On Finite Satisfiability of the Guarded Fragment with Equivalence or Transitive Guards ⋆

Emanuel Kieroński[1] and Lidia Tendera[2]

[1] Institute of Computer Science, University of Wrocław
[2] Institute of Mathematics and Informatics, Opole University

Abstract. The guarded fragment of first-order logic, GF, enjoys the finite model property, so the satisfiability and the finite satisfiability problems coincide.
We are concerned with two extensions of the two-variable guarded fragment that do not possess the finite model property, namely, GF^2 with equivalence and GF^2 with transitive guards. We prove that in both cases every finitely satisfiable formula has a model of at most double exponential size w.r.t. its length.
To obtain the result we invent a strategy of building finite models that are formed from a number of multidimensional grids placed over a cylindrical surface. The construction yields a 2NEXPTIME-upper bound on the complexity of the finite satisfiability problem for these fragments. For the case with equivalence guards we improve the bound to 2EXPTIME.

1 Introduction

In this paper we assume that a first-order signature contains only relational symbols and, possibly, equality.

The *guarded fragment*, GF, of first-order logic is defined as the least set of formulas such that: (i) every atomic formula belongs to GF; (ii) GF is closed under logical connectives $\neg, \vee, \wedge, \rightarrow$; and (iii) quantifiers are appropriately relativized by atoms, i.e. if $\varphi(\mathbf{x}, \mathbf{y})$ is a formula of GF and $\alpha(\mathbf{x}, \mathbf{y})$ is an atomic formula containing all the free variables of φ, then the formulas

$$\forall \mathbf{y}(\alpha(\mathbf{x}, \mathbf{y}) \rightarrow \varphi(\mathbf{x}, \mathbf{y})) \quad \text{and} \quad \exists \mathbf{y}(\alpha(\mathbf{x}, \mathbf{y}) \wedge \varphi(\mathbf{x}, \mathbf{y}))$$

belong to GF. The atom $\alpha(\mathbf{x}, \mathbf{y})$ is called the *guard* of the quantifier. Equalities are also allowed as guards.

The guarded fragment was introduced by Andréka et al., [1], who showed that modal logic can be embedded in GF and they also argued convincingly that GF inherits the nice properties of modal logic. The nice behavior of GF was confirmed by Grädel, [4], who proved that GF enjoys the finite model property and that the satisfiability problem for GF is complete for double exponential time and complete for exponential time, when the number of variables is bounded.

GF was later generalized to the *loosely guarded fragment*, [19], and to the *clique guarded fragment*, [5], where all quantifiers were relativized by more general formulas,

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

preserving the idea of quantification only over elements that were close together in the model. Most of the properties of GF can be generalized to those fragments. In particular, I. Hodkinson [8] showed that they all enjoy the finite model property (see also [9] for a simpler and nicer proof of the result).

For modal or description logics, a restriction of the underlying class of structures is very natural. We are particularly interested in classes with some transitive or equivalence relations that correspond e.g. to Kripke structures for multi-modal variants of the logics K4 and S5. The results for these cases can be phrased either for satisfiability of extensions of GF, or of GF itself over restricted classes of structures.

In this paper we are concerned with two extensions of GF^2, i.e. the two-variable guarded fragment with equivalence guards, $\mathrm{GF}^2+\mathrm{EG}$, and the two-variable guarded fragment with transitive guards, $\mathrm{GF}^2+\mathrm{TG}$. In $\mathrm{GF}^2+\mathrm{EG}$ ($\mathrm{GF}^2+\mathrm{TG}$) some binary relation symbols are required to be interpreted as equivalences (transitive relations), but these *special* symbols are allowed to appear only in guards. These fragments capture many expressive modal and description logics. Both are decidable for satisfiability: $\mathrm{GF}^2+\mathrm{EG}$ is NEXPTIME-complete, [12], $\mathrm{GF}^2+\mathrm{TG}$ – 2EXPTIME-complete, [16], and both contain infinity axioms, i.e. formulas that possess only infinite models.

The lack of the finite model property for the above mentioned fragments naturally leads to the question, whether their finite satisfiability problems are decidable. This question is particularly important, if one would like to use these formalism for automatic reasoners in practical applications, where the structures investigated are essentially finite.

In [18] it is shown that the finite satisfiability problem for GF with one transitive relation is decidable in double exponential time, thus it is of the same complexity as the (unrestricted) satisfiability problem. For the case of $\mathrm{GF}^2+\mathrm{EG}$, a NEXPTIME lower bound is implied by GF with one equivalence relation, that enjoys the finite (exponential) model property and is NEXPTIME-complete for satisfiability, [12].

In this paper we show that every finitely satisfiable $\mathrm{GF}^2+\mathrm{EG}$-formula and every finitely satisfiable $\mathrm{GF}^2+\mathrm{TG}$-formula has a model of at most double exponential size. In the case of $\mathrm{GF}^2+\mathrm{EG}$ our small models are constructed as multidimensional grids, in which the number of dimensions equals the number of equivalence relations. Models for $\mathrm{GF}^2+\mathrm{TG}$ are built from a number of such grids placed over a cylindrical surface and connected in a special way.

The results give also a trivial, double exponential, bound on the size of a single equivalence class in a finite model of a $\mathrm{GF}^2+\mathrm{EG}$-formula. We argue that this bound is essentially optimal, which contrasts with the general case, where it can be shown that every satisfiable $\mathrm{GF}^2+\mathrm{EG}$-sentence has a (possibly infinite) model with at most exponential equivalence classes.

The constructions yield 2NEXPTIME-upper bounds on the complexity of the finite satisfiability problem for these fragments. For the case with equivalence guards we improve the bound to 2EXPTIME. We also note that in $\mathrm{GF}^2+\mathrm{TG}$ transitive relations can be enforced in a simple way to be symmetric and reflexive, see Lemma 2 in [12]. Thus, the upper bound for $\mathrm{GF}^2+\mathrm{TG}$ gives also the upper bound for a combined variant $\mathrm{GF}^2+\mathrm{EG}+\mathrm{TG}$, in which some symbols appearing only in guards have to be interpreted as equivalences and some other as transitive relations.

It is perhaps worth mentioning that there is another interesting extension of the guarded fragment, namely GF with fixpoints, that has been shown decidable for satisfiability, [7], and of the same complexity as pure GF. However, decidability of the finite satisfiability problem for this fragment is only partially answered by Bojańczyk in [2], where decidability of the finite satisfiability problem for the modal μ-calculus with backwards modalities is shown.

We also remark that recently, [15], tight exponential complexity bounds for GF^2 with counting quantifiers for both satisfiability and finite satisfiability have been established.

2 Preliminaries

For a fragment $\mathcal{L}$ of first-order logic we denote by $\mathcal{L}^k$ the restriction of the logic to k variables. Since in this paper we are concerned with extensions of GF^2, we assume without loss of generality that signatures contain only unary and binary relation symbols.

