be evaluated in any $\mathfrak{A}_m$ for $m > n$ while the second is evaluated in $\mathfrak{A}_n$.

Lemma 8 *Every normal form φ that is satisfiable in an $\mathcal{EQ}[\tau_0, E]$ model, has an $\mathcal{EQ}[\tau_0, E]$ model with at most countably many equivalence classes each of which is bounded in size by an exponential function in $|\tau_0|$.*

4.1.2 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 $\boldsymbol{\alpha}[B]$).

1-types that are only realised in a single E-class (a simple $\forall\forall$-sentence can force a 1-type $\alpha(x)$ to be of this kind), and the E-classes that are distinguished in relation to at least one 1-type in this manner, clearly play a special role.

Definition 9 An equivalence class B in $\mathfrak{A}$ is called *singular* if it contains all realisations of some 1-type $\alpha \in \boldsymbol{\alpha}[\mathfrak{A}]$. A 1-type $\alpha \in \boldsymbol{\alpha}[\mathfrak{A}]$ *singular* if it is realised in only singular classes.

The following Lemma is immediate.

Lemma 10 *For every non-singular $\alpha \in \boldsymbol{\alpha}[\mathfrak{A}]$:*

(i) *there exists at least one non-singular class B containing a realisation of α.*

(ii) *for every $\alpha' \in \boldsymbol{\alpha}[\mathfrak{A}]$ (singular or non-singular) there exists $\beta \in \boldsymbol{\beta}^-[\mathfrak{A}]$ with $\beta{\restriction}x = \alpha$ and $\beta{\restriction}y = \alpha'$.*

As there is an obvious injection from the set of all singular classes of $\mathfrak{A}$ into $\boldsymbol{\alpha}[\mathfrak{A}]$, we have the following bound.

Lemma 11 *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.*

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.

Let $\mathfrak{B} \subseteq \mathfrak{A}$ be a union of non-singular 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. We fix such $\mathfrak{B}$ together with some choice of such $b_\alpha \in \mathfrak{B}$ for all non-singular α.

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 from different parts. Again we proceed in several stages.

(1) For $b \in \mathfrak{B}$ 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 $\beta_k \not\in \boldsymbol{\beta}^-[c, C]$ 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 $\beta_k \not\in \boldsymbol{\beta}^-[b, A_s]$ 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. This is possible by part (ii) of Lemma 10, since all such pairs involve at least one non-singular 1-type.

The combination of the transformations that firstly yield models with few classes (section 4.1.1) and then allow us to pass to a model relaising only few of these, now proves Proposition 7.

5 Two equivalence relations

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

5.1 Examples

We start with a few examples showing that adding the second equivalence relation increases the expressive power substantially.

5.1.1 Failure of the finite model property

We show that even over $\mathcal{EQ}[\tau_0; E_1, E_2]$, for $\tau_0 = \{P, Q, R\}$ with unary P, Q, R, FO^2 does not have the finite model property. We exhibit an infinity axiom that requires an infinite number of E_1-classes and of E_2-classes. 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.

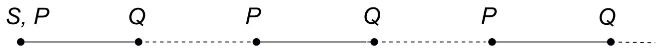

Figure 1: A model of φ.

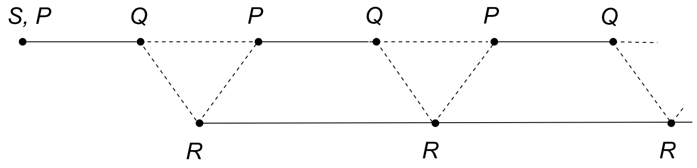

Figure 2: A model of φ'.

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

It is easy to formalize (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. Fig. 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 positions 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 as depicted in Fig. 2.

5.1.2 Large finite classes

The tree-like structure in Fig. 3 contains an E_2-class S with exactly 2^{k-1} elements of the same type along its leaves. L 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 $Q_0, \ldots, Q_{n-1}$. 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 not in L; similarly for $a \in P_{2i+1}$ and E_2-links to two successors in P_{2i+2}.

(3) distinct elements in P_{2i} are not E_1-connected; distinct elements in P_{2i+1} may be

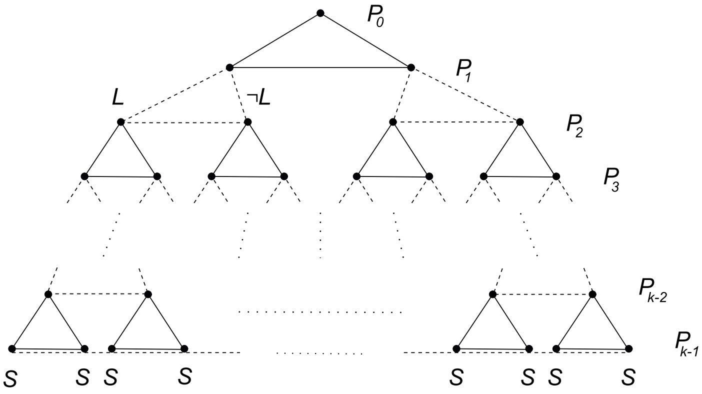

Figure 3: A large finite class.

E_2-connected only if $2i + 1 = k - 1$; this ensures that elements fulfilling the requirements from (2) are not reused.

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

(5) distinct elements from every P_{2i} may be E_2-connected only if exactly one of them is in L; similarly distinct elements from every P_{2i+1} may be E_1-connected only if exactly one of them is in L;

(6) an element satisfies S iff it is in P_{k-1}; 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 bullets from Fig. 3 as intersections, each intersection consisting of 2^n elements numbered by means of another set of unary predicates $U_1, \ldots, U_{n-1}$. 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 is presented in [17].

5.1.3 Exactly two realisations of a 1-type

In plain FO^2 we are able to enforce that some 1-type has to be realised exactly once in a structure (following [9] realisations of such types are called *kings*). In the previous example we enforced exactly 2^k realisations of some 1-type in a model, for some k. However, a king played a very important role in this construction, viz. the root element in P_0. In contrast, we now exhibit a simple satisfiable formula, whose models do not have kings but *exactly two* realisations of every 1-type. In fact, the following formula defines the structure depicted in Fig. 4 up to isomorphism. It says that

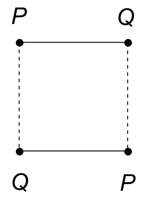

Figure 4: Two realisations of a type.

(1) P and Q form a partition of the universe and $P \times Q \subseteq E_1 \cup E_2$,
(2) there are at least two E_1-classes.
(3) P and Q each intersect every E_1-class and every E_2-class in precisely one point.

5.2 Small intersections

Let $\mathfrak{A} \in \mathcal{EQ}[\tau_0; E_1, E_2]$ be an equivalence structure. From the previous subsections we know 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 $\mathfrak{J}$ in $\mathfrak{A}$ by a new substructure $\mathfrak{J}'$ 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 $\mathfrak{J}$ are reproduced by $\mathfrak{J}'$, which will then again be the intersection of these classes. We also preserve any normal form formula in this modification.

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

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

We may apply the above process to all intersections of a countable $\mathfrak{A} \models \varphi$ in $\mathcal{EQ}[\tau_0, E_1, E_2]$ as in the construction of a model with small classes from section 4.1.1. This yields the following lemma.

Lemma 12 *Every normal form φ that is satisfiable in an $\mathcal{EQ}[\tau_0, E_1, E_2]$ model, has an $\mathcal{EQ}[\tau_0, E_1, E_2]$ model with at most countably many intersections each of which is bounded in size by an exponential function in $|\tau_0|$.*

5.3 Certificates for satisfiability

We introduce the notion of a *certificate* (of satisfiability) for a formula φ in $\mathcal{EQ}[\tau_0, E_1, E_2]$-Scott normal form. We then establish, in sections 5.4 and 5.5, that φ is satisfiable if, and only if, it possesses such a certificate (cf. Theorem 13). The size of certificates can be bounded by a triply exponential function in $|\varphi|$. We thus establish decidability of $\mathrm{SAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E_1, E_2])$ in 3-NEXPTIME ((ii) in Theorem 1).

