EPiC Series in Computing
Volume 73, 2020, Pages 334–353
LPAR23. LPAR-23: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning

The Triguarded Fragment with Transitivity*

Emanuel Kieroński and Adam Malinowski

University of Wrocław, Poland

Abstract

The triguarded fragment of first-order logic is an extension of the guarded fragment in which quantification for subformulas with at most two free variables need not be guarded. Thus, it unifies two prominent decidable logics: the guarded fragment and the two-variable fragment. Its satisfiability problem is known to be undecidable in the presence of equality, but becomes decidable when equality is forbidden. We consider an extension of the triguarded fragment without equality by transitive relations, allowing them to be used only as guards. We show that the satisfiability problem for the obtained formalism is decidable and 2-ExpTime-complete, that is, it is of the same complexity as for the analogous extension of the classical guarded fragment. In fact, in our satisfiability test we use a decision procedure for the latter as a subroutine. We also show how our approach, consisting in exploiting some existing results on guarded logics, can be used to reprove some known facts, as well as to derive some other new results on triguarded logics.

1 Introduction

The *triguarded fragment*, TGF, was introduced by Rudolph and Šimkus [15] in order to unify two seminal decidable fragments of first-order logic, the *guarded fragment*, GF, defined by Andréka, van Benthem and Németi [1], and the *two-variable* fragment, FO^2, first considered by Scott [16]. TGF is obtained as an extension of GF in which quantification for subformulas with at most two free variables need not be guarded. Alternatively, one can think about the equivalent logic GFU, *the guarded fragment with the universal role*, whose formulas are just the formulas of GF, but the admissible models interpret the distinguished binary symbol U as the universally true relation.

Clearly, compared to GF and FO^2, TGF brings a new quality as it contains formulas expressible in neither of them, like $(\forall xy)\big(P(x) \wedge Q(y) \rightarrow (\exists z)R(x,y,z)\big)$. Let us observe that it also embeds the *Gödel class*, that is the class of all prenex sentences with the quantifier prefix $\forall\forall\exists^*$. Technically, given such a sentence $(\forall xy)\big((\exists z_1 \dots z_n)\varphi(x,y,z_1,\dots,z_n)\big)$ we may need to add a dummy guard $G(x,y,z_1,\dots,z_n)$, with a fresh symbol G, for the block of existential quantifiers, but, what is crucial, the initial pair of the universal quantifiers may be left unguarded as the subformula following them has only two free variables, x and y. Since the satisfiability problem for the Gödel class *with equality* is undecidable, as shown by Goldfarb [6], this embedding implies that also satisfiability of TGF with equality is undecidable. (A more direct

*Supported by Polish National Science Centre grant No 2016/21/B/ST6/01444.

proof of this fact was given by Rudolph and Šimkus [15].) However, in the *absence of equality*, TGF becomes decidable, which makes it potentially attractive for researchers in various areas of computer science.

One of the main motivations behind FO^2 and GF is that they embed, via the so-called *standard translation*, many modal and description logics, *e.g.*, the extension of basic description logic $\mathcal{ALC}$ with role hierarchies ($\mathcal{H}$), role inverse ($\mathcal{I}$), nominals ($\mathcal{O}$) and role intersection ($\sqcap$). Offering an elegant first-order perspective for some standard description logics, FO^2 and GF extend them in two, partially orthogonal directions. In particular, GF, generalizes basic description logics to settings with relations of arbitrary arity, while FO^2 allows one to express any boolean combination of roles, including their negations. TGF naturally inherits both types of benefits (though, the absence of equality limits the potential use of nominals).

Another motivation for TGF comes from databases, where an important role is played by GF. GF was, for example, an inspiration for the fruitful notion of the *guarded tuple generating dependencies*; see Calì, Gottlob and Kifer [3]. Moving to TGF allows us to express some additional concepts important for database theory, for example we can write the formula $(\forall xy)\big(P(x) \wedge Q(y) \rightarrow R(x,y)\big)$ saying that R is the cross product of P and Q.

The idea behind TGF is not new and can be traced back already in Kazakov's PhD thesis, [11], where the fragment $\mathrm{GF}|\mathrm{FO}^2$, capturing the spirit of TGF, was defined. What is relevant, $\mathrm{GF}|\mathrm{FO}^2$ does not admit constants. Kazakov proved that the satisfiability problem for $\mathrm{GF}|\mathrm{FO}^2$ without equality is decidable and 2-ExpTime-complete using a resolution method. The idea of enhancing GF by the already-mentioned binary cross product appears in the later work by Bourhis, Morak and Pieris [2] who introduced the logic $\mathrm{GF}^{\times_2}$, in which equality-free GF formulas can be conjoined with sentences defining the cross products of pairs of unary relations. Being motivated by database applications, that work implicitly assumes a separation between ground facts (a database) and a constant-free theory. Constants in ground facts can be easily simulated by existentially quantified variables, and thus $\mathrm{GF}^{\times_2}$ can be seen as a fragment of $\mathrm{GF}|\mathrm{FO}^2$. Actually, it is not difficult to perform a reduction also in the opposite direction. The authors of [2], being unaware of Kazakov's work, prove the decidability of $\mathrm{GF}^{\times_2}$ from scratch. They obtain a tight 2-ExpTime-upper bound using the classical database concept of *chase*.

The results described in the above paragraph imply that the satisfiability problem for TGF without equality and constants is decidable in 2-ExpTime. As we already said, the formal definition of TGF appears in the later work by Rudolph and Šimkus [15]. That paper properly analyses the case *with constants*. Interestingly, the presence of constants increases the complexity, making the satisfiability problem 2-NExpTime-complete. The upper bound is obtained by the mosaic method: to verify satisfiability of a given formula it suffices to produce some relatively small number of relatively small building blocks meeting some verifiable properties; a model of the given formula can be then constructed by stitching together a (usually infinite) number of copies of those building blocks. In the same paper it is also observed that the complexity drops down to NExpTime if there is a constant bound on the arity of relation symbols.

In this paper we consider an extension of TGF without equality in which some binary symbols are required to be interpreted as *transitive relations*. Transitivity of a binary relation is often a natural requirement in applications (consider, *e.g.*, the relations *greater-than*, *later-than* or *part-of*). However, this property is not expressible in typical decidable fragments of first-order logic, including FO^2 and GF. Moreover, augmenting FO^2 or GF with simple transitivity statements for binary relations leads to undecidability, as shown respectively by Grädel, Otto and Rosen [9] and by Grädel [7]; see also [5, 12, 11] for tighter undecidability results. Clearly, this implies that also TGF with the unrestricted use of transitive relations is undecidable.

On the other hand modal and description logics cope well with transitivity: Modal logic is decidable over transitive frames, and adding transitive roles ($\mathcal{S}$) to some expressive description logics, like $\mathcal{ALCOIQ}$, does not spoil their decidability; see, *e.g.*, the PhD thesis of Tobies [18]. This phenomenon was partially explained by Ganzinger, Meyer and Veanes [5], who considered a two-variable *monadic* guarded fragment, in which non-unary relations may occur only as guards, and show its decidability in the presence of transitive relations. Their variant is sufficiently strong to embed basic modal logic or some standard description logics, including $\mathcal{ALCI}$. Its severe restrictions are moderated in a later work by Szwast and Tendera [17] who demonstrate the decidability of the *guarded fragment with transitive guards*, GF+TG, the whole guarded fragment with arbitrarily many variables, relations of arbitrary arity, equality and arbitrarily many transitive relations, just restricting the use of transitive relations to guards. The latter variant allows, *e.g.*, to express role hierarchies ($\mathcal{H}$) on non-transitive relations (one can also say that a transitive relation is contained in one not required to be transitive, but not the opposite).

As our main contribution, we lift Szwast and Tendera's result to the level of the triguarded fragment, that is, we prove the decidability and 2-ExpTime-completeness of the *triguarded fragment with transitive guards*, TGF+TG. We remark that in our proof we do not admit constants, as they are not allowed in GF+TG, decidability of which we plan to use (decidability of GF+TG with constants is, up to our knowledge, an open problem). What we get is quite a powerful logic, inheriting good motivations of GF, FO^2 and TGF and strengthening them significantly by incorporating transitive relations. For example, in our logic we can embed the description logic $\mathcal{SI}$ extended by role hierarchies on non-transitive roles and arbitrary boolean combinations of non-transitive roles.

To present our approach in a simple setting we first provide a new decidability proof for TGF without transitive relations, yielding the optimal upper complexity bound. We concentrate on the case without constant symbols, but it is possible to include them, as we discuss later. The proof goes as follows. We first convert a given formula into its Scott-like *normal form*, resembling the normal form for GF [8]. We show that to verify satisfiability of a normal form φ it suffices to guess a set of 1-types A which are going to be realized in a model of φ, a set of 2-types B containing, for every pair of 1-types from A a 2-type that completes them, and check that for any 2-type from B there is a model of a minor modification of φ, containing a realization of this 2-type. This modification of φ belongs to GF so in this step we can use any existing algorithm for GF-satisfiability. Essentially, in the proof of the correctness of our method, when constructing a model of φ we just use (many times) a model construction for GF as a black box, and interleave it with the completion step reminiscent of the completion step from the small model construction for FO^2 by Grädel, Kolaitis and Vardi [8].

In fact, the resulting decision procedure for TGF is somehow similar to the decision procedure of Rudolph and Šimkus [15]. In some sense, in the proof of its soundness we also apply the mosaic method, but our building blocks are bigger. We believe that our approach is slightly simpler conceptually, as it reuses some existing results for GF while Rudolph and Šimkus prove everything from scratch. We hope that our view is valuable, and helps to better understand the decidability of TGF. More importantly, as already advertised, our approach generalizes to the stronger logic TGF+TG. Actually, while the proof of soundness of our decision procedure for TGF+TG is slightly more involved than for TGF, the procedures themselves are almost identical, the only real difference being that the former invokes a decision subprocedure for GF+TG while the latter uses a subprocedure for GF.

After a detailed presentation of our method in case of TGF and TGF+TG (without constants) we discuss its potential applications in some other scenarios, explaining how it allows us to prove some new results or reprove some existing ones. Most importantly, we argue that our method can be extended to cover the case of TGF *with constants* leading to optimal 2-NExpTime-upper complexity bound. Moreover, in our approach we can also include a limited use of equality, namely, admit atoms of the form $x = c$ for a variable x and a constant symbol c. This additional construct is suggested by Rudolph and Šimkus [15] as a natural mean of expressing the concept of *nominals* ($\mathcal{O}$) from description logics. Rudolph and Šimkus suspected that such use of equality does not spoil the decidability and does not increase the complexity, which we confirm here. We also consider some limited use of equality in TGF+TG, with which the decidability is preserved. This will allow us to express that some transitive relations are equivalences, which extends the potential applicability of the logic.

