On Finite Satisfiability of Two-Variable First-Order Logic with Equivalence Relations

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

Lidia Tendera
Institute of Mathematics and Informatics
Opole University
Poland
tendera@math.uni.opole.pl

***Abstract*—We show that every finitely satisfiable two-variable first-order formula with two equivalence relations has a model of size at most triply exponential with respect to its length. Thus the finite satisfiability problem for two-variable logic over the class of structures with two equivalence relations is decidable in nondeterministic triply exponential time.**

We also show that replacing one of the equivalence relations in the considered class of structures by a relation which is only required to be transitive leads to undecidability. This sharpens the earlier result that two-variable logic is undecidable over the class of structures with two transitive relations.

***Keywords*-two-variable logic, satisfiability, finite models, equivalence relations**

I. Introduction

In [21] two-variable first-order logic, FO^2, was shown to be decidable over structures with one or two equivalence relations, and undecidable over structures with at least three equivalence relations. In this paper we turn towards the *finite* satisfiability problem and provide an exact bound on the size of models for finitely satisfiable FO^2 formulas with two equivalence relations. This gives a natural decision procedure for the finite satisfiability problem for this logic. Additionally, we prove that FO^2 over structures with one transitive and one equivalence relation becomes undecidable. This extends the undecidability result for FO^2 (and even for the two-variable *guarded fragment*) over structures with two transitive relations, proved in [19] and [20].

To consider the restriction of first-order logic to its two-variable fragment is a classical idea when looking for decidability [17]. Over graphs or over any relational structures, first-order logic is undecidable, while its two-variable fragment is decidable [26]. Mortimer [26] proved that FO^2 with equality has a *finite model property*, i.e. every satisfiable FO^2-formula has a finite model. Grädel, Kolaitis and Vardi [16] proved the *exponential* model property for FO^2 which together with the lower complexity bound by Lewis [23] implied that the satisfiability problem for FO^2 is NEXPTIME-complete. As FO^2 enjoys the finite model property, the same complexity bound holds for the finite satisfiability problem.

One of the main motivations to study FO^2 comes from the fact that it embeds propositional modal logic [15] which in the last decades has been used in many areas of computer science like e.g. artificial intelligence [10], [25], program verification [14], [32], [31], database theory [13], [24] and distributed computing [11], [18]. The overall success of propositional modal logic and its numerous variants and extensions is caused by very good model-theoretic and algorithmic behaviour and in particular robust decidability which persists under many extensions towards higher expressiveness.

FO^2 is also used as a representative language for a number of knowledge representation logics (description logics) [1], cf. [2]. Moreover, many extensions of computational logics that are not fragments of FO^2 can easily be embedded in some extensions of FO^2, for example, CTL and the μ-calculus can be treated as fragments of FO^2 with a fixed-point operator [33] and many powerful variants of description logics can be embedded in the extension of FO^2 with counting quantifiers, C^2, or in C^2 with transitivity [8].

In many modal and computational logics it also is natural to restrict the classes of structures into consideration. For example, modal correspondence theory associates transitivity of the accessibility relation with the modal logic K4, and equivalence relations with S5. Linear orders are crucial for linear temporal logics, whereas tree structures are considered in computational tree logics.

Recently, extensions of FO^2 on *data* words or *data* trees are studied in the context of XML reasoning. Each position in a data word and every node of a data tree carries a label from a finite alphabet and a data value from some infinite domain. It is shown by Bojańczyk et al. [6] that satisfiability for $FO^2(\sim,<,+1)$ is decidable over finite and over infinite data words, where the language contains a binary predicate $\sim$ testing the data value equality and $<$, +1 are the usual order and successor predicates. The same authors in another paper [7] have shown that satisfiability for FO^2 for data trees is decidable if the tree structure can

Supported by Polish Ministry of Science and Higher Education research project nr N206 022 31/3660 2006/2009

be accessed only through the child and the next sibling predicates (corresponding to successor relations) and the access to data values is restricted to equality tests (which is an equivalence relation). Decidability of $\mathrm{FO}^2(\sim,<,+1)$ over data trees remains open. Unfortunately, our undecidability result for FO^2 over structures with one equivalence and one transitive relation does not imply a solution of the previous problem.

There is also a growing interest in hybrid logics, i.e. modal or related logics with *nominals*, see e.g. [4], [5]. Nominals denote singleton sets, they are expressible in FO^2 and are interpreted as *kings* in FO^2 models. They occur naturally in ontologies defined in OWL - the language for publishing and sharing ontologies on the World Wide Web that is now the W3C recommendation [3]. As hybrid logics are quite powerful and easily become undecidable, restrictions of the underlying classes of structures has become one of the directions to retain decidability despite of using an expressive language. See [28] and [27] for a complete study of the complexity of hybrid logics over, respectively, transitive and equivalence structures.

It is obvious that structures considered in practical applications are often required to be finite. So, lack of the finite model property for a logic naturally leads to the question, whether its finite satisfiability problem (or related reasoning tasks) is decidable.

Decision procedures for the finite satisfiability problem for modal or related logics are usually more complex than in the arbitrary case. As argued in [33], the model theoretic reason for the good behaviour of modal logics is the tree model property. Tree models, obtained by unravelling of arbitrary models, are usually infinite and therefore are not very useful in the study of the finite satisfiability problem.

From the above mentioned fragments, the finite model property is enjoyed by pure FO^2 and FO^2 over structures with one equivalence relation. Other extensions of FO^2 contain *infinity axioms*, i.e. sentences satisfiable only over infinite models. And so, for the latter logics, the finite satisfiability problem and the satisfiability problem do not coincide.

In this paper we are concerned with the finite satisfiability problem for FO^2 over equivalence structures. As we have already mentioned, FO^2 with three equivalence relations is undecidable and FO^2 with one equivalence relation enjoys the finite model property. So, the interesting case is FO^2 over the class of structures with two equivalence relations. Below we recall an infinity axiom from [21]. Besides equivalence relation symbols E_1, E_2, it uses unary symbols P, Q and S. It is not difficult to formalize the following statements by an FO^2 formula:

$\star$ P and Q are disjoint and each E_i-class contains at most one element from P and one from Q; the E_2-class of any element of S is trivial (a singleton).

$\star$ every element of P is E_1-equivalent to one in Q; every element of Q is E_2 equivalent to one in P.

$\star$ $S \cap P \neq \emptyset$.

In any model of this formula an infinite chain of elements, starting from an element in $S \cap P$ and then joining new elements in Q and P in an alternate fashion along links in E_1 and E_2, respectively, is embedded (see Figure 1).

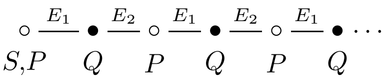

Figure 1. Infinite chain

In the proof of our result that for every FO^2 sentence φ that is finitely satisfiable over structures with two equivalence relations, one can find in this class of structures a model triply exponential with respect to $|\varphi|$, we combine new ideas with known techniques used for (finite) satisfiability for variants of two-variable logics.

