




Uniform Guarded Fragments

Reijo Jaakkola(✉) 

Tampere University, Tampere, Finland

`reijo.jaakkola@tuni.fi`

`https://reijojaakkola.github.io`

Abstract. In this paper we prove that the uniform one-dimensional guarded fragment, which is a natural polyadic generalization of guarded two-variable logic, has the Craig interpolation property. We will also prove that the satisfiability problem of uniform guarded fragment is NEXPTIME-complete.

Keywords: Guarded fragment · Interpolation · Satisfiability problem.

1 Introduction

The guarded fragment \mathcal{GF} is a well studied fragment of first-order logic \mathcal{FO} , which was introduced by Andréka, van Benthem and Némethi [1] as a generalization of modal logic. Informally speaking, \mathcal{GF} is obtained from \mathcal{FO} by requiring that all quantification must be relativised by \mathcal{FO} -atoms, which is motivated by the observation that “quantification” in modal logics is relativised by accessibility relations. Like modal logic, \mathcal{GF} behaves well both computationally and model-theoretically. In particular, it is decidable, it has a (generalized) tree-model property and it satisfies various preservation theorems [1,7].

We say that a logic \mathcal{L} has Craig interpolation property (CIP), if for every two formulas φ and ψ of \mathcal{L} we have that if $\varphi \models \psi$, then there exists a third formula — the interpolant — χ of \mathcal{L} , so that $\varphi \models \chi$, $\chi \models \psi$ and χ contains only relation symbols which occur in both φ and ψ . CIP is widely regarded as a property that a “nice” logic should have and for (reasonable logics with compactness) it implies several other desirable model-theoretic properties such as Projective Beth Definability and Robinson’s consistency theorem [1,4,13,19].

It is well-known that various modal logics have CIP [1,6,19], while \mathcal{GF} fails to have it [11]. This is somewhat surprising, given that \mathcal{GF} is a very natural generalisation of modal logic, and certainly raises the question of how the syntax of \mathcal{GF} should be modified so as to obtain a logic which does have CIP, and which also behaves well both computationally and model-theoretically. One option would be to extend further the expressive power of \mathcal{GF} , and in this direction we have the *guarded negation fragment*, which has CIP, is decidable and shares with \mathcal{GF} various desirable model-theoretic properties [2].

The other option (and the one which is more relevant for this paper) is to investigate fragments of \mathcal{GF} . In this direction we also have a positive result,

namely that \mathcal{GF}^2 — the two-variable fragment of \mathcal{GF} — has CIP [11]. Given this result, it is natural to ask whether there exists a polyadic extension of \mathcal{GF}^2 which would also have CIP, where by a polyadic extension we mean intuitively a logic which contains \mathcal{GF}^2 and can express non-trivial properties of polyadic relations. Indeed, it seems rather unlikely that there would not be such an extension, since it is well-known that there are polyadic modal logics which have CIP [1].

In [9] the *uniform one-dimensional fragment* \mathcal{UF}_1 was introduced, which is a very natural polyadic extension of the two-variable fragment \mathcal{FO}^2 of \mathcal{FO} . Roughly speaking, \mathcal{UF}_1 is obtained from \mathcal{FO} by requiring that each maximal existential (or universal) block of quantifiers leaves at most one variable free and that when forming boolean combinations of formulas with more than one free variable, the formulas need to have exactly the same set of variables. Formulas satisfying the first restriction are called *one-dimensional*, while formulas satisfying the second restriction are called *uniform*. In [16] it was proved that \mathcal{UF}_1 has the finite model property and the complexity of its satisfiability problem is NEXPTIME-complete, which is the same as for \mathcal{FO}^2 [8]. The research around \mathcal{UF}_1 and its variants has been quite active, see for instance [12,14,15,17,18].

Given that \mathcal{UF}_1 is a polyadic extension of \mathcal{FO}^2 , the guarded \mathcal{UF}_1 is a natural candidate for being a polyadic extension of \mathcal{GF}^2 with CIP. As the first main result of this paper we will prove that guarded \mathcal{UF}_1 does, in fact, have CIP. Our proof follows closely the argument given in [11] for proving that \mathcal{GF}^2 has CIP, the main technical difference being that the proof presented in [11] uses crucially the fact that in the case of \mathcal{GF}^2 we can assume live sets to have size at most two, while in our case we have to deal with live sets of arbitrary size.

Since the research around modal-like fragments of \mathcal{FO} is largely motivated by the fact that their satisfiability problems are often decidable, it is natural to also study the complexity of the satisfiability problem of the guarded \mathcal{UF}_1 , which was in fact already done in [15]. More precisely it was proved in [15] that the satisfiability problem of one-dimensional \mathcal{GF} is in NEXPTIME, while it is already NEXPTIME-hard for guarded \mathcal{UF}_1 . These results left open the problem of determining the complexity of uniform \mathcal{GF} and as the second main result of this paper we will prove that the satisfiability problem of uniform \mathcal{GF} is also in NEXPTIME (and hence it is NEXPTIME-complete).

We also emphasize that as a necessary by-product of this second technical result, we isolate the uniformity restriction imposed to formulas of \mathcal{UF}_1 as an independent syntactical restriction and provide a formal definition for it (which so far has been missing from the literature).¹ We believe that uniformity is an important and a natural syntactical restriction (at least) in the context of fragments of \mathcal{FO} . Indeed, in addition to \mathcal{UF}_1 there are several known decidable fragments of \mathcal{FO} which satisfy this restriction up to some degree, such as the one-binding fragments introduced in [20] and the ordered logic introduced in [10]. We hope that the results presented in this paper provide further motivation for the study of various uniform fragments of \mathcal{FO} .

¹ To be precise, we only define what it means for a formula to be uniform in the context of \mathcal{GF} ; however, it is easy to extend this definition for other logics.

The structure of this paper is as follows. After the preliminaries in Section 2, we define a notion of bisimulation for \mathcal{UGF}_1 and establish its basic properties in Section 3. After this we will prove that \mathcal{UGF}_1 has CIP in Section 4. In Section 5 we will establish that the complexity of the satisfiability problem of uniform \mathcal{GF} is NEXPTIME-complete. The final Section will list some new problems that the research conducted in this paper raises.

2 Preliminaries

2.1 Notation

In this paper we will work with vocabularies which do not contain constants and function symbols. We will also assume that there are no relation symbols of arity 0. We will use the Fraktur capital letters to denote structures, and the corresponding Roman letters to denote their domains. Given a model \mathfrak{A} and $C \subseteq A$, we will use $\mathfrak{A} \upharpoonright C$ to denote the restriction of \mathfrak{A} to the set C . Given two structures \mathfrak{A} and \mathfrak{B} , we will use $\mathfrak{A} \leq \mathfrak{B}$ to denote that \mathfrak{A} is a substructure of \mathfrak{B} .

Occasionally we will identify tuples $\bar{a} = (a_1, \dots, a_n)$ with sets $\{a_1, \dots, a_n\}$, which allows us to use notations such as $b \in \bar{a}$ and $\bar{a} = X$, where X is a set. Given two tuples \bar{a} and \bar{b} of the same length, we will use $\bar{a} \mapsto \bar{b}$ and $p : \bar{a} \rightarrow \bar{b}$ to denote the mapping induced by the relation $a_i \mapsto b_i$. Given a tuple $\bar{a} = (a_1, \dots, a_n)$ and a unary function f , we will use $f(\bar{a})$ to denote the tuple $(f(a_1), \dots, f(a_n))$. Given a positive integer n we will denote $[n] = \{1, \dots, n\}$. Finally, if $\bar{a} = (a_1, \dots, a_n)$ and $k \geq n$ and $\mu : [k] \rightarrow [n]$ is a surjection, we will use \bar{a}_μ to denote the tuple $(a_{\mu(1)}, \dots, a_{\mu(k)})$.

