Decidability of the Guarded Fragment with the Transitive Closure*

Jakub Michaliszyn**

Institute of Computer Science,
University Of Wroclaw,
ul. Joliot-Curie 15, 50-383 Wroclaw, Poland

Abstract. We consider an extension of the guarded fragment in which one can guard quantifiers using the transitive closure of some binary relations. The obtained logic captures the guarded fragment with transitive guards, and in fact extends its expressive power non-trivially, preserving the complexity: we prove that its satisfiability problem is 2EXPTIME-complete.

1 Introduction

The guarded fragment of first-order logic, GF, introduced in [1], is a well-known generalisation of modal logics. The main idea is to allow only a restricted form of quantification, simulating local nature of the modal operators $\Diamond$, $\Box$. GF retains a lot of nice properties of modal logics, including (a generalisation of) the tree model property, the finite modal property and the robust decidability. This makes it a promising starting point for logics for reasoning about programs and hardware.

Many extensions and variants of GF have been extensively investigated last years. One important direction is considering satisfiability of GF over restricted classes of structures, in which some distinguished binary symbols are interpreted as transitive relations (or, alternatively phrased, satisfiability of GF extended by positive statements about transitivity of some binary relations). It appeared ([7], [3]) that allowing transitive symbols to appear in arbitrary positions in formulas leads quickly to undecidability. However, if we restrict the usage of transitive symbols to guards only, the satisfiability problem becomes decidable and 2EXPTIME-complete ([11], [9]). The lower bound can be proved even in the presence of only two variables. This two-variable version, $[\mathrm{GF}^2 + \mathrm{TG}]$, captures (and non-trivially extends) modal logics K4, S4 and S5.

Transitivity of some binary relations is a desirable property in many reasoning tasks. However, to reason about programs it would be nice to have some way of expressing recursion. One idea is to extend GF by least and greatest fixed point operators. It was done in [5]. This way we obtain a powerful logic embedding

* Part of the M.Sc. thesis, under the supervision of Emanuel Kieroński

** Partially supported by Polish Ministry of Science and Higher Education research project N206 022 31/3660, 2006/2009.

modal μ-calculus with backward modalities. Another idea is to add to GF some form of the transitive closure operator. In this paper we consider [GF^2+$^+$], an extension of the two-variable guarded fragment, in which the transitive closure operator can be applied to some atomic formulas. We note that augmenting fragments of first-order logic even by a week form of the transitive closure leads quickly to undecidability (see, e. g., [6]). Thus we have to be careful: analogously to the variant with transitive relations we restrict the usage of the transitive closure operator $^+$ to symbols appearing only in guards. More formally, the signature has a distiguished subset Σ_+, containing binary symbols, which may be used only in guards, either individually or under the transitive closure operator. We note that a variant, in which symbols from Σ_+ are used additionally outside guards, but not under the transitive closure, is undecidable.

A tempting idea is to go towards an extension of GF which would be strong enough to embed propositional dynamic logic, PDL [2]. However, this is not easy. For example, a variant allowing generalised guards of the form $T \odot S(x, y)$, simulating the composition of actions, is undecidable, even if the symbols which can be used in compositions are allowed only in guards, and there is no transitive closure operator [8].

[GF^2+$^+$] easily simulates [$\mathrm{GF}^2 + \mathrm{TG}$]: in a formula of [$\mathrm{GF}^2 + \mathrm{TG}$] instead of a transitive relation T we can simply use a transitive closure of a relation T'. However [GF^2+$^+$] is strictly more expressive than [$\mathrm{GF}^2 + \mathrm{TG}$]. For example, consider the formula $\exists x(S(x) \wedge \exists y(xT^+y \wedge R(y)))$, stating that from some point in S there exists a T-path to a point in R. This is clearly not a first-order property and thus it cannot be expressed in [$\mathrm{GF}^2 + \mathrm{TG}$].

We prove that the satisfiability problem for [GF^2+$^+$] is decidable in 2EXPTIME, exactly as [$\mathrm{GF}^2 + \mathrm{TG}$]. Similarly to the case of [$\mathrm{GF}^2 + \mathrm{TG}$] the proof is based on a tree-like model property. However, there are some serious complications, due to the fact that [GF^2+$^+$] can speak both about direct successors and about reachable elements. For example, in contrast to [$\mathrm{GF}^2 + \mathrm{TG}$], for a given element its witnesses cannot always be its direct successors.

The paper is organised as follows. In Section 2 we give some basic definitions and introduce a normal form of formulas. In Section 3 we prove a useful result on the two-variable logic FO^2, which will be an important tool in our proof. In Section 4 we show that every satisfiable [GF^2+$^+$] formula in normal form has a model of a special, tree-like shape. In Section 5 we outline an alternating algorithm checking the existence of such a special model for a given normal form sentence.