Following Rudolph and Šimkus [15], we present most of our results in the equivalent setting of the logic GFU. We organize the paper as follows. Section 2 contains some preliminaries. In Sections 3 and 4 we present our decidability proof for TGF (GFU) and its extension to TGF+TG (GFU+TG), respectively. In Section 5 we discuss some other possible applications of our method. In Section 6 we conclude the paper.

2 Conventions and some model theory

Models are denoted by Gothic letters $\mathfrak{M}, \mathfrak{N}, \ldots$ while their universes are denoted with the corresponding Roman letters $M, N, \ldots$. If $\mathfrak{M}$ is a model with domain M, we write $\overline{a} \subseteq M$ when $\overline{a} = \langle a_1, \ldots, a_n \rangle$ is a tuple of elements of M. For such a tuple, $|\overline{a}| = n$ denotes the length of $\overline{a}$. If $\overline{a} = \langle a_1, \ldots, a_n \rangle$ and $\overline{b} = \langle b_1, \ldots, b_m \rangle$ are two such tuples, by $\overline{a}\overline{b}$ we mean the concatenation of those tuples, *i.e.*, $\overline{a}\overline{b} = \langle a_1, \ldots, a_n, b_1, \ldots, b_m \rangle$. $(\exists x)$ stands for the usual quantification, but $(\exists \overline{x})$ is an abbreviation for $(\exists x_1) \ldots (\exists x_n)$ where $\overline{x} = \langle x_1, \ldots, x_n \rangle$ is a tuple of variables.

As a general rule, we implicitly allow emptiness. Hence unless explicitly excluded, we allow the empty tuple and denote it as $\langle\rangle$. A predicate can be of arity 0, so it is meaningful to write $\mathfrak{M} \models P(\langle\rangle)$, but we can use a shorter form $\mathfrak{M} \models P$ instead. A finite conjunction can be the empty conjunction, which is a universally true sentence denoted as $\top$. Similarly, the empty disjunction is $\perp$ which is always a false sentence. When we write $(\exists \overline{x})\, \theta(\overline{x}, \overline{y})$, we allow that $\overline{x}$ be an empty tuple of variables, in which case we call the quantifier vacuous, and such a formula is notationally equivalent to $\theta(\overline{y})$. The exception is that, as usual, we don't allow empty models.

We work in a finite relational first-order language $\mathcal{L}$, using the following sets of formulas:

- $\mathcal{F}$ is the set of all formulas in $\mathcal{L}$,
- $\mathcal{F}_{\neq}$ is the set of all formulas in $\mathcal{L}$ that do not contain the equality symbol,
- $\mathcal{F}^{\mathrm{qf}}$ is the set of all quantifier free formulas in $\mathcal{L}$,
- $\mathcal{F}^{\mathrm{qf}}_{\neq} = \mathcal{F}^{\mathrm{qf}} \cap \mathcal{F}_{\neq}$ is the set of all quantifier free formulas in $\mathcal{L}$ without equality.

If there is more than one language at hand, we write $\mathcal{F}(\mathcal{L})$, $\mathcal{F}_{\neq}(\mathcal{L})$, etc. for disambiguation.

Definition 2.1. Assume $\mathfrak{N}$ is a model, M is a set and $f : M \to N$ is a function. We define the structure $\mathfrak{M}$ on the domain M, called the *pullback of* $\mathfrak{N}$ *through* f, such that for any predicate $P(\overline{x})$ of $\mathcal{L}$ and tuple $\overline{a} \subseteq M$, we have $\mathfrak{M} \models P(\overline{a}) \iff \mathfrak{N} \models P(f(\overline{a}))$.

Remark 2.2. Assume $\mathfrak{M}, \mathfrak{N}$ are models and $f : M \to N$ is a surjection such that $\mathfrak{M}$ is the pullback of $\mathfrak{N}$ through f. Then for any formula $\varphi(\overline{x})$ in $\mathcal{F}_{\neq}$ and any tuple $\overline{a} \subseteq M$ we have $\mathfrak{M} \models \varphi(\overline{a}) \iff \mathfrak{N} \models \varphi(f(\overline{a}))$. In particular, $\mathfrak{M}$ and $\mathfrak{N}$ satisfy the same sentences in $\mathcal{F}_{\neq}$.

Proof. We proceed by induction on the complexity of φ. W.l.o.g. we assume that φ does not use universal quantification and disjunction.

- If $\varphi(\overline{x})$ is a predicate, the statement holds by definition.
- The induction steps involving conjunction and negation are trivial.
- Suppose the statement holds for $\psi(\overline{x}, y)$ and let $\varphi(\overline{x}) = (\exists y)\, \psi(\overline{x}, y)$. Fix a tuple $\overline{a} \subseteq M$. If $\mathfrak{M} \models \varphi(\overline{a})$, then there is some $b \in M$ such that $\mathfrak{M} \models \psi(\overline{a}, b)$. By induction hypothesis, $\mathfrak{N} \models \psi(f(\overline{a}), f(b))$, hence $\mathfrak{N} \models \varphi(f(\overline{a}))$.

  Now suppose $\mathfrak{N} \models \varphi(f(\overline{a}))$ and pick $b' \in N$ with $\mathfrak{N} \models \psi(f(\overline{a}), b')$. Since f is surjective, there is some $b \in M$ such that $f(b) = b'$, so $\mathfrak{N} \models \psi(f(\overline{a}), f(b))$. By induction hypothesis we have $\mathfrak{M} \models \psi(\overline{a}, b)$ so $\mathfrak{M} \models \varphi(\overline{a})$. □

Definition 2.3. Let $n \in \mathbb{N}$. A *complete (atomic) n-type* $p(x_1, \ldots, x_n)$ is a maximal logically consistent set of literals in $\mathcal{F}^{\mathrm{qf}}$ in n variables $x_1, \ldots, x_n$. So a complete n-type $p(x_1, \ldots, x_n)$ is uniquely determined by the following choices:

- an equivalence relation $\sim$ on the set $\{1, \ldots, n\}$ such that for each $1 \leqslant i, j \leqslant n$, we have $p(\overline{x}) \models x_i = x_j$ if $i \sim j$ and $p(\overline{x}) \models x_i \neq x_j$ otherwise,
- for every predicate $Q(x_1, \ldots, x_k)$ and function $\sigma : \{1, \ldots, k\} \to \{1, \ldots, n\}/{\sim}$, either $p(\overline{x}) \models Q(x_{i_1}, \ldots, x_{i_k})$ for all sequences $1 \leqslant i_1, \ldots, i_k \leqslant n$ where each $i_j \in \sigma(j)$, or $p(\overline{x}) \models \neg Q(x_{i_1}, \ldots, x_{i_k})$ for all such sequences.

Moreover, each choice as above corresponds to some complete n-type. Since every type $p(\overline{x})$ is finite, it may be identified with the formula $\bigwedge p(\overline{x}) \in \mathcal{F}^{\mathrm{qf}}$ which is the conjunction of all formulas in $p(\overline{x})$.

Definition 2.4. Assume $\mathfrak{M}$ is a model and $\overline{a} = (a_1, \ldots, a_n) \subseteq M$ is a tuple. The *type of $\overline{a}$ in $\mathfrak{M}$*, or the type *realized* by $\overline{a}$ in $\mathfrak{M}$, written $\mathrm{tp}^{\mathfrak{M}}(\overline{a})$, is the complete n-type consisting of all literals $\varphi(\overline{x}) \in \mathcal{F}^{\mathrm{qf}}$ such that $\mathfrak{M} \models \varphi(\overline{a})$. When $\mathfrak{M}$ is clear from the context, we write $\mathrm{tp}(\overline{a})$ instead of $\mathrm{tp}^{\mathfrak{M}}(\overline{a})$. If $a \in M$ is an element whose 1-type is realized in $\mathfrak{M}$ by no other element then a is called a *king*.

Lemma 2.5. *Assume φ is a satisfiable sentence in $\mathcal{F}_{\neq}$. Then φ has a model without kings.*

Proof. Let $\mathfrak{N} \models \varphi$. Pick a set M and a function $f : M \to N$ such that $f^{-1}[\{n\}]$ has at least two elements for each $n \in N$. Define $\mathfrak{M}$ as the pullback of $\mathfrak{N}$ through f. By Remark 2.2, $\mathfrak{M} \models \varphi$ and clearly if $a \in M$ realizes some 1-type $\alpha(x)$ and $b \in M \setminus \{a\}$ is an element such that $f(b) = f(a)$, then b also realizes $\alpha(x)$. □

3 Decidability of the triguarded fragment

In this section we present our technique by reproving the decidability of the triguarded fragment. But first, let us recall the definition of the guarded fragment and introduce a normal form for its formulas similar to that used by Grädel in [7].

Definition 3.1. The family of GF-*formulas* is the smallest family of formulas such that:

(i) Any formula in $\mathcal{F}^{\mathrm{qf}}_{\neq}$ is a GF-formula.

(ii) GF-formulas are closed under logical connectives $\neg, \vee, \wedge, \to$ and $\leftrightarrow$.

(iii) If $\varphi(\overline{x}, \overline{y})$ is a GF-formula and $G(\overline{x}, \overline{y})$ is an atom containing all the free variables of φ then $(\exists\overline{y})\big(G(\overline{x}, \overline{y}) \wedge \varphi(\overline{x}, \overline{y})\big)$ and $(\forall\overline{y})\big(G(\overline{x}, \overline{y}) \rightarrow \varphi(\overline{x}, \overline{y})\big)$ are GF-formulas.

An atom $G(\overline{x}, \overline{y})$ relativizing a quantifier as in point (iii) of the above definition is called *the guard* of this quantifier. We assume that GF-formulas may not have any occurrence of equality. This is different from the definition of GF in [7]. The reason to not allow equality here is that we later define TGF (GFU) in terms of GF, and TGF with equality is undecidable.

