The Guarded Fragment with Nested Equivalences

Oskar Fiuk
Institute of Computer Science, University of Wrocław, Poland

Abstract

The Guarded Fragment (GF) is a well-established decidable fragment of first-order logic. We study an extension of GF with *nested equivalence relations*, namely a family of distinguished binary predicates $E_1, E_2, \ldots$ interpreted as equivalence relations such that E_{k+1} is coarser than E_k for every k. We show that the equality-free GF with nested equivalence relations enjoys the finite model property and has a decidable satisfiability problem. Moreover, we establish tight complexity bounds for satisfiability: TOWER-completeness in general, and $(K+2)$-EXPTIME-completeness when the number of distinguished predicates is fixed to K. Finally, we show that satisfiability becomes undecidable if either the nesting condition is dropped (already with two equivalence relations) or equality is admitted (already with a single equivalence relation).

2012 ACM Subject Classification Theory of computation → Finite Model Theory

Keywords and phrases guarded fragment, finite model property, nested equivalence relations

Digital Object Identifier 10.4230/LIPIcs.LICS.2026.44

Related Version *Full Version*: https://arxiv.org/abs/2605.15072

Funding Polish National Science Center, grant No. 2021/41/B/ST6/00996.

Acknowledgements The author thanks Emanuel Kieroński for helpful discussions, and the reviewers of LICS 2026 for their valuable suggestions.

1 Introduction

Many reasoning tasks reduce to satisfiability problems over domains whose elements carry data values drawn from potentially infinite or very large ranges. To model situations in which these values can be compared at different levels of granularity, one may introduce a family of increasingly coarser equivalence relations: two elements are related by the k-th relation if their associated values agree at the k-th level of precision, where precision decreases monotonically as k increases. We refer to such families of equivalence relations as *nested*.

For example, Geographical Information Systems often organise data using progressively broader categories: city, state, region, and country. This hierarchy is captured by four equivalence relations E_1, E_2, E_3, and E_4, where E_1 relates locations within the same city, E_2 within the same state, E_3 within the same region, and E_4 within the same country. These equivalence relations are naturally nested: addresses located in the same city are also in the same state, those in the same state are also in the same region, and so on.

Similar hierarchies appear in numerous other contexts: Data Storage (organised into files and folders), Network Management (organised into networks and subnetworks), Dependency Maintenance (organised into modules and submodules), E-Commerce (organised into products, categories, and subcategories). Given the naturalness and pervasiveness of hierarchically structured data, it is of broad interest to develop reasoning frameworks capable of modelling such data using nested equivalence relations.

Previous work. Logics with nested equivalences have so far been studied predominantly within the two-variable regime: the Two-Variable Fragment (FO^2) and related formalisms.

One can reason about nested equivalence relations in expressive description logics; such as Horrocks and Sattler's [16] logic $\mathcal{SHI}$, that is, the standard $\mathcal{ALC}$ logic enriched with transitive roles ($\mathcal{S}$), role hierarchies ($\mathcal{H}$), and role inverses ($\mathcal{I}$). Björklund and Bojańczyk [5] investigated FO^2 over words with *nested data*, that is, structures interpreting a linear order, nested equivalence relations, and further unary predicates. Recent work of Fiuk, Kieroński, and Michielini [13] advanced this line of research by establishing decidability for several variants of FO^2 with nested equivalence relations; see Section 6 for a detailed comparison.

Although FO^2 is a prominent decidable fragment, its expressive power becomes severely limited once one needs to reason about more than two elements simultaneously. To the best of our knowledge, the only decidable fragment supporting unboundedly many variables together with nested equivalence relations that has been systematically studied is an extension of the Unary Negation Fragment, denoted as $\mathrm{UNFO} + \mathcal{SOH}$ [12].

Our contribution. In this work we study the Guarded Fragment (GF) equipped with nested equivalence relations. Up to our knowledge, even the case of GF with a single equivalence relation had remained open. Previous work on GF considered equivalence relations only in restricted settings, such as *equivalence guards*; see Section 6 for discussion.

Guarded Fragment. Introduced by Andréka, van Benthem, and Németi [1] as a generalisation of modal logic, the Guarded Fragment has become, alongside FO^2, one of the central decidable fragments of first-order logic. The prominence of GF stems from its robust computational and model-theoretic properties. In particular, GF provides a logical foundation for basic description logics. For a comprehensive introduction to the Guarded Fragment, see Chapter 4 in Pratt-Hartmann's monograph [24].

▶ **Definition 1.** *The* Guarded Fragment *(*GF*) is the set of function-free first-order formulas generated by the following rules:*

- **(i)** *Every atomic formula is in* GF*.*
- **(ii)** GF *is closed under Boolean connectives.*
- **(iii)** *If $\psi(x) \in \mathrm{GF}$ has at most one free variable, then both $\exists x\, \psi(x)$ and $\forall x\, \psi(x)$ are in* GF*.*
- **(iv)** *If $\psi(\bar{x}, \bar{y}) \in \mathrm{GF}$ and $\gamma(\bar{x}, \bar{y})$ is an atom mentioning every variable from $\bar{x} \cup \bar{y}$, then both $\exists \bar{x}\, \big(\gamma(\bar{x}, \bar{y}) \wedge \psi(\bar{x}, \bar{y})\big)$ and $\forall \bar{x}\, \big(\gamma(\bar{x}, \bar{y}) \to \psi(\bar{x}, \bar{y})\big)$ belong to* GF*.*

The atom $\gamma(\bar{x}, \bar{y})$ in rule (iv) is called a guard*.*

Nested equivalences. We assume that logical signatures include a family of *distinguished* binary symbols, denoted as $E_1, E_2, E_3, \ldots$. The interpretation of these distinguished symbols is defined such that for every relevant $k \geq 1$ the following holds:

- E_k is an equivalence relation.
- E_{k+1} is coarser than E_k, that is, $\forall x, y\, \big(xE_k y \to xE_{k+1} y\big)$ is an axiom.

Beyond the distinguished symbols, signatures may also include *non-distinguished* symbols of any arity as well as constant symbols; these symbols are *free*, that is, their interpretation is not constrained. In this work, function symbols of positive arity are not considered.

For $K \in \mathbb{N}$, we denote by $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ the logic consisting of all formulas of GF whose signatures include the first K distinguished symbols $E_1, \ldots, E_K$, and whose models are restricted to structures satisfying the intended semantics of these predicates. We write $\mathrm{GF}[\mathcal{EQ}^{\subseteq}]$ for the union of $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ over all $K \geq 0$. More generally, for any fragment $\mathcal{F}$, we use the notation $\mathcal{F}[K\text{-}\mathcal{EQ}^{\subseteq}]$ and $\mathcal{F}[\mathcal{EQ}^{\subseteq}]$ as well. Finally, for $m \in \mathbb{N}$, we denote by $\mathcal{F}^m$ the fragment of $\mathcal{F}$ consisting of formulas with at most m distinct variables.

Satisfiability problem. In this work, we study the satisfiability problem for GF[$\mathcal{EQ}^{\subseteq}$] and GF[K-$\mathcal{EQ}^{\subseteq}$]: given a sentence φ, does there exist a model $\mathfrak{A}$ such that $\mathfrak{A} \models \varphi$?

We also investigate whether these fragments enjoy the *finite model property*: namely, whether every satisfiable sentence of a fragment has a finite model.

It is known that GF (without equivalences) has the finite model property, with a doubly exponential upper bound on the size of minimal models [3]. The satisfiability problem for GF is 2-ExpTime-complete; with predicates of bounded arity or finite number of variables, the complexity becomes ExpTime-complete [15]. It is also known that FO^2[$\mathcal{EQ}^{\subseteq}$] has the finite model property, with an exponential upper bound on the size of minimal models. The satisfiability problem for FO^2[$\mathcal{EQ}^{\subseteq}$] is NExpTime-complete [13]. In what follows, we present the results obtained for GF with nested equivalence relations and at least 3 variables.

Equality. We first observe that GF^3[1-$\mathcal{EQ}^{\subseteq}$] – that is, GF with three variables and a single equivalence relation – is undecidable, provided that equality is available in the syntax.

▶ **Proposition 2.** *The satisfiability and finite satisfiability problems are undecidable for the constant-free* GF^3[1-$\mathcal{EQ}^{\subseteq}$].

Proof. The satisfiability and finite satisfiability problems are undecidable for the Gödel class with identity, i.e., sentences in the form $\forall x, y\, \exists z\, \psi(x, y, z)$, where ψ is quantifier-free and without constant symbols, but may use equality (see [14, 7]).

A Gödel-class sentence $\varphi = \forall x, y\, \exists z\, \psi(x, y, z)$ embeds into GF^3[1-$\mathcal{EQ}^{\subseteq}$] as follows:

$$\forall x, y\, \big(xE_1y \to \exists z\, \big(G(x, y, z) \wedge xE_1z \wedge \psi(x, y, z)\big)\big),$$

where E_1 is the distinguished equivalence symbol and G is a fresh symbol serving as a dummy guard.

The correctness of this reduction is evident: since quantifiers are relativised by E_1, the structure on every equivalence class is a model of φ. For converse, it suffices to take any model of φ and interpret E_1 and G as full relations. ◀

In the rest of the paper, we therefore focus on the equality-free fragments of GF[$\mathcal{EQ}^{\subseteq}$] and GF[K-$\mathcal{EQ}^{\subseteq}$]. As we will see, excluding equality is sufficient to restore the finite model property and decidability of satisfiability.

Lower bound. Our first main contribution is to show that both the minimal model size and the complexity of reasoning grow drastically in GF^3[$\mathcal{EQ}^{\subseteq}$]. More precisely, each additional distinguished equivalence relation may contribute an extra exponential layer to the size of minimal models. For every $K, n \in \mathbb{N}$, we construct a satisfiable sentence $\varphi_{K,n}$ in GF^3[K-$\mathcal{EQ}^{\subseteq}$] whose length is polynomial in K and n, yet every model of $\varphi_{K,n}$ has domain size at least

$$\mathfrak{t}(K, n) = 2^{\cdot^{\cdot^{\cdot^{2^n}}}},$$

where the tower of exponentials contains K occurrences of 2.

We then use these formulas to implement succinct "large counters" in complexity lower-bound constructions, obtaining (K+1)-ExpTime-hardness of satisfiability for GF^3[K-$\mathcal{EQ}^{\subseteq}$]. Subsequently, we lift this lower bound to (K+2)-ExpTime-hardness for GF[K-$\mathcal{EQ}^{\subseteq}$], i.e., for the unrestricted-variable fragment with constants. When the number of distinguished equivalence symbols is unbounded, i.e., in GF[$\mathcal{EQ}^{\subseteq}$], the satisfiability problem becomes non-elementary; more precisely, it becomes Tower-hard.

▶ **Theorem 3.** *The satisfiability problem for the constant-free, equality-free* $\mathrm{GF}^3[\mathcal{EQ}^{\subseteq}]$ *is* TOWER*-hard.*

For every $K \geq 1$*, the satisfiability problem for the constant-free, equality-free* $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ *is* $(K{+}1)$*-*EXPTIME*-hard.*

For every $K \geq 1$*, the satisfiability problem for the equality-free* $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ *is* $(K{+}2)$*-*EXPTIME*-hard.*

Upper bound. To establish decidability of satisfiability, we first show that the finite model property holds. Moreover, we derive tight bounds on the size of minimal models.

▶ **Theorem 4.** *The equality-free* $\mathrm{GF}[\mathcal{EQ}^{\subseteq}]$ *has the finite model property.*

For every $K \geq 1$*, the equality-free* $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ *has the finite model property with a* $(K{+}2)$*-exponential upper bound on the size of minimal models. This bound decreases to* $(K{+}1)$*-exponential whenever either*

- *constants are disallowed, or*
- *the number of variables* $m \geq 3$ *is fixed.*

The finite model property alone does not, however, yield complexity upper bounds that match the lower bounds of Theorem 3. Indeed, a naïve "guess-and-verify" approach would incur nondeterministic complexity for the fragments $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ parametrised by K. We therefore design a dedicated decision procedure for satisfiability and obtain the following result, complementing Theorem 3 with matching upper bounds.

▶ **Theorem 5.** *The satisfiability problem for the equality-free* $\mathrm{GF}[\mathcal{EQ}^{\subseteq}]$ *is* TOWER*-complete.*

For every $K \geq 1$*, the satisfiability problem for the equality-free* $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ *is* $(K{+}2)$*-*EXPTIME*-complete. The complexity decreases to* $(K{+}1)$*-*EXPTIME*-complete whenever either*

- *constants are disallowed, or*
- *the number of variables* $m \geq 3$ *is fixed.*

Equivalences without nesting. Kieroński and Otto [20] showed that GF^2 is undecidable when extended with three *independent* – that is, not necessarily nested – equivalence relations, but remains decidable with at most two of them. We complement this result by showing that GF^3 becomes undecidable when extended with just two independent equivalence relations.

▶ **Theorem 6.** *The satisfiability and finite satisfiability problems are undecidable for the constant-free, equality-free* GF^3 *extended with two distinguished binary symbols* E *and* F*, each interpreted as an equivalence relation, but with no nesting conditions between them.*

Overview of technical sections.

Section 2. Introduces basic notions and definitions.

Section 3. Discusses application of our results in ontology languages. We propose a decidable extension of the description logic $\mathcal{ALCHI}$ equipped with nested equivalence roles.

Section 4. Establishes non-elementary lower bounds on the satisfiability problem (Theorem 3).
Section 5. Establishes the finite model property and the decidability of satisfiability. This section also presents the decision procedure for the fragments $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ parametrised by $K \geq 1$ (Theorems 4 and 5).

Section 6. Discusses related work.

Section 7. Concludes the paper with observations concerning slight extensions of our results.

Full version. Part of the material is delegated to the full version of the paper, available on arXiv. In particular, the undecidability result for GF^3 with two independent equivalence relations (Theorem 6) is entirely deferred there.