2 Preliminary

2.1 Logics

We work on First Order Logic (FO) with purely relational signatures, containing no constants and functional symbols. Let GF stand for the Guarded Fragment of First Order Logic defined as follows.

- every atomic formula belongs to the language of GF
- GF is closed under boolean operators
- if ψ belongs to the language of GF, $\boldsymbol{x}, \boldsymbol{y}$ are vectors of variables and $\alpha(\boldsymbol{x}, \boldsymbol{y})$ is an atomic formula that contains all variables from $\boldsymbol{x}$ and $\boldsymbol{y}$, then formulas $\forall \boldsymbol{x}.\alpha(\boldsymbol{x}, \boldsymbol{y}) \Rightarrow \psi(\boldsymbol{x}, \boldsymbol{y})$ and $\exists \boldsymbol{y}.\alpha(\boldsymbol{x}, \boldsymbol{y}) \wedge \psi(\boldsymbol{x}, \boldsymbol{y})$ belong to the language of GF

Formulas $\alpha(\boldsymbol{x}, \boldsymbol{y})$ are called *guards*. Note that $x = x$ is a special case of a guard.

Logic with the Transitive Closure In Guards ([GF+$^+$]) is an extension of GF in which some generalised guards are allowed. We divide the signature Σ into three disjoint parts $\Sigma = \Sigma_U \cup \Sigma_B \cup \Sigma_+$. Σ_+ and Σ_B are sets of binary symbols, where symbols from Σ_+ cannot appear outside guards, and Σ_U is a set of symbols with arity different than 2. For a given symbol $T \in \Sigma_+$ we can form a guard in a usual way, e. g. xTy, or by adding operator $^+$ to T, e. g. yT^+x. The semantics of the operator $^+$ is defined as usual: T^+ denotes the transitive closure of T.

We work with two-variable variants of the logics only. We denote them by FO^2, GF^2, $[\mathrm{GF}^2{+}^+]$. Without loss of generality we assume that signatures contain only unary and binary symbols.

2.2 Terminology

1-types and 2-types of elements in a structure over a signature Σ are defined in a standard way. The (atomic) 1-type of an element v is the set of atomic formulas with free variable x satisfied by v. Similarly, the 2-type of a pair of elements v, w is the set of atomic formulas with free variable x, y satisfied by v, w. We assume that 2-types are proper, i. e. contain $x \neq y$. For a 1-type t, let $t[x/y]$ stand for the set that contains the formulas from t in which each occurrence of x is replaced by y. For a 2-type t let $t|_{\mathcal{B}}$, which is called a restriction of t to family $\mathcal{B}$, stand for the set that contains exactly those atomic formulas from t that either have only one free variable or their relation symbol belongs to B. Similarly, $t_B = t|_{\{B\}}$ is called a restriction of t to the relation B.

To simplify the notation, we introduce for a binary symbol R a auxiliary symbol R^{-1}, whose intended meaning is to denote the inverse relation of R. Let, for a given set of binary relations or binary relation symbols $\mathcal{T}$, $\mathcal{T}^{-1}$ be the set $\{R^{-1} | R \in \mathcal{T}\}$.

For a given logic, the satisfiability problem of this logic is defined as follows: for a given formula ψ without free variables, is there any structure $\mathfrak{M}$ such that $\mathfrak{M} \models \psi$ (model $\mathfrak{M}$ satisfies ψ)?

2.3 Normal Forms

Definition 1. *We say that a formula $\psi \in \mathrm{FO}^2$ is in Scott normal form[10] if it is a conjunction of formulas in the following forms:*

(i) $\exists x.\rho(x)$

(ii) $\forall xy.\delta(x, y)$
(iii) $\forall x.\exists y.\delta(x, y)$

where both $\rho(x)$ and $\delta(x, y)$ are quantifier-free.

Every formula ψ of FO^2 can be transformed, in polynomial time, to a formula φ in Scott normal form over am extended signature, such that ψ is satisfiable if and only if φ is satisfiable. Furthermore, ψ has models of the same size as φ.

Definition 2. *We say that formula $\psi \in [\mathrm{GF}^2{+}^+]$ is in normal form if it is a conjunction of formulas in following forms:*

(i') $\exists x.\alpha(x) \wedge \rho(x)$
(ii') $\forall xy.\beta(x, y) \Rightarrow \delta(x, y)$
(iii') $\forall x.\alpha(x) \Rightarrow \exists y.\beta(x, y) \wedge \delta(x, y)$

where both $\alpha(x)$ and $\beta(x, y)$ are proper guards and neither $\rho(x)$ nor $\delta(x, y)$ contains quantifiers.