Definition 3.2. Assume φ is a sentence in GF. We say φ is in GF-*normal form* if it is a finite conjunction of GF-sentences of the following two forms:

N_1. any sentence in $\mathcal{F}^{\mathrm{qf}}_{\neq}$,

N_2. $(\forall\overline{x})\big(G(\overline{x}) \rightarrow (\exists\overline{y})\big(H(\overline{x}, \overline{y}) \wedge \theta(\overline{x}, \overline{y})\big)\big)$, where G, H are guards and $\theta \in \mathcal{F}^{\mathrm{qf}}_{\neq}$.

Lemma 3.3. *Let $\varphi(\overline{x})$ be a* GF*-formula in $\mathcal{L}$. Then there is $\mathcal{L}'$ extending $\mathcal{L}$ by some finite number of fresh predicates, a formula $\varphi'(\overline{x})$ in $\mathcal{F}^{\mathrm{qf}}_{\neq}(\mathcal{L}')$ and a sentence ψ in* GF*-normal form in $\mathcal{L}'$, such that*

- *Every model $\mathfrak{M}$ in $\mathcal{L}$ has a unique expansion $\mathfrak{M}'$ to $\mathcal{L}'$ satisfying ψ,*
- $\mathfrak{M}' \models (\forall\overline{x})\big(\varphi(\overline{x}) \leftrightarrow \varphi'(\overline{x})\big)$ *for each model $\mathfrak{M}'$ in $\mathcal{L}'$ satisfying ψ.*

In particular, if φ is a sentence, then φ has a model if and only if the normal form formula $\chi = \varphi' \wedge \psi$ has a model. Moreover, $|\chi|$ is in $\mathcal{O}(|\varphi|)$ and is computable in $\mathcal{O}(|\varphi|)$ time.

Proof. The idea of the proof is standard: we recursively replace existential subformulas of φ with fresh predicates and attach at the end a GF-normal sentence asserting the equivalence between the fresh predicate and the replaced subformula. Thus we end up with a conjunction of a quantifier free sentence φ' and some sentences ψ_i of the form N_2 in the extended language.

Let us turn to the details. The proof is by induction on the complexity of φ. We only demonstrate the existence part, as the computability statement easily follows from the analysis of the construction.

- If $\varphi(\overline{x})$ is in $\mathcal{F}^{\mathrm{qf}}_{\neq}$, we take $\mathcal{L}' = \mathcal{L}$, $\varphi' = \varphi$ and $\psi = \top$.
- The induction steps involving logical connectives are straightforward.
- Suppose $\varphi(\overline{x}) = (\exists\overline{y})\big(G(\overline{x}, \overline{y}) \wedge \varphi_0(\overline{x}, \overline{y})\big)$ and find $\mathcal{L}'_0 \supseteq \mathcal{L}$, $\varphi'_0(\overline{x}, \overline{y})$ and ψ_0 corresponding to $\varphi_0(\overline{x}, \overline{y})$ as in the statement. Let $P(\overline{x})$ be a fresh predicate symbol and $\mathcal{L}' = \mathcal{L}'_0 \cup \{P\}$. Then put $\varphi'(\overline{x}) = P(\overline{x})$ and

$$\psi = \psi_0 \wedge (\forall\overline{x})\big(P(\overline{x}) \leftrightarrow (\exists\overline{y})\big(G(\overline{x}, \overline{y}) \wedge \varphi'_0(\overline{x}, \overline{y})\big)\big)$$

To see that ψ can be equivalently written in GF-normal form, we write the part after $\wedge$ as a conjunction on the following two sentences of the form N_2:

$$(\forall\overline{x})\big(P(\overline{x}) \rightarrow (\exists\overline{y})\big(G(\overline{x}, \overline{y}) \wedge \varphi'_0(\overline{x}, \overline{y})\big)\big)$$
$$(\forall\overline{x}\overline{y})\big(G(\overline{x}, \overline{y}) \rightarrow \big(G(\overline{x}, \overline{y}) \wedge \big(\neg\varphi'_0(\overline{x}, \overline{y}) \vee P(\overline{x})\big)\big)\big)$$

(In the second sentence even though the existential quantifier is vacuous, we formally still need to guard the formula occurring after it, whence the doubled guard $G(\overline{x}, \overline{y})$).

It is routine to check that the inductive statement holds in that setting. □

We will also make use of the notion of a *guarded tuple*:

Definition 3.4. Assume $\mathfrak{M}$ is a model and $\overline{c} \subseteq M$ is a tuple. We say that $\overline{c}$ is *guarded* in $\mathfrak{M}$ if the set of elements occurring in $\overline{c}$ is a singleton or this set is contained in the set of elements occurring in some tuple $\overline{d} \subseteq M$ for which there exists a predicate G such that $\mathfrak{M} \models G(\overline{d})$.

Now we turn to the *triguarded fragment*, TGF. It can be thought of as the natural common extension of GF and FO^2. A TGF-formula is obtained in the same manner as a GF-formula, except that we allow quantifying an unguarded subformula provided that it has no more than two free variables. TGF also has an equivalent logic GFU, *the guarded fragment with the universal role*, introduced in [15], which is more convenient to work with, hence the results in this section are expressed in terms of GFU.

Definition 3.5. The GFU logic is defined as follows:

- A language $\mathcal{L}$ is a GFU-*language* if it contains a distinguished binary symbol U,
- A model $\mathfrak{M}$ in a GFU-language is a GFU-*model* if U is universally true in $\mathfrak{M}$,
- A GFU-*formula* is a GF-formula in a GFU-language.

Fact 3.6 ([15], Proposition 3)**.** TGF *and* GFU *have the same expressive power.* $\square$

For example the TGF-formula $(\forall xy)\big(P(x) \wedge Q(y) \to (\exists z)R(x,y,z)\big)$ can be transformed to the equivalent GFU-formula $(\forall xy)\big(\mathsf{U}(x,y) \to (P(x) \wedge Q(y) \to (\exists z)R(x,y,z))\big)$. In the opposite direction, GFU-formulas can be translated to TGF just by appending to them the conjunct $(\forall xy)\mathsf{U}(x,y)$.

Now we set to prove the decidability of satisfiability of formulas in GFU.

Lemma 3.7. *Assume φ is a* GFU*-sentence in* GF*-normal form. Then φ has a* GFU*-model if and only if there exist a set A of complete 1-types and a set $B \neq \varnothing$ of complete 2-types such that the following conditions hold:*

(i) $\alpha(x) \models \mathsf{U}(x,x)$ *for each* $\alpha \in A$,

(ii) $\beta(x,y) \models x \neq y \wedge \mathsf{U}(x,y) \wedge \mathsf{U}(y,x)$ *for each* $\beta \in B$,

(iii) for each $\alpha_1(x), \alpha_2(y) \in A$ *the partial 2-type* $\alpha_1(x) \cup \alpha_2(y)$ *extends to some complete type* $\beta(x,y) \in B$,

(iv) for each $\beta(x,y) \in B$ *there exists a (not necessarily* GFU*-) model* $\mathfrak{M} \models \varphi$ *such that*

  - *M contains some a,b with $\mathfrak{M} \models \beta(a,b)$,*
  - *each $c \in M$ realizes some $\alpha(x) \in A$,*
  - *every guarded pair $(a,b) \subseteq M$ satisfies $\mathfrak{M} \models \mathsf{U}(a,b)$.*

Proof. Write φ in GF-normal form as $\varphi = \varphi_0 \wedge \bigwedge_{s \in S} \varphi_s$ where φ_0 is a sentence in $\mathcal{F}^{\mathrm{qf}}_{\neq}$, S is a finite set of indices,

$$\varphi_s = (\forall \overline{x})\big(G_s(\overline{x}) \to (\exists \overline{y})\big(H_s(\overline{x},\overline{y}) \wedge \theta_s(\overline{x},\overline{y})\big)\big) =: (\forall \overline{x})\, \psi_s(\overline{x}) \qquad \text{for } s \in S,$$

each G_s and H_s is a guard and $\theta_s \in \mathcal{F}^{\mathrm{qf}}_{\neq}$.

First we prove the left-to-right implication. Suppose $\mathfrak{M}$ is a GFU-model satisfying φ and let A be the set of all 1-types realized in $\mathfrak{M}$. Applying Lemma 2.5 to the sentence $\varphi \wedge (\forall x,y)\,\mathsf{U}(x,y)$, we can assume each type in A has at least two realizations in $\mathfrak{M}$. To define B, we choose for each pair of 1-types $\alpha_1(x), \alpha_2(x) \in A$ two distinct elements $a, b \in M$ realizing these types and insert the 2-type $\beta(x, y) = \mathrm{tp}(a, b)$ into B. Then clearly A and B satisfy (i) – (iv).

Now we prove the implication from right to left. The idea of the construction is as follows. We use (iv) to build a model $\mathfrak{M}_1^*$ satisfying φ. If it is, by chance, a GFU-model, then we are done. Otherwise there are some unguarded pairs satisfying $\neg\mathsf{U}$. We use (iii) to *fix* such pairs by assigning to each a 2-type from B extending their 1-types. Thus we obtain a model $\mathfrak{M}_1$ which, by (i) and (ii), is a GFU-model.

Given $s \in S$, we will say that a tuple $\bar{c}\bar{d} \subseteq \mathfrak{M}_1$ such that $\mathfrak{M}_1 \models G_s(\bar{c})$ is an s-witness (or just a witness) for $\bar{c}$ if $\mathfrak{M}_1 \models H_s(\bar{c}, \bar{d}) \wedge \theta_s(\bar{c}, \bar{d})$. Since every witness is guarded and we only changed types of unguarded tuples, all witnesses are still valid. Also the unchanged type of the empty tuple $\langle\rangle$ ensures that φ_0 still holds. So the only possible reason for $\mathfrak{M}_1$ to not be the desired model is that some pairs (a, b) became guarded by being assigned a type $\beta_{a,b} \in B$ and might need new witnesses.

For each such pair (a, b) we use (iv) to find a new model $\mathfrak{M}^{a,b}$ in which φ holds, so in particular, (a, b) has the required witnesses for all $s \in S$. Then we consider a common superstructure $\mathfrak{M}_2^*$ of $\mathfrak{M}_1$ and all $\mathfrak{M}^{a,b}$ defined so that there are no relations across these models. Then again, $\mathfrak{M}_2^*$ satisfies everything except it might not be a GFU-model because of some unguarded pairs. So we fix these pairs obtaining a model $\mathfrak{M}_2$ and iterate the construction infinitely many times. The union of all models along the way has the desired properties.