Proviso. For the rest of this section, we shall work with two signatures: besides the basic vocabulary $\tau = \tau_0 \cup \{E_1, E_2\}$ we consider an extension $\bar{\tau} = \bar{\tau}_0 \cup \{E_1, E_2\}$ for $\bar{\tau}_0 = \tau_0 \cup S_1 \cup S_2$ by two sets S_1 and S_2 of new unary predicates to be used as marker colours for certian special E_1 or E_2 equivalence classes. References to sets like $\boldsymbol{\beta}^{++}$, $\boldsymbol{\beta}^{+-}$, etc., however, will always still be with respect to the basic vocabulary τ.

5.3.1 Certificates: the general format

Let φ be a formula in $\mathcal{EQ}[\tau_0, E_1, E_2]$-Scott normal form. A certificate of satisfiability for φ is a finite description of a (possibly infinite) $\mathcal{EQ}[\bar{\tau}_0, E_1, E_2]$-model of φ. Formally, a certificate is a tuple $X = (S_1, S_2, \bar{\Delta}, \mathfrak{K}, \mathfrak{C}, \mathfrak{L}_0, \mathbb{P}^{+-}, \mathbb{P}^{-+}, \mathbb{P}^{--}, \mathbb{T})$, such that:

(a) S_1 and S_2 are disjoint sets of new unary predicates;

(b) $\bar{\Delta}$ is a set of isomorphism types of intersections, with respect to the extended vocabulary $\bar{\tau}$;

(c) $\mathfrak{K} \subseteq \mathfrak{C} \subseteq \mathfrak{L}_0$ are finite $\mathcal{EQ}[\bar{\tau}_0, E_1, E_2]$-structures, at least $\mathfrak{L}_0$ is non-empty;

(d) $\mathbb{P}^{+-}$, $\mathbb{P}^{-+}$, $\mathbb{P}^{--}$ are partial functions from $\bar{\Delta}$ to finite sets of $\mathcal{EQ}[\bar{\tau}_0, E_1, E_2]$-structures;

(e) $\mathbb{T}$ is a finite set of $\mathcal{EQ}[\bar{\tau}_0, E_1, E_2]$-structures.

Structures from $\mathbb{T}$ and from sets returned by $\mathbb{P}^{+-}$, $\mathbb{P}^{-+}$, $\mathbb{P}^{--}$ are called *pattern structures*. In particular, elements from values of $\mathbb{P}^{+-}$ and $\mathbb{P}^{-+}$ are called *pattern classes*. We require all objects involved in a certificate, i.e., S_1, S_2, $\bar{\Delta}$, $\mathfrak{K}$, $\mathfrak{C}$, $\mathfrak{L}_0$, $\mathbb{T}$, all sets returned by functions $\mathbb{P}^{+-}, \mathbb{P}^{-+}, \mathbb{P}^{--}$ and all pattern structures to be finite.

5.3.2 Models from certificates: informal description

Before presenting the detailed list of requirements that charcterise valid certificates, we informally describe the components of a certificate and outline the process by which a model of the given sentence φ can be constructed from a certificate.

We mostly work on the level of intersections rather than individual elements. Unary predicates $U \in S_1 \cup S_2$ will eventually mark entire E_1 or E_2-classes; at the level of intersections as the basic building blocks provided in a certificate, each intersection type $\bar{\delta} \in \bar{\Delta}$ is either fully contained in U or disjoint from U. We correspondingly denote by $s(\bar{\delta})$ the set of unary U which mark (the whole of) the intersection type $\bar{\delta}$, and put $s_i(\bar{\delta}) := s(\bar{\delta}) \cap S_i$ to distinguish S_1-markers from S_2-markers. In fact, each $\bar{\delta}$ will be marked with at most one marker from S_1 and one from S_2, so that the $s_i(\bar{\delta})$ are either empty or singletons. We use corresponding notation for realisations, i.e., actual intersections J of type $\bar{\delta}$; for instance, $s(J)$ also stands for $s(\bar{\delta})$ if J is of type $\bar{\delta}$.

For an isomorphism type $\bar{\delta} \in \bar{\Delta}$, we denote by $\bar{\delta}{\restriction}\tau$ the isomorphism type of the τ-reduct. Similarly, for an atomic type γ with respect to $\bar{\tau}$, we denote by $\gamma \restriction \tau$ its restriction to τ.

We say that a structure is *built from* (not necessarily distinct) isomorphism types of intersections $\bar{\delta}_1, \ldots, \bar{\delta}_k$ if it consists of exactly k intersections of preceisely these types (this still leaves a lot of variablity, in terms of the *connections*, i.e., 2-types, between these parts). Recall that 2-types in $\boldsymbol{\beta}^{+-}$, $\boldsymbol{\beta}^{-+}$ and $\boldsymbol{\beta}^{--}$ describe links between distinct intersections (that belong to the same E_1-class, E_2-class, or neither, respectively), while 2-types in $\boldsymbol{\beta}^{++}$ live within one and the same intersection only.

We say that intersections J and J' in a structure $\mathfrak{A}$ are in a *free relation* if they belong to distinct E_1- and E_2-classes, i.e., if $\boldsymbol{\beta}[J, J'] \subseteq \boldsymbol{\beta}^{--}$.

The set $\bar{\Delta}$ contains only types of intersections allowed by φ: $\bar{\Delta}$ must agree with the $\forall\forall$ constraints in φ and must provide all the required $\boldsymbol{\beta}^{++}$ witnesses for their elements (these witnesses can only come from within the same intersection). Only such types of intersections will be considered in the following.

The unary predicates from S_i are used as markers for *special* E_i-classes. Roughly speaking, a class is regarded as special if it comprises intersections of a type which can be realised either only in this one class or only some specific number of times in the whole model. All special classes will be represented in $\mathfrak{L}_0$ and in fact even completed w.r.t. all $\forall\exists$ requirements that call for witnesses within the class. We say, for instance, that an E_1-class is *completed* if all $\boldsymbol{\beta}^{+-}$ requirements are witnessed within this class; similarly for E_2-classes and $\boldsymbol{\beta}^{-+}$ requirements. Note that the fact that some E_i-class is completed at some stage of the construction of a model does not mean that it has been completed in the sense that it cannot acquire further members at a later stage; it just means that as far as the internal view of this (maybe still partial) class is concerned, there is no need to extend it further. It is in this sense that all special classes are completed in $\mathfrak{L}_0$.

$\mathfrak{L}_0$ serves as the initial part of the model to be constructed. The following distinguished substructures come as parts of $\mathfrak{L}_0$: $\mathfrak{K} \subseteq \mathfrak{L}_0$ consists of those intersections whose E_1 and E_2-classes are both special; and $\mathfrak{C}$, where $\mathfrak{K} \subseteq \mathfrak{C} \subseteq \mathfrak{L}_0$, is such that every intersection in $\mathfrak{K}$ has its E_1- and E_2-classes completed in $\mathfrak{C}$ and has witnesses for all $\boldsymbol{\beta}^{--}$ requirements in $\mathfrak{C}$. Furthermore the relationship between $\mathfrak{C}$ and $\mathfrak{L}_0$ is such that $\mathfrak{L}_0$ completes all classes in $\mathfrak{C}$ and provides witnesses for all $\boldsymbol{\beta}^{--}$ requirements for elements of $\mathfrak{C}$.

The remaining part of the model is constructed in a regular fashion, with the help of the pattern structures provided by the certificate. The format of pattern functions $\mathbb{P}^{+-}$, $\mathbb{P}^{-+}$ and $\mathbb{P}^{--}$ is such that $\mathbb{P}^{+-}(\bar{\delta}) = \{\mathfrak{B}\}$ is a singleton set containing a completed E_1-class with a distinguished intersection of type $\bar{\delta}$; $\mathbb{P}^{-+}(\bar{\delta}) = \{\mathfrak{B}\}$ similarly contains a single completed E_2-class with a distinguished intersection of type $\bar{\delta}$.