2 Preliminaries

We denote the set of natural numbers with 0 by $\mathbb{N}$. For $k \in \mathbb{N}$, the notation $[k]$ stands for the set $\{1, \ldots, k\}$, with the convention that $[0] = \varnothing$. More generally, we use interval notation $[a, b] \subseteq \mathbb{N}$ to denote the set $\{a, a+1, \ldots, b\}$ whenever $a \leq b$, and the empty set $\varnothing$ whenever $a > b$. We denote by 2^S the powerset of a set S. If $E \subseteq S \times S$ is an equivalence relation, then $E[a]$ is the equivalence class of $a \in S$, and S'/E is the quotient set of a subset $S' \subseteq S$.

For $k \in \mathbb{N}$ and $n \in \mathbb{N}$, the notation $\mathfrak{t}(k, n)$ stands for *tetration*, defined inductively: $\mathfrak{t}(0, n) := n$ and $\mathfrak{t}(k+1, n) := 2^{\mathfrak{t}(k,n)}$. Whenever we say that y (a quantity) is *k-exponential* in x (a parameter), then we really mean that $y \leq \mathfrak{t}(k, p(x))$ for some *fixed* polynomial p.

First-Order Logic. We assume general familiarity with First-Order Logic (FO). The logical symbols are $=, \bot, \top, \vee, \wedge, \neg, \to, \leftrightarrow$, and the quantifiers $\forall, \exists$. Formulas may also use non-logical symbols: relation symbols (from a countably infinite set), constant symbols (also from a countably infinite set), and variables (again countably many). Every relation symbol comes with associated arity. We do not allow function symbols of positive arity.

Let φ be a first-order formula. The *length* of φ, denoted $|\varphi|$, is the total number of symbols it contains, where each occurrence of a symbol – be it a variable, relation symbol, or constant – contributes 1. We use $\mathrm{fv}(\varphi)$ to denote the set of *free* variables of φ; and write $\varphi(\bar{x})$ to denote $\mathrm{fv}(\varphi) \subseteq \bar{x}$. A *signature* σ of φ is the finite set of relation and constant symbols that appear in φ. For $K \geq 1$, if $\varphi \in \mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$, then we assume that $\{E_1, \ldots, E_K\} \subseteq \sigma$ but $E_\ell \notin \sigma$ for all $\ell > K$, where $E_1, E_2, \ldots$ is the family of distinguished binary predicates.

We use Fraktur letters such as $\mathfrak{A}, \mathfrak{B}, \ldots$ to denote structures, and the corresponding Roman letters $A, B, \ldots$ for their domains. A *σ-structure* $\mathfrak{A}$ interprets the symbols from σ: a relation symbol R as a relation $R^{\mathfrak{A}} \subseteq A^k$ with k denoting the arity of R; and a constant symbol c as an element $c^{\mathfrak{A}} \in A$. If $B \subseteq A$, we write $\mathfrak{A} \restriction B$ for the *restriction* of $\mathfrak{A}$ to the subdomain B; to remain a σ-structure, B must include all interpretations of constants.

3 Application in Ontology Languages

To illustrate the relevance of our results, we discuss a case study inspired by real-life access control policies, a topic naturally arising in knowledge representation and reasoning.

Example (Rule-Based Access Control). Consider an IT system involving administrators, users, and documents. Each entity is affiliated with exactly one organisation, and every organisation is partitioned into departments. Administrators may manage documents within their departments, and users may request to download documents. To enforce data confidentiality, the following access policy is adopted: *A user may download a document only if an administrator who manages the document has granted the user permission, provided that the user and the administrator belong to the same organisation.*

The described scenario can be formalised in $\mathrm{GF}[2\text{-}\mathcal{EQ}^{\subseteq}]$ as follows. Let E_1 relate entities belonging to the same department, and let E_2 relate entities within the same organisation, where E_2 is coarser than E_1. The policy is then captured by the following axioms:

$$\forall a, u \, \big(\mathrm{grant}(a, u) \to \big(\mathrm{Admin}(a) \wedge \mathrm{User}(u)\big)\big) \tag{1}$$

$$\forall a, d \, \big(\mathrm{manages}(a, d) \to \big(\mathrm{Admin}(a) \wedge \mathrm{Doc}(d)\big)\big) \tag{2}$$
$$\forall u, d \, \big(\mathrm{access}(u, d) \to \big(\mathrm{User}(u) \wedge \mathrm{Doc}(d)\big)\big) \tag{3}$$
$$\forall a, u \, \big(\mathrm{grant}(a, u) \to a E_2 u\big) \tag{4}$$
$$\forall a, d \, \big(\mathrm{manages}(a, d) \to a E_1 d\big) \tag{5}$$
$$\forall u, d \, \big(\mathrm{access}(u, d) \to \exists a \, \big(\mathrm{grant}(a, u) \wedge \mathrm{manages}(a, d)\big)\big) \tag{6}$$

In (6) the quantifier $\exists a$ is left unguarded; however, w.r.t. satisfiability, one can introduce a fresh ternary symbol Guard and put there $\mathrm{Guard}(u, d, a)$ to serve as a dummy guard.

Extending Description Logics. We adopt a common terminology of *description logics* (DLs); for a gentle introduction to the topic, see, e.g., [2].

The sentences (1)–(6) use only unary and binary predicates. Nevertheless, sentence (6) is expressible neither in FO^2 nor in standard description logics from the $\mathcal{ALC}$ family, including expressive formalisms such as $\mathcal{SHI}$ that can define nested equivalence relations. To integrate such properties into the DL framework, we propose an *ad-hoc* extension of the logic $\mathcal{ALCHI}$.

The standard constructs of $\mathcal{ALCHI}$ on top of $\mathcal{ALC}$ are *role inclusions* ($\mathcal{H}$) and *role inverses* ($\mathcal{I}$) – both already expressible in the equality-free GF^2. The unique features added here are *nested equivalence roles* and *pairwise existential dependencies*.

Nested equivalence roles. The set of role names contains a distinguished family $E_1, E_2, \ldots$. In every interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, these roles are interpreted so that:

- for every $k \geq 1$, $E_k^{\mathcal{I}}$ is an equivalence relation over $\Delta^{\mathcal{I}}$;
- for every $k \geq 1$, $E_{k+1}^{\mathcal{I}}$ is coarser than $E_k^{\mathcal{I}}$, i.e., $E_k^{\mathcal{I}} \subseteq E_{k+1}^{\mathcal{I}}$.

Pairwise existential dependencies. Let r, p, q be role names or their inverses. A *pairwise existential dependency* is an axiom of the form

$$r \sqsubseteq \exists(p, q) \tag{7}$$

An interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ satisfies (7) whenever, for every pair $(a, b) \in r^{\mathcal{I}}$, there exists some $c \in \Delta^{\mathcal{I}}$ such that $(a, c) \in p^{\mathcal{I}}$ and $(b, c) \in q^{\mathcal{I}}$. Intuitively, such axioms link pairs of r-related individuals to a common witness via p and q. This construct goes beyond FO^2 as well as the core description logics of the $\mathcal{ALC}$ family. This additional expressivity is necessary to capture many real-life properties such as (6).

Having the discussed above features, the sentences (1)–(6) can be expressed concisely as:

$$\begin{array}{llllll} \mathrm{grant} & \sqsubseteq \mathrm{Admin} \times \mathrm{User} & \mathrm{manages} & \sqsubseteq \mathrm{Admin} \times \mathrm{Doc} & \mathrm{access} & \sqsubseteq \mathrm{User} \times \mathrm{Doc} \\ \mathrm{grant} & \sqsubseteq E_2 & \mathrm{manages} & \sqsubseteq E_1 & \mathrm{access} & \sqsubseteq \exists(\mathrm{grant}^{-1}, \mathrm{manages}^{-1}) \end{array}$$

From our results it follows that the knowledge-base consistency problem for the outlined extension of $\mathcal{ALCHI}$ is decidable: given a knowledge base $\mathcal{K}$ (i.e., a finite set of axioms), the task is to decide whether there exists an interpretation $\mathcal{I}$ such that $\mathcal{I} \models \mathcal{K}$ (i.e., $\mathcal{I}$ satisfies every axiom from $\mathcal{K}$). In addition to decidability, we obtain that general and finite consistency coincide: the interpretation $\mathcal{I}$ can be chosen to be finite.

4 Non-Elementary Lower Bound

In this section, we first show that the satisfiability problem for the constant-free, equality-free $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ is $(K{+}1)$-ExpTime-hard (Subsections 4.1 and 4.2). We then extend the lower-bound construction to obtain $(K{+}2)$-ExpTime-hardness for equality-free $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ with constants and unboundedly many variables (Subsection 4.3). This establishes Theorem 3.

4.1 Implementing Large Counters

Fix $K, n \in \mathbb{N}$. For each $k \in [K]$, an integer in the range $[0, \mathfrak{t}(k,n) - 1]$ is called a *k-numeral*.

As a first step in the lower-bound construction, we define a sentence $\varphi_{K,n}$ in the constant-free, equality-free $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ that describes a counter over $\mathfrak{t}(K,n)$ elements. The sentence $\varphi_{K,n}$ is given as the conjunction $\varphi^{(1)}_{K,n} \wedge \varphi^{(2)}_{K,n} \wedge \cdots \wedge \varphi^{(K)}_{K,n}$, realising the following idea.

- Using n unary predicates, the sentence $\varphi^{(1)}_{K,n}$ assigns to each element a 1-numeral, and enforces that every E_1-class contains 2^n elements (at least one for each 1-numeral).
- Treating domain elements as bits, the sentence $\varphi^{(2)}_{K,n}$ assigns to each E_1-class a 2-numeral, and requires every E_2-class to contain 2^{2^n} distinct E_1-classes (one for each 2-numeral).
- Treating entire E_1-classes as bits, the sentence $\varphi^{(3)}_{K,n}$ assigns to each E_2-class a 3-numeral, and requires every E_3-class to contain $2^{2^{2^n}}$ distinct E_2-classes (one for each 3-numeral).
- Proceeding inductively, for each $k \geq 2$, the sentence $\varphi^{(k)}_{K,n}$ assigns a k-numeral to every E_{k-1}-class, and enforces every E_k-class to contain $\mathfrak{t}(k,n)$ distinct E_{k-1}-classes.
- Consequently, every E_K-class contains $\mathfrak{t}(K,n)$ distinct E_{K-1}-classes, each representing a different K-numeral.
- Finally, the length of $\varphi_{K,n}$ is polynomial in n and K; in fact, $|\varphi_{K,n}| = \mathcal{O}(n+K)$.

The described idea of "nested counters" is certainly not new, as similar ones have already appeared, e.g., in Stockmeyer's proof concerning FO on words [29], or in the non-elementary lower bound for the Fluted Fragment [25]. The non-trivial part, however, is to express required properties using just the quite limited syntax of GF^3.

Signature. The signature of $\varphi_{K,n}$ consists of the following predicates.

- The distinguished nested equivalence symbols $E_1, \ldots, E_K$.
- An additional symbol E_0 to be used as a sentinel. A priori it is a free symbol; however, we later axiomatise it to behave like the finest equivalence relation. Thus, we form an extended family of nested equivalence relations $E_0, E_1, \ldots, E_K$.
- Unary symbols BB_i and BC_i for $i \in [0, n-1]$ to represent bit and carry values of 1-numerals, respectively.
- Unary symbols B_k and C_k for $k \in [2, K]$ to represent bit and carry values of k-numerals, respectively.
- Binary symbols S_k for $k \in [K]$ to connect (elements of) E_{k-1}-classes that represent circularly successive k-numerals.
- Unary symbols Z_k for $k \in [K]$ to decorate (elements of) E_{k-1}-classes that represent the k-numeral 0.
- Binary symbols Q_k for $k \in [K]$ to connect (elements of) E_{k-1}-classes that represent equal k-numerals.
- Additional symbols serving auxiliary roles, e.g., as guards; not listed here.

Interpretation. The sentence $\varphi_{K,n}$ will be constructed so that every model $\mathfrak{A} \models \varphi_{K,n}$ will satisfy the following conditions.

First, the sentinel symbol E_0 shall be interpreted as an equivalence relation which is finer than $E_1^{\mathfrak{A}}$. In particular, all below quotients via E_0 are well-defined.

For every $k \in [K]$ and every E_k-class $\mathcal{C} \in A/E_k^{\mathfrak{A}}$, we shall have

$$\left|\mathcal{C}/E_{k-1}^{\mathfrak{A}}\right| = \mathfrak{t}(k,n).$$

To formally state the intended interpretation of the predicates S_k, Z_k, and Q_k, we introduce a family of functions $\mathfrak{n}_k \colon A \to [0, \mathfrak{t}(k,n){-}1]$ for $k \in [K]$ (i.e., the kth function assigns k-numerals to domain elements). These functions are defined inductively:

- The first function $\mathfrak{n}_1$ on element $a \in A$ is given by

$$\mathfrak{n}_1(a) := \sum\nolimits_{i=0}^{n-1} b_i \cdot 2^i,$$

  where $b_i = 1$ if $\mathfrak{A} \models BB_i(a)$; and otherwise $b_i = 0$.
- For $k \in [2, K]$, the function $\mathfrak{n}_k \colon A \to [0, \mathfrak{t}(k,n) - 1]$ is given by mapping $a \in A$ to

$$\mathfrak{n}_k(a) := \sum\nolimits_{i=0}^{\mathfrak{t}(k-1,n)-1} b_i \cdot 2^i,$$

  where $b_i = 1$ if $\mathfrak{A} \models B_k(b)$ for some $b \in E_{k-1}^{\mathfrak{A}}[a]$ with $\mathfrak{n}_{k-1}(b) = i$; and otherwise $b_i = 0$.

Now, for every $k \in [K]$ and every pair $(a,b) \in E_K^{\mathfrak{A}}$, the following shall hold:

- $\mathfrak{A} \models E_{k-1}(a,b)$ iff $(a,b) \in E_k^{\mathfrak{A}}$ and $\mathfrak{n}_k(a) = \mathfrak{n}_k(b)$.
- $\mathfrak{A} \models S_k(a,b)$ iff $\mathfrak{n}_k(a) + 1 \equiv \mathfrak{n}_k(b) \mod \mathfrak{t}(k,n)$.
- $\mathfrak{A} \models Z_k(a)$ iff $\mathfrak{n}_k(a) = 0$.
- $\mathfrak{A} \models Q_k(a,b)$ iff $\mathfrak{n}_k(a) = \mathfrak{n}_k(b)$.