We recall first a result from [21] that, over the class of structures with two equivalence relations (E_1 and E_2), every finitely satisfiable FO^2 sentence has a model with exponentially bounded *intersections* (i.e., equivalence classes of the common refinement $E = E_1 \cap E_2$ of E_1 and E_2). Then, reviving some ideas from [22], we show that every equivalence class in such a model can be appropriately simulated (from local point of view) by its *completed subclass* of only double exponential size. In the above step we apply in particular the adaptation of the *small substructures lemma* ([21]) to the level of intersections. As a next step, in any *regular* model (i.e., a model with exponential intersections and distinguished small completed subclasses in every class) we identify a bounded number of *special* classes, which define an initial fragment of the required small model. Some properties of the remaining part of the model are captured by a system of linear inequalities whose small solution establishes the domain of the constructed model. The structure of the model is defined by applying some combinatorics and graph theoretic properties.

The construction yields a triply-exponential upper bound on the size of finite models. This bound is optimal since in [21] it is shown how to construct for a given n a (finitely satisfiable) formula of length polynomial in n whose all models have at least $2^{2^{2^n}}$ elements. Our result gives naturally a 3NEXPTIME-upper complexity bound for the finite satisfiability problem for FO^2 over the class of structure with two equivalence relations. We note that it is still higher than the best lower bound, 2EXPTIME [20], and that the tight complexity bound is left open, for both the satisfiability and the finite satisfiability problem.

It is worth mentioning that the presented construction differs significantly from the constructions used for unrestricted satisfiability for FO^2 with two equivalence relations [21] and for finite satisfiability of the two-variable guarded fragment with equivalence or transitive relations [22]. A regular (infinite) model for FO^2 even though it is far from being a tree, is essentially constructed in a tree-like manner: in a given model its central part (root) is identified and, raughly speaking, the remaining part of the model is unrevealed, with some back connection to the root created. In the case of the guarded fragment we have a very convenient property that an elemens may be copied as may times as required. In constrast, in FO^2 one can enforce that some types of elements, intersections or classes have to be realized some specific number of times (even triply exponential in the size of the formula, cf. [21]) in the entire model.

Undecidability of FO^2 over structures with one equivalence and one transitive relations is obtained in a rather standard fashion, by axiomatization of a rich class of grid-like structures that suffice to encode undecidable domino tilings [34].

II. Equivalence structures and normal form

Without loss of generality we restrict our attention to signatures containing only unary and binary symbols. We assume that a signature is of the form $\tau = \tau_0 \dot\cup \{E_1, E_2\}$, with E_1, E_2 distinguished binary symbols. By $\mathcal{EQ}[\tau_0; E_1, E_2]$ we denote the class of all *equivalence structures* in the vocabulary $\tau = \tau_0 \dot\cup \{E_1, E_2\}$, i.e., the class of all $\tau_0 \dot\cup \{E_1, E_2\}$ structures $\mathfrak{A} = (\mathfrak{A}_0, (E_1^{\mathfrak{A}}, E_2^{\mathfrak{A}}))$ that interpret E_1, E_2 as equivalence relations.

We use superscripts $+$ and $-$ to indicate which equivalences/non-equivalences a given formula stipulates. For instance, for a quantifier free formula $\chi = \chi(x, y) \in \mathrm{FO}^2[\tau_0]$ we let

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

We say that an FO^2 sentence φ is in $\mathcal{EQ}[\tau_0, E_1, E_2]$-normal form if

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

for quantifier-free $\chi \in \mathrm{FO}^2[\tau_0; E_1, E_2]$, $\chi_i \in \mathrm{FO}^2[\tau_0]$ and $s_i \in \{+,-\} \times \{+,-\}$.

For a normal form φ and $s \in \{+,-\} \times \{+,-\}$ we denote by φ_s the conjunction of all $\forall x \exists y\ \chi_i^{s_i}$ with $s_i = s$. We also let $\varphi_{univ} = \forall x \forall y\ \chi$. Thus

$$\varphi = \varphi_{univ} \wedge \varphi_{++} \wedge \varphi_{+-} \wedge \varphi_{-+} \wedge \varphi_{--}.$$

The following theorem can be proved in a standard fashion.

Theorem 1: For every formula φ over $\tau = \tau_0 \cup \{E_1, E_2\}$ there exist polynomially computable $\widetilde{\tau}_0 \supseteq \tau_0$ and a formula $\widetilde{\varphi}$ in $\mathcal{EQ}[\widetilde{\tau}_0, E_1, E_2]$-normal form, which is (finitely) satisfiable over $\mathcal{EQ}[\widetilde{\tau}_0; E_1, E_2]$ if and only if φ is (finitely) satisfiable over $\mathcal{EQ}[\tau_0; E_1, E_2]$. Moreover, every model of φ can be expanded to a model of $\widetilde{\varphi}$ and the restriction of a model of $\widetilde{\varphi}$ to τ is a model of φ.

The process of building a model for a normal form formula φ proceeds by providing *witnesses* for all elements, without violating φ_{univ} constraints. A witness for an element $a \in A$ and a conjunct $\forall x \exists y \chi_i^s$ is an element $b \in A$, such that $\mathfrak{A} \models \chi_i^s(a, b)$. Note that, for an element a, its witnesses for φ_{--} conjuncts are located outside its equivalence classes; for φ_{+-} conjuncts – inside its E_1-class but outside its E_2-class; for φ_{-+} conjuncts – inside its E_2-class but outside its E_1-class; and for φ_{++} conjuncts – inside its both E_1-class and E_2-class. Witnesses for formulas from φ_{--} are sometimes called *free witnesses*.

III. Completed subclasses

The equivalence classes of the relation $E_1 \cap E_2$ are called *intersections*. We recall a lemma which was proved in Section 5.3 of [21].

Lemma 2: Every sentence φ in $\mathcal{EQ}[\tau_0, E_1, E_2]$-normal form, (finitely) satisfiable over the class $\mathcal{EQ}[\tau_0, E_1, E_2]$ has in this class a (finite) model whose all intersections have size bounded exponentially in $|\tau_0|$.

From this point we will only work with models with such small intersections. Moreover our constructions will proceed on the level of intersections rather than individual elements. It means, that intersections will be treated as basic building blocks for structures.

For a given structure $\mathfrak{A}$ we denote by $\Delta(\mathfrak{A})$ the set of the isomorphism types of intersections in $\mathfrak{A}$. Note that due to our assumption on the size of intersections, the size of $\Delta(\mathfrak{A})$ is bounded doubly exponentially in the size of the signature.

We say that an intersection type $\delta \in \Delta(\mathfrak{A})$ is *royal* for E_i-classes if in every E_i-class in $\mathfrak{A}$ there is at most one intersection of type δ.[1]

We say that B_0 is a *subclass* of an E_i-class B if $B_0 \subseteq B$ and every intersection in $\mathfrak{B}_0$ is also an intersection in $\mathfrak{B}$, i.e., $\mathfrak{B}_0$ contains only whole intersections of $\mathfrak{B}$.

For a given structure $\mathfrak{A}$ and a normal form sentence φ we say that a subclass B_0 of an E_1-class B in $\mathfrak{A}$ is *completed* (for φ) if $\mathfrak{B}_0 \models \varphi_{univ} \wedge \varphi_{+-} \wedge \varphi_{++}$. Similarly a subclass B_0 of an E_2-class B is completed if $\mathfrak{B}_0 \models \varphi_{univ} \wedge \varphi_{-+} \wedge \varphi_{++}$. In other words, an E_1-subclass B_0 is completed if all of its elements have all the required φ_{++}- and φ_{+-}-witnesses inside B_0. Analogously for E_2-subclasses.