We fix a signature σ. An *atomic k-type* is a maximal consistent set of atomic or negated atomic formulas over σ in k variables. Let $\boldsymbol{\alpha}$ be the set of all atomic 1-types, and $\boldsymbol{\beta}$ the set of all atomic 2-types over σ.

Let $\mathfrak{A}$ be a σ-structure with universe A, and let $a, b \in A$, $a \neq b$. We denote by $\alpha(a)$ the unique atomic 1-type realized in $\mathfrak{A}$ by the element a, and by $\beta(a, b)$ – the unique atomic 2-type realized in $\mathfrak{A}$ by the pair (a, b). We note that the notation $\alpha(a)$ and $\beta(a, b)$ does not explicitly show from which structure the elements a and b are taken, but, in this paper, the structure will always be clear from the context.

2.1 Simple observations on FO^2 models

Since equivalence classes in models of $\mathrm{GF}^2+\mathrm{EG}$ formulas behave similarly to whole models of FO^2 formulas, we start with some simple definitions and lemmas for FO^2.

Let $M = 3|\boldsymbol{\beta}|^3$ be fixed for the rest of this paper. Note that the value of M is exponential in $|\sigma|$.

Definition 1 (Scott normal form). *We say that an* FO^2*-formula φ is in* normal form *if* $\varphi = \forall x \forall y \psi_0(x, y) \wedge \bigwedge_i \forall x \exists y \psi_i(x, y)$*, where ψ_i are quantifier-free.*

Definition 2. (i) *An* exact counting type *is a function $\theta : \boldsymbol{\alpha} \to \mathbb{N}$. We say that a structure $\mathfrak{A}$ has an exact counting type θ if for every $t \in \boldsymbol{\alpha}$, $\theta(t)$ is the number of realizations of t in $\mathfrak{A}$.*

(ii) *Let $n > 0$. An n*-counting type *is a function $\theta : \boldsymbol{\alpha} \to \{0, 1, \ldots, n\}$. We say that a structure $\mathfrak{A}$ has an n-counting type θ if for every $t \in \boldsymbol{\alpha}$, $\theta(t)$ is the number of realizations of t in $\mathfrak{A}$, if this number is less than n, and $\theta(t) = n$ otherwise.*

We say that a counting type θ' *safely extends* a counting type θ if for every $t \in \boldsymbol{\alpha}$, $\theta'(t) = \theta(t)$ if $\theta(t) = 0$ or $\theta(t) = 1$, and $\theta'(t) \geq \theta(t)$ otherwise.

The following lemma establishes two useful properties of counting types and the relevance of M. Part (ii) is a simple application of the "small substructures lemma" from [13]. The proof is omitted due to space limits.

Lemma 1. (i) *Let φ be an* FO^2*-formula in normal form, $\mathfrak{A} \models \varphi$ and let θ be the exact counting type of $\mathfrak{A}$. Then for every exact counting type θ' safely extending θ, there exists a model $\mathfrak{A}' \models \varphi$ whose exact counting type is θ'.*

(ii) *Let φ be an* FO^2*-formula in normal form, $\mathfrak{A} \models \varphi$, θ be the M-counting type of $\mathfrak{A}$. Then, there exists a model $\mathfrak{A}' \models \varphi$ whose exact counting type is θ.*

2.2 Normal forms

Let $S_1, S_2, \ldots, S_k$ denote the *special* binary relation symbols of σ, i.e. those that are required to be interpreted as equivalence or transitive relations, and $R_1, R_2, \ldots$ – the remaining binary symbols in σ. A σ-structure is *special*, if all the special relation symbols are interpreted in the required way.

When the special relation symbols are allowed to appear in guards only, it is useful to consider the following normal form for GF^2-formulas.

Definition 3. *We say that a* GF^2*-formula φ with special guards is in normal form if it is a conjunction of formulas of the following form:*

(i) $\exists x(p(x) \wedge \psi(x))$,
(ii) $\forall x \forall y(S_i xy \rightarrow \psi(x,y))$,
(iii) $\forall x \forall y(q(x,y) \rightarrow \psi(x,y))$,
(iv) $\forall x(p(x) \rightarrow \exists y(q(x,y) \wedge \psi(x,y)))$,
(v) $\forall x(p(x) \rightarrow \exists y(S_i xy \wedge S_i yx \wedge \psi(x,y)))$,
(vi) $\forall x(p(x) \rightarrow \exists y(S_i xy \wedge \neg S_i yx \wedge \psi(x,y)))$,
(vii) $\forall x(p(x) \rightarrow \exists y(S_i yx \wedge \neg S_i xy \wedge \psi(x,y)))$,

where each ψ is quantifier-free and does not contain special relation symbols, $p(x)$ is atomic, $q(x,y)$ is an atom $R_i(x,y)$ or $R_i(y,x)$. We denote by φ_{sym} the fragment of φ consisting of all the conjuncts of types (i), (ii), (iii), (iv) and (v), and by φ_{eq} the fragment of φ consisting of all the conjuncts of types (i), (ii), (iii) and (v). For a given i, we denote by $\varphi_{eq}^{S_i}$ the fragment of φ_{eq} consisting of those conjuncts of type (ii) and (v), which have S_i in their guards, and all conjuncts of type (iii). Additionally, let φ_{univ} be the conjunction of all implications $q(x,y) \rightarrow \psi(x,y)$ from formulas of type (iii).

The (finite) satisfiability problem for GF^2-formulas with special guards over special structures can be reduced to the (finite) satisfiability problem for disjunctions of exponential number of linear size GF^2-formulas in normal form over special structures. See for example [18] for a proof of a similar result.

We note that if the special relation symbols are required to be interpreted as equivalence relations, then in the normal form we do not need to consider conjuncts of the forms (vi) and (vii) and the conjuncts of the form (v) may be simplified to $\forall x(p(x) \rightarrow \exists y(S_i xy \wedge \psi(x,y)))$.

We remark that, although in conjuncts of the forms (v)-(vii) the special symbols appear not only as simple guards, the normal form is more useful than the standard one where, instead of conjuncts of the forms (v)-(vii), we have just the form: $\forall x(p(x) \rightarrow \exists y(S_i xy \wedge \psi(x,y)))$.

In next sections we use E_is to denote the special relations symbols that are to be interpreted as equivalence relations, and T_is to denote the transitive relation symbols.

3 Equivalence Guards

In this section we show that every finitely satisfiable GF^2+EG-formula has a model of a bounded size.

If we restrict the number of equivalence symbols to one, then the finite satisfiability coincides with satisfiability, since, [13], even whole FO^2 with one equivalence relation has the finite (exponential) model property. However, in the presence of two equivalence symbols, there are satisfiable GF^2+EG formulas whose all models are infinite [13].