Lemma 1. *Every formula φ of $[\mathrm{GF}^2{+}^+]$ can be effectively transformed to a set of formulas Δ of $[\mathrm{GF}^2{+}^+]$ over extended signature in normal form such that*

- *φ is satisfiable if and only if $\bigvee \Delta$ is satisfiable*
- *$|\Delta| = O(2^{|\varphi|})$, $\Sigma' = O(|\varphi|)$ and for each $\psi \in \Delta$ we have $|\psi| = O(|\varphi| \log |\varphi|)$*
- *Δ can be computed in exponential time, and every $\psi \in \Delta$ can be computed in polynomial time.*

The proof of this lemma is identical to the proof of Lemma 3.2 from paper [11] about normal form for $[\mathrm{GF}^{\mathrm{k}} + \mathrm{TG}]$, because that proof does not depend on guards. Additionally we assume that a conjunct in form (i') appears exactly once in the whole formula.

We say that w is a *witness* for an element v, if for some conjunct of the form (iii) the formula $\delta(v, w)$ is satisfied or for some conjunct in the form (iii') formulas $\alpha(v)$ and $\beta(v, w) \wedge \delta(v, w)$ are satisfied.

3 Exponential Model Property for Strongly-connected FO^2

In this section we show a lemma about FO^2, which in fact is an extension of the exponential model property [4]. This lemma will then become a crucial tool in our construction.

We say that a structure $\mathfrak{M}$ is *T-strongly-connected* if the digraph obtained from $\mathfrak{M}$ by removing all edges except T is strongly-connected. Note that a structure is T-strongly-connected if and only if the transitive closure of T contains every pair of different elements of this structure.

Lemma 2. *Let φ be a sentence from FO^2 in Scott normal form over signature Σ, with a distinguished binary symbol T, and let $\mathfrak{M}$ be a T-strongly-connected model of φ. Then φ has a T-strongly-connected model $\mathfrak{M}'$ of size bounded by $2^{4|\Sigma|+5}$, such that*

- *$\mathfrak{M}'$ contains all 1-types realized in $\mathfrak{M}$*
- *for every point v from $\mathfrak{M}'$ of 1-type t_v and every 1-type t_w such that there is a 2-type t in $\mathfrak{M}$ that contains t_v, $t_w[x/y]$ there is a point w from $\mathfrak{M}'$ such that the pair v, w has the 2-type t.*

Proof. Let $\Sigma_U = \{U_1, U_2, \ldots, U_u\}$ be a set of unary relation symbols, $\Sigma_B = \{T, B_1, B_2, \ldots, B_b\}$ be a set of binary relation symbols and $\Sigma = \Sigma_B \cup \Sigma_U$. Let us fix a sentence from FO^2 in Scott normal form over a signature Σ

$$\varphi = \exists x \rho(x) \wedge \forall xy \phi(x, y) \wedge \bigwedge_{i=1}^{k} \forall x \exists y \psi_i(x, y)$$

such that $\mathfrak{M}$ is a T-strongly-connected model of φ with universe M.

We will now build a T-strongly-connected model $\mathfrak{M}'$ with universe M', where $|M'| \leq 2^{4|\Sigma|+5}$. This construction can be seen as an extension of the construction for the exponential model for FO^2 [4]. However we have to work much harder to preserve strong-connectivity of the model.

Definitions 1.

- *Kings* are these points which have a unique 1-type in the model.
- An R-path $v_1, v_2, \ldots, v_k$ is a substructure generated by pairwise different elements $v_1, v_2, \ldots, v_k$ such that for each $i < k$ we have $\mathfrak{M} \models v_i R v_{i+1}$.
- A path $s, v_1, \ldots, v_n, s'$ is *non-royal*, if none of the elements $v_1, \ldots, v_n$ is a king.
- A *shortcut* of a path $s, \ldots, v_{pi}, v_i, \ldots, v_j, v_{nj}, \ldots, s'$, where vertices v_i and v_j have the same 1-types, is the path $s, \ldots, v_{pi}, v_i, v_{nj}, \ldots, s'$, where v_i and v_{nj} are connected in the same way as v_j with v_{nj}.
- We say that a path s is a *compression* of a path s' if s has no shortcut and there exists a sequence $r_1, r_2, \ldots, r_n$ where $r_1 = s'$, $r_n = s$ and r_{i+1} is a shortcut of r_i for $1 \leq i \leq n$.

Note that if the signature is finite, then every path can be compressed to a path of length bounded exponentially in the size of the signature.

The universe M' contains the following parts: the royal palace V_k, the court V_d and three cities V_1, V_2 and V_3. Their construction proceeds as follows:

1. We insert into the royal palace V_k copies (i. e. elements of the same 1-type) of all kings from $\mathfrak{M}$. If none of the kings satisfies ρ, then we add to the royal palace one copy of a point that satisfies ρ in $\mathfrak{M}$. We preserve the connections between these elements from $\mathfrak{M}$.
2. The court contains all witnesses for kings. More precisely, for each royal 1-type t_r, non-royal type t_n and 2-type t that contains t_r and $t_n[x/y]$, and appears in $\mathfrak{M}$, we insert into the court V_d a new point of 1-type t_n and connect it with the point of type t_r from royal palace as in t.
3. We build three cities. First, for each city we add $2^{2|\Sigma|}$ copies of every non-royal point from $\mathfrak{M}$.