2.2 Types and Tables

The following definitions are standard in the context of \mathcal{UF}_1 and were first introduced in [16]. Let σ be a vocabulary. Given a set $X = \{x_1, \dots, x_n\}$ of distinct variables and a k -ary relation $R \in \sigma$, we say that an atomic formula $R(x_{i_1}, \dots, x_{i_k})$ is an X -atom over σ , if $X = \{x_{i_1}, \dots, x_{i_k}\}$. If α is an X -atom, then α and $\neg\alpha$ are both X -literals over σ . A 1-type over σ is a maximal satisfiable set of $\{x\}$ -literals over σ . We identify 1-types π with conjunctions of their elements

$$\bigwedge \pi(x)$$

A k -table is a tuple $\langle \rho, \pi_1, \dots, \pi_k \rangle$, where each π_ℓ is a 1-type over σ , while ρ is a maximal satisfiable set of $\{x_1, \dots, x_k\}$ -literals over σ . We identify k -tables $\langle \rho, \pi_1, \dots, \pi_k \rangle$ with conjunctions

$$\bigwedge \rho(x_1, \dots, x_k) \wedge \bigwedge_{1 \leq \ell \leq k} \pi_\ell(x_\ell).$$

Let \mathfrak{A} be a σ -model. Given a 1-type π over σ , we say that $a \in A$ realizes π if π is the unique 1-type so that $\mathfrak{A} \models \pi[a]$; we denote by $\text{tp}_{\mathfrak{A}}^\sigma[a]$ the (unique) 1-type

π over σ which is realized by a in \mathfrak{A} . For *distinct* elements $a_1, \dots, a_k \in A$ we will use $\text{tp}_{\mathfrak{A}}^\sigma[a_1, \dots, a_k]$ to denote the (unique) k -table over σ which is realized by the tuple (a_1, \dots, a_k) .

2.3 Syntax of Uniform Fragments of \mathcal{GF}

Given a vocabulary σ , we define $\mathcal{GF}[\sigma]$ to be the smallest set \mathcal{F} which satisfies the following requirements.

- \mathcal{F} contains all the atomic formulas over σ , which includes also equalities between variables.
- If $\varphi, \psi \in \mathcal{F}$, then $\neg\varphi \in \mathcal{F}$ and $(\varphi \wedge \psi) \in \mathcal{F}$.
- If $\psi(\bar{x}) \in \mathcal{F}$, where each free variable of ψ occurs in the tuple \bar{x} , then

$$\exists \bar{y}(\alpha(\bar{x}) \wedge \psi(\bar{x})) \in \mathcal{F},$$

where $\bar{y} \subseteq \bar{x}$ and α is an atomic formula over σ .

If the vocabulary σ is irrelevant or known from the context, then we will simply use \mathcal{GF} to denote $\mathcal{GF}[\sigma]$.

Next we will give a formal definitions for the syntactical notions of one-dimensionality and uniformity. We will start by making the technical remark that we will define recursively the set of subformulas $\text{Sf}(\varphi)$ of $\varphi \in \mathcal{GF}$ otherwise in a standard way, except that for formulas of the form $\varphi := \exists \bar{y}(\alpha(\bar{x}) \wedge \psi(\bar{x}))$, we define $\text{Sf}(\varphi)$ to be

$$\{\exists \bar{y}(\alpha(\bar{x}) \wedge \psi(\bar{x}))\} \cup \text{Sf}((\alpha(\bar{x}) \wedge \psi(\bar{x}))).$$

In other words, we treat each maximal sequence of existential quantification as a single logical operator.

Definition 1. *Let $\varphi \in \mathcal{GF}$ be a formula. We say that φ is one-dimensional, if every subformula of φ of the form*

$$\exists \bar{y}(\alpha(\bar{x}) \wedge \psi(\bar{x}))$$

has at most one free variable. In other words each maximal sequence of (guarded) existential quantification leaves at most one variable free.

Next we will define what it means for a formula of \mathcal{GF} to be uniform. The precise definition turns out to be somewhat technical, and we will start with the following auxiliary definition.

Definition 2. *Let X be a (possibly empty) set of variables and let σ be a vocabulary. A relative X -atom over σ is a formula ψ of $\mathcal{GF}[\sigma]$ which satisfies one of the following conditions.*

1. ψ is a sentence.
2. ψ has a one free variable which belongs to X .

3. ψ is of the form $x = y$, where $x, y \in X$.
4. ψ is an X -atom over σ .
5. ψ is of the form $\exists \bar{z}(\alpha(\bar{x}) \wedge \psi(\bar{x}))$ and the set of its free variables is precisely X .

With the aid of this definition we are in a position where we can define the notion of uniformity formally.

Definition 3. Let $\varphi \in \mathcal{GF}[\sigma]$ be a formula. We say that φ is uniform, if every subformula ψ of φ is a boolean combination of relative X -atoms, where X is the set of free variables of ψ .

Remark 1. Consider a uniform quantifier-free formula $\psi(x_1, \dots, x_k)$ of $\mathcal{GF}[\sigma]$. Let \mathfrak{A} be a σ -model and let (a_1, \dots, a_k) be a tuple of not necessarily distinct elements. Then whether or not

$$\mathfrak{A} \models \psi(a_1, \dots, a_k)$$

holds depends only on the table of (c_1, \dots, c_ℓ) , where (c_1, \dots, c_ℓ) is an arbitrary enumeration of the set of distinct elements of (a_1, \dots, a_k) .

The definition of uniformity is somewhat technical, but the following examples should clarify the intuition behind it.

Example 1. Let $\sigma = \{S, R, P\}$, where S is a ternary relation symbol, R is a binary relation symbol and P is a unary relation symbol. The formula

$$\exists x \exists y (P(x) \wedge R(x, y) \wedge S(x, y, y) \wedge R(y, x) \wedge P(y))$$

is both uniform and one-dimensional. On the other hand the formula

$$\exists x \exists y (\exists z (S(x, y, z) \wedge P(z)) \wedge R(x, y) \wedge S(x, y, x))$$

is uniform but not one-dimensional. Finally, the formula

$$\exists x \exists y \exists w (R(x, y) \wedge \exists z S(x, w, z))$$

is neither one-dimensional nor uniform.

Example 2. The standard translation of polyadic modal logic into \mathcal{FO} results in formulas of the form

$$\exists x_1 \dots \exists x_k (R(x_0, x_1, \dots, x_k) \wedge \bigwedge_{1 \leq \ell \leq k} \psi_\ell(x_\ell))$$

which are uniform and one-dimensional [5].

We will use \mathcal{UGF} to denote the set of formulas of \mathcal{GF} which are uniform and \mathcal{UGF}_1 to denote the set of formulas of \mathcal{GF} which are both uniform and one-dimensional. Throughout this paper we will use $\varphi(x_1, \dots, x_n)$, where all the variables in the tuple (x_1, \dots, x_n) are distinct, to denote a formula of either \mathcal{UGF}_1 or \mathcal{UGF} such that either $\{x_1, \dots, x_n\}$ is precisely the set of free variables of φ or φ has at most one free variable which belongs to $\{x_1, \dots, x_n\}$ or φ is of the form $x_i = x_j$, where $1 \leq i, j \leq n$.

2.4 Interpolation

We start by recalling the definition of the Craig interpolation property.

Definition 4. *Given a logic \mathcal{L} , we say that \mathcal{L} has the Craig interpolation property (CIP), if for every $\varphi \in \mathcal{L}[\sigma]$ and $\psi \in \mathcal{L}[\tau]$ we have that $\varphi \models \psi$ implies that there exists an interpolant $\chi \in \mathcal{L}[\sigma \cap \tau]$ for this entailment, i.e., a sentence for which $\varphi \models \chi$ and $\chi \models \psi$ hold.*

It is well-known that the full \mathcal{GF} fails to have CIP. The known examples of sentences which demonstrate this can be used to make the following observation.

Proposition 1. *The one-dimensional \mathcal{GF} does not have CIP.*

Proof. Consider the following sentences, which are simple variants of the formulas used in [13].

$$\varphi := \exists x \exists y \exists z (G(x, y, z) \wedge R(x, y) \wedge R(y, z) \wedge R(z, x))$$

$$\psi := \forall x \forall y (R(x, y) \rightarrow (A(x) \leftrightarrow \neg A(y)))$$

Notice that both of these sentence are one-dimensional. Now one can show, using essentially the same argument as the one used in Example 1 in [13], that there is no interpolant for the implication $\varphi \models \neg\psi$.