3.1 Example

Let us now observe, that finite models for GF^2+EG have different properties from infinite models. In the following example we show how to construct a family of finitely satisfiable formulas $\{\varphi_n\}$ with only two equivalence relation symbols, such that every finite model of φ_n contains at least one equivalence class of size at least doubly exponential in n, and $|\varphi_n|$ is polynomial in n. This is in contrast to the (unrestricted) satisfiability: in [12] it was observed that every satisfiable GF^2+EG-formula φ has a model, in which all equivalence classes have size at most exponential in $|\varphi|$. The following example is a refinement of the example from [18], where several transitive relations were used.

Example 1. Let us assume that E_1 and E_2 have to be interpreted as equivalence relations. We construct a finitely satisfiable formula φ_n, such that its every model contains some number of full binary trees of depth 2^n, whose every leaf requires a root in its E_2-class. Trees will have to be disjoint, so there will always be 2^{2^n} leaves per one root, which will guarantee the existence of at least one large E_2-class.

Except E_1 and E_2, we use only a number of unary relation symbols. Symbols $P_0, \ldots, P_{n-1}$ encode in each element a number from $\{0, \ldots 2^n - 1\}$ which is the depth of the element in the tree (the number of the level to which the element belongs). Let us denote by $\mathcal{L}_i$ the i-th level, i.e. the set of elements a, with the encoded number i. Symbol R indicates roots, symbol L is used to distinguish left sons from right sons. The formula φ_n consists of conjuncts stating:

- There exists an element satisfying R.
- Every element satisfying R belongs to $\mathcal{L}_0$.
- For each element a from $\mathcal{L}_{2i}$ there exist two elements in $\mathcal{L}_{2i+1}$ connected to a by E_1. One of them satisfies L, the other $\neg L$.
- For each element a from $\mathcal{L}_{2i+1}$ (for $2i + 1 < 2^n - 1$) there exist two elements in $\mathcal{L}_{2i+2}$ connected to a by E_2. One of them satisfies L, the other $\neg L$.
- If two distinct elements belong to a level $\mathcal{L}_{2i}$ for some i, then they are not connected by E_1.
- If two distinct elements belong to a level $\mathcal{L}_{2i+1}$ (for $2i + 1 < 2^n - 1$), then they are not connected by E_2.
- Every element in $\mathcal{L}_{2^n-1}$ is connected by E_2 to an element satisfying R.

It is not difficult to formulate the above sentences in GF^2+EG, and to see, that each $\mathfrak{A} \models \varphi_n$ has a desired large E_2-class.

3.2 System of linear inequalities

Let φ be a $\mathrm{GF}^2 + \mathrm{EG}$-formula in normal form over a signature σ, where $E_1, E_2, \ldots, E_k$ are the equivalence relation symbols.

We say that an n-counting type θ is *admissible for* φ and E_i, if there exists a structure $\mathfrak{C}$ whose exact counting type is θ and

$$\mathfrak{C} \models \varphi_{eq}^{E_i} \wedge \forall xy E_i xy \wedge \bigwedge_{j \neq i} \forall xy (x \neq y \rightarrow \neg E_j xy).$$

Let $\Theta = (\Theta^{E_1}, \Theta^{E_2}, \ldots, \Theta^{E_k})$ be a tuple of sets of M-counting types. We say that an atomic type $t \in \boldsymbol{\alpha}$ *appears* in Θ^{E_i} if there exists $\theta \in \Theta^{E_i}$ such that $\theta(t) > 0$. We say that t appears in Θ if t appears in Θ^{E_i} for some i. We say that $t \in \boldsymbol{\alpha}$ is *royal* for E_i-classes, if t appears in Θ^{E_i} and, for every $\theta \in \Theta^{E_i}$, we have $\theta(t) \leq 1$.

Let $\boldsymbol{\alpha}_0 \subseteq \boldsymbol{\alpha}$ and let $\overline{\rho} = \rho_1, \rho_2, \ldots, \rho_k$ be a sequence of subsets of $\boldsymbol{\alpha}_0$. Let $\Theta(\varphi, \boldsymbol{\alpha}_0, \overline{\rho}) = (\Theta^{E_1}, \Theta^{E_2}, \ldots, \Theta^{E_k})$ be the tuple of maximal sets of M-counting types such that: (i) $\boldsymbol{\alpha}_0$ is the set of one-types appearing in $\Theta(\varphi, \boldsymbol{\alpha}_0, \overline{\rho})$, (ii) ρ_i is the set of royal types for E_i-classes, and (iii) every element of Θ^{E_i} is admissible for φ and E_i.

The set $\boldsymbol{\alpha}_0$ corresponds to the set of one-types realized in a structure in which elements of ρ_i are royal for E_i-classes. Elements of Θ^{E_i} correspond to M-counting types of restrictions of the structure to equivalence classes of E_i.

With $\Theta(\varphi, \boldsymbol{\alpha}_0, \overline{\rho})$ we associate a system of linear inequalities Γ_Θ. For each i and each $\theta \in \Theta^{E_i}$ we use a variable $X_\theta^{E_i}$, whose purpose is to count the number of E_i-classes with type θ, and, for each $t \in \boldsymbol{\alpha}_0$, we use a variable $Y_t^{E_i}$ to count the lower bound of the number of elements of type t in E_i-classes.

For every $t \in \boldsymbol{\alpha}_0$ and every i, we put to Γ_Θ the following (in)equalities:

- $$Y_t^{E_i} = \sum_{\theta \in \Theta^{E_i}} \theta(t) \cdot X_\theta^{E_i}, \tag{E0}$$
- if $t \in \rho_i$, then we write:
$$Y_t^{E_i} \geq 1, \tag{E1a}$$
- if $t \notin \rho_i$, then we write:
$$\sum_{\theta \in \Theta^{E_i}, \theta(t) > 1} X_\theta^{E_i} \geq 1, \tag{E1b}$$
- if $t \in \rho_i$, then, for every $j \neq i$, we write:
$$Y_t^{E_i} \geq Y_t^{E_j}. \tag{E2}$$

Please note that if t is royal for E_i-classes then the number $Y_t^{E_i}$ is exact, and if t is royal for both E_i-classes and E_j-classes ($i \neq j$) then we write two inequalities of the form (E2) that give $Y_t^{E_i} = Y_t^{E_j}$.

Moreover, we add to Γ_Θ the following inequalities:

- for every conjunct δ of the form (i), $\delta = \exists x(p(x) \wedge \psi(x))$,we write

$$\sum_{t \in \boldsymbol{\alpha}_0, t \models p(x) \wedge \psi(x)} Y_t^{E_1} > 0, \tag{E3}$$

- for every conjunct δ of the form (iv), $\delta = \forall x(p(x) \to \exists y(q(x,y) \wedge \psi(x,y)))$, for every $t \in \boldsymbol{\alpha}_0$ such that $t \models p(x)$, we write:

$$\sum_{t' \in \boldsymbol{\alpha}_0, t \to \delta t'} Y_{t'}^{E_1} > 0, \tag{E4}$$

where $t \to \delta t'$, if for some $s(x,y) \in \boldsymbol{\beta}$ such that $t(x) \subseteq s(x,y)$, $t'(y) \subseteq s(x,y)$, we have $\neg E_i xy \in s(x,y)$, for all i, and $s(x,y) \models \varphi_{univ} \wedge q(x,y) \wedge \psi(x,y)$.