Note that every E_K-class is supposed to describe an independent K-exponential counter. In particular, we do not impose any constraints on pairs $(a,b) \notin E_K^{\mathfrak{A}}$.

Convention. In the presentation that follows, we occasionally write formulas that do not match Definition 1 (the syntax of GF), but are equivalent to guarded ones modulo simple rewritings. We also omit guards for existential quantifiers, since one may always introduce fresh symbols to act as dummy guards.

Base conjunct. We now construct the first sentence $\varphi_{K,n}^{(1)}$.

To work with 1-numerals, i.e., integers in $[0, 2^n{-}1]$, we introduce auxiliary formulas implementing cyclic successor (i.e., successor modulo 2^n), zero test, and equality test.

Recall the standard algorithm for incrementing a number m modulo a power of two. We first locate the least significant zero bit of m, say at position i, and then flip every bit $j \leq i$. If all bits of m are already equal to 1, then every bit is flipped. Equivalently, the increment operation can be seen as a bitwise XOR with the *carry bits*. For example, if $m = 1010111$, then the carry bits are 0001111, and the bitwise XOR produces the successor 1011000.

We implement increments as described above. Axiomatise first the carry bits:

$$\forall x \left(BC_0(x) \wedge \bigwedge\nolimits_{i=0}^{n-2} \big(BC_{i+1}(x) \leftrightarrow \big(BB_i(x) \wedge BC_i(x) \big) \big) \right) \tag{8}$$

Define then a formula testing whether two elements represent circularly successive 1-numerals:

$$\varphi_S(x_1, x_2) \;:=\; \bigwedge\nolimits_{i=0}^{n-1} \big(BB_i(x_2) \leftrightarrow \big(BB_i(x_1) \oplus BC_i(x_1) \big) \big)$$

The formulas for zero and equality tests are straightforward:

$$\varphi_Z(x) \;:=\; \bigwedge\nolimits_{i=0}^{n-1} \neg BB_i(x) \quad \text{and} \quad \varphi_Q(x_1, x_2) \;:=\; \bigwedge\nolimits_{i=0}^{n-1} \big(BB_i(x_1) \leftrightarrow BB_i(x_2)\big)$$

We now axiomatise the symbols S_1, Z_1, and Q_1; and enforce the existence of successors:

$$\forall x_1, x_2 \, \big(x_1 E_K x_2 \to \big(S_1(x_1, x_2) \leftrightarrow \varphi_S(x_1, x_2)\big)\big) \tag{9}$$
$$\forall x \, \big(Z_1(x) \leftrightarrow \varphi_Z(x)\big) \tag{10}$$
$$\forall x_1, x_2 \, \big(x_1 E_K x_2 \to \big(Q_1(x_1, x_2) \leftrightarrow \varphi_Q(x_1, x_2)\big)\big) \tag{11}$$
$$\forall x \, \exists y \, \big(S_1(x, y) \wedge x E_1 y\big) \tag{12}$$

Since S_1 is cyclic, inside every E_1-class all 2^n elements will be eventually generated, regardless of the initial value.

We make the sentinel symbol E_0 act as the finest equivalence relation:

$$\forall x_1, x_2 \, \big(x_1 E_0 x_2 \to x_1 E_1 x_2\big) \tag{13}$$
$$\forall x_1, x_2 \, \big(x_1 E_1 x_2 \to \big(x_1 E_0 x_2 \leftrightarrow Q_1(x_1, x_2)\big)\big) \tag{14}$$

Define $\varphi^{(1)}_{K,n}$ as the conjunction of sentences (8)–(14).

Inductive conjuncts. Let $k \in [2, K]$. Suppose that the sentences $\varphi^{(1)}_{K,n}, \ldots, \varphi^{(k-1)}_{K,n}$ have already been defined. We now construct the next sentence $\varphi^{(k)}_{K,n}$, which axiomatises the predicates S_k, Z_k, and Q_k. As a result, each E_{k-1}-class is associated with a k-numeral, that is, an integer in the range $[0, \mathfrak{t}(k, n) - 1]$.

Bits and carry bits. We treat the entire E_{k-2}-classes as individual bits, where the respective bit values are given by the predicate B_k. We require that elements of every E_{k-2}-class are assigned consistent bit values:

$$\forall x_1, x_2 \, \big(x_1 E_{k-2} x_2 \to \big(B_k(x_1) \leftrightarrow B_k(x_2)\big)\big) \tag{15}$$

We next axiomatise the carry bits. First, the carry bit of the least significant position is always on:

$$\forall x \, \big(Z_{k-1}(x) \to C_k(x)\big) \tag{16}$$

We then propagate the carry until we encounter the first bit off or reach the last bit position; since S_{k-1} is cyclic, the latter is detected by $\neg Z_{k-1}(x_2)$:

$$\forall x_1, x_2 \, \Big(\big(x_1 E_{k-1} x_2 \wedge S_{k-1}(x_1, x_2) \wedge \neg Z_{k-1}(x_2)\big) \to \big(C_k(x_2) \leftrightarrow \big(C_k(x_1) \wedge B_k(x_1)\big)\big)\Big) \tag{17}$$

Successors. We next ensure that whenever S_k connects two elements, then their E_{k-1}-classes indeed represent consecutive k-numerals.

First we ensure the following property: whenever S_k connects a pair of elements across two E_{k-1}-classes, say $\mathcal{C}$ and $\mathcal{C}'$, then necessarily S_k connects a pair of elements across every E_{k-2}-classes $\mathcal{D} \in \mathcal{C}/E_{k-2}$ and $\mathcal{D}' \in \mathcal{C}'/E_{k-2}$. This property translates to:

$$\forall x_1, x_2 \, \big(S_k(x_1, x_2) \to \exists y \, \big(S_{k-1}(x_1, y) \wedge x_1 E_{k-1} y \wedge S_k(y, x_2)\big)\big) \tag{18}$$
$$\forall x_1, x_2 \, \big(S_k(x_1, x_2) \to \exists y \, \big(S_{k-1}(x_2, y) \wedge x_2 E_{k-1} y \wedge S_k(x_1, y)\big)\big) \tag{19}$$

We next apply the bitwise XOR with the respective carry bits:

$$\forall x_1, x_2 \, \big(\big(S_k(x_1, x_2) \wedge Q_{k-1}(x_1, x_2) \big) \rightarrow \chi(x_1, x_2) \big) \tag{20}$$

where $\chi(x_1, x_2) := B_k(x_2) \leftrightarrow \big(B_k(x_1) \oplus C_k(x_1) \big)$.

So far, whenever S_k links any two elements, then their E_{k-1}-classes indeed represent successive k-numerals. Though, we did not actually enforced any S_k-connections. To fix this, we require that whenever S_k does not link some two elements, then their E_{k-1}-classes do not represent successive k-numerals, i.e., there is a position where the respective bit and carry-bit values disagree. With 4 variables, this property is expressed by the sentence:

$$\forall x_1, x_2 \, \Big(\big(x_1 E_K x_2 \wedge \neg S_k(x_1, x_2) \big) \rightarrow \\ \exists y_1, y_2 \, \big(x_1 E_{k-1} y_1 \wedge x_2 E_{k-1} y_2 \wedge Q_{k-1}(y_1, y_2) \wedge \neg \chi(y_1, y_2) \big) \Big)$$

Using a fresh binary symbol W_k, we rewrite the above sentence into a pair of three-variable sentences:

$$\forall x_1, x_2 \, \big(\big(x_1 E_K x_2 \wedge \neg S_k(x_1, x_2) \big) \rightarrow \exists y_1 \, \big(x_1 E_{k-1} y_1 \wedge W_k(y_1, x_2) \big) \big) \tag{21}$$

$$\forall y_1, x_2 \, \big(W_k(y_1, x_2) \rightarrow \exists y_2 \, \big(x_2 E_{k-1} y_2 \wedge Q_{k-1}(y_1, y_2) \wedge \neg \chi(y_1, y_2) \big) \big) \tag{22}$$

We next postulate the existence of cyclic successors:

$$\forall x \, \exists y \, \big(S_k(x, y) \wedge x E_k y \big) \tag{23}$$

Zeros. We now decorate with Z_k every E_{k-1}-class that represents the k-numeral 0. First, when Z_k holds, every B_k is off:

$$\forall x_1, x_2 \, \big(\big(x_1 E_{k-1} x_2 \wedge Z_k(x_1) \big) \rightarrow \neg B_k(x_2) \big) \tag{24}$$

For converse, if Z_k does not hold, then some B_k is on:

$$\forall x \, \big(\neg Z_k(x) \rightarrow \exists y \, \big(x E_{k-1} y \wedge B_k(y) \big) \big) \tag{25}$$

At this point, every E_k-class decomposes into at least $\mathsf{t}(k, n)$ inner E_{k-1}-classes. The E_{k-1}-classes representing circularly successive k-numerals are linked by S_k, provided that they are contained in the same E_K-class. Finally, the E_{k-1}-classes representing the k-numeral 0 are marked by Z_k. It remains therefore to axiomatise the kth level equality predicate Q_k.

Equality. The predicate Q_k shall connect E_{k-1}-classes representing the same k-numeral, provided that these classes are contained in the same E_K-class. We use the fact that a common S_k-successor exists precisely when the k-numerals agree:

$$\forall x_1, x_2 \, \Big(x_1 E_K x_2 \rightarrow \exists y \, \big(S_k(x_1, y) \wedge \big(Q_k(x_1, x_2) \leftrightarrow S_k(x_2, y) \big) \big) \Big) \tag{26}$$

Having the predicate Q_k axiomatised, we enforce that in every E_k-class each k-numeral is represented by a unique E_{k-1}-class:

$$\forall x_1, x_2 \, \big(\big(x_1 E_k x_2 \wedge Q_k(x_1, x_2) \big) \rightarrow x_1 E_{k-1} x_2 \big) \tag{27}$$

Finally, define $\varphi^{(k)}_{K,n}$ as the conjunction of (15)–(27).

4.2 Complexity Lower Bound

Using the sentences $\varphi_{K,n}$ constructed in Subsection 4.1, we encode accepting runs of alternating Turing machines operating in K-exponential space. The desired lower bound for the constant-free, equality-free $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ follows then from the result of Chandra, Kozen, and Stockmeyer [10], stating that K-AExpSpace $=$ $(K{+}1)$-ExpTime.

Alternating machines. Let $K \in \mathbb{N}$. Suppose that $\mathbf{M}$ is an alternating Turing machine that operates in a space bound $\mathfrak{t}(K, p(n))$ for some polynomial p.

The machine $\mathbf{M}$ is defined by the finite set of *states* S with a distinguished *initial state* $s_0 \in S$; the finite *tape alphabet* Γ; the *transition functions* δ_1 and δ_2; and a partition of S into *universal*, *existential*, *accepting*, and *rejecting* states.

Let $m = \mathfrak{t}(K, p(|x|))$ denote the space bound for the input x. A *configuration* of $\mathbf{M}$ is a complete description of the tape with m cells, the head position, and the current state. The *run* of $\mathbf{M}$ on input x is represented as a directed tree T with a distinguished root r. Each node u of T is labelled by a configuration $C(u)$, where $C(r)$ is the *initial* configuration. Every non-leaf node has exactly two children. For each edge $v \to u$, the configuration $C(u)$ is obtained from $C(v)$ by a single deterministic transition step: the left child applies the transition rule δ_1, and the right child applies δ_2. W.l.o.g. only the initial configuration $C(r)$ has state s_0 and a sequence of configurations $C(v)$ on any path from the root to a leaf is without repetitions; in particular, by the space bound, the tree T is finite.

Leaves of T carry accepting or rejecting states, while internal nodes carry existential or universal ones. We define *positivity* inductively: a leaf is positive if its state is accepting; a node with an existential state is positive if at least one of its children is positive; a node with a universal state is positive if all its children are positive. Finally, the machine *accepts* x if the root r is positive.

In the following, we encode the run of $\mathbf{M}$ on an input string x: we construct a $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ sentence $\Theta_{\mathbf{M},x}$ such that $\mathbf{M}$ accepts x iff $\Theta_{\mathbf{M},x}$ is satisfiable.

Signature. The signature of $\Theta_{\mathbf{M},x}$ consists of the following symbols:

- The distinguished symbols $E_1, \dots, E_K$, where every E_K-class models a single configuration with the inner E_{K-1}-classes representing individual tape cells.
- The binary symbol S_K and the unary symbol Z_K, which are shared with the sentence $\varphi_{K,n}$ (where $n := p(|x|)$) from Subsection 4.1.
- Unary symbols P_a for $a \in \Gamma$ and P_s for $s \in S$.
- Unary symbols $P_>$ and $P_<$, marking the head being located to the left and to right, respectively.
- Binary symbols T_i, T_i^L, T_i^R for $i \in \{1, 2\}$ for modelling δ_i-transitions, where T_i links the jth cell of a configuration to the jth cell of a successive configuration, T_i^L links to the $(j{-}1)$th cell, and T_i^R links to the $(j{+}1)$th cell.
- A symbol F to mark positive nodes in the computation tree.

Configurations. Every E_K-class shall represent a single configuration of $\mathbf{M}$: the inner E_{K-1}-classes correspond to the individual tape cells; the sentence $\varphi_{K,n}$ from Subsection 4.1 enforces that the correct number of such cells is present and arranges them in a linear order. The symbol written in each cell is indicated by the unary predicates P_a (for $a \in \Gamma$), which are made E_{K-1}-universal. The current state s and the head position are encoded analogously: exactly one E_{K-1}-class is marked with P_s, identifying the cell scanned by the head.