$\mathbb{P}^{--}(\bar{\delta})$ is a finite set of structures $\mathfrak{D}$, each consisting of two intersections in a free relation, one of them of type $\bar{\delta}$. The purpose of $\mathbb{P}^{+-}(\bar{\delta})$, $\mathbb{P}^{-+}(\bar{\delta})$ and $\mathbb{P}^{--}(\bar{\delta})$ is to provide extensions of an intersection of type $\bar{\delta}$ to completed E_1 and E_2-classes and to fulfil its $\boldsymbol{\beta}^{--}$ requirements. The pattern structures from $\mathbb{T}$, finally, also consists of two intersections in a free relation; these serve to allocate connections between pairs of intersections that are in a free relationship and that have not already been connected through requirements during their introduction.

These extensions are adjoined in an inductive process generating a chain of structures $\mathfrak{L}_0 \subseteq \mathfrak{L}_1 \subseteq \cdots$. The desired model for φ will be the τ-reduct of the limit $\mathfrak{L} := \bigcup_{k\in\mathbb{N}} \mathfrak{L}_k$.

In the single extension step from $\mathfrak{L}_k$ to $\mathfrak{L}_{k+1}$ we build classes and provide witnesses for intersections in $\mathfrak{L}_k$, insofar as these were not already provided in $\mathfrak{L}_k$. The patterns provided by $\mathbb{P}^{+-}$, $\mathbb{P}^{-+}$ and $\mathbb{P}^{--}$ only specify connections between individual intersections in $\mathfrak{L}_k$ and their respective extensions by intersections added to $\mathfrak{L}_{k+1}$. Some further connections are made for new intersections marked with some $U \in S_1 \cup S_2$: these tell us how to join intersections to special classes represented in $\mathfrak{L}_0$.

To complete $\mathfrak{L}_{k+1}$ we specify the remaining connections (2-types). Firstly, we join intersections marked with symbols from $S_1 \cup S_2$ to appropriate classes in $\mathfrak{L}_0$, then set the required links with witnesses in $\mathfrak{K}$, and, finally, we use patterns from $\mathbb{T}$ to settle the remaining connections.

5.3.3 Certificates: validity criteria

We collect a number of constraints, which have to be fulfilled by valid certificates, chosen such that the model construction outlined above can be performed without conflicts and yields a model of the given normal form sentence φ.

(1) All intersections in $\mathfrak{L}_0$ and in pattern structures have isomorphism types from $\bar{\Delta}$.

(2) $\mathfrak{L}_0$ and all pattern structures respect the $\forall\forall$ part of φ.

(3) Every intersection J of type $\bar{\delta} \in \bar{\Delta}$ fulfils all $\boldsymbol{\beta}^{++}$ requirements.

(4) For $\bar{\delta} \in \bar{\Delta}$ and $U \in S_1 \cup S_2$: either $\bar{\delta} \models \forall x U x$ or $\bar{\delta} \models \forall x \neg U x$.

(5) For every $\bar{\delta} \in \bar{\Delta}$ and $i = 1, 2$, we have $|s_i(\bar{\delta})| \leq 1$ (recall that $s_i(\bar{\delta}) = s(\bar{\delta}) \cap S_i$).

(6) For intersections J, J' in $\mathfrak{L}_0$:
if J and J' are members of the same E_i-class then $s_i(J) = s_i(J')$;
if J and J' are not members of the same E_i-class, then $s_i(J) \cap s_i(J') = \emptyset$.

(7) $|s(\bar{\delta})| = 2$ (i.e., $s_i(\bar{\delta}) \neq \emptyset$ for both of $i = 1, 2$) if, and only if, $\bar{\delta}$ is realised exactly once in $\mathfrak{K}$, not realised in $\mathfrak{L}_0$ outside $\mathfrak{K}$, and does not occur in any of the pattern classes produced by $\mathbb{P}^{+-}$ and $\mathbb{P}^{-+}$.

(8) If for some $\bar{\delta}$ and U, $s(\bar{\delta}) = \{U\}$, then there exists a completed class $\mathfrak{B}$ in $\mathfrak{L}_0$ marked with U; moreover, if $\bar{\delta}$ is a type of a non-distinguished intersection J in some pattern structure, then $\mathfrak{B}$ contains at least two realisations of $\bar{\delta}$.[2]

(9) Intersections from $\mathfrak{C}$ have completed E_1- and E_2-classes in $\mathfrak{L}_0$, and all $\boldsymbol{\beta}^{--}$ requirements at elements in $\mathfrak{C}$ are fulfilled in $\mathfrak{L}_0$.

(10) Every E_i-class that is not completed in $\mathfrak{L}_0$ contains at most one intersection in $\mathfrak{L}_0$.

(11) For an intersection $J \subseteq L_0 \setminus C$ of type $\bar{\delta}$ such that $\mathbb{P}^{--}(\bar{\delta})$ contains a structure $\mathfrak{D}$ built from $\bar{\delta}$ and some $\bar{\delta}'$ with $|s(\bar{\delta}')| = 2$, there is an isomorphism between $\mathfrak{D}$ and $\mathfrak{L}_0 \restriction (J \cup J')$, where $\mathfrak{J}' \subseteq \mathfrak{K}$ is the unique realisation of $\bar{\delta}'$ in $\mathfrak{L}$.

[2] This allows us to safely add additional intersections of this type $\bar{\delta}$ (in particular a copy of J) to $\mathfrak{B}$.

(12) $\mathbb{P}^{+-}(\bar{\delta})$ is defined whenever $s_1(\bar{\delta}) = \emptyset$. In this case, $\mathbb{P}^{+-}(\bar{\delta})$ is a singleton set consisting of a completed E_1-class with a distinguished intersection of type $\bar{\delta}$ (and referred to as a pattern E_1-class below). Similarly for $\mathbb{P}^{-+}(\bar{\delta})$, $s_2(\bar{\delta})$ and pattern E_2-classes.

(13) $s_i(\bar{\delta}) = \emptyset$ for all intersection types $\bar{\delta}$ that occur in pattern E_i-classes. [3]

(14) For each $U \in S_2$, any pattern E_1-class contains at most one intersection marked with U; similarly for pattern E_2-classes and $U \in S_1$. [4]

(15) $\mathbb{P}^{--}(\bar{\delta})$ is defined whenever $|s(\bar{\delta})| < 2$. In this case, $\mathbb{P}^{--}(\bar{\delta})$ is a finite set of structures, each built from a distinguished intersection of type $\bar{\delta}$ and another intersection of some type $\bar{\delta}'$ such that $s(\bar{\delta}) \cap s(\bar{\delta}') = \emptyset$, in a free relation. Every type $\bar{\delta}'$ such that $s(\bar{\delta}') = 2$ can be used to built at most one of the structures in $\mathbb{P}^{--}(\bar{\delta})$ for this $\bar{\delta}$.

(16) Every $\mathbb{P}^{--}(\bar{\delta})$ provides witnesses for all $\boldsymbol{\beta}^{--}$ requirements at elements in $\bar{\delta}$. More precisely, the structure obtained by gluing all the member structures in $\mathbb{P}^{--}(\bar{\delta})$ in their distinguished realisations of $\bar{\delta}$, fulfils all $\boldsymbol{\beta}^{--}$ requirements at all elements in this realisation of $\bar{\delta}$.

(17) For every $\bar{\delta}, \bar{\delta}' \in \bar{\Delta}$, such that $s(\bar{\delta}) \cap s(\bar{\delta}') = \emptyset$ there exists a structure in $\mathbb{T}$ built from $\bar{\delta}, \bar{\delta}'$ in a free relation.

The proof of the following theorem will be divided into two parts, presented in the following two sections.

Lemma 13 *Let φ be a formula in $[\tau_0, E_1, E_2]$-normal form. Then φ is satisfiable if and only if it has a certificate of satisfiability enjoying properties (1)–(17) whose size if bounded by a function f, at most triply exponential in $|\varphi|$.*

5.4 Constructing a regular model from a valid certificate

Let $X = (S_1, S_2, \bar{\Delta}, \mathfrak{K}, \mathfrak{C}, \mathfrak{L}_0, \mathbb{P}^{+-}, \mathbb{P}^{-+}, \mathbb{P}^{--}, \mathbb{T})$ be a certificate for φ. We construct a model $\mathfrak{A} \models \varphi$ as a limit of a countable chain of finite structures $\mathfrak{L}_0 \subseteq \mathfrak{L}_1 \subseteq \ldots \subseteq \mathfrak{L}_k \subseteq \mathfrak{L}_{k+1} \subseteq \cdots$, such that for every $k \geq 0$, intersections from L_k have both their classes completed in $\mathfrak{L}_{k+1}$ ($\boldsymbol{\beta}^{+-}$ and $\boldsymbol{\beta}^{-+}$ requirements fulfilled) and all $\boldsymbol{\beta}^{--}$ requirements fulfilled by witnesses in $\mathfrak{L}_{k+1}$. Let $M_k = L_k \setminus L_{k-1}$ for $k > 0$ and $M_0 = L_0 \setminus C$.