Observe that the number of inequalities is at most exponential and the number of variables is at most doubly exponential in the size of the signature.

3.3 Model construction

Lemma 2. *Let φ be a $\mathrm{GF}^2{+}\mathrm{EG}$-formula φ in normal form. Then, φ is finitely satisfiable if and only if there exist $\boldsymbol{\alpha}_0$, $\overline{\rho} = \rho_1, \ldots, \rho_k$, with $\rho_i \subseteq \boldsymbol{\alpha}_0$, such that the system Γ_Θ constructed for $\Theta(\varphi, \boldsymbol{\alpha}_0, \overline{\rho})$ has a nonnegative integer solution.*

Proof ($\Rightarrow$) Let $\mathfrak{A}$ be a finite model of φ. Let $\boldsymbol{\alpha}_0$ be the set of 1-types realized in $\mathfrak{A}$. Let A/E_i be the set of equivalence classes of E_i in $\mathfrak{A}$. Let Θ^{E_i} be the set of all M-counting types of the structures $\mathfrak{A}\restriction K$, for $K \in A/E_i$. Let ρ_i be the set of all royal types for E_i-classes.

Let $\theta \in \Theta^{E_i}$ be the M-counting type of $\mathfrak{A}\restriction K$ for some K. Observe that $\mathfrak{A}\restriction K$ is a model of the FO^2-formula $\varphi_{eq}^{E_i} \wedge \forall xy E_i xy$. It can be simply modified to a model of $\varphi_{eq}^{E_i} \wedge \forall xy E_i xy \wedge \bigwedge_{j \neq i} \forall xy(x \neq y \to \neg E_j xy)$, by replacing in every atomic 2-type atoms $E_j xy$ by $\neg E_j xy$ and atoms $E_j yx$ by $\neg E_j yx$. The last formula can be easily rewritten to FO^2 normal form. Now we apply part (ii) of Lemma 1, obtaining the desired exact counting type of the model. So, θ is admissible for φ and E_i.

It is not difficult to see that $X_\theta^{E_i} = r_\theta^{E_i}$, where $r_\theta^{E_i}$ is the number of those E_i-classes in $\mathfrak{A}$, whose M-counting types equal θ, and $Y_t^{E_i} = \sum_{\theta \in \Theta^{E_i}} \theta(t) \cdot r_\theta^{E_i}$ is a nonnegative integer solution of Γ_Θ.

($\Leftarrow$) Let the system Γ_Θ constructed for $\Theta(\varphi, \boldsymbol{\alpha}_0, \overline{\rho}) = (\Theta^{E_1}, \ldots \Theta^{E_k})$ with appropriate $\boldsymbol{\alpha}_0$ and $\overline{\rho}$ has a nonnegative integer solution: $X_\theta^{E_i} = r_\theta^{E_i}$ and $Y_t^{E_i} = r_t^{E_i}$.

First, we show how to build a model $\mathfrak{A} \models \varphi_{eq}$. Then, a model $\mathfrak{A}' \models \varphi$, will be obtained by taking an appropriate number of copies of $\mathfrak{A}$ and setting the 2-types in a simple way.

For every $t \in \boldsymbol{\alpha}_0$, let $n_t = \max_i \{r_t^{E_i}\}$. Observe that by (E0) and (E1a) or (E1b), $n_t > 0$. The structure $\mathfrak{A}$ is constructed from some number of copies of the *base set* of elements. The base set consists of exactly n_t realizations of type t, for every $t \in \boldsymbol{\alpha}_0$. Let $\{a_0, \ldots a_{l-1}\}$ be the elements of the base set. We put copies of elements from the base set into the nodes of a k-dimensional grid (recall that k is the number of equivalence relation symbols): at location $(x_1, \ldots, x_k)$, for $0 \leq x_i < l$, we put an element whose 1-type equals the 1-type of a_s, where $s = (x_1 + \ldots + x_k) \bmod l$. Note that the set of elements located on each line of the grid is exactly a copy of the base set. In Fig. 1 we picture the desired arrangement of elements for $l = 5$ and $k = 2$.

$$
\begin{array}{ccccccccc}
a_4 & — & a_0 & — & a_1 & — & a_2 & — & a_3 \\
| & & | & & | & & | & & | \\
a_3 & — & a_4 & — & a_0 & — & a_1 & — & a_2 \\
| & & | & & | & & | & & | \\
a_2 & — & a_3 & — & a_4 & — & a_0 & — & a_1 \\
| & & | & & | & & | & & | \\
a_1 & — & a_2 & — & a_3 & — & a_4 & — & a_0 \\
| & & | & & | & & | & & | \\
a_0 & — & a_1 & — & a_2 & — & a_3 & — & a_4
\end{array}
$$

Fig. 1. Arrangement of elements on the two-dimensional grid.

To settle properly the 2-types of elements in $\mathfrak{A}$ we first observe that for every i, the base set can be divided into $\sum_{\theta \in \Theta^{E_i}} r_\theta^{E_i}$ disjoint parts, so that for each part P there exists an exact counting type θ', safely extending some $\theta \in \Theta^{E_i}$, such that the number of elements of every atomic 1-type t in P equals $\theta'(t)$. The desired division is obtained as follows. We create $r_\theta^{E_i}$ parts for every $\theta \in \Theta^{E_i}$. To each of this parts we put exactly $\theta(t)$ elements of type t, for every $t \in \boldsymbol{\alpha}_0$. Note, that, because of the choice of the numbers n_t, we have enough copies of elements of every type t. After this step we may have some elements remaining. Observe, that due to inequalities of the form (E2), none of the types of the remaining elements is royal for E_i. All the remaining elements of type t are joined to a part which contains at least two elements of t. Note that such a part exists due to inequalities of the form (E1b).

Now, horizontal lines of the grid are divided as above to form equivalence classes of E_1, vertical lines – to form classes of E_2, lines going in the third dimension – to form classes of E_3, and so on. To define the structure on each of the classes we use the fact that each variable $X_\theta^{E_i}$ in the system corresponds to an M-counting type θ^{E_i} admissible for φ and E_i and, (if necessary) part (i) of Lemma 1. We complete the definition of $\mathfrak{A}$ by setting all 2-types for pairs of elements not belonging to the same class to *binary-free types*, i.e. 2-types containing $\neg Rxy$ and $\neg Ryx$ for every binary R (including equivalence symbols).

To construct $\mathfrak{A}'$ we take three sets of copies of $\mathfrak{A}$, each consisting of h elements, where h is the number of conjuncts of φ of the form (iv); we can suppose $h > 0$. This step is reminiscent of the construction of the small model for FO^2 [6].