Since a class is its own subclass, we can in a natural way speak about completed/non-completed classes. It is clear that all equivalence classes in a model of φ are completed.

In the following fact we observe that to every completed class B we can add a non-royal intersection of a type which

[1] Realizations of types royal for E_i-classes play role of *kings* inside E_i-classes, cf. [16].

is realized in B, in such a way that the extended class remains completed.[2]

Fact 3: Let $\mathfrak{B}$ be a completed E_i-class in a model $\mathfrak{A}$ of a normal form sentence φ. Let δ be a type of an intersection realized in $\mathfrak{B}$, which is not royal for E_i-classes. Then $\mathfrak{B}$ can be extended to a completed class $\mathfrak{B}'$ containing an additional intersection of type δ.

Proof: Let J' be an intersection in B of type δ. We obtain B' by extending B by a new intersection J of type δ. For every intersection J'' of B, $J'' \neq J'$, set the connection between J and J'', i.e. the 2-types realized by pairs of elements from, respectively, J and J'', isomorphically to the connection between J' and J''. This ensures all the required witnesses for J inside B'. Find any E_i-class in $\mathfrak{A}$ with two realizations of δ (such a class exists since δ is not royal for E_i-classes) and set the connection between J and J' isomorphically to the connection between those realizations. Together with the previous step this ensures that 2-types connecting elements of J with elements from B are consistent with φ_{univ}. ■

In the next Lemma we show that a model of φ can be modified in such a way that in every class B we can distinguish its small completed subclass containing realizations of all intersection types appearing in $\mathfrak{B}$.

Lemma 4: Let φ be an FO^2 sentence in $\mathcal{EQ}[\tau_0, E_1, E_2]$-normal form, (finitely) satisfiable over $\mathcal{EQ}[\tau_0, E_1, E_2]$. Then there exists a natural number M, doubly exponential in $|\tau_0|$, and a (finite) model $\mathfrak{A} \models \varphi$ in $\mathcal{EQ}[\tau_0, E_1, E_2]$ with intersections bounded exponentially in $|\tau_0|$, whose every E_i-class B contains a completed subclass B' with at most M intersections, realizing all intersection types from B.

This Lemma is similar in spirit to Lemma 1, part (ii) from [22]. In the proof we use a natural adaptation of the *small substructures lemma* from [21] to the level of intersections. The adapted small substructures lemma allows us, for a given class B, to build a small (doubly exponential with respect to the signature) completed class B', realizing all (and only) isomorphism types of intersections from B, in such a way that the number of realizations of every type δ in B' is less or equal to the number of realizations of δ in B. Thus we can distinguish in B its small subclass B'', make it isomorphic to B', and then join all the intersections from $B \setminus B''$ to B'' using Fact 3, without changing their connections to the remaining part of the model. We skip the details of the proof.

IV. Small model property

In this section we sometimes use the following notation: $\widetilde{i}$ denotes 1 for $i = 2$ and 2 for $i = 1$. This is convenient when speaking about relations E_1 and E_2.

[2]Note however that we do not say how to connect the additional intersection to the remaining part of the structure $\mathfrak{A}$. Thus, we do not build an entire model of φ.

Let φ be an FO^2 sentence in $\mathcal{EQ}[\tau_0, E_1, E_2]$-normal form, finitely satisfiable over $\mathcal{EQ}[\tau_0, E_1, E_2]$. We show that φ has a model of size bounded triply exponentially in $|\tau_0|$. To make the presentation more readable we assume first that φ_{--} is empty.

Let M and $\mathfrak{A}$ be as in Lemma 4. Let $\Delta = \Delta(\mathfrak{A})$. Let $subcl$ be a function which for every E_i-class B in $\mathfrak{A}$ returns a small completed subclass of B whose existence is postulated in Lemma 4. We construct a new small model $\mathfrak{A}'$. In $\mathfrak{A}'$ we use only intersection types from Δ, which allows us not to bother about conjuncts from φ_{++} – they will be automatically satisfied.

The proof consists in providing completed E_1- and E_2-classes for all intersections in the model. We first distinguish in $\mathfrak{A}$ its fragment containing intersections of such types δ which are realized in a *small* number of classes. Completed subclasses of this fragment will form the initial part of our new model $\mathfrak{A}'$. Then we add copies of small completed subclasses of some other classes from $\mathfrak{A}$ to $\mathfrak{A}'$. The number of these copies will be given by a solution to a system of linear inequalities, which will allow us to solve the most difficult task: providing second classes for royal intersections. Non-royal intersections will be joined to some other classes using Fact 3.

A. Special classes

Let $K = 3M|\Delta|$. We distinguish a set $S = S_1 \cup S_2$ of *special* classes, consisting of two parts S_1 and S_2 of, respectively, special E_1- and special E_2-classes. S_i is defined by the procedure which starts with $S_i = \emptyset$ and repeats the following step as long as the fixed point is not obtained:

- If there exists $\delta \in \Delta$, such that the number of E_i-classes not belonging to S_i, containing a realisation of δ, is less than K, then add all these classes to S_i.

The following observations are straightforward.

Fact 5: The number of classes in S_i is bounded by $K|\Delta|$, which is doubly exponential in τ_0.

Fact 6: If an intersection type δ is realized in a non-special E_i-class then it is realized in at least K non-special E_i-classes.

Fact 7: If J_1, J_2 are intersections in $\mathfrak{A}$ of (not necesserily different) types δ_1, δ_2, respectively, not contained in the same special class then there exist intersections J_1', J_2' in $\mathfrak{A}$, realizing δ_1, δ_2 in a free connection, i.e., for every $a_1 \in J_1', a_2 \in J_2'$ we have $\mathfrak{A} \models \neg E_1 a_1 a_2 \wedge \neg E_2 a_1 a_2$.

For further purposes we denote by $\Delta_N^{E_i}$ the subset of Δ containing those types of intersections which are realized in non-special E_i-classes. Note that they can also appear in special E_i-classes.

B. Initial part of $\mathfrak{A}'$

We choose a small substructure $\mathfrak{C}$ of $\mathfrak{A}$ which after minor modifications will become an initial part $\mathfrak{C}'$ of $\mathfrak{A}'$. Let S' be the set of those classes in $\mathfrak{A}$ that have at least one intersection in $subcl(B)$ for some $B \in S$. Obviously $S \subseteq S'$. Note that the size of S' is still doubly exponentially bounded in $|\tau_0|$. We define C to consists of all intersections from $subcl(B)$ for all $B \in S'$.

Now we slightly reorganise $\mathfrak{C}$ obtaining $\mathfrak{C}'$. Firstly, we make all the counterparts of classes from S' in $\mathfrak{C}'$ completed:[3] if an intersection $J \subseteq C$ is a fragment of $B \setminus subcl(B)$ for some class $B \in S'$ then, if necessary, we change its connections to $subcl(B)$ using Fact 3, to ensure that elements from J have all the required witnesses inside $subcl(B)$. Secondly, if a pair of intersections J, J' in C, not belonging to any special classes and not belonging to the same class from S' is connected by an equivalence relation, then we change this connection to a free connection, consistent with φ_{univ}. An appropriate pattern for such a connection can be found by Fact 7.