4. For each point v' from the court or one of the cities, we find in $\mathfrak{M}$ a T-path s_1 from a point that has the same type as v' to the point k_1 from the royal palace, and a path s_2 from the point k_2 from the royal palace to a point that has the same type as v', where both s_1 and s_2, do not contain kings except k_1 and k_2, respectively. We add the compression of the paths s_1 and s_2 to the V_1, if v was from the court, or to the city of v', if v' was from a city, by replacing extreme points with v' and the copies of k_1 and k_2. Similarly, for each pair (v', w') from the royal palace which has not been yet connected by a T-path, we find in $\mathfrak{M}$ a non-royal T-path s from the copy o v' to the copy of w' and add a copy of a compression of s to the court. Finally, for each newly-added element u' that is connected directly (i.e. by some relation, not only by the transitive closure of some relation) with a king of 1-type t_k, we find in $\mathfrak{M}$ an element u which has the same 1-type as u' and is connected with the king of the 1-type t_k in the same way as u'. We set connections between u' and V_K as between u and all kings from $\mathfrak{M}$.
5. We ensure that all non-royal points have non-royal witnesses. For each non-royal point v' from $\mathfrak{M}'$ of 1-type t_s, each non-royal 1-type t_r and each 2-type t from $\mathfrak{M}$ which contains t_s and $t_r[x/y]$, we copy connections from t to $\mathfrak{M}'$ in the following way. If v' was in V_d or V_3, then we connect it with points from V_1, if it was in V_1 then with V_2, and if it was from V_2 then with V_3. Each city contains $2^{2|\Sigma|}$ elements of 1-type t_r, so for each 2-type we can choose a different element.
6. We ensure that all non-royal points have the requested witnesses among kings: for each non-royal element w' from $\mathfrak{M}'$, if a connection between w' and the royal palace is not set yet, then we find in $\mathfrak{M}$ an element of the same 1-type and copy connections between this point and kings to $\mathfrak{M}'$. Note that we always can find such a pair, because all points in $\mathfrak{M}'$ are copied from $\mathfrak{M}'$ and, moreover, if some 1-type appears in $\mathfrak{M}$ only once, then it appears also only once in V_K and it does not appear in the court or in any city.
7. For each pair v', w' of points from $\mathfrak{M}'$, if the connection between v' and w' is not already set, we copy some connection beetwen points of the same 1-types from $\mathfrak{M}$. Again, a proper connection can be found because of special treatment of kings (as in step (6)).

Note that $|V_k| \leq 2^{|\Sigma|}$, $|V_d| \leq 2^{2|\Sigma|}|V_k| + 2^{|\Sigma|}|V_k|^2$ and $|V_i| \leq (1+2^{|\Sigma|}) \cdot 2^{2|\Sigma|} + 2^{|\Sigma|}|V_d|$ for $i \in \{1, 2, 3\}$, so $|M| \leq 2^{4|\Sigma|+5}$. Let us observe that:

- The formula $\exists x \rho(x)$ is satisfied in $\mathfrak{M}'$, because a point that satisfies this formula was added in step 1.
- All 1-types and 2-types from $\mathfrak{M}'$ appear also in $\mathfrak{M}$; thus, ϕ is satisfied.
- Every point has all witnesses it needs, so each ψ_j is satisfied in $\mathfrak{M}'$.
- The royal palace is a T-strongly-connected subgraph (because of step 4) and moreover each courtier and citizen is on a T-path from the royal palace to the royal palace because of paths added in step 4.

Therefore, $\mathfrak{M}'$ satisfies φ and is T-strongly-connected. The construction also implies that in $\mathfrak{M}'$ there appear all 1-types from $\mathfrak{M}$ are realized. $\square$

Theorem 1. *Let φ be a an FO^2 formula over a signature Σ with a distinguished symbol T, and let $\mathfrak{M}$ be a T-strongly-connected model of this formula. Then φ has a T-strongly-connected model of size exponential in $|\Sigma|$.*

This Theorem is a straightforward consequence of Lemma 2 and the observation about Scott normal form.

4 Ramified model property for $[\mathrm{GF}^2{+}^{+}]$

The proof of a special model property of $[\mathrm{GF}^2{+}^{+}]$ has the common outline with the idea presented in [11] for $[\mathrm{GF}^2 + \mathrm{TG}]$. In both proofs we at first prove that each satisfiable formula has a model of a tree-like shape. This is done in the following way:

1. Take a satisfiable formula in normal form and a model $\mathfrak{M}$ of this formula.
2. Take one point from the model $\mathfrak{M}$ and start building the new model $\mathfrak{M}'$ from this point.
3. If this point is in some clique whose edges are defined by a one of transitive relation (a transitive closure of a relation from Σ_+), take this clique, compress it and add it to $\mathfrak{M}'$.
4. If the current point needs a witness outside its cliques, find a path to this witness in $\mathfrak{M}$, compress it and add it $\mathfrak{M}'$.
5. Process recursively the newly-added points (as in points (3) - (5)).