We remark that, in the context of fragments of \mathcal{FO} , CIP is usually defined for *formulas* instead of sentences (as we have defined it). We could have also formulated it for formulas, but we decided to work with sentences for simplicity.

3 Bisimulation for \mathcal{UGF}_1

Given two models \mathfrak{A} and \mathfrak{B} , and tuples $\bar{c} \in A^n$ and $\bar{d} \in B^n$ we will use

$$(\mathfrak{A}, \bar{c}) \equiv_{\sigma} (\mathfrak{B}, \bar{d})$$

to denote the fact that for every $\varphi(x_1, \dots, x_n) \in \mathcal{UGF}_1$ we have that

$$\mathfrak{A} \models \varphi(c_1, \dots, c_n) \iff \mathfrak{B} \models \varphi(d_1, \dots, d_n).$$

The purpose of this section is to define a corresponding notion of bisimulation for \mathcal{UGF}_1 which captures the above equivalence relation. We will start by defining a suitable notion of partial isomorphism.

Definition 5. *Let \mathfrak{A} and \mathfrak{B} be models, and let $X := \{a_1, \dots, a_n\} \subseteq A$ and $Y \subseteq B$. A bijection $p : X \rightarrow Y$, is called a uniform partial σ -isomorphism between \mathfrak{A} and \mathfrak{B} , if*

$$\text{tp}_{\mathfrak{A}}^{\sigma}[a_1, \dots, a_n] = \text{tp}_{\mathfrak{B}}^{\sigma}[p(a_1), \dots, p(a_n)].$$

Quantification in \mathcal{GF} over a model \mathfrak{A} is restricted to *live* subsets of \mathfrak{A} , i.e., subsets of \mathfrak{A} which are either singletons or are *contained* in a single tuple $\bar{a} \in R^{\mathfrak{A}}$, for some $R \in \sigma$. In the case of \mathcal{UGF}_1 we will need the following modified version of the notion of live set, which takes into account the requirement that our formulas are uniform.

Definition 6. *Let \mathfrak{A} be a model and let $X \subseteq A$. We say that X is σ -live, if either $|X| \leq 1$ or there exists $R \in \sigma$ and $(a_1, \dots, a_n) \in R^{\mathfrak{A}}$ so that $X = \{a_1, \dots, a_n\}$.*

We are now ready to define the notion of bisimulation for \mathcal{UGF}_1 .

Definition 7. *Let \mathcal{Z} be a non-empty set of uniform partial σ -isomorphism between two structures \mathfrak{A} and \mathfrak{B} . Let $\bar{c} \in A^n$ and $\bar{d} \in B^n$ be tuples. We say that \mathcal{Z} is a uniform guarded σ -bisimulation between (\mathfrak{A}, \bar{c}) and (\mathfrak{B}, \bar{d}) , if for every $p : X \rightarrow Y \in \mathcal{Z}$ the following conditions hold:*

(cover) *There exists $h \in \mathcal{Z}$ with $\bar{c} = \text{dom}(h)$ so that $h(\bar{c}) = \bar{d}$.*

(forth) *For any $a \in X$ and a σ -live set $X' \subseteq A$, with $a \in X'$, there exists $q : X' \rightarrow Y' \in \mathcal{Z}$ so that*

$$p(a) = q(a).$$

(back) *For any $b \in Y$ and a σ -live set $Y' \subseteq B$, with $b \in Y'$, there exists $q : X' \rightarrow Y' \in \mathcal{Z}$ so that*

$$p^{-1}(b) = q^{-1}(b).$$

If there exists a guarded σ -bisimulation between (\mathfrak{A}, \bar{c}) and (\mathfrak{B}, \bar{d}) , then we denote this by $(\mathfrak{A}, \bar{a}) \sim_{\sigma} (\mathfrak{B}, \bar{b})$.

In what follows we will often refer to uniform guarded bisimulations simply as guarded bisimulations. The following two lemmas establish that our notion of bisimulation is correct, the first of which can be proved in a standard manner by using induction.

Lemma 1. *Let \mathfrak{A} and \mathfrak{B} be models, and let $\bar{c} \in A^n$ and $\bar{d} \in B^n$ be tuples so that $(\mathfrak{A}, \bar{c}) \sim_{\sigma} (\mathfrak{B}, \bar{d})$. Then $(\mathfrak{A}, \bar{c}) \equiv_{\sigma} (\mathfrak{B}, \bar{d})$.*

For the proof of the second lemma we need to recall the definition of ω -saturated model. A *elementary n -type* over a vocabulary σ is a consistent set of first-order formulas (not necessarily quantifier-free) with free variables in $\{x_1, \dots, x_n\}$. Given a σ -model \mathfrak{A} , we say that it is *ω -saturated*, if for every tuple $\bar{a} \in A^n$ of elements of A we have that each elementary n -type over the extended vocabulary $\sigma \cup \{a_1, \dots, a_n\}$, where each a_i denotes a constant to be interpreted as the element a_i , which is finitely consistent with the \mathcal{FO} -theory of (\mathfrak{A}, \bar{a}) , is realized in (\mathfrak{A}, \bar{a}) . It is well-known that every σ -model, where σ is finite and relational, has an ω -saturated elementary extension [3].

Lemma 2. *Let \mathfrak{A} and \mathfrak{B} be two ω -saturated models, and let $\bar{c} \in A^n$ and $\bar{d} \in B^n$ be tuples so that $(\mathfrak{A}, \bar{c}) \equiv_{\sigma} (\mathfrak{B}, \bar{d})$. Then $(\mathfrak{A}, \bar{c}) \sim_{\sigma} (\mathfrak{B}, \bar{d})$.*

Proof. Consider the following set

$$\mathcal{Z} := \{p : \bar{a} \rightarrow \bar{b} \mid (\mathfrak{A}, \bar{a}) \equiv_{\sigma} (\mathfrak{B}, \bar{b})\}.$$

We claim that \mathcal{Z} is a guarded σ -bisimulation between (\mathfrak{A}, \bar{c}) and (\mathfrak{B}, \bar{d}) . We first note that by assumption $\bar{c} \mapsto \bar{d} \in \mathcal{Z}$, and hence \mathcal{Z} satisfies (cover). \mathcal{Z} also clearly consists of uniform partial σ -isomorphism between \mathfrak{A} and \mathfrak{B} . What remains to be proved is that \mathcal{Z} also satisfies (forth) and (back). Since these two cases are analogous, we will concentrate on (forth).

Let $p : \bar{a} \rightarrow \bar{b} \in \mathcal{Z}$, $a \in X$ and $X' := \{c_1, \dots, c_m\} \subseteq A$ be a σ -live set so that $a \in X'$. For simplicity we will assume that $a = c_1$. Consider now the following elementary m -type

$$\Sigma := \{\varphi(p(a), x_2, \dots, x_m) \in \mathcal{UGF}_1[\sigma \cup \{p(a)\}] \mid \mathfrak{A} \models \varphi(a, c_2, \dots, c_m)\}.$$

We claim that Σ is realized in $(\mathfrak{B}, p(a))$. Since \mathfrak{B} is ω -saturated, it suffices to show that each finite subset of Σ is realized in $(\mathfrak{B}, p(a))$. Let

$$\psi_1(p(a), x_2, \dots, x_m), \dots, \psi_r(p(a), x_2, \dots, x_m) \in \Sigma.$$

Since X' is σ -live, there exists an atomic formula $\alpha(x_1, \dots, x_m)$ over σ with the property that

$$\mathfrak{A} \models \exists x_2 \dots \exists x_m (\alpha(a, x_2, \dots, x_m) \wedge \bigwedge_{1 \leq i \leq r} \psi_i(a, x_2, \dots, x_m)).$$

Note that Definition 6 guarantees that this is indeed a formula of $\mathcal{UGF}_1[\sigma]$. Since $(\mathfrak{A}, \bar{a}) \equiv_{\sigma} (\mathfrak{B}, \bar{b})$, we know that

$$\mathfrak{B} \models \exists x_2 \dots \exists x_m (\alpha(p(a), x_2, \dots, x_m) \wedge \bigwedge_{1 \leq i \leq r} \psi_i(p(a), x_2, \dots, x_m)).$$

Thus $\{\psi_1(p(a), x_2, \dots, x_m), \dots, \psi_r(p(a), x_2, \dots, x_m)\}$ is satisfiable in $(\mathfrak{B}, p(a))$, and hence Σ is satisfiable in $(\mathfrak{B}, p(a))$, say by the tuple $(p(a), d_2, \dots, d_m)$. Now $\bar{c} \mapsto \bar{d} \in \mathcal{Z}$ is the mapping we were after.