We stress the obtained properties in the following observation:

Fact 8: If an intersection J in $\mathfrak{C}'$ of type δ has not its E_i-class completed in $\mathfrak{C}'$ then:

(i) J is not a part of any special class and thus δ is realized in at least K non-special E_1-classes and at least K non-special E_2-classes in $\mathfrak{A}$.

(ii) J is not E_i-connected to any other intersection in $\mathfrak{C}'$.

C. Colours

In this section we introduce sets F^{E_1} and F^{E_2} of *colours* for, respectively, E_1- and E_2-classes, and mark with them some intersections in $\mathfrak{A}$. An intersection can be coloured by at most one colour from F^{E_1} and at most one from F^{E_2}. Intersections from the same E_i-class cannot have different F^{E_i}-colours. For $i = 1, 2$, $F^{E_i} = F_S^{E_i} \dot\cup F_N^{E_i}$, where $F_S^{E_i}$ is the set of colours for special classes and $F_N^{E_i}$ is the set of colours for non-special classes. As we will see later the purposes of $F_S^{E_i}$ and $F_N^{E_i}$ are different.

For every special E_i-class $B \in S_i$ there is a unique colour $f_B \in F_S^{E_i}$. The sets $F_N^{E_1}$ and $F_N^{E_2}$ contain $M(M-1)+1$ unique colours each. We colour intersections in $\mathfrak{A}$ in the following way:

Colours for special classes. For $B \in S_i$, all intersections belonging to $B \cap subcl(B')$ for some class B' (B' may be equal to B) are coloured by $f_B \in F_S^{E_i}$.

Colours for non-special classes. For $i = 1, 2$, define a graph whose vertices are non-special E_i-classes of $\mathfrak{A}$, and there is an edge between classes B, B' if and only if there exists a non-special $E_{\tilde{i}}$-class B'' and intersections $J, J' \subseteq subcl(B'')$, such that $J \subseteq subcl(B)$ and $J' \subseteq subcl(B')$. Associate a colour $c(B)$ from $F_N^{E_i}$ with every non-special E_i-class B in such a way that adjacent classes have different colours. Such a colouring exists since the degree of a vertex in the defined graph is bounded by $M(M-1)$. Finally colour all intersections in $subcl(B)$ with $c(B)$ for every non-special class B.

[3] This is not necessarily the case in $\mathfrak{C}$, since, for instance, an intersection in C may belong to two special classes B_1, B_2, but not belong to $subcl(B_1)$ or $subcl(B_2)$, and thus it may not have all the φ_{+-}- or φ_{-+}-witnesses inside C.

Below we collect some observations about the defined colouring.

Fact 9: (i) Two intersections belonging to $subcl(B)$ for a non-special E_i-class B are not coloured by the same colour from $F^{E_{\tilde{i}}}$.

(ii) An intersection whose type is royal for E_i-classes is coloured by a colour from F^{E_i}. Thus an intersection whose type is royal for both E_1- and E_2-classes is coloured by two colours.

(iii) An intersection J whose type is royal for E_i-classes, and $J \subseteq subcl(B)$ for an $E_{\tilde{i}}$-class B, is coloured by two colours.

To see (ii) and (iii) note that an intersection J whose type is royal for E_i-classes has to belong to $subcl(B)$ for its E_i-class B, by the definition of the function $subcl$.

Remark. Colours from $F_N^{E_1} \cup F_N^{E_2}$ are relevant only for intersections that are royal for E_1- or E_2-classes. We coloured with them entire subclasses $subcl(B)$ to simplify the presentation.

D. System of linear inequalities

A *coloured type* of an intersection J in $\mathfrak{A}$ is a triple $\overline{\delta} = (\delta, c_1, c_2)$, where $\delta \in \Delta$ is the isomorphism type of J and $c_i \in F^{E_i} \cup \{null\}$ is the colour J obtained in the colouring process in the previous subsection ($c_i = null$ means that J was not a coloured by F^{E_i}). Let $\overline{\Delta}$ be the set of the coloured types of intersections realized in $\mathfrak{A}$. Note that $|\overline{\Delta}|$ is bounded doubly exponentially in the size of the signature.

A *counting type* of a class B in $\mathfrak{A}$ is the function from $\overline{\Delta} \to \{0, \ldots, M\}$ which for a given coloured intersection type $\overline{\delta} \in \overline{\Delta}$ says how many times it is realized in $subcl(B)$. Let Θ_i be the set of the counting types of non-special E_i-classes in $\mathfrak{A}$. Note that $|\Theta_i|$ is bounded triply exponentially in the size of the signature.

For every $\theta \in \Theta_i$ we introduce a variable $X_\theta^{E_i}$ counting the number of E_i-classes of type θ. For $i = 1, 2$ and every $\delta \in \Delta_N^{E_i}$, we write an inequality stating that δ is realized in at least K non-special E_i-classes:

$$(\mathbf{I}_\delta^{\mathbf{E_i}}): \sum_{\{\theta \in \Theta_i:\ \text{for some } \overline{\delta} = (\delta, c_1, c_2)\ \theta(\overline{\delta}) > 0\}} X_\theta^{E_i} \geq K.$$

Let $k_\theta^{E_i}$ be the number of E_i-classes of type θ in $S' \setminus S$. For every θ and $i = 1, 2$, if $k_\theta^{E_i} > 0$ we state that the number of classes of type θ is at least $k_\theta^{E_i}$:

$$(\mathbf{I}_\theta^{\mathbf{E_i}}): X_\theta^{E_i} \geq k_\theta^{E_i},$$

moreover, if $k_\theta^{E_i} = 1$ and there is only one E_i-class of the counting type θ in the whole model $\mathfrak{A}$, we strengthen the above inequality to

$$(\mathbf{E}^{\mathbf{E_i}}_\theta) : X^{E_i}_\theta = 1.$$

For every $\overline{\delta} = (\delta, c_1, c_2)$, such that $c_1 \in F^{E_1}_N$, $c_2 \in F^{E_2}_N$, we write an equation stating that $\overline{\delta}$ is realized exactly the same number of times in non-special E_1- and E_2-classes (note that by Fact 9, part (i), such $\overline{\delta}$ may be realized at most once in a class):

$$(\mathbf{E}_{\overline{\delta}}) : \sum_{\{\theta\in\Theta_1:\theta(\overline{\delta})=1\}} X^{E_1}_\theta = \sum_{\{\theta\in\Theta_2:\theta(\overline{\delta})=1\}} X^{E_2}_\theta.$$

Fact 10: The constructed system of inequalities and equalities has a non-negative integer solution.

Proof: To obtain a solution it is enough to take the number of realizations of $\theta \in \Theta_i$ in $\mathfrak{A}$ as $X^{E_i}_\theta$. ■

E. Small model construction

Building on the classical result [30] and following [12], we first observe that our system of inequalities has a *small* solution.

Lemma 11 (Calvanese): Let Γ be a system of k linear inequalities in n unknows, and let the coefficients and constants that appear in the inequalities be in $\{-a, -a+1, \ldots, a-1, a\}$. If Γ admits a nonnegative integer solution, then it also admits one in which the values assigned to the unknows are all bounded by $(n+k)\cdot(k\cdot a)^{2k+1}$.