For the construction of the structures $\mathfrak{L}_k$ we proceed inductively. The starting point is the structure $\mathfrak{L}_0$ as given in the certificate. Note that by (9) all intersections in C have both their classes completed and all $\boldsymbol{\beta}^{--}$ requirements witnessed in $\mathfrak{L}_0$.

In the inductive process generating the $\mathfrak{L}_k$ we want to maintain the following additional properties, for all k and every intersection $J \subseteq M_k$:

(A0) the extended type $\bar{\delta}$ of J is a member of $\bar{\Delta}$ and $|s(\bar{\delta})| < 2$.

(A1) if J is of type $\bar{\delta}$ and if $\mathbb{P}^{--}(\bar{\delta})$ contains a structure $\mathfrak{D}$ built from intersection types $\bar{\delta}, \bar{\delta}'$, such that $|s(\bar{\delta}')| = 2$ and δ' is realised in $\mathfrak{K}$ by J', then $\mathfrak{L}_k \restriction (J \cup J')$ is isomorphic to $\mathfrak{D}$.

[3] $s_i(\bar{\delta})$ markers say to which special classes an intersection should be adjoined; we do not want to add intersections from pattern E_i-classes to special E_i-classes.

[4] We do not allow two distinct intersections from a pattern class to be joined to the same special class.

(A2) if J is of type $\bar{\delta}$ and if $U \in s(\bar{\delta})$, then J is connected to the class marked by U, whose initial fragment is represented in $\mathfrak{L}_0$, in such a way that this class, after joining J, remains completed.

(A3) if J does not have its E_i-class completed in $\mathfrak{L}_k$, then it is not E_i-connected to any other intersection in $\mathfrak{L}_k$.

Note that M_0 satisfies (A0) by (7), (A1) by (11), (A2) by (6) and (8), and (A3) by (10).

Given $\mathfrak{L}_k$, we now indicate how to obtain its extension $\mathfrak{L}_{k+1}$ while maintaining (A0)–(A3). This extension step consists of defining classes and providing $\boldsymbol{\beta}^{--}$ witnesses in M_{k+1} for elements from M_k, and defining the remaining 2-types.

Completing classes. Let $J \subseteq M_k$ be of type $\bar{\delta}$. If $s_1(\bar{\delta}) \neq \emptyset$, then by (A2) its E_1-class is completed.

Otherwise $s_1(\bar{\delta}) = \emptyset$ and, by (12), $\mathbb{P}^{+-}(\bar{\delta})$ contains an E_1-class $\mathfrak{D}$, built from $\bar{\delta}$ and some $\bar{\delta}_1, \dots, \bar{\delta}_k$. If the E_1-class of J is not already completed in $\mathfrak{L}_k$, then we add new intersections $J_1, \dots, J_k$ in $\mathfrak{L}_{k+1}{\restriction}M_{k+1}$ and make the substructure $J \cup J_1 \cup \dots J_k$ isomorphic to $\mathfrak{D}$. Note that this is possible without conflicts, since, by (A3), J is not E_1-connected to any other intersection in $\mathfrak{L}_k$. Property (12) guarantees that all $\boldsymbol{\beta}^{+-}$ requirements of $J, J_1, \dots J_k$ are now fulfilled, i.e., that the new class has been completed.

We proceed similarly with E_2-classes: if $s_2(\bar{\delta}) \neq \emptyset$ then the E_2-class of J is completed by (A2), if $s_2(\bar{\delta}) = \emptyset$, we complete its E_2-class, using $\mathbb{P}^{-+}(\bar{\delta})$.

This step (and the pre-condition (A2)) guarantees that every element from M_k has both its classes completed in $\mathfrak{L}_{k+1}$.

Adjoining $\boldsymbol{\beta}^{--}$ witnesses. Let $J \subseteq M_k$ be of type $\bar{\delta}$. By (A0) we have $|s(\bar{\delta})| < 2$. By (15), $\mathbb{P}^{--}(\bar{\delta})$ is defined. For every $\mathfrak{D} \in \mathbb{P}^{--}(\bar{\delta})$ built from intersection types $\bar{\delta}$ and $\bar{\delta}'$, if $s(\bar{\delta}') < 2$, then put a new intersection J' of type $\bar{\delta}'$ in M_{k+1} making the substructure $\mathfrak{L}_{k+1}{\restriction}(J \cup J')$ isomorphic to $\mathfrak{D}$. Note that we do not (and cannot) add new copies of intersection types with $s(\bar{\delta}') = 2$, since these are only to be realised in $\mathfrak{K}$. By (A1) and property (16) we thus fulfil all $\boldsymbol{\beta}^{--}$ requirements of J.

All newly added intersections in M_{k+1} satisfy (A0). To see this note that they either come from pattern classes produced by $\mathbb{P}^{+-}$, $\mathbb{P}^{-+}$ or from pattern structures produced by $\mathbb{P}^{--}$. In the former case the required property is guaranteed by (8), in the latter by the special treatment of intersections marked by two colours in the procedure adjoining $\boldsymbol{\beta}^{--}$ witnesses.

In the remaining steps we also need to set conditions (A1)–(A3) right for intersections in M_{k+1}.

(A1): $\boldsymbol{\beta}^{--}$ witnesses in K. Let $J \subseteq M_{k+1}$. Let $\bar{\delta}$ be the type of J. For every $\mathfrak{D} \in \mathbb{P}^{--}(\bar{\delta})$ built from $\bar{\delta}, \bar{\delta}'$, with $s(\bar{\delta}') = 2$, find the unique $J' \subseteq K$ with $\bar{\delta}(J') = \bar{\delta}'$ and make the substructure $\mathfrak{L}_{k+1}{\restriction}(J \cup J')$ isomorphic to $\mathfrak{D}$. This ensures property (A1) for M_{k+1}.

(A2): joining intersections to special classes. Let J be an intersection in M_{k+1} of type $\bar{\delta}$ marked, for instance, with $s_1(\bar{\delta}) = \{U\}$. By (A0) we have $|s(\bar{\delta})| < 2$, so $s(\bar{\delta}) = s_1(\bar{\delta}) = \{U\}$. Then, by (8), there is an E_1-class $\mathfrak{B}$ in $\mathfrak{L}_0$ marked by U and with at least two realisations J_1 and J_2 of $\bar{\delta}$. We join J to the E_1-class $\mathfrak{B}$ in the following way. For every $J' \subseteq B$, $J' \neq J_2$ we make the connection between J and J' isomorphic to the connection between J_2 and J'; the connection between J and J_2 is made isomorphic to the connection between J_1 and J_2. Since $\mathfrak{B}$ was completed, which means that J_2 had all its $\boldsymbol{\beta}^{+-}$ requirements fulfilled in $\mathfrak{B}$, so does J. Note that the above procedure can be performed without conflicts by (13), (14), and (5). This step guarantees property (A2) for $\mathfrak{L}_{k+1}$.

(A3): the remaining connections. Consider a pair of intersections $J \subseteq M_{k+1}$, $J' \subseteq L_{k+1}$ for which we have not specified a connection so far. Let $\bar{\delta}$, $\bar{\delta}'$ be their extended types. Observe that $s(\bar{\delta}) \cap s(\bar{\delta}') = \emptyset$, since otherwise the intersections would have been made members of the same class. So by (17) there is a structure $\mathfrak{D} \in \mathbb{T}$, built from $\bar{\delta}$ and $\bar{\delta}'$. We make the substructure $\mathfrak{L}_{k+1}\restriction(J \cup J')$ isomorphic to $\mathfrak{D}$. Note that in this step we do not add any equivalence connections, since all structures in $\mathbb{T}$ consist of pairs of intersections in free relation. Thus we maintain property (A3).