The formal proof begins now. We inductively define an increasing sequence of GFU-models $\langle\mathfrak{M}_n : n < \omega\rangle$ which satisfies the following conditions for each n:

(1) if $n > 0$, $\mathfrak{M}_{n-1}$ is a substructure of $\mathfrak{M}_n$,

(2) each element of M_n realizes some 1-type $\alpha(x) \in A$,

(3) if $n > 0$, for each tuple $\bar{c} \subseteq M_{n-1}$ we have $\mathfrak{M}_n \models \bigwedge_{s\in S} \psi_s(\bar{c})$

(4) for each tuple $\bar{c} \subseteq M_n$ if $\mathfrak{M}_n \models \neg \bigwedge_{s\in S} \psi_s(\bar{c})$, then there exist $a, b \in M_n$ such that $\mathrm{tp}^{\mathfrak{M}_n}(a, b) \in B$ and $\bar{c} \in \{a, b\}^k$, where k is the length of $\bar{c}$.

Pick any $\beta_0(x, y) \in B$ and define a GFU-model $\mathfrak{M}_0$ with domain $M_0 = \{a_0, b_0\}$ so that $\mathfrak{M}_0 \models \beta_0(a_0, b_0)$. It is easy to check that $\mathfrak{M}_0$ satisfies all the stipulated conditions. In particular, (iv) indirectly implies that $\beta_0{\restriction}x$ and $\beta_0{\restriction}y$ (the restrictions of β_0 to atoms containing only the specified variables) belong to A, which shows (2). Another indirect consequence of (iv) is that $\beta_0(x, y) \models \varphi_0$, hence $\mathfrak{M}_0 \models \varphi_0$, which we will use later.

Now fix $n > 0$ and suppose that structures $\langle\mathfrak{M}_i : i < n\rangle$ satisfying the inductive conditions have already been defined. The scheme of the induction step is as follows. First we will extend $\mathfrak{M}_{n-1}$ to an auxiliary model $\mathfrak{M}_n^*$ and prove some intermediate properties. Then we will obtain $\mathfrak{M}_n$ by modifying the structure of $\mathfrak{M}_n^*$, without changing the domain. Finally we will check that $\mathfrak{M}_n$ has the desired properties.

For any pair $a, b \in M_{n-1}$ realizing some 2-type $\beta_{a,b}(x, y) \in B$ use assumption (iv) to find a model $\mathfrak{M}^{a,b}$ satisfying the conditions stated there with respect to $\beta_{a,b}(x, y)$. In particular $\mathfrak{M}^{a,b}$ contains a pair of elements realizing $\beta_{a,b}(x, y)$, so we may assume that $M^{a,b} \cap M_{n-1} = \{a, b\}$ and $\{a, b\}$ is a common substructure of those models. We also assume that for two such pairs $a, b, a', b' \in M_{n-1}$ the domains of $\mathfrak{M}^{a,b}$ and $\mathfrak{M}^{a',b'}$ have the least possible intersection, *i.e.*, $M^{a,b} \cap M^{a',b'} = \{a, b\} \cap \{a', b'\} \subseteq M_{n-1}$.

The domain of $\mathfrak{M}_n^*$ is the union of M_{n-1} and $M^{a,b}$ over all pairs as above.[1] Now we define the structure of $\mathfrak{M}_n^*$ so that $\mathfrak{M}_{n-1}$ and each $\mathfrak{M}^{a,b}$ is a substructure of $\mathfrak{M}_n^*$. Hence for any tuple

[1]Note that if both $M^{a,b}$ and $M^{b,a}$ have been defined for some distinct $a, b \in M_{n-1}$, we include both of them in the union. We might as well include just one of them and the proof remains valid.

$\bar{c} \subseteq M_n^*$ contained in a single model M' considered in the union, we have $\mathrm{tp}^{\mathfrak{M}_n^*}(\bar{c}) = \mathrm{tp}^{\mathfrak{M}'}(\bar{c})$. Note that if $\bar{c}$ is contained in two such models, then it is in fact contained in some pair $\{a,b\} \subseteq M_{n-1}$ and the type of $\bar{c}$ in both models is the same as in $\mathfrak{M}_{n-1}$ so no inconsistency arises. If $\bar{c} \subseteq M_n^*$ is not contained in any single model, we declare $\mathfrak{M}_n^* \models \neg P(\bar{c})$ for all predicates of $\mathcal{L}$ of the same arity as the length of $\bar{c}$.

We claim $\mathfrak{M}_n^*$ has the following properties:

(a) $\mathfrak{M}_n^* \models \bigwedge_{s \in S} \varphi_s$,

(b) $\mathfrak{M}_n^* \models \mathsf{U}(a,b)$ for each guarded pair $(a,b) \subseteq M_n^*$,

(c) $\mathrm{tp}(c) \in A$ for each $c \in M_n^*$.

To check (a), fix $s \in S$ and a tuple $\bar{c} \subseteq M_n^*$ satisfying $\mathfrak{M}_n^* \models G_s(\bar{c})$. Since $\bar{c}$ is guarded in $\mathfrak{M}_n^*$, it is contained in a single model from the union defining M_n^*. If $\bar{c} \subseteq M^{a,b}$ for some $a,b \in M_{n-1}$, then there is a witness for $\bar{c}$ in $\mathfrak{M}^{a,b}$ because $\mathfrak{M}^{a,b} \models \varphi_s$. Otherwise $\bar{c} \subseteq M_{n-1}$ and by (4) there are two cases: either $\bar{c}$ has a witness in $\mathfrak{M}_{n-1}$ or $\bar{c} \in \{a,b\}^k$ for some $a,b \in M_{n-1}$ realizing in $\mathfrak{M}_{n-1}$ a type $\beta_{a,b}(x,y) \in B$. In the first case we are done, in the second $\bar{c} \subseteq M^{a,b}$ so again $\bar{c}$ has a witness in $\mathfrak{M}^{a,b}$.

To prove (b), take any guarded pair $(a,b) \subseteq M_n^*$. Then (a,b) is contained in one of the models from the union defining M_n^*. If this model is $\mathfrak{M}_{n-1}$, which is a GFU-model, then the conclusion is clear. Otherwise it is some $\mathfrak{M}^{a,b}$, which satisfies in particular the last bullet from (iv), hence the conclusion is clear as well. Finally, (c) follows from (2) and the properties of each $\mathfrak{M}^{a,b}$ listed in (iv).

Now we define a new structure $\mathfrak{M}_n$ on the domain $M_n = M_n^*$. For each pair $a,b \in M_n^*$ such that $\mathfrak{M}_n^* \models \neg\mathsf{U}(a,b)$, by (c) we can use (iii) to find some $\beta_{a,b}(x,y) \in B$ extending $\mathrm{tp}(a)(x) \cup \mathrm{tp}(b)(y)$. Let $\mathfrak{M}_n$ be the model obtained from $\mathfrak{M}_n^*$ by assigning to each pair (a,b) as above the new type $\beta_{a,b}(x,y)$.[2] This is possible because if (a,b) is such a pair, then $a \neq b$ by (i) and (c). It remains to check that $\mathfrak{M}_n$ is a GFU-model satisfying (1) – (4). We will use the following observations:

($\dagger_1$) for each $\bar{c} \subseteq M_n$ and predicate P, if $P(\bar{c})$ has different truth value in $\mathfrak{M}_n^*$ and $\mathfrak{M}_n$, then there are $a,b \in M_n^*$ such that $\mathfrak{M}_n^* \models \neg\mathsf{U}(a,b)$ and $\bar{c} \in \{a,b\}^k \setminus \{a^k, b^k\}$,

($\dagger_2$) for each $\bar{c} \subseteq M_n^*$ if $\mathfrak{M}_n^* \models \mathsf{U}(a,b)$ for all $a,b \in \bar{c}$ then $\mathrm{tp}^{\mathfrak{M}_n}(\bar{c}) = \mathrm{tp}^{\mathfrak{M}_n^*}(\bar{c})$,

($\dagger_3$) for each $\bar{c} \subseteq M_n^*$ guarded in $\mathfrak{M}_n^*$ we have $\mathrm{tp}^{\mathfrak{M}_n}(\bar{c}) = \mathrm{tp}^{\mathfrak{M}_n^*}(\bar{c})$,

where ($\dagger_1$) follows directly from the definition, ($\dagger_2$) is implied by ($\dagger_1$) and ($\dagger_3$) is true by ($\dagger_2$), (b) and (c) + (i).

To check that $\mathfrak{M}_n$ is a GFU-model, fix $a,b \in M_n$. If $\mathfrak{M}_n^* \models \mathsf{U}(a,b)$, then (a,b) is guarded in $\mathfrak{M}_n^*$, so $\mathfrak{M}_n \models \mathsf{U}(a,b)$ by ($\dagger_3$). Otherwise $\mathrm{tp}^{\mathfrak{M}_n}(a,b) = \beta_{a,b}$ or $\beta_{b,a}$, either way $\mathfrak{M}_n \models \mathsf{U}(a,b)$ by (ii).

Finally we check that $\mathfrak{M}_n$ satisfies properties (3) and (4), as the other two are easy. Let $s \in S$ and fix $\bar{c} \subseteq M_n$ such that $\mathfrak{M}_n \models G_s(\bar{c})$. We consider two cases. If $\mathfrak{M}_n^* \models G_s(\bar{c})$, then by (a) there is $\bar{d} \subseteq M_n^*$ such that $\mathfrak{M}_n^* \models H_s(\bar{c},\bar{d}) \wedge \theta_s(\bar{c},\bar{d})$. In particular $\bar{c}\bar{d}$ is guarded in $\mathfrak{M}_n^*$, so ($\dagger_3$) implies $\mathfrak{M}_n \models H_s(\bar{c},\bar{d}) \wedge \theta_s(\bar{c},\bar{d})$, as required. In the second case we have $\mathfrak{M}_n^* \models \neg G_s(\bar{c})$, so by ($\dagger_1$) we have $\bar{c} \in \{a,b\}^k \setminus \{a^k, b^k\}$ for some $a,b \in M_n^*$ such that $\mathfrak{M}_n^* \models \neg\mathsf{U}(a,b)$. But then $\mathrm{tp}^{\mathfrak{M}_n}(a,b) = \beta_{a,b}$ or $\beta_{b,a}$, so (4) holds with a and b possibly exchanged. If we additionally assume that $\bar{c} \subseteq M_{n-1}$, then from ($\dagger_2$) and the fact that $\mathfrak{M}_{n-1}$ is a GFU-model, we have $\mathrm{tp}^{\mathfrak{M}_n^*}(\bar{c}) = \mathrm{tp}^{\mathfrak{M}_n}(\bar{c})$. So under that assumption the first of the two cases above must hold, hence we get (3).