In our system a, k are bounded doubly exponentially, and n triply exponentially in $|\tau_0|$. Thus it has a non-negative integer solution in which all variables have values bounded triply exponentially in $|\tau_0|$. Let $X^{E_i}_\theta := r^{E_i}_\theta$ be such a solution. The solution will determine the cardinality of our new model.

In the process of building the new model we use coloured types of intersections. The new model $\mathfrak{A}'$ consists of its initial part $\mathfrak{C}'$ as defined in section IV-B and, for every counting type θ and $i = 1, 2$, $r^{E_i}_\theta - k^{E_i}_\theta$ disjoint copies of completed E_i-classes, each of them containing $\theta(\overline{\delta})$ realizations of $\overline{\delta}$ (an E_i-class of type θ is made isomorphic to $subcl(B)$ for some class B in $\mathfrak{A}$ of type θ; colours are also retained).

At this point every intersection in the partially defined $\mathfrak{A}'$ has at least one completed class. In the further process we take care of providing the other completed class for all intersections. This is obtained by identifying some intersections in non-special E_i-classes with intersections of the same coloured types in non-special $E_{\widetilde{i}}$-classes and by joining some intersections from non-special E_i-classes to special or non-special $E_{\widetilde{i}}$-classes.

Identifying intersections. Consider a coloured type $\overline{\delta}$ of the form $\overline{\delta} = (\delta, c_1, c_2)$ for $c_1 \in F^{E_1}_N$ and $c_2 \in F^{E_2}_N$. At this point the only intersections of this type which may have both classes completed belong to two classes from $S' \setminus S$. Let J be an intersection in $\mathfrak{A}'$ of coloured type $\overline{\delta}$ in a non-special E_1-class whose E_2-class has not been defined yet. We find an intersection J' of coloured type $\overline{\delta}$ in a non-special E_2-class, such that the E_1-class of J' is not defined yet, and identify J with J'. Observe that even if J belongs to $\mathfrak{C}'$ then by Fact 8, part (ii) its non-completed class in $\mathfrak{C}'$ consists of just of one intersection, so we do not have to join any additional intersection to the class of J'; similarly if J' belongs to $\mathfrak{C}'$.

Since the equality $(\mathbf{E}_{\overline{\delta}})$ states that $\overline{\delta}$ is realized exactly the same number of times in non-special E_1- and E_2-classes, after this step all intersections coloured with two non-special colours have both their classes completed.

Note, that the identification step is safe in that it cannot identify two intersections from an E_1-class B with two intersections from an E_2-class B' (which would be incorrect, since two classes have obviously at most one intersection in common) - this is guaranteed by Fact 9, part (i).

Joining intersections to special classes. Let J of the coloured type $\overline{\delta} = (\delta, f_B, null)$, for some $f_B \in F^{E_1}_S$, be an intersection in $A' \setminus C'$ whose E_1-class has not been defined yet. Let it belong to a completed E_2-class of type θ in $\mathfrak{A}'$. Note that in $\mathfrak{A}$ the special E_1-class B has at least two intersections of type δ. Assume to the contrary that it has only one, say J_1. Then, by the definition of the function $subcl$ in Lemma 4, J_1 belongs to $subcl(B)$ and has its E_2-class B' completed in $\mathfrak{C}'$. B' has type θ, since J_1 is the only intersection in $\mathfrak{A}$ of the coloured type $\overline{\delta}$. In this case the equality $(E^{E_2}_\theta)$ states that there is only one realization of θ in $\mathfrak{A}'$. Contradiction. Thus δ is not royal for E_1-classes. We extend the class B, whose initial part was defined in $\mathfrak{C}'$ by J, using Fact 3. We proceed analogously with intersections whose coloured types are of the form $(\delta, null, f_B)$, joining them to special E_2-classes.

After this step none of the special classes will be further extended and all intersections coloured by a colour from $F^{E_1}_S \cup F^{E_2}_S$ have both their classes completed.

Joining intersections to non-special classes. For $i = 1, 2$ we define a division

$$\{U^{E_i}_\delta\}_{\delta\in\Delta^{E_i}_N} \cup \{V^{E_i}_\delta\}_{\delta\in\Delta^{E_i}_N} \cup \{W^{E_i}_\delta\}_{\delta\in\Delta^{E_i}_N}$$

of non-special E_i-classes containing at least one intersection not coloured by $F^{E_{\widetilde{i}}}_S$.[4] Each of the sets $U^{E_i}_\delta$, $V^{E_i}_\delta$, $W^{E_i}_\delta$ contains at least M E_i-classes, each of them with an intersection of type δ. Colours of intersections are not relevant in the remaining part of this paragraph.

The desired division of the set of non-special E_i-classes can be obtain in a greedy manner: enumerate the types $\delta_1, \delta_2, \ldots, \delta_k \in \Delta^{E_i}_N$. Choose M non-special E_i-classes with intersections of type δ_1 and put them to $U^{E_i}_{\delta_1}$. Do the same for $V^{E_i}_{\delta_1}$ and $W^{E_i}_{\delta_1}$. After this step there is $3M$ E_i-classes used. Continue to δ_2, choosing appropriate fresh non-special classes for $U^{E_i}_{\delta_2}$, $V^{E_i}_{\delta_2}$, $W^{E_i}_{\delta_2}$, then to δ_3 and so on.

[4] If a non-special E_i-class has only intersections coloured by $F^{E_{\widetilde{i}}}_S$-colours then all intersections in this class have its $E_{\widetilde{i}}$-classes already defined.

The choice of the number $K = 3M|\Delta|$ guarantees that we have enough appropriate non-special classes. Eventually we have usually some non-special E_i-classes remained. A class B is put to $U^{E_i}_\delta$ for some δ realized in $\mathfrak{B}$.

We let $U^{E_i} := \{U^{E_i}_\delta\}_{\delta \in \Delta^{E_i}_N}$, $V^{E_i} := \{V^{E_i}_\delta\}_{\delta \in \Delta^{E_i}_N}$ and $W^{E_i} := \{W^{E_i}_\delta\}_{\delta \in \Delta^{E_i}_N}$.

The general idea is that, for $i = 1, 2$, intersections from classes from U^{E_i} are joined to classes from $V^{E_{\bar{i}}}$, those form V^{E_i} to $W^{E_{\bar{i}}}$ and those from W^{E_i} to $U^{E_{\bar{i}}}$.

Let B be a non-special E_i-class in U^{E_i}. We repeat the following procedure for all its intersections. Let J be an intersection of B of type δ whose $E_{\bar{i}}$-class has been not defined yet. Note that in this case δ is not royal for $E_{\bar{i}}$-classes, since otherwise J would be coloured by two colours (by Fact 9, part (iii)) and its second class would have been defined in the step of identification. Thus, using Fact 3, we can safely join J to an $E_{\bar{i}}$-class B' in $V^{E_{\bar{i}}}_\delta$. When choosing B' we have to ensure that none of the intersections of B has been earlier identified with an intersection of B' and none of the intersections of B has been joined to B'. It is always possible because we have at least M $E_{\bar{i}}$ classes in $V^{E_{\bar{i}}}_\delta$, and M is the maximal number of intersections in B whose second class is not completed. We repeat this procedure for all non-special classes in U^{E_i}.