This finishes the construction of $\mathfrak{L}_{k+1}$. We take $\mathfrak{L} = \bigcup_{k \in \mathbb{N}} \mathfrak{L}_k$. We conclude this section by noting that $\mathfrak{L} \models \varphi$:

- all connections between intersections in $\mathfrak{L}$ are isomorphic to some connections from $\mathfrak{L}_0$ or from pattern structures. By (2), $\mathfrak{L}$ satisfies the $\forall\forall$ part of φ.
- by (1) and (3), every element has all the required $\boldsymbol{\beta}^{++}$ witnesses.
- $\boldsymbol{\beta}^{--}$, $\boldsymbol{\beta}^{+-}$ and $\boldsymbol{\beta}^{-+}$ requirements are witnessed by elements in $\mathfrak{L}_0$ for elements of C and by elements of L_{k+1} for elements of L_k.

5.5 Extracting a certificate from a model

Let $\mathfrak{A} \models \varphi$. By Lemma 12 we may assume that all intersections in $\mathfrak{A}$ are of size exponential in $|\varphi|$. Let Δ be the set of isomorphism types of these intersections. We construct a certificate $X = (S_1, S_2, \bar{\Delta}, \mathfrak{K}, \mathfrak{C}, \mathfrak{L}_0, \mathbb{P}^{+-}, \mathbb{P}^{-+}, \mathbb{P}^{--}, \mathbb{T})$ of satisfiability for φ.

5.5.1 Special classes

We distinguish a set S of equivalence classes and a set $\Delta_0 \subseteq \Delta$ of isomorphism types of intersections in $\mathfrak{A}$. Δ_0 and S are obtained by saturating intial instantiations of these sets recursively in order to achieve suitable closure conditions. We present a recursive process for their definition that will also facilitate the proof of an exponential size bound on S in Lemma 14 below.

Initialisation:

(0) if a single class B contains all intersections of some type $\delta \in \Delta$ then add δ to Δ_0 and B to S; in particular, if δ is realised only once in the model, then both its classes become special in this step.

(0') if two intersecting classes B, B' contain all realisations of some $\delta \in \Delta$, but if there is no single class containing all realisations of δ, then add δ to Δ_0 and B, B' to S.

Recursively:

(1) if all realisations of some δ are contained in classes from S, then add δ to Δ_0.

(2) if some E_1-class $B \in S$ contains an intersection J of type $\delta \in \Delta_0$ such that J is the only realisation of δ in B whose E_2-class B' is not in S, then add B' to S; similarly with the roles of E_1 and E_2 exchanged.

The members of S are called *special classes*, members of Δ_0 are *special types* and realisations of special types are *special intersections*. Note that a special intersection belongs to at least one special class and that special classes may contain also non-special intersections.

Let us turn back to our examples and see which classes in the models depicted above are special. In Fig. 1 the only special classes are the E_1- and E_2-classes of the leftmost element (the latter is a singleton). In Fig. 2 we find one more special class, viz. the lower row of elements satisfying R. In both examples the classes are special already according to (0). In Fig. 4 all four classes are special, according to (0').

The model from Fig. 3 is more interesting in that it illustrates a non-trivial saturation process. Initially, by step (0), only the E_2-class formed by the bottom row of elements in S, both classes of the root (its E_2-class is trivial) and two E_2-classes containing elements from P_1 are special.[5] By repeated applications of steps (1) and (2), however, all classes and intersections successively become special.

This shows that the number of special classes may necessarily be triply exponential in terms of formula size. The following Lemma establishes a corresponding upper bound.

Lemma 14 *The size of S is at most exponential in $|\Delta|$ (triply exponential in $|\tau_0|$).*

Proof Let $t = |\Delta|$. Step (1) can be executed at most t times. Let us denote by s_k the number of classes in S after the k-th iteration of step (1), and by s_0 the number of classes in S after steps (0) and (0'). Clearly $s_0 \leq 4t$ initially.[6] Between the k-th and $(k+1)$-st iteration of (1), step (2) can be executed no more than ts_k times, since every $(\delta, B) \in \Delta_0 \times S$ can produce at most one new member B' of S (and if B' contains an intersection J whose type is a member of the current Δ_0, then the second class of J is already in S). Therefore $s_{k+1} \leq s_k + ts_k \leq (1+t)s_k$, and it follows that $s_t \leq (1+t)^t 4t$, which is exponential in t. □

The extended signature. Let S_1 be a new set of unary predicates containing for each special E_1-class B a marker U_B. Similarly let S_2 consist of colours for special E_2-classes. Recall that $\bar{\tau}_0 := \tau_0 \cup S_1 \cup S_2$.

[5] Note that the elements in P_1 have different types, since only one of them is in L.

[6] We bound s_0 by $4t$ rather than $2t$, because sometimes an intersection type can produce four special classes, cf. Fig. 4.

We expand the $\mathcal{EQ}[\tau_0, E_1, E_2]$-structure $\mathfrak{A}$ to an $\mathcal{EQ}[\bar{\tau}_0, E_1, E_2]$-structure $\mathfrak{A}'$ in the obvious way: $U_B^{\mathfrak{A}'} := B$ for all $B \in S$.

We let $\bar{\Delta}$ be the set of isomorphism types of intersections in $\mathfrak{A}'$ (i.e., certain uniformly coloured expansions of the intersection types in Δ).

Note that conditions (3), (4) and (5) for certificates are fulfilled.

The following Lemma will allow us to choose a set $\mathbb{T}$ in accordance with condition (17).

Lemma 15 *Let $\bar{\delta}, \bar{\delta}' \in \bar{\Delta}$ be such that $s(\bar{\delta}) \cap s(\bar{\delta}') = \emptyset$. Then $\mathfrak{A}$ has intersections J, J' in a free relation such that J is of type $\delta = \bar{\delta}{\restriction}\tau$ and J' of type $\delta' = \bar{\delta}'{\restriction}\tau$.*

Proof Let J be any intersection realising $\bar{\delta}$ in $\mathfrak{A}'$. Thus J is of type δ in $\mathfrak{A}$. Assume no realisations of δ' in $\mathfrak{A}$ were in a free relation with J. Then all realisations of δ' in $\mathfrak{A}$ (and thus also of $\bar{\delta}'$ in $\mathfrak{A}'$) belong to the E_1-class B_1 of J, or all belong to the E_2-class B_2 of J, or all belong to the union of these two classes, $B_1 \cup B_2$. In each of these cases, $\delta' \in \Delta_0$, by (0) or (0').

In the first case, $B_1 \in S$ and thus $U_{B_1} \in s(\bar{\delta}')$. Since $J \subseteq U_{B_1}$, also $U_{B_1} \in s(J) = s(\bar{\delta})$ and $s(\bar{\delta}) \cap s(\bar{\delta}') \neq \emptyset$, contradicting the assumption. The second case is symmetric to the first.

In the third case $B_1, B_2 \in S$, and thus either $U_{B_1} \in s(\bar{\delta}')$ or $U_{B_2} \in s(\bar{\delta}')$ (note that in this case all realizations of $\bar{\delta}'$ in $\mathfrak{A}'$ belong only to one of B_1, B_2) and we get a contradiction in a similar way. $\square$

We can now define $\mathbb{T}$ as follows. For every $\bar{\delta}, \bar{\delta}' \in \bar{\Delta}$, such that $s(\bar{\delta}) \cap s(\bar{\delta}') = \emptyset$, let J, J' be a pair of intersections in a free relation in $\mathfrak{A}$, realising $\delta = \bar{\delta}{\restriction}\tau$ and $\delta' = \bar{\delta}'{\restriction}\tau$, as guranteed by Lemma 15. Let $\mathfrak{D} = \mathfrak{A}{\restriction}(J \cup J')$. We add to $\mathbb{T}$ the structure $\mathfrak{D}'$, obtained from $\mathfrak{D}$ by marking J with symbols from $s(\bar{\delta})$ and J' by symbols from $s(\bar{\delta}')$. Thus we fulfil condition (17).