Remark 2. Using the two previous lemmas one prove in a standard manner that \mathcal{UGF}_1 is the maximal fragment of \mathcal{FO} which is invariant under uniform guarded bisimulation, see for example [2].

4 Proof that \mathcal{UGF}_1 has CIP

In this section we will prove that \mathcal{UGF}_1 has CIP. We will start with the following lemma.

Lemma 3. *Let σ and τ be signatures, and let $\varphi \in \mathcal{UGF}_1[\sigma]$ and $\psi \in \mathcal{UGF}_1[\tau]$. Suppose that there is no $\chi \in \mathcal{UGF}_1[\sigma \cap \tau]$ with the property that $\varphi \models \chi$ and $\chi \models \psi$. Then there is a σ -model \mathfrak{A} and a τ -model \mathfrak{B} with the property that $\mathfrak{A} \models \varphi$, $\mathfrak{B} \not\models \psi$ and $\mathfrak{A} \equiv_{\sigma \cap \tau} \mathfrak{B}$.*

Proof. Essentially the same argument as the one used in the proof of Theorem 4.1 in [2] gives the result.

To give a high level overview of the rest of the proof, suppose that the assumption of Lemma 3 holds for sentences φ and ψ , which implies in particular that there are models \mathfrak{A} and \mathfrak{B} so that $\mathfrak{A} \sim_{\sigma \cap \tau} \mathfrak{B}$. Now, what we want to prove is that $\varphi \wedge \neg\psi$ is satisfiable. To do this, we will follow a standard approach in modal logic [2,11] by constructing an *amalgam* \mathfrak{U} which has the property that $\mathfrak{U} \sim_{\sigma} \mathfrak{A}$ and $\mathfrak{U} \sim_{\tau} \mathfrak{B}$. In particular, it will be a model of $\varphi \wedge \neg\psi$, since $\mathfrak{A} \models \varphi$ and $\mathfrak{B} \models \neg\psi$.

Suppose now that $\mathfrak{A} \sim_{\sigma \cap \tau} \mathfrak{B}$ and let \mathcal{Z} be a guarded $(\sigma \cap \tau)$ -bisimulation which witnesses it. Given a pair (\bar{a}, \bar{b}) we will use $(\bar{a}, \bar{b}) \in \mathcal{Z}$ to denote the fact that there exists $p \in \mathcal{Z}$ with the property that $\bar{a} = \text{dom}(p)$ and $p(\bar{a}) = \bar{b}$. In other words the relation $a_i \mapsto b_i$ induces a uniform partial $(\sigma \cap \tau)$ -isomorphism which belongs to \mathcal{Z} .

Before describing the construction of \mathfrak{U} , we need to introduce some additional notation. Given two tuples \bar{a} and \bar{b} of the same length, we will let $(\bar{a} \otimes \bar{b})$ denote the following tuple:

$$((a_1, b_1), \dots, (a_n, b_n))$$

Given $(\bar{a} \otimes \bar{b})$, we say that it is *left-good*, if for every $1 \leq i < j \leq n$ we have that if $a_i = a_j$, then $b_i = b_j$.² Similarly we say that $(\bar{a} \otimes \bar{b})$ is *right-good*, if for every $1 \leq i < j \leq n$ we have that if $b_i = b_j$, then $a_i = a_j$. Finally we say that $(\bar{a} \otimes \bar{b})$ is *good* if it is left-good and right-good. Note that if $(\bar{a} \otimes \bar{b})$ is of length n , $k \geq n$ and $\mu : [k] \rightarrow [n]$ is a surjection, then we have that if $(\bar{a} \otimes \bar{b})$ is left-good, then so is $(\bar{a} \otimes \bar{b})_{\mu}$. Analogous observation of course holds for right-good and good.

As the domain of the amalgam \mathfrak{U} we will take the set $U = \{(a, b) \in A \times B \mid (a, b) \in \mathcal{Z}\}$, while the interpretations of relation symbols will be defined as follows. First, for every $R \in \sigma \cap \tau$ we define that

$$(\bar{a} \otimes \bar{b}) \in R^{\mathfrak{U}} \text{ iff } \bar{a} \in R^{\mathfrak{A}} \text{ and } (\bar{a}, \bar{b}) \in \mathcal{Z}$$

Then, for every $R \in (\sigma \setminus \tau)$ we define that $(\bar{a} \otimes \bar{b}) \in R^{\mathfrak{U}}$ iff $\bar{a} \in R^{\mathfrak{A}}$ and one of the following conditions holds:

- $(\bar{a}, \bar{b}) \in \mathcal{Z}$.
- $(\bar{a} \otimes \bar{b})$ is left-good and \bar{a} is not $(\sigma \cap \tau)$ -live.

Similarly, for every $R \in (\tau \setminus \sigma)$ we define that $(\bar{a} \otimes \bar{b}) \in R^{\mathfrak{U}}$ iff $\bar{b} \in R^{\mathfrak{B}}$ and one of the following conditions holds:

- $(\bar{a}, \bar{b}) \in \mathcal{Z}$.
- $(\bar{a} \otimes \bar{b})$ is right-good and \bar{b} is not $(\sigma \cap \tau)$ -live.

This concludes the construction of \mathfrak{U} . This construction is similar to the one given in [11] with the exception that we require tuples that are not $(\sigma \cap \tau)$ -live to be either right-good or left-good.

² In other words, if $(\bar{a} \otimes \bar{b})$ is left-good, then the projection $(\bar{a} \otimes \bar{b}) \mapsto \bar{a}$ is an injection.

We now define

$$\mathcal{Z}_1 := \{(\bar{a} \otimes \bar{b}) \mapsto \bar{a} \mid (\bar{a} \otimes \bar{b}) \text{ is } \sigma\text{-live in } \mathfrak{U}\}$$

and

$$\mathcal{Z}_2 := \{(\bar{a} \otimes \bar{b}) \mapsto \bar{b} \mid (\bar{a} \otimes \bar{b}) \text{ is } \tau\text{-live in } \mathfrak{U}\}$$

Note that if $(\bar{a} \otimes \bar{b})$ is σ -live, then by construction it is also left-good (and an analogous observation obviously holds for τ -live tuples in \mathfrak{U}).

Lemma 4. \mathcal{Z}_1 consists of uniform partial σ -isomorphism between \mathfrak{U} and \mathfrak{A} , and \mathcal{Z}_2 consists of uniform partial τ -isomorphism between \mathfrak{U} and \mathfrak{B} .

Proof. We will only consider the case of \mathcal{Z}_1 , since the case of \mathcal{Z}_2 is analogous. Let $(\bar{a} \otimes \bar{b}) \mapsto \bar{a} \in \mathcal{Z}_1$, where the length of $(\bar{a} \otimes \bar{b})$ is n . We will separately check that this mapping preserves 1-types and n -ary atomic formulas.

