Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Guarded Fragments Meet Dynamic Logic: The Story of Regular Guards

Bartosz Bednarczyk[1,2], **Emanuel Kieroński**[1]
[1]Institute of Computer Science, University of Wrocław
[2]Knowledge-Based Systems Group, TU Wien
{bartek, kiero}@cs.uni.wroc.pl

Abstract

We study the *Guarded Fragment with Regular Guards* (RGF), which combines the expressive power of the Guarded Fragment (GF) with Propositional Dynamic Logic with Intersection and Converse (ICPDL). Our logic generalizes, in a uniform way, many previously-studied extensions of GF, including (conjunctions of) transitive or equivalence guards, transitive or equivalence closure and more. We prove 2ExpTime-completeness of the satisfiability problem for RGF, showing that RGF is not harder than ICPDL or GF. Shifting to the query entailment problem, we provide undecidability results that significantly strengthen and solidify earlier results along those lines. We conclude by identifying, in a natural sense, the maximal ExpSpace-complete fragment of RGF.

1 Introduction

The *Guarded Fragment* (GF) of first-order logic (FO), introduced by Andréka et al. (1998), generalizes modal and description logics (DLs) to higher-arity relational vocabularies. Over the past 25 years, GF has become the canonical first-order fragment that balances expressive power with attractive model-theoretic properties, such as the finite model property (Grädel 1999), preservation theorems (Otto 2012), and robust decidability under various extensions involving fixed-point operators (Grädel and Walukiewicz 1999) or query languages (Bárány et al. 2014). Since classical (polyadic) multimodal and description logic formulæ embed naturally into GF via standard translations, this fragment serves as a versatile logical framework central to both theoretical studies and applications in knowledge representation and databases.

However, not all widely-used families of modal and description logics (DLs) are expressible within the scope of GF, as it cannot express properties such as transitivity or equivalence of relations. Consequently, translating transitive description logics like those from the $\mathcal{S}$ family of DLs or modal logics interpreted over equivalence frames, including S5, into the guarded fragment is not directly possible. To overcome this limitation, Ganzinger et al. (1999) initiated the study of *semantically-constrained guards*, an extension of GF allowing certain relations——confined to guards—to be interpreted with additional semantic constraints, notably transitivity or equivalence. This direction spurred intensive research, yielding several positive results, notably the 2ExpTime-completeness of GF extended with (conjunctions of) transitive guards (consult the works of Szwast&Tendera (2004), Kazakov (2006) and Kieroński&Rudolph (2021)), as well as the two-variable fragment of GF augmented by transitive or equivalence closures of binary guards, established by Michaliszyn and his co-authors (2009; 2017). Check Tendera's survey (2017) for a comprehensive overview. On the negative side, natural extensions of GF *with equality* ($\mathrm{GF}_{\approx}$), intended to capture popular description logics from the $\mathcal{SR}$ family, turned out to be undecidable. Examples include $\mathrm{GF}_{\approx}$ with exponentiation (regular expressions that are compositions of the same letter) (Michaliszyn 2008) or associative compositional axioms (Kazakov 2006). The decidability status of these logics without equality $\approx$ is still open. Consequently, there is no known decidable extension of GF with semantically-constrained guards capturing propositional dynamic logic (PDL) and its generalizations such as the $\mathcal{Z}$ family (Calvanese et al. 2009) of DLs, PDL with intersection and converse (ICPDL) (Göller et al. 2009), or higher-arity DLs such as $\mathcal{DLR}_{\mathsf{reg}}$ (Calvanese et al. 2008).

Our contributions. We introduce and study a novel logic called RGF, which extends (equality- and constant-free) GF by allowing ICPDL-programs as guards (Definition 3.1). Our main result establishes that $\mathsf{Sat}(\mathrm{RGF})$ is 2ExpTime-complete, matching the complexity of plain GF and ICPDL. This also lifts decidability to several logics where it was previously known only in their two-variable case. Our proof employs a *fusion* technique, reducing $\mathsf{Sat}(\mathrm{RGF})$ to instances of the satisfiability problem in plain GF and two-variable RGF, solving the latter via a careful encoding into ICPDL. We further address two questions: (i) Is the query entailment problem decidable for RGF? (ii) Is there an expressive fragment of RGF of complexity lower than 2ExpTime? We answer question (i) negatively, showing undecidability of conjunctive query entailment even for two-variable fluted GF with a single transitive guard, substantially strengthening prior results of Gottlob et al. (2013) in three ways: our logic is more restricted (belongs to the so-called fluted fragment), our queries do not use disjunction (we use conjunctive queries rather than the unions thereof), and our proof is also applicable to the finite-model scenario (which remained open). This transfers to the classical $\mathcal{ALC}$ extended with unqualified existential restrictions with intersection of the form $\exists(\mathrm{r} \cap \mathrm{s}).\top$, role-inclusion axioms

having the form $\mathrm{r} \subseteq \mathrm{s} \cup \mathrm{t}$, and a single transitivity statement. For (ii), we conduct a thorough case analysis, pinpointing when subfragments of RGF admit lower complexity than 2ExpTime. We conclude that a novel forward variant of GF extended with transitive closure is the largest (in some natural sense) ExpSpace-complete subfragment of RGF.

Our motivations. We explain our motivations behind the study of the GF with regular guards in the form of a Q&A.

☛ Why the guarded fragment (GF)?

Because GF is *the canonical* extension of modal and description logics to the setting of higher-arity relations (Grädel 1998), heavily investigated in the last 25 years. GF is not only well-behaved both computationally (Grädel 1999) and model-theoretically (Grädel and Otto 2014), but is also robust under extensions like fixed points (Grädel and Walukiewicz 1999) or semantically-constrained guards (Tendera 2017). It was studied also in the setting of knowledge representation in multiple recent papers (Lutz et al. 2024, Jung et al. 2021, Figueira et al. 2020, Zheng et al. 2020, Bourhis et al. 2017).

☛ Do we generalize any previously studied logics?