5.5.2 Regular classes

All classes whose isomorphism types are used in patterns $\mathbb{P}^{+-}$, $\mathbb{P}^{-+}$ or in $\mathfrak{L}_0$ are regular versions of classes from $\mathfrak{A}'$. We describe the process of constructing a regular counterpart B' of a class B of $\mathfrak{A}'$, relative to some tuple $(J_1, \ldots, J_k)$ of distinguished intersections contained in B which are treated as parameters.

We explicitly consider the case of an E_1-class B. Let $J_{k+1}, \ldots, J_l$ be all the intersections in $B \setminus (J_1 \cup \ldots \cup J_k)$ whose isomorphism types with respect to $\bar{\tau}$ are realised exactly once in $\mathfrak{B}$. For all the isomorphism types from $\bar{\Delta}$ 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.

First, an 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 $J_1 \cup \ldots \cup J_l$ together with a (minimal) set of intersections J such that for every $a \in J_i$, $1 \leq i \leq l$, and $\beta \in \boldsymbol{\beta}^{+-}[a, B \setminus (J_1 \cup \ldots \cup J_l)]$ there exists some $b \in J \subseteq C$ such that $\mathrm{tp}_{\mathfrak{B}}(a, b){\restriction}\tau = \beta$.

For the universe of B' we now put

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

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

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

We allocate the remaining 2-types inside B' in cyclic fashion as follows. The basic idea follows the model construction from [9].

For $a \in C \setminus (J_1 \cup \ldots \cup J_l)$ and $\beta \in \boldsymbol{\beta}^{+-}[a, B \setminus (J_1 \cup \ldots \cup J_l)]$ select some b of the right 1-type in D_0 and set $\mathrm{tp}_{\mathfrak{B}'}(a,b)$ such that $\mathrm{tp}_{\mathfrak{B}'}(a,b)\restriction\tau = \beta$.

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 $J_1, \ldots, J_k$ or in $D_{i'}$ for $i' = (i+1) \bmod 3$, and set $\mathrm{tp}_{\mathfrak{B}'}(a',b')$ such that $\mathrm{tp}_{\mathfrak{B}'}(a',b')\restriction\tau = \beta$.

This procedure can 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 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 $\boldsymbol{\beta}^{+-}[\mathfrak{B}]$.

Note that this procedure provides witnesses for all $\boldsymbol{\beta}^{+-}$ or $\boldsymbol{\beta}^{-+}$ requirements in E_1- or E_2-classes, so that the resulting regular classes are all completed. Thus we gurantee condition (12).

5.5.3 Removing special markers from a substructure

In definition of some components of the certificate we will need a procedure, which removes special markers from isomorphism types of some intersections of a given substructure.

Below we describe this procedure for a $\bar{\tau}$-substructure $\mathfrak{D}$ with distinguished intersections $J_1, \ldots, J_k$. Isomorphism types of $J_1, \ldots J_k$ are retained. For any other intersection J of type $\bar{\delta}$, if $\delta = \bar{\delta}\restriction\tau \notin \Delta_0$ change its extended type to δ. Please note that for such δ we have $\delta \in \bar{\Delta}$, i.e. in $\mathfrak{A}$ there is a realization of δ with both classes non-special.

The operation of this procedure may be described also as removing special markers from all non-special intersections (if they are not distinguished) in a given substructure.

5.5.4 Patterns

$\mathbb{P}^{+-}(\bar{\delta})$ is not defined unless $s(\bar{\delta}) \cap S_1 = \emptyset$. If $s(\bar{\delta}) \cap S_1 = \emptyset$, then let $\mathfrak{B}$ be a non-special E_1-class in which a realization $\mathfrak{J} \subseteq \mathfrak{B}$ of $\bar{\delta}$ is distinguished. Make the regular version of this class as described in 5.5.2 and then modify it using the procedure from subsection 5.5.3. $\mathbb{P}^{+-}(\bar{\delta})$ is set to be a singleton set containing the obtained E_1-class. Similarly for $\mathbb{P}^{-+}$ w.r.t. E_2-classes. Condiditions (12) and (13) are thus satisfied. For condition (14) observe that, for instance, for $\mathbb{P}^{+-}(\bar{\delta})$, each $U \in S_2$ marks at most one intersection in $\mathfrak{B}$ and the construction of regular classes cannot duplicate such intersections, so the same is true of $\mathfrak{B}'$.

$\mathbb{P}^{--}$ is defined precisely for $\bar{\delta}$ such that $|s(\bar{\delta})| < 2$. For such $\bar{\delta}$, $\mathbb{P}^{--}(\bar{\delta})$ is set to be a finite set of substructures, each of them consisting of two intersections in free relation. More precisely, we choose any realisation J of $\bar{\delta}$ and select a minimal set of intersections $J_1, \dots, J_k$ such that for every $a \in J$ and $\beta \in \boldsymbol{\beta}^{--}[a, \mathfrak{A}]$ there is some $b \in J_1 \cup \dots \cup J_k$ with $\mathrm{tp}_{\mathfrak{A}}(a, b) = \beta$. For every substructure $\mathfrak{A}' \restriction J \cup J_i$ with distinguished J we apply the procedure from subsection 5.5.3 and add it to $\mathbb{P}^{--}(\bar{\delta})$. This choice of $\mathbb{P}^{--}(\bar{\delta})$ guarantees conditions (15) and (16).

5.5.5 The initial part $\mathfrak{L}_0$

We define the structure $\mathfrak{L}_0$ and substructures $\mathfrak{K} \subseteq \mathfrak{C} \subseteq \mathfrak{L}_0$.

Let $\mathfrak{K}$ be the substructure of $\mathfrak{A}'$ induced by the set of intersections whose types $\bar{\delta}$ have $|s(\bar{\delta})| = 2$, i.e., the set of those intersections both of whose classes are special ($\mathfrak{K}$ may be empty). Condition (7) is satisfied, since we do not use types $\bar{\delta}$ with $s(\bar{\delta}) = 2$ in pattern classes produced by $\mathbb{P}^{+-}$ or $\mathbb{P}^{-+}$.

The universe C of $\mathfrak{C}$ contains K and 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 $\mathrm{tp}_{\mathfrak{A}}(a, b) \restriction \tau = \beta$. C thus provides witnesses for $\boldsymbol{\beta}^{--}$ requirements at all elements from K. Additionally, for every special class B, none of whose intersections has been added to C so far, we choose an arbitrary intersection of B and put it into C. We take $\mathfrak{C}$ to be the induced substructure of $\mathfrak{A}'$. (C may be empty, if there are no special classes).

In $\mathfrak{L}_0$ we complete classes of intersections from C and provide witnesses for remaining $\boldsymbol{\beta}^{--}$ requirements at elements of C. For every intersection J in C whose E_1-class is not completed in $\mathfrak{L}_0$ yet, we take its E_1-class B in $\mathfrak{A}'$ and construct its regular version $\mathfrak{B}'$ as described in section 5.5.2, treating the intersections $J_1, \dots, J_k$ from $B \cap C$ (including J) as distinguished parameters. If B is not a special class then we apply additionally the procedure from subsection 5.5.3 to $\mathfrak{B}'$ (again, treating $J_1, \dots, J_k$ as distinguished).

We adjoin the intersections in $B' \setminus (J_1 \cup \dots \cup J_k)$ to $\mathfrak{L}_0$ in such a way as to create an isomorphic copy of $\mathfrak{B}'$ in $\mathfrak{L}_0$. With E_2-classes we proceed analogously. This ensures condition (9). For condition (8) we note that, if a class $\mathfrak{B}$ is marked with $U \in S$, then every type $\bar{\delta}$, such that $s(\bar{\delta}) = \{U\}$ which is used in some pattern structure as a non-distinguished intersection is realised at least twice in $\mathfrak{B}$: otherwise, by step (2) of the procedure collecting special classes and intersections, its other equivalence class would also be special (note that in this case $\delta = \bar{\delta} \restriction \tau \in \Delta_0$) and $s(\bar{\delta}) = 2$. The construction of the regular version of $\mathfrak{B}$ retains at least two such realisations.