Let $1 \leq i \leq n$ and suppose that

$$((a_i, b_i), \dots, (a_i, b_i)) \in R^{\mathfrak{U}},$$

where $R \in \sigma$. By construction we know that $(a_i, \dots, a_i) \in R^{\mathfrak{A}}$. Suppose then that

$$(a_i, \dots, a_i) \in R^{\mathfrak{B}}.$$

Since by definition of U we have that $(a_i, b_i) \in \mathcal{Z}$, we can conclude that

$$((a_i, b_i), \dots, (a_i, b_i)) \in R^{\mathfrak{U}}.$$

Thus (a_i, b_i) and a_i have the same 1-types over σ .

We will then verify that the mapping preserves n -ary atomic formulas. Let $R \in \sigma$ be a k -ary relation, where $k \geq n$, and let $\mu : [k] \rightarrow [n]$ be a surjection. We need to show that $(\bar{a} \otimes \bar{b})_\mu \in R^{\mathfrak{U}}$ iff $\bar{a}_\mu \in R^{\mathfrak{A}}$. Again, the left to right direction follows immediately from the definition of \mathfrak{U} , so we will concentrate on the direction from right to left. First we note that if \bar{a} is not $(\sigma \cap \tau)$ -live, then we are done, since then also \bar{a}_μ is not $(\sigma \cap \tau)$ -live.

Thus we can assume that \bar{a} is $(\sigma \cap \tau)$ -live. Now, due to the definition of \mathcal{Z}_1 , we know that $(\bar{a} \otimes \bar{b})$ is σ -live in \mathfrak{U} . Hence, by definition of \mathfrak{U} , and the fact that \bar{a} is $(\sigma \cap \tau)$ -live, we know that $(\bar{a}, \bar{b}) \in \mathcal{Z}$, which is the same as $(\bar{a}_\mu, \bar{b}_\mu) \in \mathcal{Z}$. Now we can deduce, due to the definition of \mathfrak{U} , that $(\bar{a} \otimes \bar{b})_\mu \in R^{\mathfrak{U}}$. This, together with the fact that $(\bar{a} \otimes \bar{b}) \mapsto \bar{a}$ preserves 1-types over σ , allows us to conclude that $\text{tp}_{\mathfrak{U}}^\sigma[\bar{a} \otimes \bar{b}] = \text{tp}_{\mathfrak{A}}^\sigma[\bar{a}]$.

Lemma 5. \mathcal{Z}_1 is a guarded σ -bisimulation between \mathfrak{U} and \mathfrak{A} , and \mathcal{Z}_2 is a guarded τ -bisimulation between \mathfrak{U} and \mathfrak{B} .

Proof. Again, we will only consider the case of \mathcal{Z}_1 , since the case of \mathcal{Z}_2 is analogous. Due to Lemma 4 we just need to verify (back) and (forth) conditions. Let $(\bar{a} \otimes \bar{b}) \mapsto \bar{a} \in \mathcal{Z}_1$, where the length of \bar{a} and \bar{b} is n .

- (forth) Let $(a_i, b_i) \in X$ and let $X' \subseteq U$ be a σ -live set so that $(a_i, b_i) \in X'$. Since X' is σ -live, we know that it is of the form $\{(c_1, d_1), \dots, (c_m, d_m)\}$, with $(\bar{c} \otimes \bar{d})$ being left-good. Now $(\bar{c} \otimes \bar{d}) \mapsto \bar{c} \in \mathcal{Z}_1$ is the required mapping.
- (back) Let $a_i \in Y$ and let $Y' \subseteq A$ be a σ -live set so that $a_i \in Y$. For concreteness, suppose that $Y' = \{c_1, \dots, c_m\}$. Consider first the case that Y' is not $(\sigma \cap \tau)$ -live in \mathfrak{A} . For every $2 \leq i \leq m$ we will pick an element d_i so that $(c_i, d_i) \in \mathcal{Z}$. Note that such elements exists since each singleton is a live element. By construction $\{(c_1, d_1), \dots, (c_m, d_m)\}$ is σ -live, and hence $(\bar{c} \otimes \bar{d}) \mapsto \bar{c} \in \mathcal{Z}_1$ is the required mapping we were after. Suppose then that Y' is $(\sigma \cap \tau)$ -live in \mathfrak{A} . Since $(a_i, b_i) \in U$, we know that $(a_i, b_i) \in \mathcal{Z}$. Since \mathcal{Z} is a guarded $(\sigma \cap \tau)$ -bisimulation, there exists a set $\{d_1, \dots, d_m\} \subseteq B$ so that $(\bar{c}, \bar{d}) \in \mathcal{Z}$ and $(a_i, b_i) \in (\bar{c} \otimes \bar{d})$. In particular $(\bar{c} \otimes \bar{d})$ is σ -live in \mathfrak{U} , and hence $(\bar{c} \otimes \bar{d}) \mapsto \bar{d} \in \mathcal{Z}_1$, which is the mapping we were after.

Theorem 1. \mathcal{UGF}_1 has Craig interpolation property.

Proof. Let $\varphi \in \mathcal{UGF}_1[\sigma]$ and $\psi \in \mathcal{UGF}_1[\tau]$ be sentences so that $\varphi \models \psi$, but there is no interpolant for this entailment. By lemma 3 there exists a σ -model \mathfrak{A} and a τ -model \mathfrak{B} such that $\mathfrak{A} \models \varphi$, $\mathfrak{B} \not\models \psi$ and $\mathfrak{A} \equiv_{\sigma \cap \tau} \mathfrak{B}$. Take ω -saturated elementary extensions $\hat{\mathfrak{A}}$ and $\hat{\mathfrak{B}}$ of \mathfrak{A} and \mathfrak{B} . Since $\hat{\mathfrak{A}} \equiv_{\sigma \cap \tau} \hat{\mathfrak{B}}$, by lemma 2 we have that $\hat{\mathfrak{A}} \sim_{\sigma \cap \tau} \hat{\mathfrak{B}}$. Using the construction presented in this section there exists a $(\sigma \cup \tau)$ -model \mathfrak{U} with the property that $\mathfrak{U} \sim_{\sigma} \hat{\mathfrak{A}}$ and $\mathfrak{U} \sim_{\tau} \hat{\mathfrak{B}}$. Thus $\mathfrak{U} \models \varphi \wedge \neg\psi$, i.e. $\varphi \wedge \neg\psi$ is consistent, which is a contradiction with the assumption that $\varphi \models \psi$.

5 Complexity of uniform \mathcal{GF}

In this section we will prove that the complexity of the satisfiability problem of uniform \mathcal{GF} is in NEXPTIME. Since it was proved in [15] that the complexity of the satisfiability problem of \mathcal{UGF}_1 is NEXPTIME-hard, this upper bound is sharp.

5.1 Scott normal form

As usual, we will start by arguing that we can restrict our attention to sentences which are in a certain normal form. The normal form that we will use here has a somewhat awkward form, but the proof of Lemma 6 should clarify why we chose to use it.

Definition 8. Let φ be a sentence of \mathcal{UGF} . We say that φ is in normal form, if it has the following shape

$$\bigwedge_{t \in T} \exists z \lambda_t(z) \wedge \bigwedge_{i \in I} \forall \bar{x} (\alpha_i(\bar{x}) \rightarrow \exists \bar{y} (\beta_i(\bar{x}, \bar{y}) \wedge \psi_i(\bar{x}, \bar{y}))) \\ \wedge \bigwedge_{j \in J} \forall \bar{x} (\kappa_j(\bar{x}) \rightarrow (\theta_j(\bar{x}) \rightarrow \forall \bar{y} (\gamma_j(\bar{x}, \bar{y}) \rightarrow \psi_j(\bar{x}, \bar{y})))),$$

where T, I, J are non-empty (finite) sets, $\lambda_t, \alpha_i, \beta_i, \kappa_j$ and γ_j are atomic formulas and ψ_i, θ_j and ψ_j are quantifier-free formulas.

Remark 3. In the definition of the normal form we do not require that the tuples \bar{y} are necessarily non-empty, i.e., we allow formulas of the form $\forall \bar{x}(\alpha_i(\bar{x}) \rightarrow \psi_i(\bar{x}))$ in our normal forms. However, we do require that the tuples \bar{x} are non-empty, and hence we do not allow formulas of the form $\exists \bar{y}(\beta_i(\bar{y}) \wedge \psi_i(\bar{y}))$, where the length of \bar{y} is more than one.