A typical configuration C is represented schematically as

$$\begin{pmatrix} a_1 & a_2 & \dots & a_{i-1} & a_i & a_{i+1} & \dots & a_{m-1} & a_m \\ < & < & \dots & < & s_i & > & \dots & > & > \end{pmatrix}$$

Here $a_j \in \Gamma$ denotes the tape symbol in cell j, $s_i \in S$ is the machine state, and the markers $<$ and $>$ indicate whether a cell is located to the left or to the right of the head.

Let $\mathfrak{n}_K$ be the function assigning K-numerals as in Subsection 4.1. Formally, an E_K-class $\mathcal{C}$ models the configuration C if, for every element $b \in \mathcal{C}$, the following conditions hold:

- for each $a \in \Gamma$, $\mathfrak{A} \models P_a(b)$ iff $a = a_{\mathfrak{n}_K(b)}$;
- for each $s \in S$, $\mathfrak{A} \models P_s(b)$ iff $\mathfrak{n}_K(b) = i$ and $s = s_i$;
- for each $\bowtie \in \{<, >\}$, $\mathfrak{A} \models P_{\bowtie}(b)$ iff $\mathfrak{n}_K(b) \bowtie i$.

We now express that every E_K-class models some configuration.

We first enforce that each element carries some tape symbol and some state marker:

$$\forall x \bigvee\nolimits_{a \in \Gamma} P_a(x) \quad \wedge \quad \forall x \bigvee\nolimits_{s \in S \cup \{<,>\}} P_s(x) \tag{28}$$

To guarantee that every E_{K-1}-class encodes exactly one tape cell, we express that no two distinct elements of the same E_{K-1}-class may carry different tape symbols or different state markers:

$$\forall x_1, x_2 \left(x_1 E_{K-1} x_2 \rightarrow \bigwedge\nolimits_{a,a' \in \Gamma \,:\, a \neq a'} \big(\neg P_a(x_1) \vee \neg P_{a'}(x_2)\big) \right) \tag{29}$$

$$\forall x_1, x_2 \left(x_1 E_{K-1} x_2 \rightarrow \bigwedge\nolimits_{s,s' \in S \cup \{<,>\} \,:\, s \neq s'} \big(\neg P_s(x_1) \vee \neg P_{s'}(x_2)\big) \right) \tag{30}$$

Next, we propagate the markers $<$ and $>$ along the successor relation S_K to enforce that exactly one cell carries the state symbol (that is, the cell corresponding to the head position):

$$\forall x_1, x_2 \left(\Big(S_K(x_1, x_2) \wedge \neg Z_K(x_2) \wedge \bigvee\nolimits_{s \in S \cup \{>\}} P_s(x_1) \Big) \rightarrow P_>(x_2) \right) \tag{31}$$

$$\forall x_1, x_2 \left(\Big(S_K(x_1, x_2) \wedge \neg Z_K(x_2) \wedge \bigvee\nolimits_{s \in S \cup \{<\}} P_s(x_2) \Big) \rightarrow P_<(x_1) \right) \tag{32}$$

Above, Z_K is used to detect whether S_K cycles back to zero, or connects two elements successive without cyclicity (i.e., $S_K(x_1, x_2) \wedge \neg Z_K(x_2)$ holds when the cells are consecutive).

Finally, we constrain the boundary cells and rule out adjacent $<$ and $>$ markers. That is, the leftmost cell cannot carry $>$; the rightmost cannot carry $<$; and $<$ cannot be immediately followed by $>$ (i.e., there must be the head between them):

$$\forall x_1, x_2 \left(\big(S_K(x_1, x_2) \wedge Z_K(x_2) \big) \rightarrow \big(\neg P_<(x_1) \wedge \neg P_>(x_2) \big) \right) \tag{33}$$

$$\forall x_1, x_2 \left(\big(S_K(x_1, x_2) \wedge \neg Z_K(x_2) \big) \rightarrow \big(\neg P_<(x_1) \vee \neg P_>(x_2) \big) \right) \tag{34}$$

Here, Z_K is again used to detect the boundary of the tape.

Transitions. Let $\mathcal{C}$ and $\mathcal{C}'$ be two E_K-classes modelling configurations C and C', respectively. We axiomatise that whenever a relation T_i links (elements of) $\mathcal{C}$ to (elements of) $\mathcal{C}'$, then C' must be the δ_i-successor of C.

First, for every tape position $j \in [0, m-1]$, there must exist elements $a \in \mathcal{C}$ and $b \in \mathcal{C}'$ such that $\mathfrak{n}_K(a) = j$, $\mathfrak{n}_K(b) = j$, and $\mathfrak{A} \models T_i(a, b)$. Analogous constraints are imposed for the shifted relations T_i^L and T_i^R: for T_i^L we require $\mathfrak{n}_K(b) = j - 1$, and for T_i^R we require $\mathfrak{n}_K(b) = j + 1$; the conditions are modulo m.

We enforce these constraints using the following formulas:

$$\forall x_1, x_2 \, \big(T_i(x_1, x_2) \to \big(Z_K(x_1) \leftrightarrow Z_K(x_2)\big)\big) \tag{35}$$
$$\forall x_1, x_2 \, \big(T_i(x_1, x_2) \to \exists y \, \big(S_K(x_1, y) \wedge x_1 E_K y \wedge T_i^L(y, x_2)\big)\big) \tag{36}$$
$$\forall x_1, x_2 \, \big(T_i(x_1, x_2) \to \exists y \, \big(S_K(x_2, y) \wedge x_2 E_K y \wedge T_i^R(x_1, y)\big)\big) \tag{37}$$
$$\forall x_1, x_2 \, \big(T_i^L(x_1, x_2) \to \exists y \, \big(S_K(x_2, y) \wedge x_2 E_K y \wedge T_i(x_1, y)\big)\big) \tag{38}$$
$$\forall x_1, x_2 \, \big(T_i^R(x_1, x_2) \to \exists y \, \big(S_K(x_1, y) \wedge x_1 E_K y \wedge T_i(y, x_2)\big)\big) \tag{39}$$

We require the existence of δ_i-successive configurations:

$$\bigwedge\nolimits_{i \in \{1,2\}} \forall x \, \Big(\Big(\bigvee\nolimits_{s \in S \,:\, s \text{ is universal or existential}} P_s(x)\Big) \to \exists y \, T_i(x, y)\Big) \tag{40}$$

Having the above formulas, it is routine to express transition functions of $\mathbf{M}$. For instance, to express cell content staying the same for positions not under the head, we write:

$$\forall x_1, x_2 \, \Big(\Big(T_i(x_1, x_2) \wedge \bigwedge\nolimits_{s \in S} \neg P_s(x_1)\Big) \to \bigwedge\nolimits_{a \in \Gamma} \big(P_a(x_1) \leftrightarrow P_a(x_2)\big)\Big) \tag{41}$$

Accepting condition. To express the accepting condition, we decorate E_K-classes modelling positive nodes in the computation tree with the unary predicate F.

First, require the symbol F to be E_K-universal:

$$\forall x_1, x_2 \, \big(x_1 E_K x_2 \to \big(F(x_1) \leftrightarrow F(x_2)\big)\big) \tag{42}$$

Then propagate the positivity condition through the computation tree:

$$\forall x \, \big(P_{s_0}(x) \to F(x)\big) \quad \wedge \quad \forall x \, \Big(\Big(\bigvee\nolimits_{s \in S \,:\, s \text{ is rejecting}} P_s(x)\Big) \to \neg F(x)\Big) \tag{43}$$
$$\forall x \, \Big(\Big(F(x) \wedge \bigvee\nolimits_{s \in S \,:\, s \text{ is universal}} P_s(x)\Big) \to \bigwedge\nolimits_{i \in \{1,2\}} \exists y \, \big(T_i(x, y) \wedge F(y)\big)\Big) \tag{44}$$
$$\forall x \, \Big(\Big(F(x) \wedge \bigvee\nolimits_{s \in S \,:\, s \text{ is existential}} P_s(x)\Big) \to \bigvee\nolimits_{i \in \{1,2\}} \exists y \, \big(T_i(x, y) \wedge F(y)\big)\Big) \tag{45}$$

To conclude the reduction, it remains to express the initial condition: the tape content of a configuration with the initial state s_0 shall contain the input string x. Yet, this final step is relatively straightforward, so we leave details to the reader.

4.3 Lifting Lower Bound

We now extend the lower-bound construction to obtain the $(K{+}2)$-ExpTime-hardness of the satisfiability problem for $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ (with constants and unboundedly many variables).

Recall that in Subsection 4.1, the formulas φ_S, φ_Z, and φ_Q were used to encode numbers in the range $[0, 2^n - 1]$, relying on n unary predicates BB_i to represent individual bits.

The key insight underlying the $(K{+}2)$-ExpTime-hardness of $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ is that numbers from the much larger range $[0, 2^{2^n} - 1]$ can already be encoded at the level of individual elements: let BB be a predicate of arity $n + 1$, and let 0 and 1 be constants identified with the usual numbers 0 and 1. An element $a \in A$ represents a number $m \in [0, 2^{2^n} - 1]$ if, for every sequence $j_0, \ldots, j_{n-1} \in \{0, 1\}$, we have $\mathfrak{A} \models BB(a, j_{n-1}, \ldots, j_0)$ precisely when the k-th bit of m is equal to 1, where $k = \sum_{i=0}^{n-1} j_i \cdot 2^i$. Clearly, this condition cannot be expressed directly for all sequences $j_0, \ldots, j_{n-1}$, as a central requirement is that the resulting sentences shall have length polynomial in n. We therefore employ a succinct encoding technique, whose details are deferred to the full version. A closely related construction appears in the proof of 2-NExpTime-completeness for the Triguarded Fragment [28].

Now, let $\varphi_{K,n} = \varphi^{(1)}_{K,n} \wedge \cdots \wedge \varphi^{(K)}_{K,n}$ be as constructed in Subsection 4.1. We replace the sentences (9)–(11) occurring in $\varphi^{(1)}_{K,n}$ with their succinct counterparts, yielding an exponential increase in the representable counting range. All remaining components of the lower-bound construction from Subsections 4.1 and 4.2 remain unchanged. As a result, we can encode runs of alternating Turing machines operating in $(K{+}1)$-exponential space, rather than K-exponential space. This yields $(K{+}2)$-ExpTime-hardness for the equality-free $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$.

5 Upper Bound

Having established the lower bounds, we now turn to upper bounds (Theorems 4 and 5). To avoid technical distraction unrelated to nested equivalence relations, the results in this section are established for the constant-free, equality-free $\mathrm{GF}^3[\mathcal{EQ}^{\subseteq}]$. In the full version, we provide proofs for the general case, that is, for the equality-free $\mathrm{GF}[\mathcal{EQ}^{\subseteq}]$ with constants and an unbounded number of variables.

5.1 Normal Form and Atomic Types

▶ **Assumption 7.** *For simplicity, in this section we work with the constant-free, equality-free* $\mathrm{GF}^3[\mathcal{EQ}^{\subseteq}]$. *Given a sentence* φ *of this fragment, we assume that it is in* normal form*:*

$$\Psi \quad \wedge \quad \bigwedge\nolimits_{t\in\mathcal{I}} \forall x_1, x_2 \left(G_t(x_1, x_2) \to \exists y\, \psi_t(x_1, x_2, y)\right),$$

where each G_t *is a binary guard, each* ψ_t *is quantifier-free, and* Ψ *is a conjunction of purely universal prenex sentences:*

$$\bigwedge\nolimits_{t\in\mathcal{J}} \forall x_1, \ldots, x_{k_t} \left(F_t(x_1, \ldots, x_{k_t}) \to \chi_t(x_1, \ldots, x_{k_t})\right),$$

where $1 \le k_t \le 3$, *each* F_t *is a guard of arity* k_t, *and each* χ_t *is quantifier-free.*

The guards G_t *and* F_t *are allowed to be distinguished symbols.*

Using standard techniques (see, e.g., [15, 24]), given an arbitrary GF^3-sentence, one can compute in polynomial time a sentence in normal form which is satisfiable over the same domains. This reduction remains sound also for the fragment $\mathrm{GF}^3[\mathcal{EQ}^{\subseteq}]$. Hence, for establishing decidability and the finite model property for $\mathrm{GF}^3[\mathcal{EQ}^{\subseteq}]$, Assumption 7 is w.l.o.g.

Types. Let σ be a constant-free signature. For $k \in \mathbb{N}$, let $\mathrm{Lit}_k(\sigma)$ denote the set of all literals over relation symbols from σ and variables $X_k := \{x_1, \ldots, x_k\}$; literals with "=" are not included. A *k-type* over σ is a maximal consistent subset of $\mathrm{Lit}_k(\sigma)$, that is, for every atom γ over σ and X_k, precisely one of γ, $\neg\gamma$ is present. A *type* is a k-type for some $k \ge 1$.

Given a σ-structure $\mathfrak{A}$ and k distinct elements $a_1, \ldots, a_k \in A$, we write $\mathrm{tp}^{\mathfrak{A}}[a_1, \ldots, a_k]$ for the k-type *realised* by $\langle a_1, \ldots, a_k\rangle$ in $\mathfrak{A}$. Formally, $\mathrm{tp}^{\mathfrak{A}}[a_1, \ldots, a_k]$ is the set of literals $\gamma \in \mathrm{Lit}_k(\sigma)$ such that $\mathfrak{A}, f \models \gamma$, where f sends $x_i \mapsto a_i$ for all $i \in [k]$.

It is convenient to view k-types themselves as σ-structures over the canonical domain $\{x_1, \ldots, x_k\}$. We allow structure-like notation also for types; for instance, if τ is a 3-type, we may write $\mathrm{tp}^{\tau}[x_1, x_3]$ for the 2-type induced by τ on the variables x_1 and x_3.

5.2 Finite-Index Nesting Property

Our main technical tool is the *finite-index nesting property*: namely, if a sentence φ is satisfiable, then it admits a model in which the finest distinguished relation E_1 sits inside E_2 with *finite index*, that is, every E_2-class decomposes into only finitely many E_1-classes.