Let $A' = A \times \{0, 1, \ldots, h-1\} \times \{0, 1, 2\}$. We set $\alpha((a, i, j)) = \alpha(a)$ and $\beta((a, i, j), (b, i, j)) = \beta(a, b)$. Now we have to provide witnesses for all elements in A' for conjuncts of φ of the form (iv). This is done in a circular way: elements from $A \times \{0, 1, \ldots, h-1\} \times \{j\}$ find witnesses in $A \times \{0, 1, \ldots, h-1\} \times \{j'\}$, where

$j' = j + 1 \bmod 3$. Let us take for example an element (a, i, j). Let $\delta_0, \dots, \delta_{h-1}$ be all conjuncts of φ of the form (iv), $\delta_m = \forall x(p(x) \rightarrow \exists y(q(x, y) \wedge \psi(x, y)))$. Let $t = \alpha(a)$. For $0 \leq m < h$ we choose $t'_m \in \boldsymbol{\alpha}_0$, and $s_m \in \boldsymbol{\beta}$, whose existence is postulated by inequality (E4). By an appropriate inequality of the form $(E1a)$ or $(E1b)$, t'_m is realized in $\mathfrak{A}$, say, by an element b_m. We set $\beta((a, i, j), (b_m, m, j + 1 \bmod 3)) := s_m$. This circular scheme guarantees no conflict in setting 2-types, since a 2-type for a pair of elements in $\mathfrak{A}'$ is defined at most once. We complete the structure by setting 2-types for all remaining pairs of elements. Let $((a, i, j), (a', i', j'))$ be such a pair. Let $t = \alpha(a)$, $t' = \alpha(b)$. Let s be the unique 2-type containing $t(x)$, $t(y)$ and negations of all binary atoms. We set $\beta((a, i, j), (a', i', j')) := s$. ☐

3.4 Size of models and complexity of finite satisfiability

We use the following lemma from [3], built on the classical result [14], to obtain an upper bound on the size of the model constructed in part $\Leftarrow$ of the proof of Lemma 2.

Lemma 3 (Calvanese). *Let Γ be a system of m linear inequalities in n unknowns, and let the coefficients and constants that appear in the inequalities be in $\{-a, \dots, a-1, a\}$. If Γ admits a nonnegative integer solution, then it also admits one in which the values assigned to the unknows are all bounded by $(n + m) \cdot (m \cdot a)^{2m+1}$.*

In our system Γ_Θ the parameters a and m are at most exponential, and n – at most doubly exponential in the size of the signature. For a given φ in normal form we may consider the signature consisting only of relation symbols appearing in φ. Henceforth, by Lemma 3, if Γ_Θ has a solution, it has also one with values at most doubly exponential in $|\varphi|$. Thus, in the proof of Lemma 2 we may choose a base set consisting of a double exponential number of elements, obtaining finally a double exponential bound on the size of the whole model. In the consequence, we can state the following theorem.

Theorem 1. *Every finitely satisfiable* $\mathrm{GF}^2 + \mathrm{EG}$*-formula φ has a model of size at most doubly exponential in $|\varphi|$.*

The obtained bound on the size of finite models is essentially optimal – recall the formula from Example 1, which has only models of at least double exponential size.

Theorem 1 immediately gives a 2NEXPTIME decision procedure for the finite satisfiability problem for $\mathrm{GF}^2 + \mathrm{EG}$: guess a structure of at most double exponential size and check, if it is a model for the input sentence. We show that we can do better.

Theorem 2. *The finite satisfiability problem for* $\mathrm{GF}^2 + \mathrm{EG}$ *is decidable in deterministic double exponential time.*

Proof By Lemma 2, to check if a normal form $\mathrm{GF}^2 + \mathrm{EG}$-sentence φ has a finite model it suffices to check if there exists $\Theta(\varphi, \boldsymbol{\alpha}_0, \overline{\rho})$ such that Γ_Θ has a nonnegative integer solution.

To do this we enumerate all possible $\boldsymbol{\alpha}_0$ and $\overline{\rho}$, write the system Γ_Θ for $\Theta = \Theta(\varphi, \boldsymbol{\alpha}_0, \overline{\rho})$, check admissibility of each M-counting type in Θ and check if Γ_Θ has a nonnegative integer solution.

Observe that Γ_Θ consists of inequalities of the form $\sum_i c_i x_i \geq b$, with $b \geq 0$. So Γ_Θ has a nonnegative integer solution if and only if it has a nonnegative rational solution. (To get an integer solution, it suffices to take the rational solution and multiply it by the product of all the denominators.)

Since linear programming is in PTIME, [10], we can check the existence of a nonnegative rational solution of Γ_Θ in time doubly exponential w.r.t. $|\varphi|$.

All these gives a deterministic procedure working in double exponential time. $\square$

We note that the best lower bound for the finite satisfiability problem for $\mathrm{GF}^2+\mathrm{EG}$ is NEXPTIME. This bound is obtained for the fragment with one equivalence relation that enjoys the finite model property [12].

4 Transitive Guards

In this section we show that every finitely satisfiable $\mathrm{GF}^2+\mathrm{TG}$ formula φ in normal form has a model of size at most doubly exponential in $|\varphi|$. In the proof we essentially use Theorem 1 and the model construction from the previous section for solutions of systems of linear inequalities for appropriate tuples of M-counting types.

4.1 Basic notions and outline of the construction

Let $\mathfrak{A}$ be a finite model of φ. We construct a small model $\mathfrak{B} \models \varphi$. To follow the construction it is not harmful to think that in $\mathfrak{A}$ every pair of distinct elements is a member of at most one transitive relation. The small model $\mathfrak{B}$ will also have this property.

To construct $\mathfrak{B}$ we first distinguish in $\mathfrak{A}$ classes of elements forming *T-cliques*, for every transitive T. If we remove from $\mathfrak{A}$ nonsymmetric transitive connections then, in the obtained structure $\mathfrak{A}'$, transitive relations behave like equivalence relations and transitive cliques behave like equivalence classes. Obviously, $\mathfrak{A}' \models \varphi_{sym}$. Then, for every M-counting type θ of a T-class in $\mathfrak{A}'$, we apply the construction from the proof of part $\Leftarrow$ of Lemma 2 to $\mathfrak{A}'$ and produce a bounded model of φ_{sym} containing at least a T-class of M-counting type θ. This way, we obtain building blocks for the construction of $\mathfrak{B}$.

To be able to provide witnesses for conjuncts of the forms (vi) and (vii), we enrich the counting types of classes of $\mathfrak{A}$ with two subsets of 1-types, $\mathcal{A}$ and $\mathcal{B}$, corresponding to 1-types of elements located in $\mathfrak{A}$ *above*, respectively *below*, elements of the class. Additionally, we get out off $\mathfrak{A}$ a finite bounded pattern for providing nonsymmetric witnesses for each enriched counting type, called *a graph of witnesses*.