If φ is a sentence of \mathcal{UGF} in normal form, then we refer to its conjuncts of the form

$$\forall \bar{x}(\alpha_i(\bar{x}) \rightarrow \exists \bar{y}(\beta_i(\bar{x}, \bar{y}) \wedge \psi_i(\bar{x}, \bar{y})))$$

as the *existential requirements* and we will use φ_i^{\exists} to denote them. Given a model \mathfrak{A} , an existential requirement φ_i^{\exists} and $\bar{a} \in \alpha_i^{\mathfrak{A}}$ we say that a tuple \bar{c} is a *witness* for φ_i^{\exists} and \bar{a} if

$$\mathfrak{A} \models \beta_i(\bar{a}, \bar{c}) \wedge \psi_i(\bar{a}, \bar{c}).$$

Conjuncts of the form

$$\forall \bar{x}(\kappa_j(\bar{x}) \rightarrow (\theta_j(\bar{x}) \rightarrow \forall \bar{y}(\gamma_j(\bar{x}, \bar{y}) \rightarrow \psi_j(\bar{x}, \bar{y}))))$$

will be referred to as the *universal requirements* and we will use φ_j^{\forall} to denote them.

Using standard renaming techniques one can establish the following.

Lemma 6. *There is a polynomial nondeterministic procedure, taking as its input a sentence $\varphi \in \mathcal{UGF}[\sigma]$ and producing a sentence $\varphi' \in \mathcal{UGF}[\sigma']$ in normal form, where $\sigma' \supset \sigma$, such that*

1. *if $\mathfrak{A} \models \varphi$ for some model \mathfrak{A} , then there is a run of the procedure producing a normal form φ' such that $\mathfrak{A}' \models \varphi'$ for some expansion \mathfrak{A}' of \mathfrak{A} to the vocabulary σ' ,*
2. *if the procedure has a run producing φ' and $\mathfrak{A}' \models \varphi'$, for some \mathfrak{A}' , then the σ -reduct \mathfrak{A} of \mathfrak{A}' satisfies φ .*

Proof. We will essentially follow the proof of lemma 1 in [15], with some small technical modifications. Let $\varphi \in \mathcal{UGF}[\sigma]$ be a sentence, which w.l.o.g contains only existential quantification. Let ψ be the innermost formula of φ which starts with a block of existential quantifiers. If ψ is a sentence, we will nondeterministically either replace it with \perp or \top and add ψ or $\neg\psi$ (depending on our guess) as a conjunct to the resulting formula. Suppose then that ψ is a formula of the form

$$\exists \bar{y}(\alpha(\bar{x}, \bar{y}) \wedge \psi(\bar{x}, \bar{y})).$$

Since φ was a sentence, ψ occurs in a scope of another formula of the form

$$\exists \bar{z}(\alpha'(\bar{x}) \wedge \psi'(\bar{x})),$$

where $\bar{z} \subseteq \bar{x}$. Let α' be the guard of the innermost such formula. We will now replace φ with the following formula

$$\begin{aligned} & \varphi[\psi(\bar{x})/R(\bar{x})] \wedge \forall \bar{x}(R(\bar{x}) \rightarrow \exists \bar{y}(\alpha(\bar{x}, \bar{y}) \wedge \psi(\bar{x}, \bar{y}))) \\ & \wedge \forall \bar{x}(\alpha'(\bar{x}) \rightarrow (\neg R(\bar{x}) \rightarrow \forall \bar{y}(\alpha(\bar{x}, \bar{y}) \rightarrow \neg \psi(\bar{x}, \bar{y})))), \end{aligned}$$

where $\varphi[\psi(\bar{x})/R(\bar{x})]$ is the sentence obtained from φ by replacing the previously mentioned subformula $\psi(\bar{x})$ with the atomic formula $R(\bar{x})$ which has a fresh relation symbol R . It is straightforward to verify that the resulting sentence is equi-satisfiable with φ .

Now one can repeat the above procedure until one is left with a sentence of the form

$$\bigwedge_{t \in T} \exists \bar{x}(\alpha_t(\bar{x}) \wedge \psi_t(\bar{x})) \wedge \bigwedge_{i \in I} \varphi_i^{\exists} \wedge \bigwedge_{j \in J} \varphi_j^{\forall},$$

where each φ_i^{\exists} is an existential requirement, while each sentence φ_j^{\forall} is an universal requirement. Now one can replace each conjunct $\exists x_1 \dots \exists x_n(\alpha(x_1, \dots, x_n) \wedge \psi_t(x_1 \dots x_n))$ with a sentence of the form

$$\exists x \lambda_t(x) \wedge \forall x_1(\lambda_t(x_1) \rightarrow \exists x_2 \dots \exists x_n(\alpha_t(x_1, \dots, x_n) \wedge \psi_t(x_1, \dots, x_n))),$$

where λ_t is a fresh unary relation symbol. The resulting sentence is clearly equi-satisfiable with the original sentence and furthermore it is in normal form.

5.2 Satisfiability Witnesses

A standard technique in proving that the complexity of the satisfiability problem of a given fragment of \mathcal{FO} is in NEXPTIME is to show that each satisfiable sentence of this fragment has a finite model of size at most exponential with respect to the length of the sentence [8,12,15,16]. However, in the case of \mathcal{UGF} it seems to be easier to show that we can associate to each of its sentences φ a different type of certificate, which is still at most exponential with respect to the length of the sentence, and which can be used to construct a (potentially infinite) model for φ .

Definition 9. Let $\varphi \in \mathcal{UGF}[\sigma]$ be a sentence in normal form, P be a set of 1-types over σ and $\pi \in P$. A pair (\mathfrak{A}, c) , where $c \in A$, is called a (P, π) -witness for φ , if it satisfies the following requirements.

1. $\text{tp}_{\mathfrak{A}}[c] = \pi$.
2. For every $a \in A$ we have that $\text{tp}_{\mathfrak{A}}[a] \in P$.
3. For every existential requirement φ_i^{\exists} and for every tuple \bar{a} which contains c we have that if $\mathfrak{A} \models \alpha_i(\bar{a})$, then there exists a witness for φ_i and \bar{a} .
4. For every universal requirement φ_j^{\forall} and for every tuple \bar{a} which contains c we have that if $\mathfrak{A} \models \kappa_j(\bar{a}) \wedge \theta_j(\bar{a})$, then for every tuple \bar{b} we have that

$$\mathfrak{A} \models \gamma_j(\bar{a}, \bar{b}) \rightarrow \psi_j(\bar{a}, \bar{b}).$$

Here the intuition is that a (P, π) -witness (\mathfrak{A}, c) is a local certificate; it certifies that we can provide witnesses for tuples which contain the element c . The main idea now is that if we have a (P, π) -witness for each $\pi \in P$, then we can use them to construct a proper model for φ .

Definition 10. *Let $\varphi \in \mathcal{UGF}[\sigma]$ be a sentence in normal form. A set of 1-types P over σ is a witness for φ , if it satisfies the following two requirements.*

1. *For every conjunct $\exists z \lambda_t(z)$ there exists $\pi \in P$ so that $\lambda_t(x) \in \pi$.*
2. *For every $\pi \in P$ there exists a (P, π) -witness for P .*

The following lemmas prove that an existence of a witness for φ is equivalent with the satisfiability of φ .

Lemma 7. *Let $\varphi \in \mathcal{UGF}$ be a sentence in normal form. If φ is satisfiable, then there exists a witness for it.*

Proof. Suppose that $\mathfrak{A} \models \varphi$. As the set of 1-types P we can take the set

$$\{\text{tp}_{\mathfrak{A}}[a] \mid a \in A\}.$$

Clearly for every conjunct $\exists z \lambda_t(z)$ there exists a suitable 1-type in P . Towards verifying the second requirement let $\pi \in P$ and let $c \in A$ be an element which realizes π . Then (\mathfrak{A}, c) is clearly a (P, π) -witness for φ .