Yes, many of them. First, as GF encodes (via the standard translation, see *e.g.* Section 2.6.1 of Baader's textbook 2017) multi-modal and description logics (Grädel 1998), our logic also encodes (via an analogous translation) ICPDL and its subfragments such as $\mathcal{ALC}_{\mathsf{reg}}$ or $\mathcal{SRI}$ (Horrocks et al. 2006). There also exists a natural translation from (counting-free fragment of) $\mathcal{DLR}_{\mathsf{reg}}$ (Calvanese et al. 2008) to RGF. Second, there is a long tradition of studying GF extended with semantically-constrained guards (Tendera 2017), *i.e.* distinguished relations (available only as guards) interpreted as *transitive* (Ganzinger et al. 1999) or equivalence (Kieroński 2005) relations, or as transitive (Michaliszyn 2009) or equivalence (Kieroński et al. 2017) closures of another relation (that may also appear only as a guard). As one can simulate transitive (or equivalence) relation R with S^+ (resp. $(\mathrm{S} \cup \mathrm{S}^-)^*$) for a fresh relation S, our logic strictly extends all of the mentioned logics. Moreover, the mentioned papers concerning transitive and equivalence closures only focused on the extensions of GF^2, and hence our logic lifts them (without $\approx$) to the case of full GF and provides the tight complexity bound. Other ideas concern GF with exponentiation (regular expressions that are composition of the same letter) (Michaliszyn 2008) or associative compositional axioms (Kazakov 2006) (*i.e.* axioms of the form $\mathrm{R} \circ \mathrm{S} \subseteq \mathrm{T}$ for predicates R, S, T occurring only in guards). Both of them can be easily simulated in our framework. Finally, GF with conjunctions of transitive relations in guards (Kazakov 2006) can be expressed in RGF by employing $\cap$ operator. All of this makes our logic a desirable object of study.

☛ Are there any closely related but incomparable logics?

The closest logic is the Unary Negation Fragment (Segoufin and ten Cate 2013) with regular-path expressions (Christoph Jung et al. 2018) $\mathrm{UN}_{\mathsf{reg}}$, together with its very recent[1] generalizations with transitive closure operators (Figueira and Figueira 2025) and guarded negation (Nakamura 2025). All of these logics share 2ExpTime complexity of their satisfiability problem, but their expressive powers are incomparable. Indeed, RGF is not able to express conjunctive queries, while the other logics cannot express that R^*-reachable elements are B-connected. Yet another related logic is $\mathrm{GNFP}^{\mathsf{up}}$ by Benedikt et al. 2016, which extends the guarded negation fragment (Bárány, ten Cate, and Segoufin 2015) with fixed-point operators with unguarded parameters. The syntax of $\mathrm{GNFP}^{\mathsf{up}}$ is complicated, but the logic seems to embed ICPDL. Unfortunately, according to our understanding, such an encoding requires non-constant "pdepth" of the resulting formulæ, leading to a non-elementary complexity of the logic. The expressive powers of $\mathrm{GNFP}^{\mathsf{up}}$ and RGF are again incomparable and the separating examples are as before.

Extra examples, complete proofs, omitted definitions, conclusions, and almost all the content from Section 6 are delegated to the extended version of this paper (arXiv 2025).

[1] The papers by Figueira brothers and by Nakamura (still under review) appeared on arXiv during the preparation of this article.

2 Technical Background

Throughout the paper, we assume familiarity with the basics of finite model theory, formal languages, and complexity.

Models. We work with structures over a countably-infinite *equality-* and *constant-free* relational signature $\Sigma \coloneqq \Sigma_{\mathrm{FO}} \cup \Sigma_{\mathcal{R}}$, where all predicates in $\Sigma_{\mathcal{R}}$, called *regular predicates*, are binary. Fraktur letters denote structures; corresponding Roman letters denote their domains. For $\sigma \subseteq \Sigma$, a *σ-structure* interprets predicates in $\Sigma \setminus \sigma$ as $\emptyset$. A *sentence* is a formula without free variables. For a formula φ, we write $\varphi(\bar{x})$ to indicate that its free variables lie in $\bar{x}$. For a structure $\mathfrak{A}$ and a tuple $\bar{\mathrm{a}}$, we write $\mathfrak{A} \models \varphi[\bar{\mathrm{a}}]$ to indicate that $\varphi(\bar{x})$ holds in $\mathfrak{A}$ under the assignment $\bar{x} \mapsto \bar{\mathrm{a}}$. If $\mathfrak{A} \models \varphi$, we call $\mathfrak{A}$ a *model* of φ. A sentence is *(finitely) satisfiable* if it has a (finite) model. The *(finite) satisfiability problem* $(\mathsf{F})\mathsf{Sat}(\mathrm{L})$ for a logic L asks if an input L-sentence has a (finite) model. We say that L has the *finite model property* (FMP) if all satisfiable L-sentences admit finite models. Note that in this case, the finite and general satisfiability problems for L coincide.

Types. An *atomic k-type* over σ is a $\subseteq$-maximal satisfiable set of literals over σ with variables in $x_1, \ldots, x_k$. A k-tuple $\bar{\mathrm{a}}$ *realizes* a k-type γ in $\mathfrak{A}$ if $\mathfrak{A} \models \gamma[\bar{\mathrm{a}}]$; this unique type is denoted $\mathrm{tp}^{\mathfrak{A}}_k(\bar{\mathrm{a}})$. The sets of 1- and 2-types over the signature of φ are denoted $\boldsymbol{\alpha}_\varphi$ and $\boldsymbol{\beta}_\varphi$, with sizes exponential and doubly-exponential in $|\varphi|$. With superscripts $\cdot^{\mathfrak{A}}$ and $\cdot^{\mathrm{FO}}$ we restrict such sets to types realized in $\mathfrak{A}$ and the ones restricted to the Σ_{FO}-atoms. For a 2-type β, let β^{-1} be obtained by swapping its variables, and let $\beta \restriction x_i$ be the 1-type formed from literals involving only x_i. With a k-type γ over a signature σ we associate its *canonical structure* $\mathfrak{A}_\gamma$ over the domain $\{1, 2, \ldots, k\}$ with the interpretation of l-ary predicates P defined as: $\mathfrak{A}_\gamma \models \mathrm{P}[i_1, \ldots, i_l]$ iff $\mathrm{P}(x_{i_1}, \ldots, x_{i_l}) \in \gamma$. For a formula $\varphi(\bar{x})$ we write $\gamma \models \varphi$ to indicate that $\mathfrak{A}_\gamma \models \forall \bar{x}\, \varphi$.