The most important differences between the proofs are in points (3) and (4). The reason why the part (3) of the proof is more difficult in the case of $[\mathrm{GF}^2{+}^{+}]$ is that the connections defining cliques are not atomic relation as in $[\mathrm{GF}^2 + \mathrm{TG}]$, but paths. To overcome this difficulty we use the result of Section 3. The part (4) is more complicated because in $[\mathrm{GF}^2{+}^{+}]$ some witnesses cannot allways be direct successors of an element.

It is important to underline why we take care about cliques. Our point is to construct a model which looks like a tree. However, in $[\mathrm{GF}^2{+}^{+}]$ logic we can write a sentence ψ which is a conjunction of following formulas:

– $\forall x(S_0(x) \Rightarrow \forall y(xR^+y \Rightarrow S_0(y) \Rightarrow x = y))$ (if a point v satisfies S_0, then there is no point reachable from v by relation R that satisfies S_0, except, possibly, v)
– every point satisfies exactly one of the relations $S_0, \ldots, S_{n-1}$
– there exists a point that satisfies S_0
– $\forall x.S_i(x) \Rightarrow \exists y(xRy \wedge S_{(i+1) \mod n}(y))$ for each $i < n$

It is easy to see that every model of ψ contains a cycle of length n. In fact, we can obtain a cycle of length 2^n, using $S_0, \ldots, S_{n-1}$ to encode a binary number and request that every successor of a point v encodes the value greater by 1 modulo 2^n. As we see, the structure must sometimes have fragments that do not look like a tree. We will see that the size of each such fragment can be bounded exponentially.

4.1 Construction of a ramified model for [GF^2+$^+$]

Let $\varphi = \exists x \rho(x) \wedge \bigwedge_{i=1}^{j} \forall xy \delta_i(x,y) \wedge \bigwedge_{i=1}^{k} \forall x.\alpha_i(x) \Rightarrow \exists y.\psi_i(x,y)$ be a fixed formula in normal form from [GF^2+$^+$] over the signature $\Sigma = \Sigma_U \cup \Sigma_B \cup \Sigma_+$, where $\Sigma_+ = \{T_1, \ldots, T_n\}$.

First, we define some properties of connections:

Definitions 2.
- Let t be a 2-type. We say that t is *k-positive* if there are exactly k different relations $R \in \Sigma_+$ such that t satisfies $xRy \vee yRx$.
- We say that a structure is *1-positive*, if every 2-type that appears in this structure is either 1-positive or 0-positive.
- An *extended 1-type* of a point v in a structure $\mathfrak{A}$ is the set that contains the 1-type of v and the pairs $\langle R, t\rangle$, where $R \in \Sigma_+ \cup \Sigma_+^{-1}$ and for some $w \neq v$ of 1-type t formula vR^+w is satisfied in a structure $\mathfrak{A}$.

Note that a restriction of a 2-type t to a binary relation B, $t|_B$, is at most 1-positive. Now we define cliques and some operations on them.

Definitions 3.
- An R^+-clique in a structure $\mathfrak{M}$ is a subset K of elements from $\mathfrak{M}$, such that for each $v, w \in K$ we have $\mathfrak{M} \models vR^+w$.
- The maximal R^+-clique that contains v is denoted by R^+-*clique*(v).
- A *path of R^+-cliques* $\langle C_1, v_1^{in}, v_1^{out}\rangle$, $\langle C_2, v_2^{in}, v_2^{out}\rangle$, $\ldots$, $\langle C_k, v_k^{in}, v_k^{out}\rangle$ is a substructure generated by pairwise disjoint cliques $C_1, C_2, \ldots, C_k$, where for each $1 \leq i \leq k$ the clique C_i is the maximal R^+-clique containing distinguished vertices v_i^{in} and v_i^{out}, and for $i < k$, $M \models v_i^{out} R v_{i+1}^{in}$.
- We say that a vertex v is in the clique-distance m from w if the shortest (i. e. of the minimal length) path of R^+-cliques from R^+-clique(v) to R^+-clique(w) has the length m.
- A path of R^+-cliques $\langle C_1, v_1^{in}, v_1^{out}\rangle, \ldots, \langle C_z, v_z^{in}, v_z^{out}\rangle, \langle C_d, v_d^{in}, v_d^{out}\rangle, \ldots$ is a *shortcut* of a path of cliques $\langle C_1, v_1^{in}, v_1^{out}\rangle, \ldots, \langle C_z, v_z^{in}, v_z^{out}\rangle, \langle C_p, v_p^{in}, v_p^{out}\rangle, \ldots, \langle C_d, v_d^{in}, v_d^{out}\rangle, \ldots$ if v_p^{in} and v_d^{in} has the same 1-types.
- A *compression of a path of R^+-cliques* s is a minimal path of R^+-cliques obtained by an iterated shortcutting of s.