Lemma 8. *Let $\varphi \in \mathcal{UGF}$ be a sentence in normal form. If there exists a witness for φ , then it is satisfiable.*

Proof. For simplicity we will assume that φ contains exactly one conjunct of the form $\exists z \lambda_t(z)$. Let P be a witness for φ . Thus for every $\pi \in P$ there exists a pair (\mathfrak{A}^π, c) which is a (P, π) -witness for φ . Our goal is to use these witnesses to construct a sequence of models

$$\mathfrak{A}_1 \leq \mathfrak{A}_2 \leq \mathfrak{A}_3 \leq \dots$$

so that their union is a model of φ .

Let $\pi \in P$ be a 1-type so that $\pi \models \lambda_t$. As the model \mathfrak{A}_1 we will take the model which contains a single element with 1-type π . Suppose then that we have defined \mathfrak{A}_n in such a way that each 1-type realized in \mathfrak{A}_n belongs to P . To define the model \mathfrak{A}_{n+1} we will proceed as follows. Given $a \in \mathfrak{A}_n$, we will use W_a to denote the set $A^\pi - \{c\}$, where A^π refers to the domain of the model in the (P, π) -witness (\mathfrak{A}^π, c) of $\pi := \text{tp}_{\mathfrak{A}_n}[a]$. Without loss of generality we will assume that the sets W_a are pairwise disjoint. Now we will define \mathfrak{A}_{n+1} as follows.

- The domain of the model is

$$A_n \cup \bigcup_{a \in A_n} W_a^*$$

- $\mathfrak{A}_{n+1} \upharpoonright A_n$ is defined to be isomorphic with A_n .

- For each $a \in A_n$ and for each $\{c_1, \dots, c_m\} \subseteq W_a$, we define that

$$\text{tp}_{\mathfrak{A}_{n+1}}[a, c_1, \dots, c_m] := \text{tp}_{\mathfrak{A}^\pi}[c, c_1, \dots, c_m],$$

where π is the 1-type of a .

- For every tuple (a_1, \dots, a_m) and a m -ary relation R for which we have not yet defined whether (a_1, \dots, a_m) belongs to $R^{\mathfrak{A}_{n+1}}$, we will simply define that it does not belong to it.

The last step guarantees that if a tuple, which contains more than one element, is live in \mathfrak{A}_{n+1} , then it was already alive in one of the models \mathfrak{A}^π . It is straightforward to verify that the union of the models $(\mathfrak{A}_n)_{n < \omega}$ is indeed a model of φ .

5.3 Complexity of UGF

Although the size of a witness for φ is clearly only exponential with respect to $|\varphi|$, we do not yet have any upper bounds on the time it takes to verify that it really is a witness for φ . The following lemma gives us such a bound.

Lemma 9. *Let $\varphi \in UGF$ be a sentence in normal form and let σ denote the vocabulary of φ . Let P be a set of 1-types over σ and $\pi \in P$. If there exists a (P, π) -witness for φ , then there exists one in which the size of the model is at most $2^{|\varphi|^{O(1)}}$.*

Proof. Let (\mathfrak{A}, c) be a (P, π) -witness for φ and let $m = \max\{ar(R) \mid R \in \sigma\}$. Note that $m \leq |\varphi|$. Our goal is to construct a sequence

$$\mathfrak{B}_1 \leq \dots \leq \mathfrak{B}_m$$

of models so that (\mathfrak{B}_m, c) is a (P, π) -witness for φ and $|B_m| \leq 2^{|\varphi|^{O(1)}}$. As the model \mathfrak{B}_1 we will take the model which contains a single element with 1-type π ; let e denote this element.

Before moving forward, we will introduce one auxiliary definition. Let $\bar{a} = (a_1, \dots, a_n)$ and $\bar{b} = (b_1, \dots, b_n)$ be tuples of elements from two models \mathfrak{A} and \mathfrak{B} . Let $\{c_1, \dots, c_m\}$ denote the set of distinct elements in \bar{a} . We say that \bar{a} and \bar{b} are *similar*, if the mapping $p : \bar{a} \rightarrow \bar{b}$, which was the mapping induced by the relation $a_i \mapsto b_i$, is a bijection and furthermore

$$\text{tp}_{\mathfrak{A}}[c_1, \dots, c_n] = \text{tp}_{\mathfrak{B}}[p(c_1), \dots, p(c_n)].$$

Suppose now that we have defined \mathfrak{B}_k , where $k < m$, and in such a way that for each σ -live tuple \bar{b} for which $\text{tp}_{\mathfrak{B}_k}[\bar{b}]$ has been defined, there exists a similar tuple \bar{a} which consists of elements of \mathfrak{A} . Given an existential requirement φ_i^{\exists} of φ and a tuple $\bar{b} \in \alpha_i^{\mathfrak{B}_k}$, which contains the element e , we say that \bar{b} is a *i -defect* if there exists no witness for φ_i^{\exists} and \bar{b} in the model \mathfrak{B}_k . By construction, for each i -defect \bar{b} we can find a tuple \bar{a} of elements of \mathfrak{A} so that \bar{b} and \bar{a} are similar. In particular $\bar{a} \in \alpha_i^{\mathfrak{A}}$, and hence there exists a witness \bar{c} for φ_i^{\exists} and \bar{a} in \mathfrak{A} ; let $W_{\bar{b}, i}$

denote the set of elements in \bar{c} which were not contained in \bar{a} . Without loss of generality we will assume that the sets $W_{\bar{b},i}$ are pairwise disjoint. Now we will define \mathfrak{B}_{k+1} as follows.

- The domain of the model is

$$B_k \cup \bigcup_{i \in I_{\bar{b}} \text{ an } i\text{-defect}} W_{\bar{b},i}$$

- $\mathfrak{B}_{k+1} \upharpoonright B_k$ is defined to be isomorphic with \mathfrak{B}_k .
- For each i -defect \bar{b} and a set $W_{\bar{b},i} = \{c_1, \dots, c_n\}$ we define that

$$\text{tp}_{\mathfrak{B}_{k+1}}[d_1, \dots, d_r, c_1, \dots, c_n] = \text{tp}_{\mathfrak{A}}[p(d_1), \dots, p(d_r), c_1, \dots, c_n],$$

where (d_1, \dots, d_r) enumerates all the elements occurring in \bar{b} and $p: \bar{b} \rightarrow \bar{a}$.

- For every tuple (b_1, \dots, b_n) and a n -ary relation R for which we have not yet defined whether (b_1, \dots, b_n) belongs to $R^{\mathfrak{B}_{k+1}}$, we will simply define that it does not belong to it.

This completes the construction of the models $\mathfrak{B}_1, \dots, \mathfrak{B}_m$. To bound the size of \mathfrak{B}_m , we first note that $|B_{k+1}| \leq |B_k| + |\varphi| |D_k|$, where D_k denotes the number of defects in \mathfrak{B}_k . By construction, for every defect (d_1, \dots, d_r) of \mathfrak{B}_k the set $\{d_1, \dots, d_r\}$ is a σ -live set which is not contained in \mathfrak{B}_ℓ , for any $\ell < k$. If $k = 1$, then the number of such σ -live sets is one, and if $k > 1$, then the number of such σ -live sets is D_{k-1} . Since each σ -live set is of size at most $|\varphi|$, there are at most $|\varphi| |\varphi|^{|\varphi|} D_{k-1} = 2^{|\varphi|^{O(1)}} D_{k-1}$ defects in \mathfrak{B}_k , i.e., $D_k \leq 2^{|\varphi|^{O(1)}} D_{k-1}$. Since $m \leq |\varphi|$, we have that $D_k \leq 2^{|\varphi|^{O(1)}}$, for any $k < m$, and hence $|B_m| \leq 2^{|\varphi|^{O(1)}}$.