Lemma 8 establishes that the finite-index nesting property holds for the constant-free, equality-free $\mathrm{GF}^3[\mathcal{EQ}^{\subseteq}]$. It also provides an upper bound on the *index* of E_1 in every E_2-class, that is, on the number of inner E_1-classes sitting in an E_2-class.

▶ **Lemma 8.** *Let $K \geq 2$, and let φ be a constant-free, equality-free sentence of $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ in normal form as in Assumption 7. Let $\boldsymbol{\alpha}$ denote the set of all 1-types over the signature of φ. If φ is satisfiable, then it has a model $\mathfrak{A} \models \varphi$ such that*

$$|\mathcal{C}/E_1^{\mathfrak{A}}| \;\leq\; 3 \cdot 2^{|\boldsymbol{\alpha}|} \qquad \text{for every} \qquad \mathcal{C} \in A/E_2^{\mathfrak{A}}.$$

The key idea is that Lemma 8 allows us to eliminate the first distinguished symbol E_1 from the logic: a finite number of equivalence classes can be fully described without relying on the externally constrained semantics of the predicate E_1. For instance, we can introduce a family of fresh unary predicates $\{U_i\}_{i=1}^n$, for an appropriate $n \in \mathbb{N}$, and decorate elements of the domain so that the following equivalence holds:

$$xE_1y \;\leftrightarrow\; \Big(xE_2y \wedge \bigwedge\nolimits_{i=1}^{n} \big(U_i(x) \leftrightarrow U_i(y)\big)\Big).$$

This transformation of models serves as a basis for a reduction from $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ to the simpler fragment $\mathrm{GF}^3[(K{-}1)\text{-}\mathcal{EQ}^{\subseteq}]$, and by iterative applications to $\mathrm{GF}^3[1\text{-}\mathcal{EQ}^{\subseteq}]$. Technical details of this reduction are developed in the subsequent Subsection 5.3, and we now focus on proving Lemma 8.

Proof of Lemma 8. Let $K \geq 2$, and let φ be a constant-free, equality-free sentence of $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ in normal form as in Assumption 7. We write σ for the signature of φ, and $\boldsymbol{\alpha}$ for the set of all 1-types over σ. The signature σ includes the first K distinguished symbols $E_1, \ldots, E_K$ and some number of non-distinguished ones.

We begin with a simple lemma about equality-free FO.

▶ **Lemma 9.** *If φ is satisfiable, then there exists a model $\mathfrak{A} \models \varphi$ with the following properties.*

1. *For every index $t \in \mathcal{I}$ and every pair $(a_1, a_2) \in A^2$ with $\mathfrak{A} \models G_t(a_1, a_2)$, there exists $a_3 \in A \setminus \{a_1, a_2\}$ such that*

   $$\mathfrak{A} \models \psi_t(a_1, a_2, a_3).$$

2. *If $\mathcal{C} \in A/E_1^{\mathfrak{A}}$ and $\alpha \in \boldsymbol{\alpha}$, then the set*

   $$\{a \in \mathcal{C} \mid \mathrm{tp}^{\mathfrak{A}}[a] = \alpha\}$$

   is either empty or countably infinite.

Proof. Let $\mathfrak{A}_0$ be a model of φ. We define the structure $\mathfrak{A}$ over the domain $A_0 \times \mathbb{N}$: for each relation symbol $R \in \sigma$, denoting its arity as k, interpret R in $\mathfrak{A}$ as the relation

$$\bigcup_{i_1,\ldots,i_k \in \mathbb{N}} \big\{\, ((a_1, i_1), \ldots (a_k, i_k)) \;\big|\; (a_1, \ldots, a_k) \in R^{\mathfrak{A}_0} \,\big\}.$$

The required properties are readily verified. ◀

To the end of this subsection, assuming that φ is satisfiable, we fix a model $\mathfrak{A} \models \varphi$ as in Lemma 9, and subsequently, construct a new model $\mathfrak{B} \models \varphi$ that will satisfy Lemma 8.

Sorts and spectra. We introduce the following definitions.

- A *class-sort* of an element $a \in A$ is the set of 1-types realised by elements that are $E_1^{\mathfrak{A}}$-related to a; formally:

$$\mathrm{sort}^{\mathfrak{A}}[a] := \big\{ \mathrm{tp}^{\mathfrak{A}}[b] \;\big|\; b \in E_1^{\mathfrak{A}}[a] \big\} \subseteq \boldsymbol{\alpha}.$$

  We also define $\mathrm{sort}^{\mathfrak{A}}[\mathcal{C}] := \mathrm{sort}^{\mathfrak{A}}[a]$, where $\mathcal{C} \in A/E_1^{\mathfrak{A}}$ and $a \in \mathcal{C}$ is an arbitrarily chosen representative. For $\Gamma \subseteq \boldsymbol{\alpha}$, the E_1-class $\mathcal{C}$ is said to *realise* Γ if $\mathrm{sort}^{\mathfrak{A}}[\mathcal{C}] = \Gamma$.
- Let $S \subseteq A$ be such that $S/E_1^{\mathfrak{A}} \subseteq A/E_1^{\mathfrak{A}}$ (i.e., S is the union of some E_1-classes). Then a *class-spectrum* of S is a function $2^{\boldsymbol{\alpha}} \to \mathbb{N}$, denoted as $\mathrm{spec}^{\mathfrak{A}}[S]$, that sends $\Gamma \subseteq \boldsymbol{\alpha}$ to

$$\mathrm{spec}^{\mathfrak{A}}[S](\Gamma) := \big|\big\{ \mathcal{C} \in S/E_1^{\mathfrak{A}} \;\big|\; \mathrm{sort}^{\mathfrak{A}}[\mathcal{C}] = \Gamma \big\}\big|.$$

Truncation. We now select a subset of the domain A that will serve as a foundation for the new model. It will contain a *small* number of E_1-classes inside every E_2-class but at the same time will be rich enough to provide all necessary witnesses. To define this subset, we employ an auxiliary operation which we call a 3-*truncation*.

Let $\mathcal{C} \in A/E_2^{\mathfrak{A}}$. A subset $T \subseteq \mathcal{C}$ is a 3-*truncation* of $\mathcal{C}$ if it satisfies the following conditions:

(i) $T/E_1^{\mathfrak{A}} \subseteq \mathcal{C}/E_1^{\mathfrak{A}}$;

(ii) for every $\Gamma \subseteq \boldsymbol{\alpha}$, we have

$$\mathrm{spec}^{\mathfrak{A}}[T](\Gamma) \;=\; \inf\big\{\, \mathrm{spec}^{\mathfrak{A}}[\mathcal{C}](\Gamma),\, 3 \,\big\}.$$

In other words, T is the union of inner E_1-classes of $\mathcal{C}$ which considers at most 3 representatives of every class-sort $\Gamma \subseteq \boldsymbol{\alpha}$. The parameter 3 is chosen because φ uses 3 variables.

For every $\mathcal{C} \in A/E_2^{\mathfrak{A}}$, we arbitrarily fix a 3-truncation for $\mathcal{C}$ and denote it as $T_3(\mathcal{C})$. The existence of $T_3(\mathcal{C})$ is straightforward, but naturally its choice might not be unique.

Finally, we define the desired subset as

$$B := \bigcup_{\mathcal{C} \in A/E_2^{\mathfrak{A}}} T_3(\mathcal{C}).$$

Similar pairs. We now introduce a *similarity* relation $\sim$ on pairs of elements that exhibit analogous characteristics when considering their 1-types and interactions through the nested equivalence relations $E_1^{\mathfrak{A}} \subseteq \cdots \subseteq E_K^{\mathfrak{A}}$.

Let $(a_1, a_2) \in A^2$ and $(b_1, b_2) \in A^2$. We write $(a_1, a_2) \sim (b_1, b_2)$ if the following holds:

(i) $a_1 = a_2$ if and only if $b_1 = b_2$.

(ii) For $i \in \{1, 2\}$, $\mathrm{tp}^{\mathfrak{A}}[a_i] = \mathrm{tp}^{\mathfrak{A}}[b_i]$.

(iii) For $i \in \{1, 2\}$, $\mathrm{sort}^{\mathfrak{A}}[a_i] = \mathrm{sort}^{\mathfrak{A}}[b_i]$.

(iv) $(a_1, a_2) \in E_1^{\mathfrak{A}}$ if and only if $(b_1, b_2) \in E_1^{\mathfrak{A}}$.

(v) For every $k \in [2, K]$, $(a_1, b_1) \in E_k^{\mathfrak{A}}$ and $(a_2, b_2) \in E_k^{\mathfrak{A}}$.

It is readily verified that $\sim$ is indeed an equivalence relation. The following lemma will be used in the construction of a new model for φ.

▶ **Lemma 10.** *Let $(a_1, a_2) \in A^2$ and $(b_1, b_2) \in B^2$. Suppose that*

$$(a_1, a_2) \sim (b_1, b_2) \text{ and } F \subseteq B \text{ is a finite set with } \{b_1, b_2\} \subseteq F.$$

Then, for any $a_3 \in A \setminus \{a_1, a_2\}$, one can find an element $b_3 \in B \setminus F$ satisfying that

$$(a_1, a_3) \sim (b_1, b_3) \quad \text{and} \quad (a_2, a_3) \sim (b_2, b_3).$$

Proof. We consider two cases.

First, suppose that either $(a_1, a_3) \in E_1^{\mathfrak{A}}$ or $(a_2, a_3) \in E_1^{\mathfrak{A}}$; w.l.o.g. assume the former. We arbitrarily select b_3 from

$$\left\{b \in E_1^{\mathfrak{A}}[b_1] \setminus F \;\middle|\; \mathrm{tp}^{\mathfrak{A}}[b] = \mathrm{tp}^{\mathfrak{A}}[a_3]\right\}.$$

The above set is non-empty, since F is assumed to be finite and yet the equivalence class $E_1^{\mathfrak{A}}[b_1]$ realises every 1-type of $\mathrm{sort}^{\mathfrak{A}}[b_1]$ $(= \mathrm{sort}^{\mathfrak{A}}[a_1])$ infinitely often (item 2 of Lemma 9).

Now, assume that neither (a_1, a_3) nor (a_2, a_3) is $E_1^{\mathfrak{A}}$-related. Let $\mathcal{C} = E_2^{\mathfrak{A}}[a_3]$, and let $\mathcal{D} = E_1^{\mathfrak{A}}[a_3]$. We claim that one can select an E_1-class $\mathcal{D}' \in T_3(\mathcal{C})/E_1^{\mathfrak{A}}$ such that

$$\mathcal{D}' \cap \{b_1, b_2\} = \varnothing \quad \text{and} \quad \mathrm{sort}^{\mathfrak{A}}[\mathcal{D}'] = \mathrm{sort}^{\mathfrak{A}}[\mathcal{D}].$$

The existence of such E_1-class $\mathcal{D}'$ follows from the definition of 3-truncation: $T_3(\mathcal{C})$ includes the same number of realisations of $\mathrm{sort}^{\mathfrak{A}}[\mathcal{D}]$ as $\mathcal{C}$ does when this number is at most 2; and otherwise it includes precisely 3 realisations of $\mathrm{sort}^{\mathfrak{A}}[\mathcal{D}]$. In either case, this suffices to avoid – when necessary – the two "bad" E_1-classes: namely, $E_1^{\mathfrak{A}}[b_1]$ and $E_1^{\mathfrak{A}}[b_2]$.

Finally, we arbitrarily select b_3 from

$$\left\{b \in \mathcal{D}' \setminus F \;\middle|\; \mathrm{tp}^{\mathfrak{A}}[b] = \mathrm{tp}^{\mathfrak{A}}[a_3]\right\}.$$

As before this set is non-empty. ◀

Model construction. We now construct a chain of finite structures

$$\mathfrak{B}_1 \subseteq \mathfrak{B}_2 \subseteq \mathfrak{B}_3 \subseteq \dots, \qquad \text{where } \mathfrak{B}_{n+1} \restriction B_n = \mathfrak{B}_n \text{ for all } n \geq 1.$$

The natural limit structure $\mathfrak{B}_\omega = \bigcup_{n \geq 1} \mathfrak{B}_n$ will be a model of φ satisfying Lemma 8. For every $n \in \mathbb{N}$, the structure $\mathfrak{B}_n$ will be defined over a subset $B_n \subseteq B$, and the inclusion mapping $\mathfrak{B}_n \hookrightarrow \mathfrak{A}$ will preserve the 1-types and the nested equivalence relations $E_1, \dots, E_K$; though not necessarily other atoms. Formally, this requirement expresses as:

- For every $b \in B_n$, $\mathrm{tp}^{\mathfrak{B}_n}[b] = \mathrm{tp}^{\mathfrak{A}}[b]$.
- For every $k \in [K]$, $E_k^{\mathfrak{B}_n} = E_k^{\mathfrak{A}} \cap B_n^2$.

To guide the construction, we will additionally define a sequence of mappings $(\rho_n)_{n \geq 1}$:

(†) Let $n \geq 1$. If $(b_1, b_2) \in B_n^2$ is such that $\mathfrak{B}_n \models G_t(b_1, b_2)$ for some index $t \in \mathcal{I}$, then the mapping ρ_n sends (b_1, b_2) to a pair $(a_1, a_2) \in A^2$ such that $(b_1, b_2) \sim (a_1, a_2)$ and
   - $\mathrm{tp}^{\mathfrak{B}_n}[b_1] = \mathrm{tp}^{\mathfrak{A}}[a_1]$ when $b_1 = b_2$;
   - $\mathrm{tp}^{\mathfrak{B}_n}[b_1, b_2] = \mathrm{tp}^{\mathfrak{A}}[a_1, a_2]$ when $b_1 \neq b_2$.

We now formulate a sufficent criterion for satisfaction of the sentence φ in $\mathfrak{B}_\omega$.