Note that for any $R \in \Sigma_+$ a single vertex is a R^+-clique.

Now we define an operation on a structure that is usefull to express its tree-likeness. Intuitively, a flattening of a model $\mathfrak{M}$ is a graph $G = \langle V, E\rangle$, such that V contains one vertex for each clique from $\mathfrak{M}$, and E connects cliques C_1 and C_2 if at least one of the following conditions holds:

- C_1 has a common vertex with C_2 in $\mathfrak{M}$
- C_1 is connected in $\mathfrak{M}$ with C_2 by some relation from Σ_+
- C_1 is connected in $\mathfrak{M}$ with C_2 by some relation from Σ_B and is not connected by the transitive closure of any of relation from Σ_+.

The formal definition is more complicated. Each vertex is in n cliques (one for each relation from Σ_+), so connecting cliques with a common vertex lead us to cliques of size n. We want to show that some flattening are trees, so we arbitrary choose relation T_1 and connect cliques with common vertex only if one of this cliques is T_1^+-clique.

Definition 3. *We say that an undirected graph $G = \langle V, E\rangle$ is a flattening of a model $\mathfrak{M}$, if*

- $V = \{R^+\text{-}clique(w) | R \in \Sigma_+ \wedge w \in M\}$
- $\{T_i^+\text{-}clique(w), T_j^+\text{-}clique(w')\}$ *is in* E *if* $T_i^+\text{-}clique(w) \neq T_j^+\text{-}clique(w')$ *and at least one of the following conditions holds:*
  - $w = w'$ *and* $i = 1$
  - $i = j$, *and* $\mathfrak{M} \models wT_iw'$ *holds*
  - $i = j = 1$, *for some* $S \in \Sigma_B \cup \Sigma_B^{-1}$ *we have* $\mathfrak{M} \models wSw'$ *and for each* $S \in \Sigma_+ \cup \Sigma_+^{-1}$ *condition* $\mathfrak{M} \models wS^+w'$ *is not satisfied*

We are ready to define a property of models that will be used to build an algorithm that checks if a given formula from $[\mathrm{GF}^2{+}^+]$ has any model.

Definition 4. *We say that a model $\mathfrak{M}$ of a formula φ is r-ramified, if $\mathfrak{M}$ is 1-positive, the size of each R^+-clique in $\mathfrak{M}$ for $R \in \Sigma_+$ is bounded by r and the flattening of $\mathfrak{M}$ is a tree.*

Theorem 2. *Every satisfiable sentence from* $[\mathrm{GF}^2{+}^+]$ *over a signature Σ has a r-ramified model for $r = 2^{4|\Sigma|+5}$, in which every point has all the required witnesses in clique-distance not greater than $2^{|\Sigma|}$.*

Proof. Let $\varphi = \exists x \rho(x) \wedge \bigwedge_{i=1}^{j} \forall xy \delta_i(x, y) \wedge \bigwedge_{i=1}^{k} \forall x.\alpha_i(x) \Rightarrow \exists y.\psi_i(x, y)$ be a fixed formula from $[\mathrm{GF}^2{+}^+]$ in normal form over the signature $\Sigma = \Sigma_U \cup \Sigma_B \cup \Sigma_+$, where $\Sigma_+ = \{T_1, \ldots, T_n\}$. Furthermore, let $\mathfrak{M}$ be a model of φ. Now we define a recursive procedure that, for a given $2^{4|\Sigma|+5}$-ramified structure $\mathfrak{M}'$, point $v \in \mathfrak{M}'$ and function $from : \mathfrak{M}' \to \mathfrak{M}$, extends $\mathfrak{M}'$ to a $2^{4|\Sigma|+5}$-ramified structure where v and every newly-added vertex have all needed witnesses. Simultaneously, it extends the function $from$. Intuitively, $from(v)$ indicates a point "similar" to w.

1. Build the cliques of v, modifying the cliques of $from(v)$, proceed as follows. For each relation $R \in \Sigma_+$ if v is not inside an R^+-clique with size greater then 1, then we take from $\mathfrak{M}$ the R^+-clique($from(v)$) and we restrict all 2-types in this clique to the family of relations $\{R\} \cup \Sigma_B$. Such a structure is an R-strongly-connected component. If this component has more then 1 vertex, then, using Lemma 2 for formula $(\exists x \top) \wedge \bigwedge_{i=1}^{j} \forall xy \delta'_i(x, y)$ and this component, where δ'_i is obtained from δ_i by replacing all guards xR^+y by $\top$, we transform this component to a structure H' with an exponential size, such that in H' all witnesses are preserved. Then we choose from H' an element w with the same 1-type as v and add H' to $\mathfrak{M}'$ by identifying v with w. Finally, for each element $u' \neq w$ from H' we find in H an element u with the same 1-type and we set $from(u') = u$.
2. We ensure that v has required witnesses outside its cliques. In order to do that, for each formula $\psi_i = \gamma_i \wedge \delta_i$, where γ_i is a guard, vertex v satisfies α_i and ψ_i is not satisfied in $\mathfrak{M}'$ yet, we choose from $\mathfrak{M}$ a point w' that is a witness of this formula for v'. Observe that w' is not in any clique with v'.