Finally, the model $\mathfrak{B}$ is formed out of a number of the building blocks located on a cylindrical surface and the nonsymmetric witnesses are provided in a regular manner using the graphs of witnesses.

A T-clique in a structure $\mathfrak{C}$ is a maximal set D of elements, such that $\forall x, y \in D.\ \mathfrak{C} \models x \neq y \rightarrow (Txy \wedge Tyx)$. Let $C/\hat{T}$ be the set of all T-cliques in $\mathfrak{C}$. For a pair of T-cliques D, D' in $\mathfrak{C}$ we write $D <_T^{\mathfrak{C}} D'$ if for all $a \in D$, $b \in D'$, $\mathfrak{C} \models Tab \wedge \neg Tba$.

Definition 4. *We say that a T-clique $D \in A/\hat{T}$ has an enriched n-counting type $\hat{\theta} = (\theta, \mathcal{A}, \mathcal{B})$ if:*

(i) $\mathfrak{A}\restriction D$ *has the* n*-counting type* θ*,*
(ii) $\mathcal{A} = \{t \in \boldsymbol{\alpha} : \exists a \in D, a' \in A.\ \alpha(a') = t \wedge \mathfrak{A} \models T(a,a') \wedge \neg T(a',a)\}$
(iii) $\mathcal{B} = \{t \in \boldsymbol{\alpha} : \exists a \in D, a' \in A.\ \alpha(a') = t \wedge \mathfrak{A} \models T(a',a) \wedge \neg T(a,a')\}$

Similarly to the previous section we are interested in enriched M-counting types for $M = 3|\boldsymbol{\beta}|^3$. Let $\hat{\Theta}^T$ be the set of enriched M-counting types, and Θ^T the set of M-counting types, realized in $\mathfrak{A}$ by T-cliques. In our new model $\mathfrak{B} \models \varphi$, for every enriched M-counting type $\hat{\theta} = (\theta, \mathcal{A}, \mathcal{B}) \in \hat{\Theta}^T$ we will put a T-clique of type $\hat{\theta} = (\theta, \mathcal{A}', \mathcal{B}')$, such that $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{B}' \subseteq \mathcal{B}$.

4.2 Structures for symmetric requirements

Lemma 4. *For every transitive symbol* T*, for every* $\hat{\theta} = (\theta, \mathcal{A}, \mathcal{B}) \in \hat{\Theta}^T$ *there exists* $\mathfrak{C} \models \varphi_{sym}$*, such that*
(i) $\mathfrak{C}$ *contains a distinguished* T*-clique of* M*-counting type* θ*,*
(ii) $|C|$ *is at most doubly exponential in* $|\varphi|$*,*
(iii) *all transitive symbols are interpreted as transitive and symmetric relations in* $\mathfrak{C}$*,*
(iv) *for every transitive* T'*, and every* M*-counting type* θ' *of a* T'*-clique in* $\mathfrak{C}$ *there exists an enriched* M*-counting type* $\hat{\theta}' = (\theta', \mathcal{A}, \mathcal{B}) \in \hat{\Theta}^{T'}$*, for some* $\mathcal{A}, \mathcal{B}$*.*

Proof Note, that a model meeting requirements (i), (iii), (iv) exists: an example is $\hat{\mathfrak{A}}$ – a modification of $\mathfrak{A}$ in which, in all 2-types, we substitute $Txy \wedge \neg Tyx$ or $Tyx \wedge \neg Txy$ with $\neg Txy \wedge \neg Tyx$, for all transitive T.

Let $\boldsymbol{\alpha}_0$ be the set of 1-types realized in $\hat{\mathfrak{A}}$, Θ^T be the set of M-counting types realized in $\hat{\mathfrak{A}}$ by T-classes, ρ_i be the set of royal types for T_i-classes in $\hat{\mathfrak{A}}$, and $\overline{\rho} = \rho_1, \rho_2, \ldots, \rho_k$. Let Γ_Θ be the system of inequalities associated with $\Theta(\varphi_{sym}, \boldsymbol{\alpha}_0, \overline{\rho})$. We extend Γ_Θ to Γ'_Θ by adding the inequality:

$$X^T_\theta > 0 \tag{E5}$$

and the equality:

$$\sum_{(T,\theta):\theta \notin \Theta^T} X^T_\theta = 0. \tag{E6}$$

Note, that Γ'_Θ has a nonnegative solution: the one corresponding to the model $\hat{\mathfrak{A}}$. Now, using Lemma 3 and the construction from the proof of part $\Leftarrow$ of Lemma 2, we obtain $\mathfrak{C}$ as desired. □

We apply Lemma 4 to every transitive T and $\hat{\theta} \in \hat{\Theta}^T$, obtaining structures $\mathfrak{C}^T_{\hat{\theta}}$. Additionally, we define a function enr on the cliques in these structures. If D is the distinguished T-clique in $\mathfrak{C}^T_{\hat{\theta}}$, then $enr(D) = \hat{\theta}$. In the other case let us assume that D is a T'-clique and let θ' be the M-counting type of D. Choose any $\hat{\theta}' = (\theta', \mathcal{A}, \mathcal{B}) \in \hat{\Theta}^{T'}$ as the value of $enr(D)$. The purpose of $enr(D)$ is to say, elements of which 1-types are allowed to be connected to D by transitive relations from above and from below.

4.3 Graphs of witnesses

In this subsection we define labeled multidigraphs which contain patterns for providing nonsymmetric transitive witnesses. Nodes of these graphs are enriched M-types realized in $\mathfrak{A}$. An edge from $\hat{\theta}$ to $\hat{\theta}'$, labeled with a triple of atomic types (t, t', s), says that each element of 1-type t from a clique of type $\hat{\theta}$ should look for a witness of 1-type t' in a clique of type $\hat{\theta}'$, and that the element and its witness should be connected by the 2-type s.

We construct independently two graphs for each transitive symbol T. The first of them says where to look for witnesses for formulas with guards of the form (vii) $\forall x(p(x) \rightarrow \exists y(Tyx \wedge \neg Txy \wedge \psi(x,y)))$, i.e. *lower witnesses*. The second – for formulas of the form (vi) $\forall x(p(x) \rightarrow \exists y(Txy \wedge \neg Tyx \wedge \psi(x,y)))$, i.e. *upper witnesses*.

We show how to build a *graph of lower witnesses* for relation T, $G^T_<$. *A graph of upper witnesses*, $G^T_>$, can be constructed similarly. Formally, $G^T_< = (\hat{\Theta}^T, E, lab)$, where lab, is a labeling of edges $lab : E \rightarrow \boldsymbol{\alpha} \times \boldsymbol{\alpha} \times \boldsymbol{\beta}$. We will construct $G^T_<$ in steps, starting with $E = \emptyset$. During the construction we define a partial function $f : \hat{\Theta}^T \rightarrow A/\hat{T}$, associating with each enriched M-counting type a T-clique in $\mathfrak{A}$. When the whole construction is completed, f will be a total function. Additionally we use a marking of vertices. Initially all vertices are unmarked.