For every intersection J from $C \setminus K$ we produce the remaining $\boldsymbol{\beta}^{--}$ witnesses in the following way: let $J_1, \dots, J_k$ be some (minimal) choice of intersections in $\mathfrak{A}'$ such that for every $a \in J$ and $\beta \in \boldsymbol{\beta}^{--}[a, \mathfrak{A}]$ there is some $b \in J_1 \cup \dots \cup J_k$ with $\mathrm{tp}_{\mathfrak{A}}(a, b) = \beta$. For every J_i which is not a member of C, we apply procedure from 5.5.3 to $\mathfrak{J}_i$ and adjoin the result $\mathfrak{J}'_i$ to L_0 with connections between J and J'_i copied from $\mathfrak{A}'$ (each individual J_i is connected to C; no connections are introduced as yet between different J_i).

If for some J in $L_0 \setminus C$ of type $\bar{\delta}$ we have $s(J) = \{U_B\}$, we join J to the special class marked with U_B as described in section 5.4 (in connection with joining intersections to special classes in the construction of a model). This procedure can be safely performed, since $\bar{\delta}$ has to be realised at least twice in $\mathfrak{B}$. This ensures condition (6) and, since all the connections not specified yet will be some free relations, also condition (10).

To complete the definition of $\mathfrak{L}_0$, we have to define 2-types for elements not belonging to C.

First, we provide $\boldsymbol{\beta}^{--}$ witnesses for elemements of $\mathfrak{L}_0 \setminus C$ in K in accordance with conditions (9) and (11). Let $J \subseteq L_0 \setminus C$. Let $\bar{\delta}$ be the isomorphism type of J. For every $\mathfrak{D} \in \mathbb{P}^{--}(\bar{\delta})$ built from $\bar{\delta}, \bar{\delta}'$, such that $|s(\bar{\delta}')| = 2$ find $J' \subseteq K$ of type $\bar{\delta}'$ and make $\mathfrak{L}_0 \restriction (J \cup J')$ isomorphic to $\mathfrak{D}$.

Secondly, consider pairs of intersections J_1 and J_2 in $\mathfrak{L}_0$ whose connection has not yet been defined. Let $\bar{\delta}_1$ and $\bar{\delta}_2$ be their extended types. It follows that $s(\bar{\delta}_1) \cap s(\bar{\delta}_2) = \emptyset$. We therefore find some $\mathfrak{D} \in \mathbb{T}$ built from $\bar{\delta}_1$ and $\bar{\delta}_2$ and can make $\mathfrak{L}_0 \restriction (J_1 \cup J_2)$ isomorphic to $\mathfrak{D}$.

If the procedure described so far yields an empty structure $\mathfrak{L}_0$ (i.e., if there are no special classes in $\mathfrak{A}$), then we set $\mathfrak{L}_0$ to be some arbitrarily chosen single intersection from $\mathfrak{A}$, with empty distinguished substructures $\mathfrak{K}$ and $\mathfrak{C}$.

This completes the definition of the certificate. Note that we use only isomorphism types of intersections from $\bar{\Delta}$, thus (1) holds true. Moreover, every connection between a pair of intersections is either copied from $\mathfrak{A}'$ or is taken from $\mathbb{T}$, which implies condition (2).

5.5.6 The size of the certificate

It remains to show that the size of the certificate obtained in the previous section is as required in Lemma 13, i.e., at most triply exponential in $|\tau|$ and thus in $|\varphi|$.

The size of intersections in $\mathfrak{A}$ is bounded by an exponential function in $|\tau|$. By Lemma 14 the number of special classes, and thus also the size of S is bounded by a triply exponential function in $|\tau|$. This gives an at most triply exponential number of isomorphism types of intersections in $\mathfrak{A}'$ and thus also in $\bar{\Delta}$ (cf. condition (5)).

In $\mathbb{T}$, we have at most one pattern structure, consisting of two intersections, for each pair from $\bar{\Delta}$.

Every pattern class produced by $\mathbb{P}^{+-}$ or $\mathbb{P}^{-+}$ has an at most triply exponential number of intersections.

Every set of structures produced by $\mathbb{P}^{--}$ consists of an exponential number of structures, each of them consisting of two intersections.

$\mathfrak{L}_0$ contains regular versions of special classes and free witnesses for some elements. It is readily checked that its size is also bounded by a triply exponential function.

Thus, the size of the entire certificate is at most triply exponential in φ.

6 Three Equivalence Relations

It is shown in [10] 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 a stronger result. Its proof uses a special adaptation of a reduction of the domino tiling problem [3], which is particularly suited for FO^2, as presented in [21] based on work in [10, 12].

A *conservative reduction* from one logic to another is a reduction that simultaneously translates both the satisfiability problem (satisfiability over all structures) and the finite satisfiability problem (satisfiability in finite structures). The existence of a conservative reduction from FO to some other logic $\mathcal{L}$ implies not just that $\mathrm{SAT}(\mathcal{L})$ and $\mathrm{FINSAT}(\mathcal{L})$ are both undecidable, but even that the complement of $\mathrm{SAT}(\mathcal{L})$ is recursively inseparable from $\mathrm{FINSAT}(\mathcal{L})$, cf. [3]. A logic (fragment of FO) admitting a conservative reduction from FO is classically referred to as a *conservative reduction class*. These notions and this observation extend naturally to the setting with additional semantic restrictions on the class of (finite) models admitted for $\mathcal{L}$. In our case, a conservative reduction to FO^2 over equivalence structures with three equivalence relations in particular implies an analogous inseparability result concerning $\mathrm{SAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E_1, E_2, E_3])$ and $\mathrm{FINSAT}(\mathrm{FO}^2, \mathcal{EQ}[\tau_0; E_1, E_2, E_3])$.

Proposition 16 FO^2 *over* $\mathcal{EQ}[\tau_0; E_1, E_2, E_3]$ *forms a conservative reduction class.*

A closer inspection of the formulae used in the basic reductions given below, or ramifications of these, show that FO^2 can be further restricted to the two-variable guarded fragment without equality and in a vocabulary consisting of just unary predicates apart from E_1, E_2, E_3.

Corollary 17 *Satisfiability and finite satisfiability in* $\mathcal{EQ}[\tau_0; E_1, E_2, E_3]$ *are undecidable even for the following fragment of* FO^2 *in vocabularies* τ_0 *consisting only of unary predicates: conjunctions of sentences of the forms*

(a) $\forall x \forall y (E_i xy \rightarrow \chi(x, y))$*, and*

(b) $\forall x (\alpha(x) \rightarrow \exists y (E_i xy \wedge \alpha'(y)))$*,*

for quantifier-free and equality-free formulae χ *and* α *and* α'*. This fragment is in particular contained in the two-variable guarded fragment (note however, that equivalence symbols are allowed in* χ*).*

Our underlying reduction of the tiling problem to FO^2 satisfiability in restricted classes of structures, which serves to establish a conservative reduction, is closely based on the framework of [10, 12] and its ramifications in [21]. We include the following basic definitions and lemmas (without proofs) from [21] for convenience.

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

$$\mathfrak{G}_{\mathbb{Z}} = (\mathbb{Z}^2, H, V),$$

$$H = \{((p, q), (p + 1, q)) \colon p, q \in \mathbb{Z}\},$$

$$V = \{((p, q), (p, q + 1)) \colon p, q \in \mathbb{Z}\},$$

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

$$\mathfrak{G}_m = (\mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}, H, V),$$

$$H = \{((p,q),(p',q)) \colon p' - p \equiv 1 \mod m\},$$
$$V = \{((p,q),(p,q')) \colon q' - q \equiv 1 \mod m\}.$$

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

In order to allow for a simultaneous reduction of the infinite and periodic tiling problems to satisfiability and finite satisfiability, one wants to work over a sufficiently rich class of grid-like structures in the following sense.

Definition 18 An infinite structure $\mathfrak{G} = (G, H, V)$ is called *grid-like* if $\mathfrak{G}_{\mathbb{N}}$ is homomorphically embeddable into $\mathfrak{G}$; a finite $\mathfrak{G}$ is grid-like if some $\mathfrak{G}_m$ is homomorphically embeddable into $\mathfrak{G}$.