(a) If $\gamma_i = xRy$ or $\gamma_i = yRx$ for some $R \in \Sigma_B \cup \Sigma_+$, then we add a copy w of a point w' to the model $\mathfrak{M}'$ and we set connections between v and w in the same way as it was in $\mathfrak{M}$ after restriction to R. Moreover, we put $from(w) = w'$.

(b) If $\gamma_i = xR^+y$ or $\gamma_i = yR^+x$ for some $R \in \Sigma_+$, then we choose from $\mathfrak{M}$ a full path that provides the fulfilment of γ_i, with all R^+-cliques that appear on this path. We do it in following way.
Assume that $\gamma_i = xR^+y$. We take a path of R^+-cliques from R^+-clique(v') to R^+-clique(w') with minimal length, we restrict every 2-types on this path to R, and then we compress this path, obtaining a path $\langle R^+\text{-clique}(v'), v', v'^{out}\rangle$, $\langle C_1, v_1^{in}, v_1^{out}\rangle$, $\ldots$, $\langle C_k, v_k^{in}, v_k^{out}\rangle$, where $C_k = R^+$-clique(w) and $v_k^{out} = w'$.
We compress every clique C_i to C'_i in the way presented in step 1, obtaining cliques in which points v'^{in}_i and v'^{out}_i have the same 1-types as v_i^{in} and v_i^{out}, respectively. Then, in the R^+-clique(v), we find a vertex v^{out} with the same 1-type as v'^{out}. We add to $\mathfrak{M}'$ path $\langle C'_1, v'^{in}_1, v'^{out}_1\rangle$, $\ldots$, $\langle C'_k, v'^{in}_k, v'^{out}_k\rangle$ and we connect v^{out} and v'^{in}_1 in the same way as v'^{out} and v'^{in}_1 were connected.
For each i we put $from(v'^{in}_i) = v_i^{in}$ and $from(v'^{out}_i) = v_i^{out}$. Moreover, for each vertex $u' \in C'_i$ such that $from(u')$ is not set yet, we find in C_i a vertex u with the same 1-type and put $from(u') = u$.
When $\gamma_i = yR^+x$, we do the same for R^{-1}.

3. We connect vertices from $\mathfrak{M}'$ by relations from Σ_B using some patterns from $\mathfrak{M}$. More precisely, for each two vertices v, w:

(a) If for some $R \in \Sigma_+$ condition $vR^+w \wedge wR^+v$ occurs, then these vertices are in the same clique and the connections are already established.

(b) If these vertices are connected by R^+ for some $R \in \Sigma_+$, the connection is asymmetric (without loss of generality we may assume that $vR^+w \wedge \neg wR^+v$ holds) and connections between these points were not established yet, then we find in $\mathfrak{M}$ a vertex w' which has the following property: in $\mathfrak{M}$ there is a R-path from $from(v)$ to w' and w' has the same 1-type as w. Such a vertex exists, because the 1-type of w belongs to the set of the 1-types that are reachable by relation R, written in extended 1-type of $from(v)$. We add connections from Σ_B between v and w in the same way as $from(v)$ and w' were connected.

(c) If these two points are not connected by the transitive closure of some relation from Σ_+, then either all connections between this points were already set, or they are not connected at all – then we set empty connection between this elements.

4. We repeat steps 1-4 for all vertices added in this stage.

We take from $\mathfrak{M}$ a point v', that satisfies ρ, add its copy v to $\mathfrak{M}'$, set $from(v) = v'$ and apply the procedure above to v. The structure built by this procedure is a model of φ. The proof is omitted in this version due page limit. $\square$

5 Algorithm

In this section we describe an alternating algorithm working in exponential space that checks if a given formula has a ramified model. From Subsection 4.1 we know that every satisfiable formula in $[\mathrm{GF}^2\!+^+]$ has a ramified model, so this algorithm resolves the satisfiability problem. At first, let us introduce some definitions:

- 1-types of elements are defined as above.
- full 1-type of an element contains the following information: 1-type of element, list of 1-types of all direct successors for each binary relation, and, for each relation R which appears under the transitive closure in the formula, a list of 1-types of vertices reachable by R, and the information about 1-types of vertices reachable by R^{-1}, except for vertices that are in R^+-clique with considered element.
- type of a clique, containing the following information: its size, full 1-types of all the vertices in the clique, information about connections between the vertices and function *promise*, which for a given 1-type t and $b \in \{-1, 1\}$ returns the length of the path of $(R^b)^+$-cliques from the current clique to a clique that contains a vertex of type t (or 0, if there is no such path).