Morphisms. With $\mathfrak{A} \hookrightarrow_{\mathfrak{f}} \mathfrak{B}$ we denote the existence of an injective homomorphism $\mathfrak{f}$ between $\mathfrak{A}$ and $\mathfrak{B}$ over a com-

mon vocabulary. If $\mathfrak{f}$ preserves 1-types, we call it *semi-strong* ($\mathfrak{A} \twoheadrightarrow_{\mathfrak{f}} \mathfrak{B}$). We omit $\mathfrak{f}$ if it is unimportant. A (union of) *conjunctive queries* (U)CQ is a (disjunction of) existentially quantified conjunction of atoms. The *(finite) entailment problem* for a logic L and a class of queries Q, asks whether an input L-sentence φ (finitely) entails an input Q-sentence q.

First-Order Fragments. We consider fragments of first-order logic FO over Σ_{FO}. In the *guarded fragment* GF, all quantification takes the form $\forall \bar{x}(\alpha \to \psi)$ and $\exists \bar{x}(\alpha \wedge \psi)$, where α (called a *guard*) is an atom featuring all the variables in $\bar{x}$ and all the free variables of ψ. E. Grädel (1999) proved:

Theorem 2.1. $\mathsf{Sat}(\mathrm{GF})$ *is* 2ExpTime*-complete and it has a procedure that given an input* GF*-sentence works in time bounded polynomially in the length of the input, exponentially in the size of the signature, and doubly exponentially in its width (maximal arity among all predicates in the input).* ◀

The *two-variable* fragment GF^2 of GF, allows for sentences featuring only the variables x_1, x_2 (re-quantification is permitted). We call φ *index-normal* if on any branch of its syntax tree, the i-th quantifier bounds precisely x_i. Such a φ is *fluted* (*forward*) if for any atom $\alpha(\bar{x})$ that occurs in the scope of a quantifier bounding x_n (but not x_{n+1}), the sequence $\bar{x}$ is a suffix (infix) of the sequence $x_1, x_2, \ldots, x_n$. To help the reader, we provide a few examples of forward formulæ:

$$\forall x_1\ (\mathrm{S}(x_1) \to \neg\forall x_2\ (\mathrm{P}(x_2) \to \mathrm{A}(x_1x_2))),$$
$$\forall x_1\ \mathrm{L}(x_1) \to \neg\exists x_2[\mathrm{P}(x_2) \wedge \forall x_3\ (\mathrm{S}(x_3) \to \mathrm{I}(x_1x_2x_3))].$$

Next, we present formulæ that are syntactically not forward. The red highlight means a mismatch in the variable ordering.

$$\forall x_1\forall x_2\forall x_3\ \mathrm{R}(x_1x_2) \wedge \mathrm{R}(x_2x_3) \to \mathrm{R}(\underline{x_1x_3}),$$
$$\forall x_1\ \mathrm{S}(x_1) \to \mathrm{R}(\underline{x_1x_1}), \qquad \forall x_1\forall x_2\ \mathrm{R}(x_1x_2) \leftrightarrow \mathrm{R}(\underline{x_2x_1}).$$

The forward GF (FGF), designed to generalize inversion-free DLs, restricts GF to forward index-normal sentences.

ICPDL. Given a logic L, we define L-*programs* π with:

$$\pi, \rho\ ::=\ \mathrm{B} \mid \bar{\mathrm{B}} \mid \pi{\circ}\rho \mid \pi{\cup}\rho \mid \pi{\cap}\rho \mid \pi^* \mid \pi^+ \mid \varphi?,$$

where $\mathrm{B} \in \Sigma_{\mathcal{R}}$ and φ is an L-formula with a sole free variable. For semantics consult Tab. 1. We call π *simple* if all its subformulæ $\varphi?$ (tests) are of the form U? for a unary U.
The *intersection width* $\mathrm{i}(\pi)$ of an L-program π is defined as:

$$\mathrm{i}(\mathrm{B}) := \mathrm{i}(\bar{\mathrm{B}}) := \mathrm{i}(\varphi?) := 1,\ \mathrm{i}(\pi^*) := \mathrm{i}(\pi^+) := \mathrm{i}(\pi),$$
$$\mathrm{i}(\pi{\cup}\rho) := \mathrm{i}(\pi{\circ}\rho) := \max(\mathrm{i}(\pi), \mathrm{i}(\rho)), \mathrm{i}(\pi{\cap}\rho) := \mathrm{i}(\pi){+}\mathrm{i}(\rho).$$

The *iwidth* $\mathrm{I}(\varphi)$ of a formula φ is 1 if φ contains no programs; otherwise it is the maximum of $\mathrm{i}(\pi)$ over all programs π in φ. Formulæ of *PDL with Intersection and Converse* (ICPDL) by Harel et al. (1982) are defined according to the grammar:

$$\varphi, \psi ::= \top \mid \bot \mid \mathrm{U} \mid \neg\varphi \mid \varphi \wedge \psi \mid \langle\pi\rangle\varphi \mid [\pi]\,\varphi,$$

for $\mathrm{U} \in \Sigma_{\mathrm{FO}}$ and ICPDL-programs π. $\mathfrak{A}$ is a *model* of φ in ICPDL if $\mathfrak{A} \models \varphi[\mathrm{a}]$ for some $\mathrm{a} \in A$. S. Göller et al. (2008, Th. 3.28; 2009, Th. 4.8) proved that:

Theorem 2.2. *There exists an exponential function* exp *such that the satisfiability of each sentence* $\varphi \in$ ICPDL *can be decided in time bounded by* $\exp\left(|\varphi|^{\mathrm{I}(\varphi)}\right)$. ◀