[2] If $\mathfrak{M}_n^* \models \neg\mathsf{U}(a,b) \wedge \neg\mathsf{U}(b,a)$, we assign either $\beta_{a,b}$ or $\beta_{b,a}$ arbitrarily.

This ends the inductive construction. Now consider $\mathfrak{M} = \bigcup_{n<\omega} \mathfrak{M}_n$. Clearly as a union of an increasing chain of GFU-models, $\mathfrak{M}$ is a GFU-model. As marked before, $\mathfrak{M}_0 \models \varphi_0$ hence $\mathfrak{M} \models \varphi_0$. Finally, $\mathfrak{M} \models \bigwedge_{s\in S} \varphi_s$ follows from (3), so the proof is complete. $\square$

We are now ready to establish the complexity of the satisfiability problem for GFU, and thus also of TGF. Before formulating the main theorem of this section, we define some complexity parameters. A formula φ is of length $n = n(\varphi)$, has $m = m(\varphi)$ variables in total and $r = r(\varphi)$ predicates of maximal arity $\ell = \ell(\varphi)$. When x is a parameter, by $p(x)$ we mean any polynomial in x. We will use the following result by Grädel:

Lemma 3.8 (Theorem 4.3, [7]). *The satisfiability problem for* GF *is in* 2-ExpTime. *More precisely, satisfiability of a given formula φ can be tested in time $\mathcal{O}(c^{rm^\ell} \cdot n)$ for some $c > 0$.*

Strictly speaking, Grädel explicitly formulates only the first part of the above Lemma, that is, states his theorem in terms of n. However, the second part of the Lemma can be easily extracted from his proof. Namely, in his paper Grädel presents an alternating procedure running in space $\mathcal{O}(r \cdot m^\ell)$ with respect to the input formula. We can thus get the desired estimation using the classical Chandra, Kozen and Stockmeyer's simulation of alternating space-bounded Turing machines by deterministic ones [4]. Let us now show the main result of this section.

Theorem 3.9. *The satisfiability of* GFU *(and hence* TGF*) is in* 2-ExpTime.

Proof. We describe a satisfiability test for sentences in GFU: Read a sentence φ in GFU and use Lemma 3.3 to get a formula χ in GF-normal form (in some possibly extended language $\mathcal{L}'$) such that φ has a model if and only if χ has a model. Following the notation of Lemma 3.7, check if there is a set A of complete 1-types (in $\mathcal{L}'$) and a set $B \neq \varnothing$ of complete 2-types of size $\leqslant |A|^2$ such that the conditions (i) - (iv) from the Lemma hold. If yes then accept; otherwise reject. Note that we can indeed restrict attention to sets B of size $\leqslant |A|^2$, since it suffices to have in B just one 2-type for every pair of 1-types from A.

Given A and B, the verification of conditions (i) - (iii) is straightforward, while checking (iv) is done by writing for every $\beta \in B$ the following formula

$$\psi_\beta := \chi \wedge (\exists xy)\beta(x,y) \wedge (\forall x) \bigvee_{\alpha\in A} \alpha(x) \wedge \bigwedge_{P\in\mathcal{L}'} (\forall \overline{x})\big(P(\overline{x}) \to \bigwedge_{1\leqslant i,j\leqslant|\overline{x}|} \mathsf{U}(x_i,x_j)\big),$$

and verifying its satisfiability by executing a subprocedure referred to by Lemma 3.8.

We now analyse the computational complexity of the procedure and simultaneously fill in the missing details. The computation of χ takes polynomial time in n. Since $|A| = \mathcal{O}(2^{r(\chi)})$ and $r(\chi) \in \mathcal{O}(n)$, the number of possible choices of A and B is doubly exponential in n. We just exhaustively consider all possibilities. For a fixed pair of sets of types A, B, the verification of conditions (i) – (iii) clearly can be done in time exponential in n. As we said, to check (iv), for each $\beta \in B$ we need to verify satisfiability of ψ_β. Note that $n(\psi_\beta) = \mathcal{O}(2^{p(n(\varphi))})$ (the size of A and thus the number of disjuncts in the third conjunct can indeed be exponential in n), but $m(\psi_\beta) \in \mathcal{O}(n(\chi))$, $r(\psi_\beta) = r(\chi)$ and $\ell(\psi_\beta) = \ell(\chi)$. So, using the test from Lemma 3.8 we get the answer in time $\mathcal{O}(c^{r(\psi_\beta)m(\psi_\beta)^{\ell(\psi_\beta)}} \cdot n(\psi_\beta)) = \mathcal{O}(c^{r(\chi)n(\chi)^{\ell(\chi)}} \cdot 2^{p(n(\chi))})$, which, taking into account that $l(\chi)$ is bounded linearly in $n(\chi)$, is doubly exponential in $n(\chi)$. Hence the total time needed for computation is double exponential in $n(\varphi)$.

The correctness of the algorithm follows directly from Lemma 3.7. $\square$

We recall that the upper bound from Thm. 3.9 is optimal as the matching lower bound holds already for GF without equality [7]. We remark that Thm. 3.9 is not new—as we said in the Introduction, essentially the same result was already proved, using different techniques, in [11] and [2]. From the above analysis we also get the following corollary.

Corollary 3.10. *The satisfiability problem for* GFU *(* TGF*) is* NExpTime*-complete when the arity of relation symbols is bounded by a constant.*

Proof. We use a modification of the satisfiability test from the proof of Thm. 3.9 in which instead of considering all possible pairs of sets of types A, B we just nondeterministically guess one such pair. Since the parameter $\ell(\varphi')$ is now bounded by a constant, satisfiability of each ψ_β can be tested in time exponential in $n(\varphi')$. The upper bound then follows. The corresponding lower bound is inherited from FO^2 without equality [13]. □

We recall that also Corollary 3.10 is not a new result. Establishing the complexity of TGF in this case was left as an open problem in [2]. This problem was solved in [15] in a richer scenario involving constant symbols.

4 The triguarded fragment with transitive guards

In this section we consider the triguarded fragment with transitive guards, TGF+TG. First, let us recall the guarded fragment with transitive guards, GF+TG from [17]. It is an extension of GF in which the predicates are divided into two types: free and transitive. Free predicates are ordinary, while transitive predicates must be interpreted as transitive binary relations, however, their occurrences in formulas are restricted to guards. Below is a formal definition:

Definition 4.1. The GF+TG logic is defined as follows:

(i) A TG-*language* is a pair $\mathcal{L} = (\mathcal{P}_F, \mathcal{P}_T)$, where $\mathcal{P}_F$ is a set of predicate symbols, $\mathcal{P}_T$ is a set of binary predicate symbols and $\mathcal{P}_F \cap \mathcal{P}_T = \varnothing$ (but $\mathcal{P}_F$ can still contain some binary predicate symbols). The symbols in $\mathcal{P}_F$ are called *free predicates* while the symbols in $\mathcal{P}_T$ are *transitive predicates.*

(ii) Let $\mathcal{L}$ be as above. A formula in $\mathcal{P}_F \cup \mathcal{P}_T$ is called a GF+TG-*formula* in $\mathcal{L}$ if it is constructed as in Definition 3.1 with the following changes:

- in (i) we disallow formulas containing occurrences of transitive predicate symbols,
- in (iii) a vacuous quantifier cannot have a transitive guard.

Also we call such formula *free* if every predicate that occurs in that formula is free.

(iii) A *complete (atomic) type in* $\mathcal{L}$ is a complete type in $\mathcal{P}_F \cup \mathcal{P}_T$, while a *complete free type* and *complete transitive type* is a complete type in $\mathcal{P}_F$ and $\mathcal{P}_T$ respectively.

(iv) A TG-*model* in $\mathcal{L}$ is a model in $\mathcal{L}$ (*i.e.* in $\mathcal{P}_F \cup \mathcal{P}_T$) in which every symbol from $\mathcal{P}_T$ is interpreted as a transitive binary relation.

(v) A guard $P(\overline{x})$ is called *transitive* if P is transitive, otherwise it is *free.*

Now TGF+TG could be defined as the natural common extension of TGF and GF+TG. But as before, we find it easier to work with GFU rather than TGF, so we proceed directly to defining the GFU+TG logic which has the power of expression equal to that of TGF+TG.

Definition 4.2. The GFU+TG logic is defined as follows:

(i) A TG-language $\mathcal{L} = (\mathcal{P}_F, \mathcal{P}_T)$ is a GFU+TG-*language* if $\mathcal{P}_F$ contains a distinguished symbol U of a binary predicate.

(ii) A GFU+TG-*formula* is a GF+TG-formula in a GFU-language.

(iii) A GFU+TG-*model in* $\mathcal{L}$ is a TG-model in $\mathcal{L}$ which is also a GFU-model.

As in the previous section, the sentences in GFU+TG can be equivalently transformed to a more convenient normal form. The proof is analogous, but the restriction of transitive guards to non-vacuous quantifiers makes it necessary to add one type of conjunct:

Definition 4.3. Assume φ is a sentence in GF+TG (*e.g.* in GFU+TG). We say φ is in GF+TG-*normal form* if it is a conjunction of GF+TG-sentences of the following three forms:

N_1. any sentence in $\mathcal{F}^{\mathrm{qf}}_{\neq}$,

N_2. $(\forall\overline{x})\big(G(\overline{x}) \to (\exists\overline{y})\big(H(\overline{x},\overline{y}) \wedge \theta(\overline{x},\overline{y})\big)\big)$, where G, H are guards and $\theta \in \mathcal{F}^{\mathrm{qf}}_{\neq}$,

N_3. $(\forall\overline{x})\big(G(\overline{x}) \to \theta(\overline{x})\big)$, where G is a guard and $\theta \in \mathcal{F}^{\mathrm{qf}}_{\neq}$.

Lemma 4.4. *An analogue of Lemma 3.3 holds for* GF+TG-*sentences.* □

We now aim to prove that the satisfiability of GFU+TG is decidable. The key lemma is similar to Lemma 3.7 and the differences that could be overlooked are emphasized in bold font.