A pair of elements $(b_1, b_2) \in B_n^2$ is *satisfied* in $\mathfrak{B}_n$ if $\mathfrak{B}_n \models G_t(b_1, b_2) \to \exists y\ \psi_t(b_1, b_2, y)$ for every $t \in \mathcal{I}$, and otherwise it is *unsatisfied*. It is readily verified that $\mathfrak{B}_\omega \models \varphi$ whenever:

1. For every $n \geq 1$, the structure $\mathfrak{B}_n$ interprets $E_1, \dots, E_K$ as nested equivalence relations.
2. For every $n \geq 1$, we have $\mathfrak{B}_n \models \Psi$.
3. Every pair $(b_1, b_2) \in B_n^2$ with $n \geq 1$ is eventually satisfied in some structure $\mathfrak{B}_\ell$ with $\ell \geq n$; although no finite global upper bound on such ℓ is required for all pairs simultaneously.

Base case. Let $\mathfrak{B}_1 := \mathfrak{A} \restriction \{a\}$, where $a \in B$ is an arbitrarily chosen element. The corresponding mapping is $\rho_1 := \{(a, a) \mapsto (a, a)\}$.

Inductive step. Let $n \geq 1$. Assume that $\mathfrak{B}_n$ and ρ_n have already been defined. If there are no unsatisfied pairs in $\mathfrak{B}_n$, we conclude our construction and declare $\mathfrak{B}_\omega := \mathfrak{B}_n$ as the final

model. Yet, suppose that an unsatisfied pair exists: in the current step choose one such pair, denote it as $(b_1, b_2) \in B_n^2$. Let $t \in \mathcal{I}$ be an index for which the pair (b_1, b_2) lacks a witness, that is, such that $\mathfrak{B}_n \models G_t(b_1, b_2)$ holds, but $\mathfrak{B}_n \models \neg\exists y\ \psi_t(b_1, b_2, y)$.

We require that pairs are chosen in a *fair way*, which means that no pairs remain omitted for infinitely many steps. It is clear that such a strategy for selecting unsatisfied pairs exists.

Now, we show how to construct the successive structure $\mathfrak{B}_{n+1}$ extending $\mathfrak{B}_n$ by one fresh element $b_3 \in B \setminus B_n$ so that $\mathfrak{B}_{n+1} \models \psi_t(b_1, b_2, b_3)$ and inductive invariants are preserved. We consider only the case $b_1 \neq b_2$, as the case $b_1 = b_2$ follows a similar line of arguments.

Let $(a_1, a_2) := \rho_n(b_1, b_2)$. By the inductive invariant of ρ_n, we have that $(a_1, a_2) \sim (b_1, b_2)$ and $\mathrm{tp}^{\mathfrak{A}}[a_1, a_2] = \mathrm{tp}^{\mathfrak{B}_n}[b_1, b_2]$. Fix a witness $a_3 \in A$ for the pair (a_1, a_2), that is, an element satisfying $\mathfrak{A} \models \psi_t(a_1, a_2, a_3)$. By item 1 of Lemma 9, the witness a_3 can be chosen to be distinct from a_1 and a_2. Applying Lemma 10, with $F := B_n$, we obtain $b_3 \in B \setminus B_n$ satisfying $(a_1, a_3) \sim (b_1, b_3)$ and $(a_2, a_3) \sim (b_2, b_3)$.

We now construct the structure $\mathfrak{B}_{n+1}$ by extending $\mathfrak{B}_{n+1}$ with the element b_3.

- Initialise $\mathfrak{B}_{n+1} := \mathfrak{B}_n$.
- The witnessing 3-type: assign $\mathrm{tp}^{\mathfrak{B}_{n+1}}[b_1, b_2, b_3] := \mathrm{tp}^{\mathfrak{A}}[a_1, a_2, a_3]$.
- 2-types induced by the distinguished equivalence relations: for every $b \in B_n \setminus \{b_1, b_2\}$ such that $(b, b_3) \in E_K^{\mathfrak{A}}$, assign $\mathrm{tp}^{\mathfrak{B}_{n+1}}[b, b_3] := \mathrm{tp}^{\mathfrak{A}}[b, b_3]$.
- Remaining facts whose truth values has been not specified are set to false.

The construction of $\mathfrak{B}_{n+1}$ is now finished. Observe that $\mathfrak{B}_{n+1} \restriction B_n = \mathfrak{B}_n$, and also that $\mathfrak{B}_{n+1}$ interprets $E_1, \ldots, E_K$ in the same way as $\mathfrak{A} \restriction B_{n+1}$ does.

We now argue that $\mathfrak{B}_{n+1} \models \Psi$. Recall that Ψ is a conjunction of purely universal prenex sentences. Consider one of its conjuncts: $\forall x_1, \ldots, x_k\, (F(x_1, \ldots, x_k) \to \chi(x_1, \ldots, x_k))$, where $1 \leq k \leq 3$, F is a guard, and χ is quantifier-free. Take any tuple $b_1, \ldots, b_k \in B_{n+1}$ such that $\mathfrak{B}_{n+1} \models F(b_1, \ldots, b_k)$. Since χ is quantifier-free, the truth of $\chi(b_1, \ldots, b_k)$ depends only on the atomic type induced by the chosen elements. Thus, the requirement $\mathfrak{B}_{n+1} \models \Psi$ reduces to showing the following: every *guarded* type realised in $\mathfrak{B}_{n+1}$ is already realised in the original model $\mathfrak{A}$, where a k-type τ is said to be guarded if it contains a positive literal γ with $\mathrm{fv}(\gamma) = \{x_1, \ldots, x_k\}$, that is, if some positive literal uses all variables of τ.

The structure $\mathfrak{B}_{n+1}$ is obtained from $\mathfrak{B}_n$ by copying atomic types from $\mathfrak{A}$ while preserving all previously assigned types. This construction cannot introduce new guarded types – in contrast to non-guarded ones. It follows that $\mathfrak{B}_{n+1} \models \Psi$.

We now extend the mapping ρ_n. For brevity, we write $G(x_1, x_2) := \bigvee_{t \in \mathcal{I}} G_t(x_1, x_2)$.

- Initialise $\rho_{n+1} := \rho_n$.
- For every $(i, j) \in \{1, 2, 3\}^2 \setminus \{1, 2\}^2$, if $\mathfrak{B}_{n+1} \models G(b_i, b_j)$, then put $\rho_{n+1}(b_i, b_j) := (a_i, a_j)$.
- For every element $b \in B_n \setminus \{b_1, b_2\}$, if $\mathfrak{B}_{n+1} \models G(b, b_3)$, then put $\rho_{n+1}(b, b_3) := (b, b_3)$; and do symmetrically for (b_3, b): if $\mathfrak{B}_{n+1} \models G(b_3, b)$, then put $\rho_{n+1}(b_3, b) := (b_3, b)$.

This concludes the inductive step. The following claim highlights the key points:

▷ Claim 11.

1. $\mathfrak{B}_{n+1} \restriction B_n = \mathfrak{B}_n$ and $B_{n+1} \setminus B_n = \{b_3\}$.
2. For every $b \in B_{n+1}$, $\mathrm{tp}^{\mathfrak{B}_{n+1}}[b] = \mathrm{tp}^{\mathfrak{A}}[b]$.
3. For every $k \in [K]$, $E_k^{\mathfrak{B}_{n+1}} = E_k^{\mathfrak{A}} \cap B_{n+1}^2$.
4. The invariant (†) of the mapping ρ_{n+1} is preserved.
5. $\mathfrak{B}_{n+1} \models \Psi$.
6. $\mathfrak{B}_{n+1} \models \psi_t(b_1, b_2, b_3)$.

Bounding index. To finish the proof of Lemma 8, it remains to bound the number of E_1-classes in every E_2-class of $\mathfrak{B}_\omega$. Since $B_\omega \subseteq B$ and $\mathfrak{B}_\omega \hookrightarrow \mathfrak{A}$ preserves both E_1 and E_2, we may alternatively bound the number of E_1-classes in every E_2-class of $\mathfrak{A} \restriction B$.

By definition of B, every $\mathcal{D} \in B/E_2^{\mathfrak{A}}$ is a 3-truncation of some $\mathcal{C} \in A/E_2^{\mathfrak{A}}$, i.e., $\mathcal{D} = T_3(\mathcal{C})$. The number of distinct class-sorts – i.e., subsets of $\boldsymbol{\alpha}$ – is $2^{|\boldsymbol{\alpha}|}$. Since at most 3 realisations of every class-sort $\Gamma \subseteq \boldsymbol{\alpha}$ are included in $T_3(\mathcal{C})$, we obtain that $|T_3(\mathcal{C})/E_1^{\mathfrak{A}}| \leq 3 \cdot 2^{|\boldsymbol{\alpha}|}$.

5.3 Elimination of Distinguished Symbols

We now show a reduction from $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ (where $K \geq 2$) to $\mathrm{GF}^3[1\text{-}\mathcal{EQ}^{\subseteq}]$.

▶ **Lemma 12.** *Let $K \geq 2$, and let φ be a constant-free, equality-free sentence of $\mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$. Then there exists a constant-free, equality-free sentence φ' of $\mathrm{GF}^3[1\text{-}\mathcal{EQ}^{\subseteq}]$ such that:*

- *The sentences φ and φ' are equisatisfiable.*
- *If φ' has a model over a domain D, then φ has also.*
- *$|\varphi'|$ is bounded by a $(K{-}1)$-exponential function of $|\varphi|$.*
- *φ' is computable from φ in $(K{-}1)$-exponential time.*

Proof of Lemma 12. Let $K \geq 2$, and let $\varphi \in \mathrm{GF}^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ be a constant-free, equality-free sentence; w.l.o.g φ is in normal form as in Assumption 7. Write σ for the signature of φ, and $\boldsymbol{\alpha}$ for the set of all 1-types over σ.

Set $n = \lceil \log_2 N \rceil$, where $N = 3 \cdot 2^{|\boldsymbol{\alpha}|}$ is the bound of Lemma 8. We axiomatise the finest distinguished symbol E_1 using fresh unary predicates $\{U_i\}_{i=1}^n$: we declare two elements E_1-equivalent if and only if they are E_2-related and agree on U_i for all $i \in [n]$.

Let φ' denote the sentence φ with conjoined axioms:

$$\forall x_1, x_2 \left(x_1 E_1 x_2 \to x_1 E_2 x_2 \right) \tag{46}$$

$$\forall x_1, x_2 \left(x_1 E_2 x_2 \to \left(x_1 E_1 x_2 \leftrightarrow \bigwedge\nolimits_{i=1}^{n} \left(U_i(x_1) \leftrightarrow U_i(x_2) \right) \right) \right) \tag{47}$$

In φ' the symbol E_1 is treated as a non-distinguished binary symbol, i.e., its interpretation is not constrained. Hence, it can be naturally viewed as a sentence of $\mathrm{GF}^3[(K{-}1)\text{-}\mathcal{EQ}^{\subseteq}]$, with the distinguished symbols being $E_2, \ldots, E_K$; this numeration of distinguished symbols does not follow our convention, but it is convenient in the following claim.

▷ Claim 13.

1. The sentences φ and φ' are equisatisfiable.
2. The σ-reduct of any model $\mathfrak{A} \models \varphi'$ satisfies φ.
3. $|\varphi'|$ is bounded by an exponential function of $|\varphi|$.

Proof. It is immediate that the σ-reduct of any model for φ' is also a model for φ, and hence satisfiability of φ' implies that of φ.

The converse direction is more delicate: let n and N be as above. Let $\mathfrak{A} \models \varphi$ be a model satisfying Lemma 8. Since there are $2^n \geq N$ possible 1-types over $\{U_i\}_{i=1}^n$, we have enough combinations to decorate the elements of $\mathfrak{A}$ with the predicates U_i so that any $E_2^{\mathfrak{A}}$-related elements get the same combination precisely when they are also related by $E_1^{\mathfrak{A}}$.

Finally, since σ is constant-free, the number of 1-types is singly exponential in $|\varphi|$; more precisely: $|\boldsymbol{\alpha}| = 2^{|\sigma|} \leq 2^{|\varphi|}$. We thus have $|\varphi'| = |\varphi| + \mathcal{O}(n) = |\varphi| + \mathcal{O}(|\boldsymbol{\alpha}|) = 2^{\mathcal{O}(|\varphi|)}$. ◁

Having Claim 13 established, we adjust the numeration of the distinguished symbols: that is, we rename E_1 to a fresh symbol and shift $E_{k+1} \mapsto E_k$ for all $k \geq 1$. In this way, we obtain a proper sentence of $\mathrm{GF}^3[(K{-}1)\text{-}\mathcal{EQ}^{\subseteq}]$. By recursively applying the above reduction $K - 1$ times, we reach the final sentence in $\mathrm{GF}^3[1\text{-}\mathcal{EQ}^{\subseteq}]$ with the required properties.

5.4 Finite Model Property

We establish the finite model property via a reduction to an auxiliary fragment denoted GFU. This fragment is just GF extended with a single distinguished binary symbol U. Models of GFU-sentences are constrained to interpret U as the universal relation. (In expressive power, GFU coincides with the Triguarded Fragment, see [28] for further details.)

To obtain the reduction to GFU^3, we first embed GF$^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ into GF$^3[(K{+}1)\text{-}\mathcal{EQ}^{\subseteq}]$: we relativise quantifiers with the coarsest distinguished symbol E_{K+1} (this is the same technique as employed in Proposition 2). The symbol E_{K+1} can now act as the universal relation, and Lemma 12 produces a sentence which can be viewed as a GFU^3-sentence. The finite model property for the constant-free, equality-free GF$^3[\mathcal{EQ}^{\subseteq}]$ follows then from the same property for the equality-free fragment of GFU^3, which is established in [22].

The described reduction yields decidability of satisfiability, but without tight complexity bounds for fragments GF$^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ parametrised by $K \geq 1$. Indeed, Lemma 12 produces GFU^3-sentences of K-exponential length, and the complexity of satisfiability for GFU^3 is NExpTime-complete [28]. In consequence, under standard complexity-theoretic assumptions, we obtain only a non-optimal $(K{+}1)$-NExpTime bound. The optimal $(K{+}1)$-ExpTime bound is established in the following subsection.