Name	Syntax of π	Semantics $\pi^{\mathfrak{A}}$ of π in a structure $\mathfrak{A}$
Test / Pred.	$\varphi?$ / B	$\{(\mathrm{a},\mathrm{a}) \mid \mathfrak{A} \models \varphi[\mathrm{a}]\}$ / Binary relation
Converse op.	$\bar{\pi}$	$\{(\mathrm{b},\mathrm{a}) \mid (\mathrm{a},\mathrm{b}) \in \pi^{\mathfrak{A}}\}$
Concat.	$\pi{\circ}\rho$	$\{(\mathrm{a},\mathrm{c}) \mid \exists \mathrm{b}.(\mathrm{a},\mathrm{b}) \in \pi^{\mathfrak{A}} \wedge (\mathrm{b},\mathrm{c}) \in \rho^{\mathfrak{A}}\}$
Union/Inter.	$\pi{\cup}\rho$ / $\pi{\cap}\rho$	$\pi^{\mathfrak{A}} \cup \rho^{\mathfrak{A}}$ / $\pi^{\mathfrak{A}} \cap \rho^{\mathfrak{A}}$
Kleene $*$/$+$	π^* / π^+	$\bigcup_{i=0}^{\infty}(\pi^i)^{\mathfrak{A}}$ / $\bigcup_{i=1}^{\infty}(\pi^i)^{\mathfrak{A}}$, where $\pi^0 := \top?$ and $\pi^{i+1} := (\pi^i) \circ \pi$.
Name	**Syntax of φ**	**Given $\mathfrak{A}$ and $\mathrm{a} \in A$, $\mathfrak{A} \models \varphi[\mathrm{a}]$ if**
Top/Bot/Pred.	$\top$ / $\bot$ / U	always / never / $\mathrm{a} \in \mathrm{U}^{\mathfrak{A}}$,
Neg./Conj.	$\neg\varphi$ / $\varphi{\wedge}\psi$	$\mathfrak{A} \not\models \varphi[\mathrm{a}]$ / $\mathfrak{A} \models \varphi[\mathrm{a}]$ and $\mathfrak{A} \models \psi[\mathrm{a}]$,
Modalities	$\langle\pi\rangle\varphi$ / $[\pi]\,\varphi$	$\mathfrak{A} \models \varphi[\mathrm{b}]$ for some/all b with $(\mathrm{a},\mathrm{b}) \in \pi^{\mathfrak{A}}$.

Table 1: Semantics of paths and formulæ of ICPDL.

3 Guarded Fragment with Regular Guards

We now introduce RGF, the *Guarded Fragment with Regular Guards*, a novel and highly-expressive logic extending GF by allowing for ICPDL-programs as binary guards. Formally:

Definition 3.1 (RGF)**.** *An* RGF*-guard* ϑ *for a formula* φ *is either an atom over* Σ_{FO} *or* $\pi(xy)$ *for some* RGF*-program* π*, such that free variables of* ϑ *include all free variables of* φ*. The set* RGF *of* RGF*-formulæ is defined with the grammar:*

$$\varphi, \varphi'\ ::=\ \mathrm{A}(\bar{x}) \mid \neg\varphi \mid \varphi \wedge \varphi' \mid \exists x\,\varphi(x) \mid \exists \bar{x}(\vartheta \wedge \varphi),$$

where $\mathrm{A} \in \Sigma_{\mathrm{FO}}$ *and* ϑ *is an* RGF*-guard for* φ*.* ◀

Standard connectives and the universal quantifier are defined as usual. For each $\mathrm{S} \subseteq \{\cdot^-, \circ, \cup, \cap, \cdot^*, \cdot^+, ?\}$ we let $\mathrm{RGF}[\mathrm{S}]$ denote RGF-formulae with operators in programs in guards restricted to S. Additionally, $\mathrm{FRGF}[\mathrm{S}]$ is the forward restriction of $\mathrm{RGF}[\mathrm{S}]$ and $\mathrm{FRGF}^2[\mathrm{S}]$ is its two-variable fragment.

Our logic generalizes a plethora of extensions of GF with semantically-constrained guards (consult the introduction for details). In particular, transitive relations in guards can be simulated in $\mathrm{RGF}[\cdot^+]$ using R^+ for a fresh binary relation R. Hence, GF+TG, the extension of GF with transitive guards, is a fragment of RGF. We explain our design choices below.

☛ Why is the signature separated, *i.e.* $\Sigma := \Sigma_{\mathrm{FO}} \cup \Sigma_{\mathcal{R}}$?
This is to ensure that binary predicates from $\Sigma_{\mathcal{R}}$ appear *only in guards*; otherwise, even the two-variable guarded fragment with transitivity is undecidable (Ganzinger et al. 1999, Th. 2).
☛ Why is the equality symbol $\approx$ excluded from Σ?
Its inclusion makes our logic undecidable. This holds already for GF^2 with assoc. compositional axioms (Kazakov 2006, Th. 5.3.1), conjunctions of transitive guards (Kazakov 2006, Th. 5.3.2), or exponentiation (Michaliszyn 2008, Th. 3.1).
☛ Why are constant symbols excluded from Σ?
They are expected to preserve decidability, especially given that Theorem 2.2 extends to Hybrid ICPDL. We leave a full analysis to the journal version of this paper.

4 Solving the Satisfiability Problem

We solve $\mathsf{Sat}(\mathrm{RGF})$ in 2ExpTime in two steps. First, Section 4.1 addresses $\mathsf{Sat}(\mathrm{RGF}^2)$ by carefully encoding it into ICPDL, producing an exponentially larger formula with only polynomial iwidth. This translation utilizes models realizing few 2-types and tables, *i.e.* lists of 2-types that can be globally realized by pairs connected via any subsets of guards.

Then, Section 4.2 handles the general case via a reduction to instances of $\mathsf{Sat}(\mathrm{RGF}^2)$ and $\mathsf{Sat}(\mathrm{GF})$, based on a *fusion* — a generalization of a method by Kieroński&Rudolph (2021) for GF with transitive guards. Its correctness hinges on the restriction that regular predicates are confined to guards.
As usual in the context of the satisfiability problem, we can w.l.o.g. restrict attention to formulæ in a handy normal form.