Lemma 4.5. *Assume φ is a* GFU+TG-*sentence in* GF+TG-*normal form. Then φ has a* GFU+TG-*model if and only if there exist a set $A \neq \varnothing$ of complete **free** 1-types and a set $B \neq \varnothing$ of complete **free** 2-types such that the following conditions hold:*

(i) $\alpha(x) \models \mathsf{U}(x,x)$ *for each* $\alpha \in A$*,*

(ii) $\beta(x,y) \models x \neq y \wedge \mathsf{U}(x,y) \wedge \mathsf{U}(y,x)$ *for each* $\beta \in B$*,*

(iii) for each $\alpha_1(x), \alpha_2(y) \in A$ *the partial **free** 2-type* $\alpha_1(x) \cup \alpha_2(y)$ *extends to some complete **free** type* $\beta(x,y) \in B$*,*

(iv) for each $\beta(x,y) \in B$ *there exists a* TG***-model*** $\mathfrak{M} \models \varphi$ *such that*

- *M contains some a, b with $\mathfrak{M} \models \beta(a,b)$,*
- *each $c \in M$ realizes some $\alpha(x) \in A$,*
- *every guarded pair $(a,b) \subseteq M$ satisfies $\mathfrak{M} \models \mathsf{U}(a,b)$.*

Proof. The core of the proof is the same as in that of Lemma 3.7, so we only discuss the parts which need a nontrivial change.

The argument for the implication from left to right remains valid, except that we apply Lemma 2.5 to the sentence $\varphi \wedge (\forall x,y)\,\mathsf{U}(x,y) \wedge \bigwedge_{T\in\mathcal{P}_T}(\forall x,y,z)\big(T(x,y) \wedge T(y,z) \to T(x,z)\big)$ so that we end up with a GFU+TG-model of φ again.

For the other direction, write φ in GF+TG-normal form as $\varphi = \varphi_0 \wedge \bigwedge_{s\in S}\varphi_s \wedge \bigwedge_{s\in S'}\varphi'_s$ where φ_0 is a sentence in $\mathcal{F}^{\mathrm{qf}}_{\neq}$, S, S' are finite sets of indices,

$$\begin{aligned}
\varphi_s &= (\forall\overline{x})\big(G_s(\overline{x}) \to (\exists\overline{y})\big(H_s(\overline{x},\overline{y}) \wedge \theta_s(\overline{x},\overline{y})\big)\big) &&=: (\forall\overline{x})\,\psi_s(\overline{x}) \qquad \text{for } s \in S,\\
\varphi'_s &= (\forall\overline{x})\big(G'_s(\overline{x}) \to \theta'_s(\overline{x})\big) &&=: (\forall\overline{x})\,\psi'_s(\overline{x}) \qquad \text{for } s \in S',
\end{aligned}$$

each G_s, G'_s, H_s is a guard and $\theta_s, \theta'_s \in \mathcal{F}^{\mathrm{qf}}_{\neq}$. We will define a sequence $\langle \mathfrak{M}_n : n < \omega\rangle$ of GFU+TG-models satisfying the following conditions for each n:

(1) – (3) as before,

(4) for each tuple $\bar{c} \subseteq M_n$ and $s \in S$, if $\mathfrak{M}_n \models \neg\psi_s(\bar{c})$, then G_s **is a free guard** and there exist $a, b \in M_n$ realizing some free type $\beta(x,y) \in B$ such that $\bar{c} \in \{a,b\}^k \setminus \boldsymbol{\{a^k, b^k\}}$, where $k = |\bar{c}|$.

(5) $\mathfrak{M}_n \models \bigwedge_{s \in S'} \varphi'_s$.

We postpone the construction of $\mathfrak{M}_0$ until the end of the proof. Now fix $n > 0$ and suppose the models $\langle \mathfrak{M}_i : i < n \rangle$ have already been defined.

For any pair $a, b \in M_{n-1}$ realizing some free 2-type $\beta_{a,b}(x,y) \in B$ let $\mathfrak{M}^{a,b}$ be a model in $\mathcal{L}$ as in (iv) such that $M^{a,b} \cap M_{n-1} = \varnothing$ and the pair in $M^{a,b}$ realizing $\beta_{a,b}$ is (a^*, b^*). Put $N^{a,b} = M^{a,b} \cup \{a,b\}$ and let $\pi^{a,b} : N^{a,b} \to M^{a,b}$ be such that $\pi^{a,b} \restriction M^{a,b}$ is the identity and $\pi(a,b) = (a^*, b^*)$. We define the model $\mathfrak{N}^{a,b}$ in the language $\mathcal{P}_F$ on the domain $N^{a,b}$ as the pullback of $\mathfrak{M}^{a,b} \restriction \mathcal{P}_F$ through $\pi^{a,b}$.

Let $\mathfrak{M}^*_n$ be the model with the domain $M_{n-1} \cup \bigcup_{(a,b)} M^{a,b}$ and structure defined as follows. For each $T \in \mathcal{P}_T$ and $c, d \in M^*_n$ we declare that $\mathfrak{M}^*_n \models T(c,d)$ holds if and only if $\{c,d\}$ is contained in M_{n-1} or some $M^{a,b}$ and $T(c,d)$ holds there. So the graph of T in $\mathfrak{M}^*_n$ is the disjoint union of the graphs of T in $\mathfrak{M}_{n-1}$ and $\mathfrak{M}^{a,b}$ over all (a,b). On the other hand, the 'free' part of the structure on $\mathfrak{M}^*_n$ comes from the amalgamation of the models $\mathfrak{M}_{n-1}$ and $\mathfrak{N}^{a,b}$ over all (a,b) as above. Note that $\mathfrak{M}_{n-1}$ and each $\mathfrak{M}^{a,b}$ is a substructure of $\mathfrak{M}^*_n$ and each $\mathfrak{N}^{a,b}$ is a substructure of $\mathfrak{M}^*_n \restriction \mathcal{P}_F$.

We will prove $\mathfrak{M}^*_n$ is a TG-model with the following properties:

(a) $\mathfrak{M}^*_n \models \bigwedge_{s \in S} \varphi_s$,

(a') $\mathfrak{M}^*_n \models \bigwedge_{s \in S'} \varphi'_s$,

(b) $\mathfrak{M}^*_n \models \mathsf{U}(a,b)$ for each guarded pair $(a,b) \subseteq M^*_n$,

(c) $\mathrm{tp}(c) \restriction \mathcal{P}_F \in A$ for each $c \in M^*_n$.

For (a'), take $s \in S'$ and assume $\bar{c} \subseteq M^*_n$ satisfies $\mathfrak{M}^*_n \models G'_s(\bar{c})$. If G'_s is a transitive guard, then by definition $\bar{c}$ is contained in a single model, $\mathfrak{M}_{n-1}$ or some $\mathfrak{M}^{a,b}$ and $G'_s(\bar{c})$ holds there. That model is a substructure of $\mathfrak{M}^*_n$ and it satisfies $\psi'_s(\bar{c})$, hence $\mathfrak{M}^*_n \models \theta'_s(\bar{c})$, as required. If G'_s is a free guard, then $\bar{c}$ is contained in a single model, $\mathfrak{M}_{n-1}$ or $\mathfrak{N}^{a,b}$, which is a substructure of $\mathfrak{M}^*_n$ with respect to the free structure. Also in this case $\psi'_s(\bar{c})$ is a free sentence which holds in that model, thus $\mathfrak{M}^*_n \models \theta'_s(\bar{c})$.

Now we check (a). Fix $s \in S$ and take $\bar{c} \subseteq M^*_n$ satisfying $\mathfrak{M}^*_n \models G_s(\bar{c})$. First suppose G_s or H_s is a transitive guard. We then claim that $\bar{c}$ is again contained in a single model $\mathfrak{M}_{n-1}$ or some $\mathfrak{M}^{a,b}$. Indeed: if G_s is transitive, it follows directly from the construction of $\mathfrak{M}^*_n$. On the other hand, if H_s is transitive, then $\bar{y}$ in φ_s must be non-empty and $|\overline{xy}| \leqslant 2$, hence $|\bar{c}| = |\bar{x}| \leqslant 1$, so the claim must hold as well. Now the said model is a substructure of $\mathfrak{M}^*_n$, so it satisfies $G_s(\bar{c})$. If that model is $\mathfrak{M}^{a,b}$, then $\mathfrak{M}^{a,b} \models \psi_s(\bar{c})$ and it is easy to get the desired conclusion. But if it is $\mathfrak{M}_{n-1}$, then $\mathfrak{M}_{n-1} \models \psi_s(\bar{c})$, for otherwise by (4), H_s would be transitive and $|\bar{c}| > 1$, which contradicts the analysis above. The conclusion follows in this case as well. Finally suppose both G_s and H_s are free. Then φ_s is free and we argue as in the original proof.

The proofs of (b) and (c) do not change and obviously $\mathfrak{M}^*_n$ is a TG-model.

Now we define the structure $\mathfrak{M}_n$ on the domain $M_n = M^*_n$ similarly as before, by assigning to each pair $(a,b) \subseteq M_n$ with $\mathfrak{M}^*_n \models \neg\mathsf{U}(a,b)$ a new free type from B extending their free 1-types, but keeping their transitive types unchanged. The properties $(\dagger_1)$ – $(\dagger_3)$ from the previous proof still hold in this situation and moreover, under the assumption of $(\dagger_1)$ we have that P is a free predicate.

$\mathfrak{M}_n$ is clearly a TG-model and we prove as before that it is a GFU-model. The properties (1), (2) are again easy to check and (3), (4) are proved the same way as before taking into account the strengthening of ($\dagger_1$) mentioned in the previous paragraph. Finally we check (5): fix $s \in S'$ and a tuple $\bar{c} \subseteq M_n$ satisfying $\mathfrak{M}_n \models G'_s(\bar{c})$. If $\mathfrak{M}^*_n \models G'_s(\bar{c})$, then ($\dagger_3$) implies $\mathrm{tp}^{\mathfrak{M}_n}(\bar{c}) = \mathrm{tp}^{\mathfrak{M}^*_n}(\bar{c})$ and also $\mathfrak{M}^*_n \models \varphi'_s$, hence $\mathfrak{M}_n \models \theta'_s(\bar{c})$. Otherwise G'_s is free, hence so is φ'_s, and $\bar{c} \in \{a,b\}^k \setminus \{a^k, b^k\}$ for some $a, b \in M_n$ with $\mathfrak{M}^*_n \models \neg\mathsf{U}(a,b)$, which again implies $\mathfrak{M}_n \models \beta(a,b)$ for some free type $\beta(x,y) \in B$. Let $\bar{z} \in \{x,y\}^k \setminus \{x^k, y^k\}$ be the tuple of variables which becomes $\bar{c}$ under the substitution $(x,y) := (a,b)$. Then $\beta(x,y) \models G'_s(\bar{z}) \to \theta'_s(\bar{z})$ by (iv), therefore $\mathfrak{M}_n \models \theta'_s(\bar{c})$, as required.