Thus what remains to be proven is that (\mathfrak{B}_m, e) is a (P, π) -witness for φ . Here the only non-trivial requirement that we need to verify is that \mathfrak{B}_m satisfies the second item in definition 9. So, let φ_i^{\exists} be an existential requirement and let $\bar{b} = (b_1, \dots, b_n) \in \alpha_i^{\mathfrak{B}_m}$ be a tuple which contains e . We can clearly assume that $n < m$. It suffices to show that \bar{b} is contained in \mathfrak{B}_k , for some $k < m$, since then by construction we know that it has a witness in \mathfrak{B}_m .

Aiming for a contradiction, suppose that \bar{b} is contained in \mathfrak{B}_m , but it is not contained in \mathfrak{B}_k for any $k < m$. By construction we know that, since \bar{b} is σ -live, we assigned a table to some tuple (b'_1, \dots, b'_r) , where (b'_1, \dots, b'_r) enumerates the set of distinct elements of (b_1, \dots, b_n) . Again, by construction we know that we assigned a table to the tuple (b'_1, \dots, b'_r) , because we wanted to provide a witness for some tuple (d_1, \dots, d_s) , which contains e and for which $\{d_1, \dots, d_s\}$ is a *strict* subset of $\{b'_1, \dots, b'_r\}$.³

Now observe that (d_1, \dots, d_s) is a σ -live tuple containing e , which is contained in \mathfrak{B}_{m-1} but is not contained in \mathfrak{B}_k for any $k < m-1$. Indeed, if it were contained in \mathfrak{B}_k , for some $k < m-1$, then by construction we would have provided a witness for it in the model \mathfrak{B}_{k+1} , i.e., (b_1, \dots, b_n) would have been contained in \mathfrak{B}_{k+1} . But now we are in a position which is the same as the one that we started in; in

³ If it were not, there would have been no need to provide a witness for it.

particular, we can repeat the above argument. After repeating the argument (at least) $(n - 1)$ -times we would end up with the conclusion that e is contained in some \mathfrak{B}_k , where $k > 1$, but it is not contained in \mathfrak{B}_1 , which would be an obvious contradiction.

Now we can prove the main theorem of this section.

Theorem 2. *The satisfiability problem of \mathcal{UGF} is NEXPTIME-complete.*

Proof. The lower bound follows from the proof of Theorem 3 in [15]. We will give an informal description of a non-deterministic procedure running in exponential time which determines whether a given sentence $\varphi \in \mathcal{UGF}$ is satisfiable. It starts by converting φ into an equi-satisfiable sentence $\varphi' \in \mathcal{UGF}$ in normal form, after which it guesses a set of 1-types P over the vocabulary of φ' and for each $\pi \in P$ a (P, π) -witness (\mathfrak{A}, c) for φ , where the size of \mathfrak{A} is at most $2^{|\varphi|^{O(1)}}$. Lemmas 6, 7, 8 and 9 guarantee that this procedure is correct. Since $|P| \leq 2^{|\varphi|}$, the algorithm runs in exponential time with respect to $|\varphi|$.

6 Conclusions

In this paper we have proved two results of quite distinct flavour on uniform guarded fragments. The first result was that although \mathcal{GF} fails to have Craig interpolation, its one-dimensional uniform fragment does have it. The second result was that the complexity of the satisfiability problem of the uniform guarded fragment is NEXPTIME-complete. The results presented in this paper suggest several new research questions, but here we will mention just two of them.

The first question is whether or not the uniform \mathcal{GF} has Craig interpolation property. While the correctness of the amalgam construction presented in Section 4 rests on the assumption of one-dimensionality, we have not been able to show that uniform \mathcal{GF} would not have Craig interpolation property. This has led the author to conjecture that the uniform \mathcal{GF} does in fact have Craig interpolation property.

The second question is whether or not uniform \mathcal{GF} has the exponential model property (note that if uniform \mathcal{GF} would have an exponential model property, then one would obtain Theorem 2 for free). As we saw in the proof of Lemma 9, the requirement of uniformity essentially prevents uniform \mathcal{GF} from enforcing long paths, and this seems to suggest that uniform \mathcal{GF} can only enforce exponentially long paths (which it can enforce, since it contains standard modal logic with the global diamond). Because of this, the author conjectures that uniform \mathcal{GF} has the exponential model property.

Acknowledgements

The author wishes to thank Bartosz Bednarczyk for several helpful discussions on interpolation and fragments of first-order logic, and for suggesting the problem

of determining the complexity of the satisfiability problem of uniform guarded fragment. The author also wishes to thank Antti Kuusisto for pointing out rather silly mistakes in the original definitions of uniformity and the uniform guarded bisimulation. Finally, the author wishes to thank the anonymous reviewers for their useful remarks which improved the presentation of this paper.

References

1. Andréka, H., Németi, I., van Benthem, J.: Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic* **27**, 217–274 (1998)
2. Bárány, V., Benedikt, M., Cate, B.T.: Some model theory of guarded negation. *The Journal of Symbolic Logic* **83**, 1307 – 1344 (2018)
3. Bell, J., Slomson, A.: *Models and Ultraproducts: An Introduction*. Dover Books on Mathematics Series, Dover Publications (2006)
4. van Benthem, J.: The many faces of interpolation. *Synthese* **164**(3), 451–460 (2008)
5. Blackburn, P., Rijke, M.d., Venema, Y.: *Modal Logic*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press (2001)
6. Gabbay, D.M.: Craig’s interpolation theorem for modal logics. In: Hodges, W. (ed.) *Conference in Mathematical Logic — London ’70*. pp. 111–127. Springer Berlin Heidelberg, Berlin, Heidelberg (1972)
7. Grädel, E.: On the restraining power of guards. *Journal of Symbolic Logic* **64**(4), 1719–1742 (1999)
8. Grädel, E., Kolaitis, P., Vardi, M.: On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic* **3**(1), 53–69 (1997)
9. Hella, L., Kuusisto, A.: One-dimensional fragment of first-order logic. In: *Advances in Modal Logic*. vol. 10, pp. 274–293 (08 2014)
10. Herzig, A.: A new decidable fragment of first order logic. In: *3rd Logical Biennial Summer School and Conference in Honour of SC Kleene* (1990)
11. Hoogland, E., Marx, M.: Interpolation and definability in guarded fragments. *Studia Logica: An International Journal for Symbolic Logic* **70**(3), 373–409 (2002)
12. Jaakkola, R.: Ordered Fragments of First-Order Logic. In: *46th International Symposium on Mathematical Foundations of Computer Science (MFCS 2021)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 202, pp. 62:1–62:14 (2021)
13. Jung, J.C., Wolter, F.: Living without beth and craig: Definitions and interpolants in the guarded and two-variable fragments. *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)* pp. 1–14 (2021)
14. Kieronski, E.: One-Dimensional Logic over Words. In: *25th EACSL Annual Conference on Computer Science Logic (CSL 2016)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 62, pp. 38:1–38:15 (2016)
15. Kieronski, E.: One-Dimensional Guarded Fragments. In: *44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 138, pp. 16:1–16:14 (2019)
16. Kieronski, E., Kuusisto, A.: Complexity and expressivity of uniform one-dimensional fragment with equality. In: *39nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2014)*. Lecture Notes in Computer Science, vol. 8634, pp. 365–376 (2014)
17. Kierónski, E., Kuusisto, A.: Uniform One-Dimensional Fragments with One Equivalence Relation. In: *24th EACSL Annual Conference on Computer Science Logic (CSL 2015)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 41, pp. 597–615 (2015)

18. Kieronski, E., Kuusisto, A.: One-Dimensional Logic over Trees. In: 42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 83, pp. 64:1–64:13 (2017)
19. Marx, M.: Interpolation in modal logic. In: Proceedings of the 7th International Conference on Algebraic Methodology and Software Technology. p. 154–163. AMAST '98, Springer-Verlag, Berlin, Heidelberg (1999)
20. Mogavero, F., Perelli, G.: Binding Forms in First-Order Logic. In: 24th EACSL Annual Conference on Computer Science Logic (CSL 2015). Leibniz International Proceedings in Informatics (LIPIcs), vol. 41, pp. 648–665 (2015)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