Definition 4.1. *An* RGF*-sentence φ is in* normal form *(NF) if $\varphi = \forall x \lambda(x) \wedge \bigwedge_i \varphi_i$ where each φ_i has one of the forms:*
$\bullet\ \forall \bar{x}\, \eta_i^{\exists}(\bar{x}) \to \exists \bar{y}\, \vartheta_i^{\exists}(\bar{x}\bar{y}) \wedge \psi_i^{\exists}(\bar{x}\bar{y})$, $\bullet\ \forall \bar{x}\, \eta_i^{\forall}(\bar{x}) \to \psi_i^{\forall}(\bar{x})$,
$\bullet\ \forall x_1\, \gamma_i^{\exists}(x_1) \to \exists x_2\, \pi_i^{\exists}(x_1x_2) \wedge \phi_i^{\exists}(x_1x_2)$,
$\bullet\ \forall x_1 \forall x_2\, \pi_i^{\forall}(x_1x_2) \to \phi_i^{\forall}(x_1x_2)$,
for (decorated) Σ_{FO}-atoms η, ϑ, γ, simple RGF*-programs π, and quantifier-free ψ, ϕ, λ over Σ_{FO} with $\bar{x} \cap \bar{y} = \emptyset$.* ◀

We refer to above types of conjuncts as $\exists^{\mathrm{FO}}$-, $\forall^{\mathrm{FO}}$-, $\exists^{\mathcal{R}}$-, and $\forall^{\mathcal{R}}$-conjuncts. We call a k-type γ FO*-compatible* with φ if $\gamma \models \lambda$ and for each $\forall^{\mathrm{FO}}$-conjunct: $\gamma \models \eta_i^{\forall}$ implies $\gamma \models \psi_i^{\forall}$.

Lemma 4.2. *Consider* $\mathrm{FRGF}[\cdot^+] \subseteq \mathrm{L} \subseteq \mathrm{RGF}$. *Then if the satisfiability problem for* L*-sentences in NF is decidable in* 2ExpTime *(resp.* ExpSpace*) then so is* $\mathsf{Sat}(\mathrm{L})$. ◀

4.1 The Two-Variable Case

Consider $\varphi \in \mathrm{RGF}^2$ in NF over a signature σ with its usual components (*cf.* Def. 4.1). Let $\mathrm{bg}(\varphi) \coloneqq \{\vartheta_1, \dots, \vartheta_K\}$ collect 2-variable guards of φ, and $\mathrm{pr}(\varphi)$ be its restriction to programs. We focus on *sparse models* of φ, *i.e.* models realizing only exponentially-many ($\leq 2^{12|\varphi|}$) 2-types from $\boldsymbol{\beta}_\varphi^{\mathrm{FO}}$. Their presence helps keep the size of our translation in control.

Lemma 4.3. *Suppose φ is (finitely) satisfiable, and let $\mathfrak{A}$ be its (finite) model φ. Then φ also has a (finite) sparse model $\mathfrak{B}$ such that: (i) $\mathfrak{A}$ and $\mathfrak{B}$ realise precisely the same 1-types, and (ii) all 2-types realised in $\mathfrak{B}$ are also realised in $\mathfrak{A}$.* ◀

Sketch. We use a filtration-based argument. Let $\mathfrak{A} \models \varphi$. Consider $\sigma^* \coloneqq \sigma \cup \{\mathrm{G}_1, \dots, \mathrm{G}_K\}$ for fresh binary G_is, and let $\mathfrak{A}^*$ be the σ^*-expansion of $\mathfrak{A}$ interpreting each G_i in $\mathfrak{A}$ as the set of these pairs (a, b) with $\mathfrak{A} \models \vartheta_i[\mathrm{a}, \mathrm{b}]$. Denoting by $\boldsymbol{\beta}^*$ the set of all 2-types over $(\sigma^* \cap \Sigma_{\mathrm{FO}})$ *realized* in $\mathfrak{A}^*$, we define an equivalence relation $\sim$ on it as follows: $\beta \sim \beta'$ if

- β and β' are equal when restricted to literals involving either only a single variable or a binary predicate; and
- for all conjuncts $\forall x_1\, \eta(x_1) \to \exists x_2\, \vartheta_i(x_1x_2) \wedge \psi_i(x_1x_2)$ of φ: $\beta \models \mathrm{G}_i(x_1x_2) \wedge \psi_i$ if and only if $\beta' \models \mathrm{G}_i(x_1x_2) \wedge \psi_i$ holds, and similarly for β^{-1} and β'^{-1} in place of β and β'.

Fix a representative for each $\sim$-class (of $\leq 2^{12|\varphi|}$ many) and denote the chosen member in $[\beta]_\sim$ by $\beta^\star$. A sparse model is obtained by replacing all 2-types β between any unordered pair of elements in $\mathfrak{A}^*$ with the corresponding 2-type $\beta^\star$. □

We reduce the satisfiability of φ to the existence of a φ*-table* $\mathcal{T} \colon (\boldsymbol{\alpha}_\varphi^{\mathrm{FO}} \times \mathcal{P}(\mathrm{pr}(\varphi)) \times \boldsymbol{\alpha}_\varphi^{\mathrm{FO}}) \to \mathcal{P}(\boldsymbol{\beta}_\varphi^{\mathrm{FO}})$ governing realizable 2-types in an intended sparse model of φ and satisfiability of $\varphi_{\mathcal{T}}$ in ICPDL governing its intended shape. We use unary U_α and binary B_β indexed by types $\alpha \in \boldsymbol{\alpha}_\varphi^{\mathrm{FO}}, \beta \in \boldsymbol{\beta}_\varphi^{\mathrm{FO}}$, intended to represent types of (pairs of) elements. Based on them, we produce $\varphi_{\mathcal{T}}$ incorporating obvious properties of sparse models, *e.g.* that every element has its 1-type and its witnesses (realizing some 2-type and its inverse) for all $\exists^{\mathrm{FO}}$- and $\exists^{\mathcal{R}}$-conjuncts of φ. The main issue is that ICPDL cannot express that any two elements are connected by some binary relation, hence models of $\varphi_{\mathcal{T}}$ have only *partially* assigned 2-types. To lift them to the actual models of φ by "completing" the missing 2-types, we employ φ-tables $\mathcal{T}$. Given its input $(\alpha_{\mathrm{a}}, \mathrm{S}, \alpha_{\mathrm{b}})$, a map $\mathcal{T}$ outputs possible 2-types β realizable by pairs (a, b) with 1-types α_{a} and α_{b} in a hypothetical model of φ, provided that a and b are connected by (at least) all the guards in S. To qualify as a φ-table, $\mathcal{T}$ must satisfy the four conditions listed below, met by any sparse model of φ.