5.5 Decision Procedure

▶ **Proposition 14.** *The satisfiability problem for the constant-free, equality-free* GF$^3[K\text{-}\mathcal{EQ}^{\subseteq}]$ *is decidable in* $(K{+}1)$*-ExpTime.*

By Lemma 12, to show the proposition, it suffices to design a 2-ExpTime procedure for GF$^3[1\text{-}\mathcal{EQ}^{\subseteq}]$. The idea employed here is inspired by Kieroński and Malinowski's approach for GFU [18], as well as Pratt-Hartmanns's procedure for deciding the pure GF [24].

The procedure is based on a criterion that expresses satisfiability of a sentence in terms of the existence of a certain set of *labelled types* – let us first introduce this notion.

Labelled types. Let σ be a constant-free signature, and let $\boldsymbol{\alpha}$ denote the set of all 1-types over σ. Suppose that τ is a k-type over σ for some $k \in \mathbb{N}$, and that $\ell\colon \mathrm{dom}(\tau) \to 2^{\boldsymbol{\alpha}}$ assigns to each variable in $\mathrm{dom}(\tau) := \{x_1, \ldots, x_k\}$, the *domain* of τ, a set of 1-types. A pair (τ, ℓ) is then called a *labelled k-type* (over σ) if the following holds:

- τ interprets E_1 as an equivalence relation.
- For every $x_i \in \mathrm{dom}(\tau)$, we have $\mathrm{tp}^{\tau}[x_i] \in \ell(x_i)$.
- For every $x_i, x_j \in \mathrm{dom}(\tau)$, if $\tau \models x_i E_1 x_j$, then $\ell(x_i) = \ell(x_j)$.

For brevity, we sometimes write just τ for (τ, ℓ) and refer to the associated labelling by ℓ_τ. A *labelled type* is simply a labelled k-type for some k.

The intended role of the labelling ℓ_τ is to capture information about 1-types realised in the E_1-classes of elements of τ: whenever $a_1, \ldots, a_k$ are distinct elements of a structure $\mathfrak{A}$, then the tuple $\langle a_1, \ldots, a_k \rangle$ *realises* the labelled k-type (τ, ℓ_τ), where

- $\tau = \mathrm{tp}^{\mathfrak{A}}[a_1, \ldots, a_k]$;
- for every $i \in [k]$, $\ell_\tau(x_i) = \{\mathrm{tp}^{\mathfrak{A}}[b] \mid b \in E_1^{\mathfrak{A}}[a_i]\}$.

To illustrate, consider a structure $\mathfrak{A}$ over a domain $A = \{1, 2, 3, 4\}$. Suppose that $E_1^{\mathfrak{A}}$ has three equivalence classes: $\{1, 3\}$, $\{2\}$, and $\{4\}$. Then the labelled 2-type of the pair $(1, 2)$ is given by $(\mathrm{tp}^{\mathfrak{A}}[1, 2], \ell)$, where ℓ maps $x_1 \mapsto \{\mathrm{tp}^{\mathfrak{A}}[1], \mathrm{tp}^{\mathfrak{A}}[3]\}$ and $x_2 \mapsto \{\mathrm{tp}^{\mathfrak{A}}[2]\}$.

Let $k, m \in \mathbb{N}$ be such that $1 \leq k \leq m$. Suppose that τ is a labelled k-type and τ' is a labelled m-type. We say that τ is a *reduct* of τ' if there exists an injective mapping $\rho\colon \mathrm{dom}(\tau) \hookrightarrow \mathrm{dom}(\tau')$ such that the following holds:

- $\tau = \mathrm{tp}^{\tau'}[\rho(x_1), \ldots, \rho(x_k)]$.
- $\ell_\tau(x_i) = \ell_{\tau'}(\rho(x_i))$ for every $x_i \in \mathrm{dom}(\tau)$.

Further, we say that τ' *canonically extends* τ if the mapping ρ is an inclusion, that is,

- $\rho(x_i) = x_i$ for every $x_i \in \mathrm{dom}(\tau)$.

Satisfiability criterion. With necessary tools in place, we formulate the satisfiability criterion.

▶ **Lemma 15.** *Let φ be a constant-free, equality-free sentence of $\mathrm{GF}^3[1\text{-}\mathcal{E}\mathcal{Q}^{\subseteq}]$, and let σ be the signature of φ. Assume that φ is in normal form as in Assumption 7, that is,*

$$\varphi \quad = \quad \Psi \quad \wedge \quad \bigwedge\nolimits_{t \in \mathcal{I}} \forall x_1, x_2 \left(G_t(x_1, x_2) \to \exists y\, \psi_t(x_1, x_2, y)\right).$$

Then φ is satisfiable if and only if there exists a non-empty set Ω of labelled types over σ such that $P_{1,\Omega}(\tau), \ldots, P_{5,\Omega}(\tau)$ hold for every $\tau \in \Omega$, where

$P_{1,\Omega}(\tau) \iff$ *every reduct of τ belongs to Ω.*

$P_{2,\Omega}(\tau) \iff \tau \models \Psi$.

$P_{3,\Omega}(\tau) \iff$ *if τ is a labelled 1-type and $\tau \models G_t(x_1, x_1)$ for some $t \in \mathcal{I}$, then there exists a labelled 2-type $\tau' \in \Omega$ canonically extending τ such that $\tau' \models \psi_t(x_1, x_1, x_2)$.*

$P_{4,\Omega}(\tau) \iff$ *if τ is a labelled 2-type and $\tau \models G_t(x_1, x_2)$ for some $t \in \mathcal{I}$, then there exists a labelled 3-type $\tau' \in \Omega$ canonically extending τ such that $\tau' \models \psi_t(x_1, x_2, x_3)$.*

$P_{5,\Omega}(\tau) \iff$ *if τ is a labelled 1-type, then for all $\alpha_1, \alpha_2 \in \ell_\tau(x_1)$, there exists a labelled 2-type $\beta \in \Omega$ such that*
$\mathrm{tp}^\beta[x_1] = \alpha_1, \quad \mathrm{tp}^\beta[x_2] = \alpha_2,$
$\beta \models x_1 E_1 x_2, \quad \ell_\beta(x_1) = \ell_\beta(x_2) = \ell_\tau(x_1).$

Intuitively, the conditions of Lemma 15 play the following roles. The predicate $P_{1,\Omega}(\tau)$ expresses closure under reducts, a property satisfied by the set of types realised in any structure. The predicate $P_{2,\Omega}(\tau)$ ensures satisfaction of the universal conjunct Ψ in any structure constructed from types in Ω. The predicates $P_{3,\Omega}(\tau)$ and $P_{4,\Omega}(\tau)$ guarantee the existence of suitable witnessing extensions required by the existential conjuncts. The predicate $P_{5,\Omega}(\tau)$ ensures that E_1-classes can be completed by providing the necessary labelled 2-types. Together, these conditions characterise the existence of a model of φ.

Completeness. Let $\mathfrak{A} \models \varphi$ be a model as in Lemma 9, and construct the set

$$\Omega \;:=\; \bigcup_{1 \leq k \leq 3} \left\{ (\tau, \ell_\tau) \text{ is a labelled } k\text{-type realised in } \mathfrak{A} \right\}.$$

One can verify that the conditions of Lemma 15 are satisfied for the constructed Ω.

Soundness. Write $\boldsymbol{\alpha}$ for the set of all 1-types over the signature σ.

Suppose that Ω is a set of labelled types over σ as in Lemma 15. We prove that φ is satisfiable by constructing a chain of finite structures

$$\mathfrak{B}_1 \subseteq \mathfrak{B}_2 \subseteq \mathfrak{B}_3 \subseteq \dots, \qquad \text{where } \mathfrak{B}_{n+1} \restriction B_n = \mathfrak{B}_n \text{ for all } n \geq 1.$$

The natural limit structure $\mathfrak{B}_\omega = \bigcup_{n\geq 1} \mathfrak{B}_n$ will be a model of φ.

To guide the construction, we define in parallel a sequence of functions $(\ell_n)_{n\in\mathbb{N}}$, where each $\ell_n \colon B_n \to 2^{\boldsymbol{\alpha}}$ assigns to every element of the domain of $\mathfrak{B}_n$ a set of 1-types. The label $\ell_n(a)$ serves as a *declaration* of which 1-types are admitted in the E_1-class of a. The following conditions shall hold during the construction:

- $\{\mathrm{tp}^{\mathfrak{B}_n}[b] \mid b \in E_1^{\mathfrak{B}_n}[a]\} \subseteq \ell_n(a)$ for all $a \in B_n$.
- $\ell_n(a) = \ell_n(b)$ for all $(a,b) \in E_1^{\mathfrak{B}_n}$.

It is not required that all 1-types from $\ell_n(a)$ are actually realised in the E_1-class of a.

We require each function ℓ_n to satisfy the following invariant:

(‡) Let $(b_1, b_2) \in B_n^2$ be such that $\mathfrak{B}_n \models G_t(b_1, b_2)$ for some $t \in \mathcal{I}$.

- If $b_1 = b_2$, then the labelled 1-type (α, ℓ_α) belongs to Ω, where

$$\alpha := \mathrm{tp}^{\mathfrak{B}_n}[b_1], \qquad \ell_\alpha(x_1) := \ell_n(b_1).$$

- If $b_1 \neq b_2$, then the labelled 2-type (β, ℓ_β) belongs to Ω, where

$$\beta := \mathrm{tp}^{\mathfrak{B}_n}[b_1, b_2], \qquad \ell_\beta(x_1) := \ell_n(b_1), \qquad \ell_\beta(x_2) := \ell_n(b_2).$$

Base case. Since Ω is non-empty and $P_{1,\Omega}$ holds, there exists a labelled 1-type $(\alpha, \ell_\alpha) \in \Omega$. Define $\mathfrak{B}_1$ to be the singleton structure with an element b realising α, and set $\ell_1(b) := \ell_\alpha(x_1)$.

Inductive step. Let $n \geq 1$. Assume that $\mathfrak{B}_n$ and ℓ_n have already been defined. Choose in a *fair way* a pair $(b_1, b_2) \in B_n^2$ such that $\mathfrak{B}_n \models G_t(b_1, b_2)$, but $\mathfrak{B}_n \models \neg\exists y\, \psi_t(b_1, b_2, y)$ for some $t \in \mathcal{I}$. We consider only the case $b_1 \neq b_2$, as the case $b_1 = b_2$ is analogous.

Let (β, ℓ_β) be the labelled 2-type defined by

$$\beta := \mathrm{tp}^{\mathfrak{B}_n}[b_1, b_2], \qquad \ell_\beta(x_1) := \ell_n(b_1), \qquad \ell_\beta(x_2) := \ell_n(b_2).$$

By the inductive invariant (‡), $(\beta, \ell_\beta) \in \Omega$. Hence, by $P_{4,\Omega}$, there exists a labelled 3-type $(\tau, \ell_\tau) \in \Omega$ that canonically extends (β, ℓ_β) and satisfies $\tau \models \psi_t(x_1, x_2, x_3)$.

We now construct the structure $\mathfrak{B}_{n+1}$ by extending $\mathfrak{B}_n$ with a fresh element b_3.

- Initialise $\mathfrak{B}_{n+1} := \mathfrak{B}_n$.
- The witnessing 3-type: assign $\mathrm{tp}^{\mathfrak{B}_{n+1}}[b_1, b_2, b_3] := \tau$. (Since τ canonically extends β, this step does not modify the 2-type $\mathrm{tp}^{\mathfrak{B}_n}[b_1, b_2]$.)
- Closure under E_1: If either (x_1, x_3) or (x_2, x_3) belongs to E_1^τ, then the corresponding transitive connections involving b_3 must be added. Assume w.l.o.g. that $(x_1, x_3) \in E_1^\tau$. For every $b \in E_1^{\mathfrak{B}_n}[b_1] \setminus \{b_1, b_2\}$, $P_{5,\Omega}$ yields a labelled 2-type $(\gamma, \ell_\gamma) \in \Omega$ such that

$$\begin{aligned} &\mathrm{tp}^\gamma[x_1] = \mathrm{tp}^{\mathfrak{B}_{n+1}}[b], \qquad && \mathrm{tp}^\gamma[x_2] = \mathrm{tp}^{\mathfrak{B}_{n+1}}[b_3], \\ &\gamma \models x_1 E_1 x_2, && \ell_\gamma(x_1) = \ell_\gamma(x_2) = \ell_\tau(x_3). \end{aligned}$$

  We then set $\mathrm{tp}^{\mathfrak{B}_{n+1}}[b, b_3] := \gamma$.

We now have $\mathfrak{B}_{n+1} \models \psi_t(b_1, b_2, b_3)$. To complete the inductive step, we declare

$$\ell_{n+1} \;:=\; \ell_n \cup \{b_3 \mapsto \ell_\tau(x_3)\}.$$

Algorithm 1 Decide satisfiability for constant-free, equality-free $\mathrm{GF}^3[1\text{-}\mathcal{EQ}^{\subseteq}]$.

Input: A normal-form sentence φ over a signature σ.
Output: Sat if φ is satisfiable, and UnSat otherwise.

$\Omega \leftarrow \{\text{all labelled } k\text{-types over } \sigma \text{ for } k \in \{1,2,3\}\}$;
while there exists $\tau \in \Omega$ violating one of $P_{1,\Omega}(\tau), \dots, P_{5,\Omega}(\tau)$ **do**
    choose such a τ;
    $\Omega \leftarrow \Omega \setminus \{\tau\}$;
if $\Omega \neq \varnothing$ **then return** Sat **else return** UnSat;

Decision procedure. To conclude the proof, it remains to show that the existence of a set Ω satisfying the conditions of Lemma 15 can be decided in 2-ExpTime. We call such a set a *satisfiability witness*. The decision procedure is given in Algorithm 1.

The algorithm is clearly sound: it always terminates, and upon termination returns Sat if and only if the resulting set Ω is a satisfiability witness.

Completeness follows from the following observation. Suppose that a satisfiability witness Ω^* exists. Then no element of Ω^* is ever removed during the execution of the algorithm. Consequently, upon termination, the algorithm produces a set $\Omega \supseteq \Omega^*$, which is itself a satisfiability witness. In particular, Ω is the unique maximal satisfiability witness.