Note that every two cliques of the same type are isomorphic. For the sake of simplification, in this section we look upon a type of a clique as a clique of this type with the attached function *promise*. This function is needed because the algorithm guesses in each stage only cliques from a direct neighbourhood, while Theorem 2 guarantees only that all needed witnesses are in clique-distance not greater than exponential in the size of the signature.

Step 1. For a given formula $\varphi = \exists x \rho(x) \wedge \bigwedge_{i=1}^{j} \forall xy \delta_i(x, y) \wedge \bigwedge_{i=1}^{k} \forall x.\alpha_i(x) \Rightarrow \exists y.\psi_i(x, y)$, the algorithm starts from guessing (i.e. existentially choosing) type of a starting clique K containing an element satisfying ρ. Then it checks if local properties of K are correct: if connections between vertices in the clique satisfy δ_i for each i, if all successors from the full 1-type of points in the clique can be connected with points from the clique in a way that satisfies each δ_i, and if the *promise* function and full 1-types of vertices are not inconsistent.
Step 2. The algorithm finds direct witnesses for each vertex in the clique by guessing types of the cliques containing witnesses and connections between guessed points and K. The algorithm checks if new cliques are locally proper. Then, for each R^+-clique K', which is connected with the previous clique by R, the algorithm checks if 1-types of every vertex from K' are included in full 1-types of vertices from K to make sure that these vertices can be connected by binary relations in a way that satisfy each δ_i. Furthermore, the algorithm checks if the sets of 1-types reachable by R from vertices in K' are subsets of the sets reachable by R from vertices in K and vice versa. Then the algorithm checks if the type of K (including the *promise* function) and guessed cliques guarantee all witnesses for each vertex from K.
Step 3. The algorithm guesses types of cliques that are on paths of cliques to some clique that contains points of 1-types guaranteed in the *promise* function

from K. For the easier control of dependencies, the algorithm guess on this stage only the first clique from this path and check if the guessed clique has less value of the *promise* function or realizes this 1-type.

Step 4. The algorithm checks the counter of stages. If the value of the counter is greater then $2^{8|\Sigma|}$, then algorithm stops and return "Yes", because then we know that some type of cliques occurred twice and another computation would be the same as previously. In the other case the algorithm increments the counter, universally chooses a clique K from the set of cliques added in this stage and goes to step 2.

The algorithm needs only exponential memory, because each type of clique has at most exponential size, so, since 2EXPTIME = AEXPSPACE, the satisfiability problem can be solved in 2EXPTIME.

The lower bound follows from the 2EXPTIME-hardness of the satisfiability problem for logic $[\mathrm{GF}^2 + \mathrm{TG}]$, presented in [9], since we can simply replace transitive relations by transitive closure of these relations.

Corollary 1. *The satisfiability problem for* $[\mathrm{GF}^2{+}^{+}]$ *is* 2EXPTIME*-complete.*

References

1. H. Andreka, I. Nemeti, J. van Benthem. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217-274, 1998.
2. M. J. Fischer, R. E. Ladner. Propositional modal logic of programs. In *Proceedings of the Ninth Annual ACM Symposium on theory of Computing (Boulder, Colorado, United States, May 04 - 04, 1977). STOC '77. ACM*, New York, NY, 286-294, 1977.
3. H. Ganzinger, C. Meyer, and M. Veanes. The two-variable guarded fragment with transitive relations. In *Proc. LICS'99. IEEE Computer Soc. Press*, 1999.
4. E. Grädel, P. G. Kolaitis, and M. Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, Mar, 1997.
5. E. Grädel, I. Walukiewicz. Guarded Fixed Point Logic. In *Proceedings of 14th IEEE Symposium on Logic in Computer Science LICS '99*, Trento, 1999.
6. E. Grädel, M. Otto, E. Rosen. Undecidability results on two-variable logics. *Lecture Notes in Computer Science Nr. 1200, Springer*, pp. 249–260, 1997.
7. E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64(4):1719-1742, 1999.
8. Y. Kazakov. Saturation-Based Decision Procedures for Extensions of the Guarded Fragment. PhD thesis, Universität des Saarlandes, Saarbrücken, March 2006.
9. E. Kieroński, The two-variable guarded fragment with transitive guards is 2EXPTIME - Hard. *Lecture Notes in Computer Science*, vol. 2620, pp. 299-312, 2003.
10. D. Scott. A decision method for validity of sentences in two variables. *J. Symb. Logic* 27, 477, 1962.
11. W. Szwast, L. Tendera. The guarded fragment with transitive guards. *Annals of Pure and Applied Logic*, 128, 227-276, 2004.