To finish the inductive construction we define the model $\mathfrak{M}_0$. Pick any $\beta(x,y) \in B$ and find a model $\mathfrak{M}^*_0$ satisfying (iv) with respect to β. Then define $\mathfrak{M}_0$ by the same procedure that was used to obtain $\mathfrak{M}_n$ from $\mathfrak{M}^*_n$. Clearly $\mathfrak{M}^*_0 \models$ (a) – (c) and therefore $\mathfrak{M}_0 \models$ (2) – (5).

Finally note that $\mathfrak{M}^*_0 \models \varphi_0$ and by ($\dagger_1$) the type of the empty tuple is the same in $\mathfrak{M}_0$ as in $\mathfrak{M}^*_0$, hence $\mathfrak{M}_0 \models \varphi_0$. The rest of the proof is as before with an additional observation that a union of a chain of TG-models is a TG-model. □

We now turn to establishing the complexity of TGF+TG. As black box we use this time the following result of Szwast and Tendera:

Lemma 4.6 (Theorem 38, [17])**.** *The satisfiability problem for* GF+TG *is in* 2-ExpTime. *More precisely, the satisfiability of a given* GF+TG*-sentence* Δ *in a specific normal form can be tested in time* $c^{\mathcal{O}\left(r^2 \cdot m^\ell \cdot (3n)^m \cdot 2^{rm}\right)} \cdot n$ *for some* $c > 0$.

Again the cited theorem only asserts the general doubly exponential bound with respect to $n(\Delta)$, but the more precise estimation can be extracted by reading the proof and using the Chandra, Kozen and Stockmeyer's simulation of alternating space-bounded Turing machines from [4].

The specific normal form mentioned in the Lemma is that from [17], which we recall here for convenience:

Definition 4.7 (Definition 1, [17])**.** A GF+TG-sentence Δ is in normal form if it is a conjunction of sentences of the following form:

(n1) $(\exists x)(\alpha(x) \wedge \psi(x))$,

(n2) $(\forall \mathbf{x})\big(\alpha(\mathbf{x}) \to (\exists y)(\beta(\mathbf{x}, y) \wedge \psi(\mathbf{x}, y))\big)$,

(n3) $(\forall \mathbf{x})\big(\alpha(\mathbf{x}) \to \psi(\mathbf{x})\big)$,

where α and β are atomic formulas (guards) and $y \notin \mathbf{x}$, all the variables listed in $\beta(\mathbf{x}, y)$ do occur in β, ψ is quantifier-free and it contains no transitive predicate letter.

The above normal form is similar to ours. A slight difference is that its blocks of existential quantifiers are always of length 1.

We are ready to show the main theorem.

Theorem 4.8. *Satisfiability of* GFU+TG *(and therefore* TGF+TG*) is in* 2-ExpTime.

Proof. Mostly, repeat the proof of Theorem 3.9. The only part that needs to be changed and reanalysed is the verification of the condition (iv) for a fixed $\beta \in B$. The modification to the algorithm is as follows: assume the program has already computed the sentence

$$\psi_\beta = \chi \wedge (\exists xy)\beta(x,y) \wedge (\forall x) \bigvee_{\alpha \in A} \alpha(x) \wedge \bigwedge_{P \in \mathcal{L}'} (\forall \bar{x})(P(\bar{x}) \to \bigwedge_{1 \leqslant i,j \leqslant |\bar{x}|} \mathsf{U}(x_i, x_j)).$$

It is then transformed into a sentence Δ in the normal form from Definition 4.7 using the standard procedure of recursively replacing single-quantified subformulas with a predicate.[3] Using Lemma 4.6 as a subroutine, verify satisfiability of Δ. The rest is as before.

Now we analyse the complexity of the new part of the algorithm. Recall that the input to our algorithm is the GFU+TG-sentence φ and χ is its reduction to our GF+TG-normal form given by Lemma 4.4. It is enough to show that the time it takes to execute the described fragment of the satisfiability test is doubly exponential in $n(\varphi)$.

The conversion of ψ_β into Δ takes (possibly nondeterministic) polynomial time in $n(\psi_\beta)$, which is exponential in $n(\varphi)$. By a routine analysis we get $n(\Delta) = \mathcal{O}(n(\psi_\beta)), m(\Delta) \leqslant m(\psi_\beta) = \max\{m(\chi), 2\} = \mathcal{O}(n(\varphi)), \ell(\Delta) = \ell(\psi_\beta) = \mathcal{O}(n(\varphi))$. Also $r(\Delta)$ is exactly $r(\psi_\beta)$ increased by the number of quantifiers in ψ_β, where $(\exists \overline{x})$ is counted $|\overline{x}|$ times, so ultimately $r(\Delta) = \mathcal{O}(r(\chi) + n(\chi)) = \mathcal{O}(n(\varphi))$. By Lemma 4.6 the verification of satisfiability of Δ takes

$$c^{\mathcal{O}\left(r(\Delta)^2 \cdot m(\Delta)^{\ell(\Delta)} \cdot (3n(\Delta))^{m(\Delta)} \cdot 2^{r(\Delta)m(\Delta)}\right)} \cdot n(\Delta)$$

time, which is doubly exponential in $n(\varphi)$.

Overall the whole algorithm works in doubly exponential time as required. □

We recall that the upper bound in the above theorem is optimal as already GF without equality is 2-ExpTime-hard [7].

5 Discussion

In this section we briefly discuss some applications of our method in slightly modified/extended scenarios, deriving some new results on triguarded logics.

5.1 Equality in formulas guarded by transitive atoms

In the proof of our central technical result, Lemma 4.5 (or of its variant for the language without transitive relations, Lemma 3.7) it is the left-to-right direction in which the absence of equalities is crucial. Consider for example the GFU conjuncts $(\exists x)P(x)$ and $(\forall xy)(\mathsf{U}xy \to (Px \wedge Py \to x = y))$. Any GFU-model $\mathfrak{M}$ which satisfies them must contain a king (cf. Def. 2.4). For a formula enforcing kings we would not be able to find appropriate sets A and B. More specifically, the set A should then contain a 1-type α realized by a king in any model of the formula and the problem would be to fulfil condition (iii) for the pair (α, α). Recall that every satisfiable formula without equality has a model without kings (Lemma 2.5), and note that for a formula (even with equality) having a model without kings there are sets A, B (*e.g.*, the sets of types realized in this model) fulfilling the conditions of Lemma 4.5.

Interestingly, the absence of equality is not important for the proof of the right-to-left direction of Lemma 4.5. A routine inspection shows that our model construction works fine even if the formulas contain equalities. Thus, using the fact that the procedure of Szwast and Tendera employed as our external subroutine works in the presence of equalities, we can use our approach to test satisfiability of formulas with equality *in models without kings*. This observation is similar to the observation by Gurevich and Shelah [10], concerning their proof of the decidability of the Gödel class without equality.

[3] A minor detail is that the work of Szwast and Tendera [17] seems not to admit relations of arity 0. So, to take this into account we can additionally nondeterministically guess the truth values of all such relations and replace their occurrences in ψ_β by $\top$ or $\bot$ depending on the guess.

One may want to try to find some syntactic conditions restricting the use of equality in order to guarantee that a TGF+TG formula, whenever satisfiable, has a model without kings. Let us propose one simple such condition, allowing us to subtly increase the expressive power of the logic. For simplicity, let us consider only formulas in normal form. It turns out that without sacrificing the decidability we can admit equalities in those conjuncts of the form N_3 whose guards are transitive. Indeed, any satisfiable normal form TGF+TG formula using equalities only in this way has a model without kings.

Lemma 5.1. *Let φ be a satisfiable normal form* GFU+TG *formula in whose conjuncts of type N_3 guarded by transitive atoms (and only in such conjuncts) equalities can be used. Then φ has a model without kings.*

Proof. Let $\mathfrak{N} \models \varphi$. Let M consists of two disjoint copies of the domain of $\mathfrak{N}$, and $f : M \to N$ be the function returning for $a \in M$ the element of N which a is a copy of. Define the structure on $\mathfrak{M}$ in such a way that its free part is the pullback of the free part of $\mathfrak{N}$ through f and $\mathfrak{M} \models T(a, b)$ iff a and b belong to the same copy of N and $\mathfrak{N} \models T(f(a), f(b))$. Clearly T is transitive in $\mathfrak{M}$. We show that $\mathfrak{M} \models \varphi$. First, note that all conjuncts of φ of the form N_1 or N_2 and those conjuncts of the form N_3 whose guard is non-transitive are satisfied in $\mathfrak{M}$ by Remark 2.2 (since they do not contain equalities). Consider a conjunct of the form N_3, $(\forall \overline{x})\big(G(\overline{x}) \to \theta(\overline{x})\big)$, with transitive G (and thus with $|\overline{x}| \leqslant 2$), possibly with occurrences of $=$ in θ. Take a tuple $\overline{a}$ such that $\mathfrak{M} \models G(\overline{a})$. Since G is transitive, by the definition of $\mathfrak{M}$ we know that the elements of $\overline{a}$ belong to the same copy of $\mathfrak{N}$, and thus $\mathfrak{M}\restriction\overline{a}$ is isomorphic to $\mathfrak{N}\restriction f(\overline{a})$. Since $\mathfrak{N} \models \theta(f(\overline{a}))$ it follows that $\mathfrak{M} \models \theta(\overline{a})$. Obviously, $\mathfrak{M}$ is without kings, as it contains at least two copies of every 1-type realized in $\mathfrak{N}$ and it realizes no other 1-types. ☐

Now, for example, we can use the trick from [12] to express that a transitive relation T is an equivalence relation: introduce an auxiliary unary predicate P_T and say: $(\forall x)(\exists y)(T(x, y) \wedge P_T(y))$, $(\forall x)(\exists y)(T(y, x) \wedge P_T(y))$, and $(\forall xy)(T(x, y) \to (P_T(x) \wedge P_T(y) \to x = y))$. One easily checks that in any model interpreting T as a transitive relation and satisfying the above formulas, T is an equivalence, and, in the opposite direction, any model interpreting T as an equivalence can be expanded to a model of those formulas by setting P_T true for precisely one element of every equivalence class of T. (Note that the formula expressing the symmetry of a transitive relation in a straightforward way, $(\forall xy)(T(x, y) \to T(y, x))$ is not in TGF+TG.)