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

The following provides a simple and sufficient local criterion for grid-likeness (cf. Lemma 2.4 in [21]). We say that H is *complete over* V in $\mathfrak{G} = (G, H, V)$ if $\mathfrak{G}$ satisfies

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

Lemma 19 *Assume that $\mathfrak{G} = (G, H, V)$ satisfies the* FO^2*-axiom $\forall x(\exists y Hxy \wedge \exists y Vxy)$. If H is complete over V, then $\mathfrak{G}$ is grid-like.*

Lemma 20 *If a rich class $\mathcal{G}$ of grid-like structures is* FO^2*-axiomatisable over $\mathcal{C}$, then* $\mathrm{SAT}[\mathrm{FO}^2, \mathcal{C}]$ *and* $\mathrm{FINSAT}[\mathrm{FO}^2, \mathcal{C}]$ *are undecidable. In fact we get a conservative reduction class, so that the complement of* $\mathrm{SAT}[\mathrm{FO}^2, \mathcal{C}]$ *is recursively inseparable from* $\mathrm{FINSAT}[\mathrm{FO}^2, \mathcal{C}]$.

See [4, 13, 3] for the undecidability and recursive inseparability of corresponding tiling problems, which serve as the natural intermediaries here, and [10] and Lemma 2.3 in [21] in particular for the last lemma.

For our present purposes it therefore suffices to show the following.

Proposition 21 *There is a rich class $\mathcal{G}$ of grid-like structures that is* FO^2*-axiomatisable over $\mathcal{EQ}[\tau_0; E_1, E_2, E_3]$. In fact there is an* FO^2*-sentence η such that*

(a) *$\mathfrak{G}_{\mathbb{N}}$ and all $\mathfrak{G}_{3m}$ for $m \geq 1$ admit expansions to $\mathcal{EQ}[\tau_0; E_1, E_2, E_3]$ models of η.*

(b) *Every $\mathcal{EQ}[\tau_0; E_1, E_2, E_3]$ model of η is grid-like.*

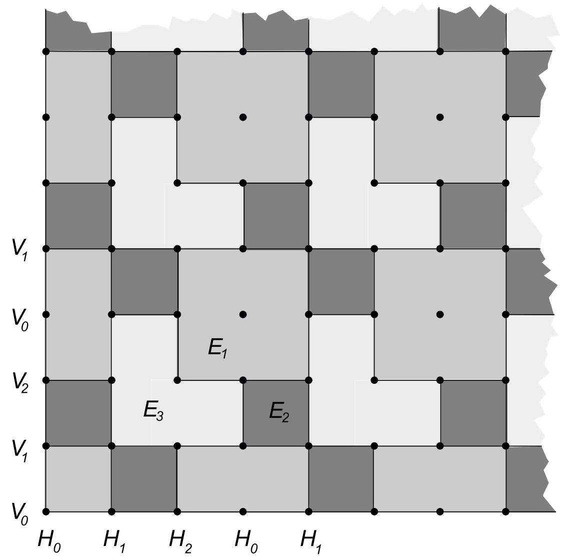

Figure 5: Expansion of $\mathfrak{G}_{\mathbb{N}}$ grid.

Proof Let us expand the standard grid $\mathfrak{G}_{\mathbb{N}}$ to the structure $\bar{\mathfrak{G}}_{\mathbb{N}}$ illustrated in Fig. 5. We capture some properties of $\bar{\mathfrak{G}}_{\mathbb{N}}$ by the sentence η and then observe that every $\mathfrak{G}_{3m}$ can also be expanded to a model of η and prove that every model of η which interprets E_1, E_2, E_3 as equivalence relations is grid-like.

The sentence η is the conjunction of the following formulae:

(1) The initial formulae

$$\exists x(H_0 x \wedge V_0 x),$$

$$\forall x(\exists y H xy \wedge \exists y V xy).$$

(2) A formula axiomatising H, which has the following shape

$$\forall xy\ (Hxy \rightarrow (\varphi_1 \vee \cdots \vee \varphi_9)),$$

where each φ_i describes one of the nine possible cases of values of unary predicates on H-related vertices and equivalences between them. For example

$$\varphi_1 \equiv E_1 xy \wedge H_0 x \wedge V_0 x \wedge H_1 y \wedge V_0 y, \quad \text{and}$$
$$\varphi_2 \equiv E_2 xy \wedge E_3 xy \wedge H_1 x \wedge V_0 x \wedge H_2 y \wedge V_0 y.$$

(3) A formula axiomatising V, which is built similarly to the one for H.

(4) A group of formulae stating that some horizontally adjacent elements that are connected by one equivalence relation are linked by H. A sample formula from this group is

$$\forall xy\ ((E_2 xy \wedge H_1 x \wedge V_1 x \wedge H_2 y \wedge V_1 y) \rightarrow Hxy).$$

It is not hard to see that every $\mathfrak{G}_{3m}$ can also be expanded to a model of η. In fact, a natural quotient of the above expansion of $\mathfrak{G}_{\mathbb{N}}$ serves this purpose.

We sketch the argument for grid-likeness of a model $\mathfrak{G} \models \eta$. By (1) $\mathfrak{G}$ satisfies $\forall x(\exists y Hxy \wedge \exists y Vxy)$ so we can use Lemma 19. Assume that for a, a', b, b'

$$\mathfrak{G} \models Hab \wedge Vaa' \wedge Vbb'.$$

We show that then $\mathfrak{G} \models Ha'b'$.

We have to consider several cases distinguished by values of the H_i and V_i on a. Let us go through one of them, for instance, $\mathfrak{G} \models H_1 a \wedge V_1 a$.

By (2) we have

$$\mathfrak{G} \models H_2 b \wedge V_1 b \wedge E_2 ab \wedge E_3 ab.$$

Similarly (3) implies

$$\mathfrak{G} \models H_1 a' \wedge V_2 a' \wedge E_2 aa' \wedge E_3 aa' \quad \text{and} \quad \mathfrak{G} \models H_2 b' \wedge V_2 b' \wedge E_3 bb'.$$

From transitivity of E_3 it follows that $\mathfrak{G} \models E_3 a'b'$. Now an appropriate formula of the form (4) guarantees $\mathfrak{G} \models Ha'b'$ which finishes the proof for this case. The remaining eight cases can be treated in the same way. ☐

Remark The technique we used in the proof was provided in [21] for extensions of FO^2. It is not difficult to adapt it to extensions of the two-variable guarded fragment GF^2 without equality. All the formulae in our proof are essentially guarded. Moreover relations H and V do not play a crucial role. They can be defined in terms of E_1, E_2 and E_3 and extra unary predicates. This means that we can show Corrolary 17. Similar result was obtained independently in [15]. For more about the satisfiability of the guarded fragment in restricted classes of models see [6, 24, 16, 17, 15, 18]

References

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

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

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

[6] H. Ganzinger, C. Meyer, M. Veanes, *The two-variable guarded fragment with transitive relations*, Proc. of 14-th IEEE Symp. on Logic in Computer Science (LICS), pages 24-34, 1999.

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

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

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

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

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

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

[13] Y. Gurevich, I. Koryakov, *Remarks on Berger's paper on the domino problem*, Siberian Mathematical Journal, 13, 1972, 319–321.

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

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

[16] E. Kieroński, *The two-variable guarded fragment with transitive guards is 2EXPTIME-hard*, FOSSACS 2003, LNCS 2620, 299-312.

[17] E. Kieroński, *Results on the guarded fragment with equivalence or transitive relations*, CSL 2005, LNCS 3634, 309-324.

[18] E. Kieroński, L. Tendera, *On finite satisfiability of the guarded fragment with equivalence or transitive guards*, LPAR 2007, LNAI 4790, 318-322.

[19] E. Kieroński, L. Tendera, *On finite decidability of two-variable first order logic with equivalence relations.*, in preparation, 2009.

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

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

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

[23] D. Scott, *A decision method for validity of sentences in two variables*, Journal of Symbolic Logic, 27, 1962, 377.

[24] W. Szwast, L. Tendera *On the decision problem for the guarded fragment with transitivity*, LICS 2001, 147-156.