To analyse the running time, observe that the procedure runs in polynomial time in the number of labelled types over at most three variables. Each labelled type is a pair (τ, ℓ_τ), where τ is a k-type with $k \leq 3$ and $\ell_\tau \colon \mathrm{dom}(\tau) \to 2^{\boldsymbol{\alpha}}$. The number of k-types with $k \leq 3$ is a priori doubly exponential in $|\varphi|$. The number of possible labellings is bounded by $(2^{|\boldsymbol{\alpha}|})^3$, which is likewise doubly exponential in $|\varphi|$; indeed, since φ is constant-free, the number of 1-types (= the size of $\boldsymbol{\alpha}$) is only singly exponential in $|\varphi|$. Hence, the total number of labelled k-types with $k \leq 3$ is doubly exponential in $|\varphi|$, and the procedure runs in deterministic doubly-exponential time. This concludes the proof of Proposition 14.

6 Related Work

Two-Variable Fragment. With one equivalence relation, FO^2 has the finite model property and NExpTime-complete satisfiability [20]. With two equivalence relations (without nesting condition), the finite model property is lost; both satisfiability and finite satisfiability are 2-NExpTime-complete [19]. With three equivalence relations (again without nesting conditions), FO^2 is undecidable [20].

The work of [5] extends the results on FO^2 with equivalences to the setting of words with *nested data*, i.e., structures equipped with a linear order, nested equivalence relations, and additional free unary predicates. It is shown there that satisfiability remains decidable when the linear order is accessible only through the successor relation, but becomes undecidable in the presence of both the order and the successor relation.

A limitation of [5] is the absence of free binary predicates. This was recently addressed in [13], where it was shown that $\mathrm{FO}^2[\mathcal{EQ}^{\subseteq}]$ enjoys the finite model property and has an NExpTime-complete satisfiability problem. Several further variants are investigated there:

- nested equivalences combined with a linear order (without successor);
- nested total preorders (i.e., equivalence classes are linearly ordered);
- nested total preorders equipped with the induced successor relation.

These extensions retain elementary complexity, namely NExpTime-completeness in the first two cases and ExpSpace-completeness in the third. Finally, [13] establishes that FO^2 with two independent families of nested equivalence relations is undecidable.

Logics for trees. An alternative perspective on nested equivalence relations is to interpret them as unranked trees of bounded depth. When only K distinguished predicates $E_1, \ldots, E_K$ are considered, the corresponding structures can be viewed as describing the leaves of unranked trees of height $K{+}1$. Domain elements correspond to leaves at level 0; the equivalence classes of E_1 form nodes at level 1; those of E_2 form nodes at level 2; and so on, up to level K, with an implicit root located at level $K{+}1$. In this view, for every $1 \leq k \leq K$, the predicate aE_kb holds precisely when a and b share a common ancestor at level k.

This interpretation differs from standard logics over trees, where all tree nodes belong to the domain and the structure is navigated via predicates such as *parent*, *child*, or *descendant*. In this setting, FO^2 was investigated in numerous variants – over both ranked and unranked trees of unbounded depth, over data trees, and over trees with counting (see, e.g., [6, 11, 4]).

Equivalence guards. GF^2 with two equivalence relations (without nesting condition) is 2-ExpTime-complete, but becomes undecidable with three equivalence relations [17, 21]. For an arbitrary number of equivalence relations, GF remains decidable provided that equivalences are restricted to occur only in guards – so-called *equivalence guards* [17, 23]. However, in this setting the expressive power of the distinguished symbols is severely limited; in particular, one cannot express that equivalence relations are nested.

Other logics. Description logics with transitive roles ($\mathcal{S}$), role hierarchies ($\mathcal{H}$), and inverse roles ($\mathcal{I}$) are expressive enough to define nested equivalence relations. Indeed, given transitive roles $T_1, T_2, \ldots$, one can axiomatise them as nested equivalence relations using the inclusions

$$\top \sqsubseteq \exists T_k.\top, \qquad T_k \sqsubseteq T_k^{-1}, \qquad T_k \sqsubseteq T_{k+1}.$$

Another formalism capable of defining families of nested equivalence relations is the Unary Negation Fragment extended with transitive roles, nominals, and role hierarchies. The work of [12] showed that this logic is 2-ExpTime-complete.

7 Concluding Remarks

We conclude with brief observations on slight extensions of our results.

Cross products. A limitation of GF is that it cannot express the following sentence:

$$\forall x, y \, \big(\big(\mathrm{Elephant}(x) \wedge \mathrm{Mouse}(y) \big) \rightarrow \mathrm{biggerThan}(x, y) \big),$$

known in the context of description logics as a *cross product* [27, 9, 8]. This sentence is, however, expressible in an extension of GF, named the *Triguarded Fragment* (TGF), introduced in [28]. TGF extends the syntax of GF (Definition 1) with the following rule:

- If $\psi(x, y) \in \mathrm{TGF}$ has two free variables, then both $\forall x \, \psi(x, y)$ and $\exists x \, \psi(x, y)$ are in TGF.

TGF captures both GF and FO^2, and in expressive power coincides with GFU (defined in Subsection 5.4). It turns out that our approach for $\mathrm{GF}[\mathcal{EQ}^{\subseteq}]$ extends to $\mathrm{TGF}[\mathcal{EQ}^{\subseteq}]$: roughly speaking, the resulting bounds mirror those of Theorem 5, but with deterministic complexity classes replaced by their nondeterministic counterparts.

Nominals. By Proposition 2, for every $K \geq 1$, the satisfiability problem for $\mathrm{GF}[K\text{-}\mathcal{EQ}^{\subseteq}]$ with equality is undecidable. Nevertheless, our proof techniques allow a restricted use of equality: the atoms $c = x$ and $c = c'$, where x is a variable and c, c' constants, pose no problem. Even such a restricted use of equality has natural applications; for example, in the standard translation of description logics with *nominals* (singleton concepts).

Undecidable extensions. Our decidability results do not extend to the Loosely Guarded Fragment (LGF) [15]. In fact, an arbitrary FO sentence embeds into LGF[1-$\mathcal{EQ}^{\subseteq}$] by employing $\bigwedge_{x,x' \in \bar{x}} E_1(x, x')$ as a loose guard for the quantifiers $\forall \bar{x}'$ and $\exists \bar{x}'$ with $\bar{x}' \subseteq \bar{x}$.

We cannot also support conjunctive queries: FO^2 embeds into GF^2[1-$\mathcal{EQ}^{\subseteq}$]; yet FO^2 with (negations of) conjunctive queries is undecidable [26].

References

1 Hajnal Andréka, István Németi, and Johan van Benthem. Modal languages and bounded fragments of predicate logic. *J. Philos. Log.*, 27(3):217–274, 1998. doi:10.1023/A:1004275029985.

2 Franz Baader, Ian Horrocks, Carsten Lutz, and Ulrike Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017. URL: http://www.cambridge.org/de/academic/subjects/computer-science/knowledge-management-databases-and-data-mining/introduction-description-logic?format=PB#17zVGeWD2TZUeu6s.97.

3 Vince Bárány, Georg Gottlob, and Martin Otto. Querying the guarded fragment. *Log. Methods Comput. Sci.*, 10(2), 2014. doi:10.2168/LMCS-10(2:3)2014.

4 Saguy Benaim, Michael Benedikt, Witold Charatonik, Emanuel Kieronski, Rastislav Lenhardt, Filip Mazowiecki, and James Worrell. Complexity of two-variable logic on finite trees. *ACM Trans. Comput. Log.*, 17(4):32, 2016. doi:10.1145/2996796.

5 Henrik Björklund and Mikolaj Bojanczyk. Shuffle expressions and words with nested data. In Ludek Kucera and Antonín Kucera, editors, *Mathematical Foundations of Computer Science 2007, 32nd International Symposium, MFCS 2007, Ceský Krumlov, Czech Republic, August 26-31, 2007, Proceedings*, volume 4708 of *Lecture Notes in Computer Science*, pages 750–761. Springer, 2007. doi:10.1007/978-3-540-74456-6_66.

6 Mikolaj Bojanczyk, Anca Muscholl, Thomas Schwentick, and Luc Segoufin. Two-variable logic on data trees and XML reasoning. *J. ACM*, 56(3):13:1–13:48, 2009. doi:10.1145/1516512.1516515.

7 Egon Börger, Erich Grädel, and Yuri Gurevich. *The classical decision problem*. Perspectives in Mathematical Logic. Springer, 1997.

8 Alexander Borgida. On the relative expressiveness of description logics and predicate logics. *Artif. Intell.*, 82(1-2):353–367, 1996. doi:10.1016/0004-3702(96)00004-5.

9 Pierre Bourhis, Michael Morak, and Andreas Pieris. Making cross products and guarded ontology languages compatible. In Carles Sierra, editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 880–886. ijcai.org, 2017. doi:10.24963/ijcai.2017/122.

10 Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981. doi:10.1145/322234.322243.

11 Witold Charatonik and Piotr Witkowski. Two-variable logic with counting and trees. *ACM Trans. Comput. Log.*, 17(4):31, 2016. doi:10.1145/2983622.

12 Daniel Danielski and Emanuel Kieronski. Finite satisfiability of unary negation fragment with transitivity. In Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen, editors, *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, Aachen, Germany, August 26-30, 2019*, volume 138 of *LIPIcs*, pages 17:1–17:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.MFCS.2019.17.

13 Oskar Fiuk, Emanuel Kieroński, and Vincent Michielini. Two-Variable Logic for Hierarchically Partitioned and Ordered Data. In *Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 316–325, October 2025. doi:10.24963/kr.2025/31.

14 Warren D. Goldfarb. The unsolvability of the godel class with identity. *J. Symb. Log.*, 49(4):1237–1252, 1984. doi:10.2307/2274274.

15 Erich Grädel. On the restraining power of guards. *J. Symb. Log.*, 64(4):1719–1742, 1999. doi:10.2307/2586808.

16 Ian Horrocks and Ulrike Sattler. A description logic with transitive and inverse roles and role hierarchies. *J. Log. Comput.*, 9(3):385–410, 1999. doi:10.1093/logcom/9.3.385.

17 Emanuel Kieronski. Results on the guarded fragment with equivalence or transitive relations. In C.-H. Luke Ong, editor, *Computer Science Logic, 19th International Workshop, CSL 2005, 14th Annual Conference of the EACSL, Oxford, UK, August 22-25, 2005, Proceedings*, volume 3634 of *Lecture Notes in Computer Science*, pages 309–324. Springer, 2005. doi:10.1007/11538363_22.

18 Emanuel Kieronski and Adam Malinowski. The triguarded fragment with transitivity. In Elvira Albert and Laura Kovács, editors, *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain, May 22-27, 2020*, volume 73 of *EPiC Series in Computing*, pages 334–353. EasyChair, 2020. doi:10.29007/z359.

19 Emanuel Kieronski, Jakub Michaliszyn, Ian Pratt-Hartmann, and Lidia Tendera. Two-variable first-order logic with equivalence closure. *SIAM J. Comput.*, 43(3):1012–1063, 2014. doi:10.1137/120900095.

20 Emanuel Kieronski and Martin Otto. Small substructures and decidability issues for first-order logic with two variables. *J. Symb. Log.*, 77(3):729–765, 2012. doi:10.2178/jsl/1344862160.

21 Emanuel Kieronski, Ian Pratt-Hartmann, and Lidia Tendera. Equivalence closure in the two-variable guarded fragment. *J. Log. Comput.*, 27(4):999–1021, 2017. doi:10.1093/logcom/exv075.

22 Emanuel Kieronski and Sebastian Rudolph. Finite model theory of the triguarded fragment and related logics. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–13. IEEE, 2021. doi:10.1109/LICS52264.2021.9470734.

23 Emanuel Kieronski and Lidia Tendera. Finite satisfiability of the two-variable guarded fragment with transitive guards and related variants. *ACM Trans. Comput. Log.*, 19(2):8:1–8:34, 2018. doi:10.1145/3174805.

24 Ian Pratt-Hartmann. *Fragments of First-Order Logic*. Oxford Logic Guides. Oxford University Press, United Kingdom, 2023.

25 Ian Pratt-Hartmann, Wieslaw Szwast, and Lidia Tendera. Quine's fluted fragment is non-elementary. In Jean-Marc Talbot and Laurent Regnier, editors, *25th EACSL Annual Conference on Computer Science Logic, CSL 2016, Marseille, France, August 29 - September 1, 2016*, volume 62 of *LIPIcs*, pages 39:1–39:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.CSL.2016.39.

26 Riccardo Rosati. The limits of querying ontologies. In Thomas Schwentick and Dan Suciu, editors, *Database Theory - ICDT 2007, 11th International Conference, Barcelona, Spain, January 10-12, 2007, Proceedings*, volume 4353 of *Lecture Notes in Computer Science*, pages 164–178, Berlin, Heidelberg, 2007. Springer. doi:10.1007/11965893_12.

27 Sebastian Rudolph, Markus Krötzsch, and Pascal Hitzler. All elephants are bigger than all mice. In Franz Baader, Carsten Lutz, and Boris Motik, editors, *Proceedings of the 21st International Workshop on Description Logics (DL2008), Dresden, Germany, May 13-16, 2008*, volume 353 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2008. URL: https://ceur-ws.org/Vol-353/RudolphKraetzschHitzler.pdf.

28 Sebastian Rudolph and Mantas Simkus. The triguarded fragment of first-order logic. In Gilles Barthe, Geoff Sutcliffe, and Margus Veanes, editors, *LPAR-22. 22nd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Awassa, Ethiopia, 16-21 November 2018*, volume 57 of *EPiC Series in Computing*, pages 604–619. EasyChair, 2018. doi:10.29007/m8ts.

29 Larry J. Stockmeyer. *The complexity of decision problems in automata theory and logic*. PhD thesis, Massachusetts Institute of Technology, USA, Cambridge, Massasuchets, USA, 1974. URL: http://hdl.handle.net/1721.1/15540.