1. Choose any unmarked vertex $\hat{\theta}$ from $\hat{\Theta}^T$. Find a T-clique D in $\mathfrak{A}$ such that the enriched M-counting type of D is $\hat{\theta}$ and there is no T-clique D' of type $\hat{\theta}$ in $\mathfrak{A}$, such that $D' <^{\mathfrak{A}}_T D$. The existence of the desired clique D is implied by finiteness of $\mathfrak{A}$. Set $f(\hat{\theta}) = D$.
2. Repeat the following steps as long as possible:
   (a) Choose an unmarked $\hat{\theta} = (\theta, \mathcal{A}, \mathcal{B}) \in \hat{\Theta}^T$ for which f is defined, let $D = f(\hat{\theta})$.
   (b) For every conjunct of the form (vii): $\forall x(p(x) \rightarrow \exists y(Tyx \wedge \neg Txy \wedge \psi(x,y)))$, and every 1-type t appearing in θ, $t = \alpha(a)$ for some $a \in D$, such that $t \models p(x)$:
      i. Find $b \in A$, such that $\mathfrak{A} \models Tba \wedge \neg Tab \wedge \psi(a,b)$.
      ii. Find $c \in A$, such that $\alpha(b) = \alpha(c)$, $\mathfrak{A} \models Tca \wedge \neg Tac$ and there is no d such that $\alpha(d) = \alpha(c)$ and $\mathfrak{A} \models Tdc \wedge \neg Tcd$. Again, such a c exists since $\mathfrak{A}$ is finite. Please note, that it is not necessary that $\mathfrak{A} \models \psi(a,c)$.
      iii. Let D' be the T-clique of c, and $\hat{\theta}'$ be its enriched M-counting type. Set $f(\hat{\theta}') = D'$. Add to E an edge e from $\hat{\theta}$ to $\hat{\theta}'$, and set $lab(e) = (\alpha(a), \alpha(b), \beta(a,b))$.
   (c) Mark $\hat{\theta}$.
3. If there is an unmarked vertex in $G^T_<$ then go to 1.

We note two properties of graphs of lower witnesses, shared also by graphs of upper witnesses.

Fact 1 *$G^T_<$ is an acyclic graph.*

Proof Assume to the contrary, that there exists a cycle $\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_l, \hat{\theta}_1$. But then $f(\hat{\theta}_2) <_T^{\mathfrak{A}} f(\hat{\theta}_1)$, $f(\hat{\theta}_3) <_T^{\mathfrak{A}} f(\hat{\theta}_2)$, $f(\hat{\theta}_1) <_T^{\mathfrak{A}} f(\hat{\theta}_l)$. Thus, from transitivity of T, $f(\hat{\theta}_1) <_T^{\mathfrak{A}} f(\hat{\theta}_1)$, which is impossible. □

Fact 2 *Every path in $G_<^T$ has length at most $|\boldsymbol{\alpha}|$.*

Proof Let $\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_l$ form a path in $G_<^T$. Assume that $\hat{\theta}_i = (\theta_i, \mathcal{A}_i, \mathcal{B}_i)$. Our construction implies that $\mathcal{B}_1 \supseteq \mathcal{B}_2 \supseteq \ldots \supseteq \mathcal{B}_l$. In fact, the choice of c in step 2(b)ii., as a farthest element of the required 1-type, guarantees, that the above set containments are strict. Since $\mathcal{B}_1 \subseteq \boldsymbol{\alpha}$ we have the desired bound on l. □

4.4 Model construction

Universe Let $K = 2|\boldsymbol{\alpha}| + 1$. Let m be the maximal degree of a vertex in all graphs of lower and upper witnesses. Note that K and m are exponential in $|\sigma|$. We define the universe of $\mathfrak{B}$ to be

$$B = \{0, \ldots, K-1\} \times \{0, \ldots, 3\} \times \{1, \ldots, m\} \times \dot{\bigcup_{T \in \sigma, \hat{\theta} \in \hat{\Theta}^T}} C_{\hat{\theta}}^T.$$

Thus, if we assume that σ consists only of symbols appearing in φ, $|B|$ is at most doubly exponential in $|\varphi|$. We can think that B consists of K rows and 4 columns, and that the intersection of each row and each column contains m copies of each $\mathfrak{C}_{\hat{\theta}}^T$. In fact, we glue row $K-1$ to row 0 obtaining thus a cylindrical surface.

Initially, we impose appropriate structures on copies of $\mathfrak{C}_{\hat{\theta}}^T$. For $a \in C_{\hat{\theta}}^T$ let us set $\alpha(i, j, l, a) := \alpha(a)$, and for $a, b \in C_{\hat{\theta}}^T$, set $\beta((i, j, l, a), (i, j, l, b)) := \beta(a, b)$.

General scheme of providing witnesses Now we provide, for all elements, witnesses for all formulas with nonsymmetric transitive guards, i.e. formulas of the form (vi) and (vii). This is done in a regular manner. Elements from row i find their lower witnesses in the row $i + 1(\text{mod } K)$, and their upper witnesses in the row $i - 1(\text{mod } K)$. Elements from column j look for their lower witnesses in column $l(j)$ and for their upper witnesses in column $u(j)$, where $l, u : \{0, 1, 2, 3\} \to \{0, 1, 2, 3\}$ are defined as follows: $l(0) = l(2) = 0$, $l(1) = l(3) = 1$, $u(0) = u(3) = 3$, $u(1) = u(2) = 2$. This strategy guarantees two important properties. First, there will be no pair of elements a, b such that a use b as a lower witness and b use a as an upper witness (since both ud and du have no fixed points). Second, none of the cliques will be used as a source for both a lower witness and an upper witness (since u and d have disjoint images).

Example of providing witnesses Let us consider explicitly the case of, say, providing lower T-witnesses for element $(5, 1, 3, a)$, where $a \in C_{\hat{\theta}'}^{T'}$. Let D be the T-clique of a in $\mathfrak{C}_{\hat{\theta}'}^{T'}$. Let $\hat{\theta} = enr(D)$. The outlined strategy says that lower witnesses should be looked for in row 6, column 1. Let $F_w = \{F_1, F_2, \ldots F_l\}$ be the set of edges in $G_<^T$ outgoing from the vertex $\hat{\theta}$. Note that $l \leq m$. Let F_i goes to $\hat{\theta}_i$, let $lab(F_i) = (\alpha_i, \alpha'_i, \beta)$.

For each $i \leq m$, such that $\alpha_i = \alpha(a)$, choose an element a' in $\{6\} \times \{1\} \times \{i\} \times C^T_{\hat{\theta}_i}$ of type α'_i. Set $\beta(a, a') := \beta$.