Then we proceed analogously with V^{E_i}, joining intersection from its classes to $W^{E_{\bar{i}}}$ and with W^{E_i}, joining intersections to $U^{E_{\bar{i}}}$. This circular scheme of joining intersections to other classes guarantees that we do not join an intersection from a class B to a class B' and then an intersection from B' to B, which would not be sound.

Setting free connections. Finally we have to set the remaining non-equivalence connections between intersections in $\mathfrak{A}'$ not belonging to the same E_1- or E_2-class, without violating φ_{univ} formula. Let J_1, J_2 be a pair of intersections in $\mathfrak{A}'$ not belonging to the same E_1- or E_2-class, of coloured types $\overline{\delta}_1 = (\delta_1, c_1, c_2)$, $\overline{\delta}_2 = (\delta_2, c_3, c_4)$, respectively. Let J'_1, J'_2 be any intersections of coloured types $\overline{\delta}_1, \overline{\delta}_2$ in $\mathfrak{A}$. Note that J'_1 and J'_2 do not belong to the same special class since in that case their coloured types would contain the same special colour and thus J_1 and J_2 would be members of the same special class in $\mathfrak{A}'$. Using Fact 7 we find intersections J''_1, J''_2 in a free connection in $\mathfrak{A}$ of types δ_1, δ_2, respectively. We set the connection between J_1 and J_2 in $\mathfrak{A}'$ isomorphically to the connection between J''_1 and J''_2 in $\mathfrak{A}$.

This finishes our construction. We thus have proved the following theorem.

Theorem 12: Every normal form formula φ with empty φ_{--}, finitely satisfiable over the class $\mathcal{EQ}[\tau_0; E_1, E_2]$, has a model in this class of size at most triply exponential in the size of $|\tau_0|$.

F. The case of free witnesses

To deal with the case of formulas which may require free witnesses, i.e., formulas with non-empty part φ_{--}, we modify the described process and combine it with some ideas from the small model construction for plain FO^2 by Grädel, Kolaitis, Vardi [16]. We outline the required changes and additions to the construction here. The detailed proof will be presented in the full version of this paper.

Special classes. The notion of special classes is slightly generalised. We add the following step which is repeated as long as the set S is changing: if all realisations of an intersection type can be *covered* by classes from S and by a *small* set T of E_1- and E_2-classes not from S, then we add all classes from T appropriatly to S_1 and S_2.

This allows us to prove the following fact: if an isomorphism type δ has a non-special realization, i.e., a realization whose both classes are non-special, then there exists a *large* set of *independent* non-special realizations of δ, where by independent we mean realizations that are, pairwise, in a free connection.

Initial part $\mathfrak{C}'$. During the construction of the initial part $\mathfrak{C}'$ of the small model we embed in it the set of *kings* and *nobles* (cf. [16]). For a special class B we add to the set of kings those intersections whose (coloured) types are realized in B *small* number of times. The set of nobles is some minimal set of intersections providing all the required free witnesses for the kings.

In $\mathfrak{C}'$ we also secure enough intersections to provide witnesses for all elements of the constructed model: we add large independent sets of non-special intersections (each with its two completed classes) of all possible types, and a sufficient number of elements in every special class B, whose type is realized in B *large* number of times.

We ensure both completed classes for all intersections from special classes in $\mathfrak{C}'$ and at least one completed class for the other intersections (as it was done in Section IV-B). We retain the connections inside the set of kings, and between kings and nobles from $\mathfrak{A}$, but in contrast to the construction from [16], we do not set free connections between nobles (and other elements inside $\mathfrak{C}'$) at this moment.

Providing free witnesses. The colouring process and the system of inequalities are as in the restricted case from the previous subsections. Both classes for all intersections are provided also in the similar way, by the procedures of identifying intersection and joining them to classes.

After this step we provide free witnesses for all elements. Kings have their witnesses ensured among nobles. All other elements can find its witnesses inside $\mathfrak{C}'$. Members of $\mathfrak{C}'$ do it in a circular way, similarly to the scheme from [16].

For example when an element looks for a witness in a non-special intersection, then it always has a lot of choices, since it can be connected by an equivalence relation only to at most two members of the independent set of intersections of a particular type.

The construction sketched above, together with Theorem 1, lead to the following result.

Theorem 13: Every FO^2 sentence finitely satisfiable over the class $\mathcal{EQ}[\tau_0; E_1, E_2]$ has a model in this class of size at most triply exponential in $|\tau_0|$.

We note that the obtained bound is essentially optimal, since as shown in [21], one can construct a family of formulas $\{\varphi_i\}_{i\in\mathbb{N}}$, such that φ_i has length $O(i)$, and its every model has size greater than $2^{2^{2^n}}$. It is not difficult to see that every φ_i in the family is finitely satisfiable.

V. UNDECIDABILITY

In this section we show that replacing one of the equivalence relations in the considered class of structures by a relation which is only required to be transitive leads to undecidability. This sharpens an earlier result that FO^2 is undecidable over the class of structures with two transitive relations [19], [20] and also corresponds to the result by [21] that FO^2 over $\mathcal{EQ}[E_1, E_2, E_3]$ is undecidable.

In fact, as shown in [21], FO^2 over $\mathcal{EQ}[E_1, E_2, E_3]$ is a conservative reduction class. We show that the same holds if we consider FO^2 over structures with one distinguished equivalence relation and one distinguished transitive relation, which yields the following theorem.

Theorem 14: The (finite) satisfiability problem for FO^2 over structures with one equivalence and one transitive relation is undecidable.

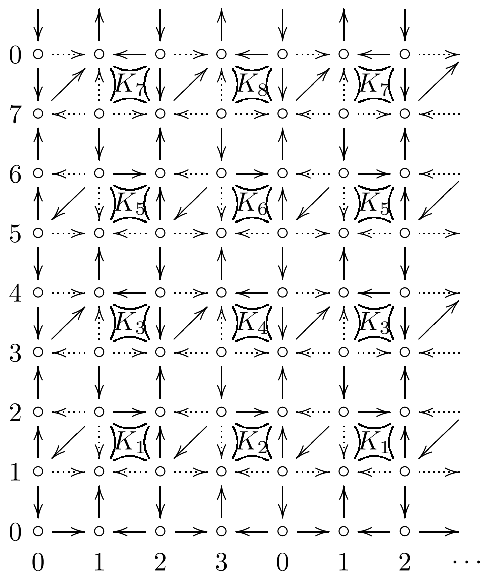

Figure 2. Expansion of a grid

As usually for two-variable logics, we expand the standard two-dimensional grid $\mathfrak{G} = (G, H, V)$ to the structure illustrated in Figure 2 whose properties we capture by a formula `Grid`. To get undecidability we show that the formula `Grid` captures a rich class of grids (cf. [9]).

In the formula `Grid`, apart from the standard horizontal and vertical grid relations H and V, we use E for the equivalence relation, T for the transitive relation and additional monadic predicates C_{ij}, $i = 0, \ldots, 3$, $j = 0, \ldots, 7$. We also refer to C_{ij}s as *colours* of the elements in the grid. Colour C_{ij} describes elements in a row numbered by i and column numbered by j, as shown in Figure 2. Non-oriented edges correspond to E-connections and oriented edges describe T-connections.