(Size) The image $\mathrm{im}(\mathcal{T})$ of $\mathcal{T}$ is bounded by $2^{12|\varphi|}$.
(Clo) If $\beta \in \mathrm{im}(\mathcal{T})$ then $\mathrm{im}(\mathcal{T})$ contains β^{-1} and the *symmetric 2-types β_α induced by* each 1-type $\alpha \in \{\beta{\restriction}x_1, \beta{\restriction}x_2\}$, namely the unique 2-types β_α satisfying $\beta_\alpha[x_2 \mapsto x_1] \equiv \alpha$.
(Com) If $\beta \in \mathcal{T}(\alpha_1, \mathrm{S}, \alpha_2)$ then $\beta{\restriction}x_1 = \alpha_1$ and $\beta{\restriction}x_2 = \alpha_2$ and the types $\alpha_1, \alpha_2, \beta$ are FO-compatible with φ.
(Sat) For all $\forall^{\mathcal{R}}$-conjuncts $\forall x_1x_2\, \pi_i^{\forall}(x_1x_2) \to \phi_i^{\forall}$ from φ and all $\beta \in \mathcal{T}(\alpha_1, \mathrm{S}, \alpha_2)$ we have $\beta \models \phi_i^{\forall}$ if $\pi_i^{\forall}(x_1x_2) \in \mathrm{S}$.

Every $\mathfrak{A}$ naturally defines its table $\mathcal{T}_{\mathfrak{A}}$, where $\mathcal{T}_{\mathfrak{A}}(\alpha_1, \mathrm{S}, \alpha_2)$ consists of all $\mathrm{tp}_2^{\mathfrak{A}}(\mathrm{a}_1, \mathrm{a}_2)$ over all $\mathrm{a}_1, \mathrm{a}_2$ from $\mathfrak{A}$ witnessing $\mathrm{tp}_1^{\mathfrak{A}}(\mathrm{a}_i) = \alpha_i$ and $\mathfrak{A} \models \vartheta[\mathrm{a}_1, \mathrm{a}_2]$ for all guards $\vartheta \in \mathrm{S}$. An $\mathfrak{A}$ satisfies $\mathcal{T}$ ($\mathfrak{A} \models \mathcal{T}$) if $\mathcal{T}_{\mathfrak{A}} \subseteq \mathcal{T}$. Clearly, if $\mathfrak{A} \models \varphi$ then $\mathcal{T}_{\mathfrak{A}}$ is a φ-table but not all $\mathfrak{A}$ satisfying φ-tables are models of φ.

We can now translate φ and $\mathcal{T}$ into $\varphi_{\mathcal{T}} \coloneqq [\star]\, (\varphi_{\mathcal{T}}^1 \wedge \ldots \wedge \varphi_{\mathcal{T}}^7)$ in ICPDL, where the formulæ $\varphi_{\mathcal{T}}^i$ with informal explainations are given below and $\star$ is the "universal modality", *i.e.*

$$\star \coloneqq \big(\bigcup_{\beta \in \mathrm{im}(\mathcal{T})} \bigcup_{\mathrm{R} \in \sigma \cap \Sigma_{\mathcal{R}}} \mathrm{B}_\beta \cup \overline{\mathrm{B}_\beta} \cup \mathrm{R} \cup \overline{\mathrm{R}}\big)^*.$$

For a program π, let $\pi_{\mathcal{T}}$ be the result of replacing each unary U in π with the union of U_α over all $\alpha \models \mathrm{U}(x_1)$ from $\boldsymbol{\alpha}_\varphi^{\mathrm{FO}}$.
1. Each element has a unique 1-type, and each pair of elements shares at *most one* 2-type (existence of a 2-type is not expressible, justifying the use of φ-tables by our reduction).

$$\varphi_{\mathcal{T}}^1 \coloneqq \forall \bigwedge_{\alpha \in \boldsymbol{\alpha}_\varphi^{\mathrm{FO}}} \mathrm{U}_\alpha \wedge \bigwedge_{\beta \neq \beta' \in \mathrm{im}(\mathcal{T})} [\mathrm{B}_\beta \cap \mathrm{B}_{\beta'}] \bot$$

2. Each element *self*-realize its unique (symmetric) 2-type.

$$\varphi_{\mathcal{T}}^2 \coloneqq \bigwedge_{\alpha \in \boldsymbol{\alpha}_\varphi^{\mathrm{FO}}} \big(\mathrm{U}_\alpha \to \langle \mathrm{B}_{\beta_\alpha} \cap \overline{\mathrm{B}_{\beta_\alpha}} \cap \top? \rangle \top\big)$$

3. Compatibility of 2-types with their inverses and induced 1-types. The big conjunction is over $\beta, \beta' \in \mathrm{im}(\mathcal{T})$: $\beta' \neq \beta^{-1}$.

$$\varphi_{\mathcal{T}}^3 \coloneqq \bigwedge \big[\mathrm{B}_\beta \cap \overline{\mathrm{B}_{\beta'}}\,\big] \bot \wedge [\neg \mathrm{U}_{\beta{\restriction}x_1}? \circ \mathrm{B}_\beta \cup \mathrm{B}_\beta \circ \neg \mathrm{U}_{\beta{\restriction}x_2}?] \bot$$

4&5. Ensuring witnesses for all $\exists^{\mathrm{FO}}$- and $\exists^{\mathcal{R}}$-conjuncts of φ.