We thus have the following strengthening of our main result.

Corollary 5.2. *The satisfiability problem for the variant of* TGF+TG *in which the set of transitive symbols $\mathcal{P}_T$ has a distinguished subset $\mathcal{P}_E$ whose elements must be interpreted as equivalence relations is decidable and 2-*ExpTime*-complete.*

The equivalence relations are natural in many applications, *e.g.*, in the field of epistemic logics, thus the ability of expressing them enhances the potential applicability of our logic.

5.2 Constants in the triguarded fragment

In our main result, the decidability of TGF+TG (GFU+TG), we do not include constants, as they are not allowed in GF+TG, decidability of which we need. We leave the decidability of TGF+TG with constants as an open problem, recalling that even the decidability of GF+TG is open. We also do not include constants when we present our method in the simpler setting of TGF (GFU). However, in this latter case this is because of the clarity of the presentation, and in fact our method can be adapted for TGF (GFU) with constants, and used to derive the optimal 2-NExpTime-upper complexity bound.

Moreover, in our approach we can also include a limited use of equality, namely, allow to use atoms of the form $x = c$ for a variable x and a constant symbol c. This additional construct is suggested by Rudolph and Šimkus [15] as a natural mean of expressing the concept of *nominals* ($\mathcal{O}$) from description logics. Rudolph and Šimkus suspect that such use of equality does not spoil the decidability nor increase the computational complexity, which we confirm here.

Theorem 5.3. *The satisfiability problem for* TGF+TG *with constants and equalities of the form $x = c$ where x is a variable and c is a constant is decidable and 2-*NExpTime*-complete.*

Let us sketch the modifications of our decidability proof for GFU covering the above extension.

First, note that Lemma 3.3, allowing us to restrict attention to normal form formulas, holds for the extended logic. In fact, no changes in the proof of this lemma are needed.

We will call a complete 1-type α *constant-free* if for every constant symbol c we have $\alpha(x) \models x \neq c$. Collecting all the ground literals of a given 1-type α we obtain a unique 0-type called the 0-type *induced* by α. This 0-type fully describes the restriction of any structure realizing α to the set of elements interpreting the constant symbols. Obviously all 1-types realized in a given structure induce the same 0-type.

To make the crucial Lemma 3.7 work in the extended scenario we modify (iii) to

(iii) for every pair of **constant-free** 1-types $\alpha_1(x), \alpha_2(y) \in A$ the partial 2-type $\alpha_1(x) \cup \alpha_2(y)$ extends to some complete type $\beta(x, y) \in B$.

and add the following three conditions

(v) all 1-types $\alpha \in A$ induce the same 0-type; call this 0-type γ_0

(vi) for every $\alpha \in A$ and every constant c we have $\mathsf{U}(x, c), \mathsf{U}(c, x) \in \alpha$

(vii) there is a model $\mathfrak{M}_0 \models \varphi \wedge \gamma_0$ such that every guarded pair $(a, b) \subseteq M_0$ satisfies $\mathfrak{M}_0 \models \mathsf{U}(a, b)$.

Note that the model $\mathfrak{M}_0$ from (vii), similarly to models whose existence is postulated by (iv), is not required to be a GFU-model, that is not necessarily all pairs of its elements are connected by U. We remark that in the case without constants the notion of 0-types also make sense; in that case a 0-type just consists of literals built out of relation symbols of arity 0. Observe that in the case without constants condition (v) is implied by (the original) condition (iii), and condition (vii) is implied by condition (iv). This may not be the case in the current scenario, as B may be empty (when no constant-free 1-types are present in A).

To show that for any normal form φ satisfiable in a GFU-model appropriate sets of types A and B exist we observe that there is a GFU-model of φ in which every constant-free 1-type is realized at least two times. To see this one can adapt the notion of the pullback (cf. Def. 2.1) to the case with constants by requiring that for every constant c the function f maps precisely one element of M to $c^{\mathfrak{N}}$. This unique element is then required to be the interpretation of c in $\mathfrak{M}$. After this adjustment, Remark 2.3 becomes true for the extended language and can be naturally used to duplicate a realization of every constant-free type in any model of φ. One easily verifies that the sets A, B of, resp., 1-types and 2-types realized in the model obtained this way satisfy the desired conditions (i)-(vii). We emphasise that this is the fragment of the proof which would not work for arbitrary equalities (just recall the formula enforcing the existence of a king from the previous subsection); it however works for equalities $x = c$.

In the opposite direction, having A and B satisfying the conditions (i)–(vii) we can seamlessly build a GFU-model $\mathfrak{M}$ of φ. We start from the model $\mathfrak{M}_0$ whose existence is postulated in (vii) and follow the lines of the proof of Lemma 3.7, with the exception that in the inductive step, as the interpretations of the constant symbols in $\mathfrak{M}^{a,b}$ we use the same elements that interpret them in $\mathfrak{M}_{n-1}$ (and hence, by induction, the same as in $\mathfrak{M}_0$), that is, that the intersection $M^{a,b} \cap M_{n-1}$ consists of a, b and the interpretations of constants from $\mathfrak{M}_{n-1}$. The structure $\mathfrak{M}_n^*$ can be then defined without conflicts because $\mathfrak{M}_{n-1}$ and each of the $\mathfrak{M}^{a,b}$ agree on constants as in all of them the 0-type γ_0 is realized.

A natural decision procedure arising from our considerations above, given an input φ, again transforms φ into normal form χ, guesses the sets A and B and verifies the required conditions. In the presence of constants the number of possible 1-types may be larger than without them, namely doubly exponential in the maximal arity of the relation symbols and hence in the size of φ. Thus, the size of the guess can be also bounded only doubly exponentially, and the procedure indeed needs nondeterministic doubly exponential time. All the required conditions except (iv) and (vii) are directly verifiable. To verify conditions (iv) and (vii) we construct GF-formulas analogous to those constructed for condition (iv) in the case without constants and then pass them to an external procedure solving their satisfiability.

The formulas constructed are this time of size doubly exponential in the size of χ, since one of their conjuncts is a big disjunction listing all the 1-types from A. However, the other complexity parameters we use in our analysis, that is the number of variables, the maximal arity of relation symbols and the number of relation symbols are linearly bounded in the size of χ. Thus, if as the external procedure we take Grädel's procedure from [7], and we repeat the complexity analysis from the proof of Thm. 3.9, we get that the procedure will return its answer in deterministic doubly exponential time in the size of χ. (We remark that in Grädel's proof constants are first eliminated and simulated by variables which results in increasing the arity of every relation symbol by the number of constants; still this is linearly bounded in the size of the input formula). Hence, overall, our procedure works in 2-NExpTime, and the only source of nondeterminism is the need of guessing the sets A and B of potentially large size.

We comment that in our approach the equalities of the form $x = c$ do not need any special care. Simply, we assume that the models whose existence is postulated by conditions (iv) and (vii) respect such equalities. They then just need to be properly handled by the external procedure; and we recall that Grädel's procedure admits equalities, in particular those of the form $x = c$.

5.3 Two-variable logic with guarded use of transitive closure

In [14] Michaliszyn proves the decidability and 2-ExpTime-completeness of the two-variable restriction of the guarded fragment in which to some distinguished binary relations one can apply the transitive closure operator. Analogously to transitive relations in our paper, those distinguished relations, as well as their transitive closures (call them both *special relations*), can be used only as guards. In Michaliszyn's variant one can simulate transitive relations (for a distinguished symbol T one can simply use only its transitive closure T^* and never mention T) but, of course, the main advantage of his logic is that it allows to express some reachability properties. Employing our approach, that is using Michaliszyn's procedure as a black box, we can lift his result to full two-variable logic (which is the same as *the two-variable triguarded fragment*) without equality with special relations appearing only as guards of quantifiers. Details of the proof of this fact are similar to those from the case of TGF+TG and we skip them here. We remark that [14] deals only with the two-variable restriction of GF, so we cannot easily derive the decidability of full TGF with transitive closure in guards.

6 Conclusion

We have proposed a new approach to prove the decidability and establish the complexity of some triguarded logics, consisting in using some existing results on the guarded logics (almost) as black boxes. We demonstrated usefulness of our method by reproving some known facts, and deriving some new results, most important of which is the decidability of TGF+TG, the triguarded fragment with transitive relations in guards.

We remark that our technique works only for the general satisfiability problem and gives no insight into the *finite* satisfiability problem, which asks for the existence, for a given formula, of its finite model. An interesting related open question is if TGF has the *finite model property*, *i.e.*, if its every satisfiable formula has a finite model. Obviously, this is certainly not the case for TGF+TG, since already GF+TG without equality does not have the finite model property: just recall the typical infinite axiom, with transitive T, $(\forall x)(\exists y)T(x,y) \wedge \neg(\exists x)T(x,x)$.

References

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

[2] P. Bourhis, M. Morak, and A. Pieris. Making cross products and guarded ontology languages compatible. In *IJCAI 2017*, pages 880–886, 2017.

[3] Andrea Calì, Georg Gottlob, and Michael Kifer. Taming the infinite chase: Query answering under expressive relational constraints. *J. Artif. Intell. Res.*, 48:115–174, 2013.

[4] A. K. Chandra, D. Kozen, and L. J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981.

[5] H. Ganzinger, Ch. Meyer, and M. Veanes. The two-variable guarded fragment with transitive relations. In *LICS*, pages 24–34, 1999.

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

[7] E. Grädel. On the restraining power of guards. *J. Symb. Log.*, 64(4):1719–1742, 1999.

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

[9] E. Grädel, M. Otto, and E. Rosen. Undecidability results on two-variable logics. *Archiv für Mathematische Logik und Grundlagenforschung*, 38(4-5):313–354, 1999.

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

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

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

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

[14] J. Michaliszyn. Decidability of the guarded fragment with the transitive closure. In *ICALP (2)*, volume 5556 of *LNCS*, pages 261–272. Springer, 2009.

[15] Sebastian Rudolph and Mantas Šimkus. The triguarded fragment of first-order logic. In *LPAR*, volume 57 of *EPiC Series in Computing*, pages 604–619, 2018.

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

[17] W. Szwast and L. Tendera. The guarded fragment with transitive guards. *Annals of Pure and Applied Logic*, 128:227–276, 2004.

[18] S. Tobies. *Complexity results and practical algorithms for logics in knowledge representation*. PhD thesis, RWTH-Aachen, Germany, 2001.