The equivalence relation E partitions the required model into four-element classes, denoted in Figure 2 as $K_1, \ldots, K_8$. Elements belonging to the same equivalence class have different colours. The transitive relation T is antisymmetric.

The formula `Grid` is a conjunction of the following formulas.

$\star$ The initial formula:

$$\exists x \exists y C_{00}x \wedge \forall x(\exists y Hxy \wedge \exists y Vxy). \tag{1}$$

$\star$ A formula axiomatizing H:

$$\forall x \forall y (Hxy \to (Row_0 \vee \ldots \vee Row_7)), \tag{2}$$

where each Row_i describes four possible cases of values of the monadic predicates and E and T connections on H-connected vertices in a row i. For example

$$\begin{aligned} Row_0 \equiv\ & Txy \wedge \neg Exy \wedge C_{00}x \wedge C_{10}y \vee \\ & Tyx \wedge Exy \wedge C_{10}x \wedge C_{20}y \vee \\ & Txy \wedge \neg Exy \wedge C_{20}x \wedge C_{30}y \vee \\ & Tyx \wedge Exy \wedge C_{30}x \wedge C_{00}y. \end{aligned}$$

$\star$ A similar formula axiomatizing V:

$$\forall x \forall y (Vxy \to (Column_0 \vee \ldots \vee Column_3)), \tag{3}$$

where, e.g., $Column_1$ is as follows

$$\begin{aligned} Column_1 \equiv\ & Txy \wedge \neg Exy \wedge C_{10}x \wedge C_{11}y \vee \\ & Tyx \wedge Exy \wedge C_{11}x \wedge C_{12}y \vee \\ & Tyx \wedge \neg Exy \wedge C_{12}x \wedge C_{13}y \vee \\ & Txy \wedge Exy \wedge C_{13}x \wedge C_{14}y \vee \\ & Txy \wedge \neg Exy \wedge C_{14}x \wedge C_{15}y \vee \\ & Tyx \wedge Exy \wedge C_{15}x \wedge C_{16}y \vee \\ & Tyx \wedge \neg Exy \wedge C_{16}x \wedge C_{17}y \vee \\ & Txy \wedge Exy \wedge C_{17}x \wedge C_{10}y. \end{aligned}$$

$\star$ The conjunction of formulas inducing the diagonal T-edges drawn in Figure 2:

$$\begin{aligned}
&\forall x(C_{01}x \rightarrow \exists y(Tyx \wedge C_{12}y)),\\
&\forall x(C_{21}x \rightarrow \exists y(Tyx \wedge C_{32}y)),\\
&\forall x(C_{03}x \rightarrow \exists y(Txy \wedge C_{14}y)),\\
&\forall x(C_{23}x \rightarrow \exists y(Txy \wedge C_{34}y)),\\
&\forall x(C_{05}x \rightarrow \exists y(Tyx \wedge C_{16}y)),\\
&\forall x(C_{25}x \rightarrow \exists y(Tyx \wedge C_{36}y)),\\
&\forall x(C_{07}x \rightarrow \exists y(Txy \wedge C_{10}y)),\\
&\forall x(C_{27}x \rightarrow \exists y(Txy \wedge C_{30}y)).
\end{aligned} \tag{4}$$

$\star$ A group of formulas saying that certain elements connected by T are in the same E-class. For example

$$\begin{aligned}
&\forall x \forall y(Tyx \wedge C_{11}x \wedge C_{21}y \rightarrow Exy),\\
&\forall x \forall y(Tyx \wedge C_{11}x \wedge C_{12}y \rightarrow Exy).
\end{aligned} \tag{5}$$

$\star$ A formula describing the partition induced by E into four-element equivalence classes.

$$\begin{aligned}
\forall x \forall y Exy \rightarrow (&\bigwedge_{i,j}(\neg(C_{ij}x \wedge C_{ij}y) \vee x = y) \wedge\\
&\bigwedge_{k \neq l} \neg(Class_k(x) \wedge Class_l(y)),
\end{aligned} \tag{6}$$

where each $Class_l(x)$, $l = 1, \ldots, 8$, describes the possible colours of elements in the class of type K_l. For example

$$\begin{aligned}
Class_1(x) &\equiv C_{11}x \vee C_{12}x \vee C_{21}x \vee C_{22}x,\\
Class_8(x) &\equiv C_{37}x \vee C_{30}x \vee C_{07}x \vee C_{00}x.
\end{aligned}$$

$\star$ Two groups of formulas saying that certain elements connected by, respectively, E or T are also connected by the grid relations. For example

$$\forall x \forall y(Exy \wedge C_{11}x \wedge C_{21}y \rightarrow Hxy) \tag{7}$$

and

$$\begin{aligned}
&\forall x \forall y(Txy \wedge C_{21}x \wedge C_{31}y \rightarrow Hxy),\\
&\forall x \forall y(Tyx \wedge C_{22}x \wedge C_{32}y \rightarrow Hxy).
\end{aligned} \tag{8}$$

We show that the expansion of the standard grid over $\mathbb{N} \times \mathbb{N}$ shown in Figure 1 is a model of the formula `Grid`. It is clear that in the model all conjuncts of the form (1)-(6) hold. To see that also conjuncts of the form (7) and (8) are satisfied, observe that every T-path in the structure is finite and of length at most 6. Moreover, any T-path connects at most three adjacent columns and at most five adjacent rows. So, the distribution of the C_{ij} colours ensures that formulas (7) and (8) cannot enforce new pairs of elements, apart from those already connected in the standard grid, to become connected by the H or V relations.

In a similar way, one can see that also every standard grid over a finite $4m \times 8m$ torus can be expanded to a model of `Grid`.

To show that the formula `Grid` captures a rich class of grids it suffices to show that every model of `Grid` is grid-like. To show that $\mathfrak{A}$ is grid-like it is enough to show that H is complete over V in $\mathfrak{A}$, i.e. $\mathfrak{A} \models \forall xyx'y'((Hxy \wedge Vxx' \wedge Vyy') \rightarrow Hx'y')$ (cf. [29], [21]).

Claim. H is complete over V in every model $\mathfrak{A}$ of `Grid`.

Assume that $\mathfrak{A} \models Hab \wedge Vaa' \wedge Vbb'$. We show that then $\mathfrak{A} \models Ha'b'$. In the proof several cases are to be considered depending on the colour of the element a. We discuss three typical ones.

Case 1. $\mathfrak{A} \models C_{10}a$. By (2) we have: $\mathfrak{A} \models Tba \wedge C_{20}b$. Formula (3) implies: $\mathfrak{A} \models Taa' \wedge C_{11}a'$ and $\mathfrak{A} \models Tb'b \wedge C_{21}b'$. Now, by transitivity of T: $\mathfrak{A} \models Tb'a'$. Using (8) we get $\mathfrak{A} \models Ha'b'$.

Case 2. $\mathfrak{A} \models C_{11}a$. By (2) we have: $\mathfrak{A} \models Tba \wedge C_{21}b \wedge Eab$. Formula (3) implies: $\mathfrak{A} \models Ta'a \wedge C_{12}a'$ and $\mathfrak{A} \models Tbb' \wedge C_{22}b'$. By (5), $\mathfrak{A} \models Eaa'$ and $\mathfrak{A} \models Ebb'$. Since E is an equivalence relation, we have: $\mathfrak{A} \models Ea'b'$. And by (7), we get $\mathfrak{A} \models Ha'b'$.