$$\varphi_{\mathcal{T}}^4 \coloneqq \bigwedge_{i, \alpha \models \eta_i^{\exists}} \mathrm{U}_\alpha \to \bigvee_{\beta \in \mathrm{im}(\mathcal{T}), \beta \models \vartheta_i^{\exists}, \beta \models \psi_i^{\exists}} \langle \mathrm{B}_\beta \cap \overline{\mathrm{B}_{\beta^{-1}}} \rangle \top$$

$$\varphi_{\mathcal{T}}^5 \coloneqq \bigwedge_{i, \alpha \models \gamma_i^{\exists}} \mathrm{U}_\alpha \to \bigvee_{\beta \in \mathrm{im}(\mathcal{T}), \beta \models \phi_i^{\exists}} \langle (\pi_i^{\exists})_{\mathcal{T}} \cap \mathrm{B}_\beta \cap \overline{\mathrm{B}_{\beta^{-1}}} \rangle \top$$

6. No pair (a, b) of elements in the intended model are B_β-connected for a 2-type β with either β or β^{-1} violating $\mathcal{T}$.

$$\varphi_{\mathcal{T}}^6 \coloneqq \bigwedge \big[(\mathrm{U}_\alpha? \circ (\bigcup \mathrm{B}_\beta) \circ \mathrm{U}_{\alpha'}?) \cap \bigcap_{\pi \in \mathrm{S}} \pi_{\mathcal{T}} \cap \bigcap_{\pi \in \bar{\mathrm{S}}} \overline{\pi_{\mathcal{T}}}\big] \bot$$

Above, $\bigwedge$ is over $\alpha, \alpha' \in \boldsymbol{\alpha}_\varphi^{\mathrm{FO}}$ (1-types of a, b), $\mathrm{S}, \bar{\mathrm{S}} \subseteq \mathrm{pr}(\varphi)$ (guards from a to b and from b to a), and $\bigcup$ ranges over $\beta \in \mathrm{im}(\mathcal{T})$ satisfying $\beta \notin \mathcal{T}(\alpha, \mathrm{S}, \alpha')$ or $\beta^{-1} \notin \mathcal{T}(\alpha', \mathrm{S}', \alpha)$.
7. There should always be an available 2-type and its inverse to assign to each pair of elements in the intended model.

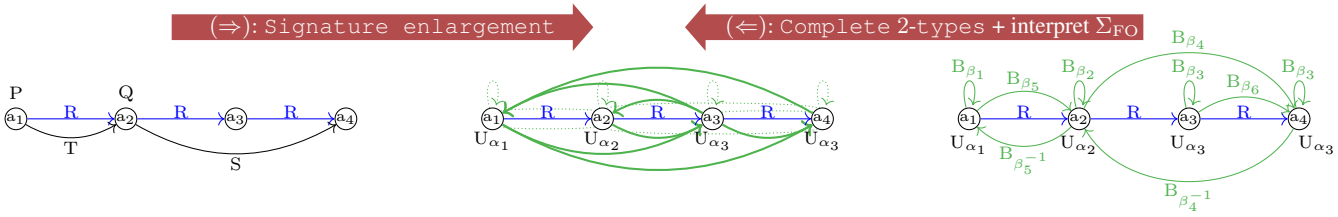

Figure 1: Structures $\mathfrak{A}$ (left), $\mathfrak{A}^+$ (middle), and $\mathfrak{A}^*$ (right) from our example illustrating the proof of Theorem 4.4.

$$\varphi^7_{\mathcal{T}} := \bigwedge \left[(\mathrm{U}_\alpha ? \circ \star \circ \mathrm{U}_{\alpha'} ?) \cap \bigcap\nolimits_{\pi \in \mathrm{S}} \pi_{\mathcal{T}} \cap \bigcap\nolimits_{\pi \in \mathrm{S}'} \overline{\pi_{\mathcal{T}}}\right] \bot$$

As in the 6th case, $\bigwedge$ iterate $\alpha, \alpha', \mathrm{S}, \mathrm{S}'$ with $\mathcal{T}(\alpha, \mathrm{S}, \alpha')$ either empty or containing no β with $\beta^{-1} \in \mathcal{T}(\alpha', \mathrm{S}', \alpha)$. ◀

Observe that $|\varphi_{\mathcal{T}}|$ is exponential in $|\varphi|$, but the iwidth of $\varphi_{\mathcal{T}}$ is bounded by $2|\mathrm{pr}(\varphi)|+2$ (hence only polynomially in $|\varphi|$). Thus, by Theorem 2.2, satisfiability of each $\varphi_{\mathcal{T}}$ can be tested in doubly-exponential time in $|\varphi|$. By enumerating possible φ-tables $\mathcal{T}$ and testing each $\varphi_{\mathcal{T}}$ for satisfiability yields the optimal algorithm (in 2ExpTime) for satisfiability of φ.

Theorem 4.4. *φ is satisfiable iff there is a φ-table $\mathcal{T}$ with a satisfiable $\varphi_{\mathcal{T}}$. Thus,* $\mathsf{Sat}(\mathrm{RGF}^2)$ *is* 2ExpTime*-compl.* ◀