In a similar way we provide lower and upper witnesses for all elements in $\mathfrak{B}$.

Transitive closure During the step of providing witnesses we defined 2-types containing nonsymmetric transitive connections between some pairs of elements. We say that there exists a T-path from T-clique D' to a T-clique D in $\mathfrak{B}$ if there exists a sequence of T-cliques $D = D_0, D_1, \ldots D_{l-1}, D_l = D'$ such that for every pair D_i, D_{i+1} we defined a lower T-witness for an element from D_i in D_{i+1} or an upper T-witness for an element from D_{i+1} in D_i.

If there exists a T-path, for some transitive T, from D to D', then we make $D' <^{\mathfrak{B}}_T D$. For every $a' \in D'$, $a \in D$, if the 2-type for a' and a is not defined, then let t' be the 1-type of a' and t – the 1-type of a. We choose any 2-type $s(x, y)$ realized in $\mathfrak{A}$, such that $s(x, y) \models Txy \wedge \neg Tyx \wedge t'(x) \wedge t(y)$, and set $\beta(a', a) := s$.

The following claim shows that nothing bad can happen during this step. Part (i) implies, that finding a required $s(x, y)$ in $\mathfrak{A}$ is always possible. Part (ii) says, that we have not created a T-path from a T-clique D to the same clique D, i.e. we have not formed a T-cycle. Part (iii) says, that setting 2-types can be done without conflicts.

Claim (OK).

(i) Assume that, after the step of providing nonsymmetric witnesses, there is a T-path from D to D' in $\mathfrak{B}$. Let $a' \in D', a \in D, t' = \alpha(a'), t = \alpha(a)$. Then there exist elements $c', c \in A$ whose 1-types are t' and t, respectively, such that $\mathfrak{A} \models Tc'c \wedge \neg Tcc'$.

(ii) Let $D_0, D_1, \ldots, D_l$ be a T-path from D_0 to D_l in $\mathfrak{B}$. Then $l \leq 2|\boldsymbol{\alpha}|$, i.e. $l < K$.

(iii) If an element $a \in B$ belongs to a T-clique D, $b \in B$ belongs to a T-clique D' and there is a T-path from D to D', then there is no pair C, C' of T'-cliques with $T \neq T'$, such that $a \in C$, $b \in C'$ and there is a T'-path from C to C', or a T'-path from C' to C.

We skip the proof of this claim. We note only that in the proof of part (iii) we use the fact that both $u^i d^j$ and $d^i u^j$ have no fixed points (for $i, j > 0$).

Remaining 2-types For every pair of elements $a, a' \in B$ for which we have not defined a 2-type yet, we set it to be the unique binary free type containing 1-types of a and a'. This step completes the construction of the structure $\mathfrak{B}$.

4.5 Complexity

The construction from the previous section allows us to state the following theorems.

Theorem 3. *Every finitely satisfiable* $\mathrm{GF}^2 + \mathrm{TG}$ *formula* φ *has a model of size at most doubly exponential in* $|\varphi|$.

Theorem 4. *The finite satisfiability problem for* $\mathrm{GF}^2 + \mathrm{TG}$ *is decidable in* 2NEXPTIME.

The best lower complexity bound we know is 2EXPTIME and can be obtained in the presence of just one transitive symbol, exactly as in [11], where (unrestricted) satisfiability problem for $\mathrm{GF}^2 + \mathrm{TG}$ was concerned.

References

1. H. Andréka, J. van Benthem, I. Németi, *Modal Languages and Bounded Fragments of Predicate Logic*, ILLC Research Report, 1996. Journal ver.: *J. Philos. Logic*, 27 (1998), No. 3, 217-274.
2. M. Bojańczyk, Two-way alternating automata and finite models, in: Proceedings of the 29th International Colloquium on Automata, Languages, and Programming, Vol. LNCS 2380, Springer, 2002, pp. 833–844.
3. D. Calvanese, *Unrestricted and finite model reasoning in class-based representation formalisms*, Ph.D. thesis, Dipartimento di Informatica e Sistemistica, Universitaa di Roma "La Sapienza", 1996.
4. E. Grädel, *On the Restraining Power of Guards*, Journal of Symbolic Logic 64:1719-1742, 1999.
5. E. Grädel. Decision procedures for guarded logics. In *16th International Conference in Artificial Intelligence*, volume LNCS 1932, pages 31–51. Springer, 1999.
6. E. Grädel, P. Kolaitis, M. Vardi, *On the Decision Problem for Two-Variable First Order Logic*, Bulletin of Symbolic Logic, 3(1):53-96, 1997.
7. E. Grädel and I. Walukiewicz. Guarded fixed point logic. In *Fourteenth Annual IEEE Symposium on Logic in Computer Science*, pages 45–54, 1999.
8. I. Hodkinson, Loosely guarded fragment of first-order logic has the finite model property, Studia Logica 70 (2002) 205–240.
9. I. Hodkinson, M. Otto, Finite conformal hypergraph covers, with two applications, Bull. Symbolic Logic 9 (2003) 387–405.
10. L. G. Khachiyan, *A polynomial algorithm in linear programming*, Soviet Mathematics Doklady, vol. 20, pages 191-194, 1979.
11. E. Kieroński, *The Two-Variable Guarded Fragment with Transitive Guards is 2EXPTIME-Hard*, Proc. of 6th Int. Conf. on Foundations of Software Science and Computational Structures (FOSSACS), LNCS 2620, pages 299-312, 2003.
12. E. Kieroński, *Results on the Guarded Fragment with Equivalence or Transitive Relations*, Proc. of CSL 2005, LNCS 3634, pages 309-324, 2005.
13. E. Kieroński, M. Otto, *Small Substructures and Decidability Issues for First-Order Logic with Two Variables*, Proc. of 20-th IEEE Symp. on Logic in Computer Science (LICS), pages 448-457, 2005.
14. Ch. Papadimitriou, *On the Complexity of Integer Programming*, Journal of ACM, vol. 28, No 4, pages, 765-786, 1981.
15. I. Pratt-Hartmann. Complexity of the guarded two-variable fragment with counting quantifiers. *Journal of Logic and Computation*, 17(1):133–155, 2007.
16. W. Szwast and L. Tendera. On the decision problem for the guarded fragment with transitivity. In *Proc. 16th IEEE Symposium on Logic in Ccomputer Science*, pages 147–156, 2001.
17. W. Szwast, L. Tendera, *The Guarded Fragment with Transitive Guards*, Annals of Pure and Applied Logic, 128, pages 227-276, 2004.
18. W. Szwast, L. Tendera, *On the Finite Satisfiability Problem For the Guarded Fragment with Transitivity*, Lect. Notes in Art. Intelligence, 3835, pages 307-321, 2005.
19. J. van Benthem. Dynamics bits and pieces. *ILLC Research Report LP-97-01*, University of Amsterdam, 1997.