Case 3. $\mathfrak{A} \models C_{21}a$. Similarly to previous cases, by (2) we have: $\mathfrak{A} \models Tab \wedge C_{31}b \wedge \neg Eab$. Formula (3) implies: $\mathfrak{A} \models Taa' \wedge C_{22}a'$ and $\mathfrak{A} \models Tb'b \wedge Ebb' \wedge C_{32}b'$. Now, by (4), for some $c \in A$, $\mathfrak{A} \models Tca \wedge C_{32}c$. By transitivity of T: $\mathfrak{A} \models Tca'$ and $\mathfrak{A} \models Tcb$. By a conjunct from (5), $\mathfrak{A} \models Ecb$. Since E is an equivalence, we have: $\mathfrak{A} \models Eb'c$ and so, by (6), $b' = c$. So, by transitivity of T, $\mathfrak{A} \models Tb'a'$. Now, using (8), we get $\mathfrak{A} \models Ha'b'$.

The remaining cases can be treated in a similar way.

We conclude noting that the grid relations H and V can be simulated using appropriate combinations of E, T and the monadic colours. So, the undecidability is retained, even when apart from E and T there are no other binary relation symbols in the signature.

VI. Conclusion

Theorem 13 implies a natural upper bound on the complexity of the finite satisfiability problem.

Theorem 15: The finite satisfiability problem for FO^2 over the class of structures with two equivalence relations is decidable in 3NEXPTIME.

We note that there is still a gap between lower and upper bounds for the complexity of (both finite and unrestricted) satisfiability of FO^2 over the class of structures with two equivalence relations. The current lower bound is 2EXPTIME [20].

The techniques developed in this paper can be adapted to the case of unrestricted satisfiability. This adaptation seems to be more complicated than the original construction from [21], but raises our hopes for an algorithm in 2NEXPTIME.

References

[1] F. Baader, H. Bürckert, B. Hollunder, A. Laux, and W. Nutt. Terminologische logiken. *KI*, 3/92:23–33, 1992.

[2] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P.F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.

[3] S. Bechhofer, F. van Harmelen, J. Hendler, I. Horrocks, D. L. McGuinness, P. F. Patel-Schneider, and L. A. Stein. OWL Web Ontology Language reference. W3C Recommendation, 10 February 2004. Available at http://www.w3.org/TR/owl-ref/.

[4] P. Blackburn. Representation, reasoning, and relational structures: a hybrid logic manifesto. *Logic Journal of the IGPL*, 3(8):339–625, 2000.

[5] P. Blackburn and J. Seligman. Hybrid languages. *Journal of Logic, Language and Information*, 4(3):251–272, 1995.

[6] M. Bojańczyk, C. David, A. Muscholl, T. Schwentick, and L. Segoufin. Two-variable logic on words with data. In *LICS '06: Proceedings of the twenty-first IEEE symposium on Logic in Computer Science*, pages 7–16, 2006.

[7] M. Bojańczyk, C. David, A. Muscholl, T. Schwentick, and L. Segoufin. Two-variable logic on data trees and xml reasoning. In *PODS '06: Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 10–19, New York, NY, USA, 2006.

[8] A. Borgida. On the relative expressiveness of description logics and predicate logics. *Artificial Intelligence*, 82:353–367, 1996.

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

[10] R. Brafman, J. C. Latombe, Y. Moses, and Y. Shoham. Knowledge as a toll in motion planning under uncertainty. In *R. Fagin, ed., 'Theoretical Aspects of Reasoning about Knowledge: Proc. Fifth Conference'*, pages 208–224. Morgan Kaufmann, 1994.

[11] M. Burrows, M. Abadi, and R. Needham. Authentication: a practical study in belief and action. In *Proc. Second Conf. on Theoretical Aspects of Reasoning about Knowledge*, pages 325–342, 1988.

[12] D. Calvanese. *Unrestricted and Finite Model Reasoning in Class-Based Representation Formalisms.* Ph.D. thesis, Dipartimento di Informatica e Sistemistica, Universit'a di Roma La Sapienza, 1996.

[13] J. M. V. Castilho, M. A. Casanova, and A. L. Furtado. A temporal framework for database specification. In *Proc. Eight Int. Conf. on Very Large Data Bases*, pages 280–291, 1982.

[14] E. M. Clarke and E. A. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proc. Workshop on Logic of Programs*, volume 131, pages 52–71. Springer-Verlag, 1981.

[15] D. Gabbay. Expressive functional completeness in tense logic. In *U. Mönnich, ed., 'Aspects of Philosophical Logic'*, pages 91–117. Reidel, 1971.

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

[17] E. Grädel and M. Otto. On Logics with Two Variables. *Theoretical Computer Science*, 224:73–113, 1999.

[18] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 32:137–161, 1985.

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

[20] E. Kieroński. Results on the guarded fragment with equivalence or transitive relations. In *Computer Science Logic*, volume 3634, pages 309–324. Springer Verlag, 2005.

[21] E. Kieroński and M. Otto. Small substructures and decidability issues for first-order logic with two variables. In *LICS*, pages 448–457, 2005.

[22] E. Kieroński and L. Tendera. On finite satisfiability of the guarded fragment with equivalence or transitive guards. In *LPAR*, volume LNAI 4790, pages 318–332. Springer Verlag, 2007.

[23] H. R. Lewis. Complexity results for classes of quantificational formulas. *J. Comp. and System Sci.*, 21:317–353, 1980.

[24] W. Lipski. On the logic of incomplete information. In *Proc. Sixth Int. Symp. on Mathematical Foundations of Computer Science*, volume 53, pages 374–381. Springer Verlag, 1977.

[25] J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *D. Michie, ed. 'Machine Intelligence'*, volume 4, pages 463–502. Edinburgh University Press, 1969.

[26] M. Mortimer. On languages with two variables. *Zeitschr. f. Logik und Grundlagen d. Math.*, 21:135–140, 1975.

[27] M. Mundhenk and T. Schneider. The complexity of hybrid logics over equivalence relations. International Workshop on Hybrid Logic, 2007.

[28] M. Mundhenk, T. Schneider, T. Schwentick, and V. Weber. Complexity of hybrid logics over transitive frames. *CoRR*, abs/0806.4130, 2008.

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

[30] C. H. Papadimitriou. On the complexity of integer programming. *J. of the ACM*, 28(2):765–769, 1981.

[31] A. Pnueli. The temporal logic of programs. In *Proc. 18th IEEE Symposium on Foundations of Computer Science*, pages 46–57, 1977.

[32] V. R. Pratt. Semantical considerations on floyd-hoare logic. In *Proc. 17th IEEE Symposium on Foundations of Computer Science*, pages 109–121, 1976.

[33] M. Y. Vardi. Why is modal logic so robustly decidable? *DIMACS Series in Discrete Mathematics and Theoretical Computer Science,*, 31:149–184, 1997.

[34] H. Wang. Proving theorems by pattern recognision. *Bell Syst. Tech. J.*, 40:1–41, 1961.