Sketch. For a (sparse) $\mathfrak{A} \models \varphi$, let $\mathfrak{A}^+$ be its expansion that interpret each U_α (resp.B_β) as the set of elements (resp. pairs) satisfying α (resp. β). Clearly $\mathfrak{A}^+ \models \varphi_{\mathcal{T}}$ for $\mathcal{T} := \mathcal{T}_{\mathfrak{A}}$. Conversely, take a connected $\mathfrak{A}^* \models \varphi_{\mathcal{T}}$ for some $\mathcal{T}$. We "complete" $\mathfrak{A}^*$ by resolving all typeless pairs, *i.e.* $(\mathrm{a}_1, \mathrm{a}_2)$ such that $\mathfrak{A} \not\models \mathrm{B}_\beta[\mathrm{a}_1, \mathrm{a}_2]$ for any β. For each such pair, let α_i be the unique 1-type with $\mathfrak{A}^* \models \mathrm{U}_{\alpha_i}[\mathrm{a}_i]$. If some β satisfies $\mathfrak{A}^* \models \mathrm{B}_\beta[\mathrm{b}, \mathrm{a}]$, we add (a, b) to $(\mathrm{B}_{\beta^{-1}})^{\mathfrak{A}^*}$. Otherwise, fix a $\beta \in \mathcal{T}(\alpha_1, \mathrm{S}, \alpha_2)$, where $\mathrm{S} \subseteq \mathrm{pr}(\varphi)$ collects all guards ϑ satisfied by (a, b), and then include (a, b) in $(\mathrm{B}_\beta)^{\mathfrak{A}^*}$. Repeat until no typeless pair remain. $\mathfrak{A}^*$ becomes a model $\mathfrak{A}^+$ of φ after interpreting Σ_{FO} minimally to satisfy $\mathfrak{A}^+ \models \mathrm{B}_\beta[\mathrm{a}, \mathrm{b}]$ iff $\mathfrak{A}^* \models \beta[\mathrm{a}, \mathrm{b}]$ for all $\beta \in \mathrm{im}(\mathcal{T})$ and all $\mathrm{a}, \mathrm{b} \in A^*$. □

To illustrate the proof of Theorem 4.4, consider φ (over the signature $\sigma := \{\mathrm{P}, \mathrm{Q}, \mathrm{S}, \mathrm{T}, \mathrm{R}\}$ with R being the only regular predicate) composed of the following four conjuncts:
- $\forall x_1\, \mathrm{Q}(x_1) \to \exists x_2\, \mathrm{R}^+(x_1 x_2) \wedge \mathrm{S}(x_1 x_2)$,
- $\forall x_1\, \mathrm{Q}(x_1) \to \exists x_2\, \bar{\mathrm{R}}(x_1 x_2) \wedge \mathrm{P}(x_2)$,
- $\forall x_1\, \mathrm{P}(x_1) \to \exists x_2\, \mathrm{T}(x_1 x_2) \wedge \mathrm{Q}(x_2)$,
- $\forall x_1 \forall x_2\, \mathrm{R}(x_1 x_2) \to \neg \mathrm{S}(x_1 x_2)$.

When starting from $\mathfrak{A} \models \varphi$ (left part of Fig. 1), we construct its expansion $\mathfrak{A}^+ \models \varphi_{\mathcal{T}_{\mathfrak{A}}}$, in which (i) each a_i satisfies U_α, where $\alpha := \mathrm{tp}^{\mathfrak{A}}_1(\mathrm{a}_i)$, and (ii) for each pair $(\mathrm{a}_i, \mathrm{a}_j)$, the tuple $(\mathrm{a}_i, \mathrm{a}_j)$ satisfies B_β and $(\mathrm{a}_j, \mathrm{a}_i)$ satisfies $\mathrm{B}_{\beta^{-1}}$, where $\beta := \mathrm{tp}^{\mathfrak{A}}_2(\mathrm{a}_i, \mathrm{a}_j)$ (dotted and solid green arrows). For the other direction, we start with $\mathfrak{A}^* \models \varphi_{\mathcal{T}}$ (right part of Fig. 1). Here:

- Some U_α are assigned to the a_i due to $\varphi^1_{\mathcal{T}}$: α_1 contains $\mathrm{P}(x_1)$, α_2 contains $\mathrm{Q}(x_1)$ (both contain no other positive atoms), and α_3 contains no positive atoms.
- Some B_β are assigned to the pairs $(\mathrm{a}_i, \mathrm{a}_i)$ due to $\varphi^2_{\mathcal{T}}$: β_1, β_2, β_3 are the trivial 2-types without positive binary atoms, compatible with the 1-types of the corresponding a_i.
- β_4 contains $\mathrm{S}(x_1, x_2)$ (and no other positive binary atoms); B_{β_4} is assigned to $(\mathrm{a}_2, \mathrm{a}_4)$ and $\mathrm{B}_{\beta_4^{-1}}$ to $(\mathrm{a}_4, \mathrm{a}_2)$ due to $\varphi^4_{\mathcal{T}}$, providing the first witness for a_2.
- β_5 contains $\mathrm{T}(x_1, x_2)$ (and no other positive binary atoms); B_{β_5} is assigned to $(\mathrm{a}_1, \mathrm{a}_2)$ and $\mathrm{B}_{\beta_5^{-1}}$ to $(\mathrm{a}_2, \mathrm{a}_1)$ due to $\varphi^5_{\mathcal{T}}$, providing the witness for a_1.
- The second witness for a_2 is provided by $\varphi^4_{\mathcal{T}}$, which requires some B_β to be assigned to $(\mathrm{a}_2, \mathrm{a}_1)$ and $\mathrm{B}_{\beta^{-1}}$ to $(\mathrm{a}_1, \mathrm{a}_2)$. By $\varphi^3_{\mathcal{T}}$, this β must be equal to β_5^{-1}.
- β_6 is the 2-type, compatible with the 1-type α_3 of its endpoints, containing no positive binary atoms. Assigning B_{β_6} to $(\mathrm{a}_3, \mathrm{a}_4)$ is not enforced by $\varphi_{\mathcal{T}}$, but it is not forbidden.

We first append $(\mathrm{a}_4, \mathrm{a}_3)$ to (the interpretation of) $\mathrm{B}_{\beta_6^{-1}}$. Then, for each pair not yet contained in any B_β, take a $\beta \in \mathrm{im}(\mathcal{T})$ compatible with its 1-types, and assign B_β to this pair; for instance, one may choose 2-types containing no positive atoms. (The newly added connections are represented by solid green arrows in Fig. 1). This structure becomes a model $\mathfrak{A}^+ \models \varphi$ after assigning the 1-types and 2-types over the original signature σ of φ in accordance with the U_α and B_β.

4.2 The Multi-Variable Case

To lift the result from the previous section to the general case, we generalize the approach of Kieroński&Rudolph (2021), originally developed for the finite satisfiability problem in GF+TG. Fix an RGF-sentence φ in NF with its usual components (*cf.* Def. 4.1). The core idea is to decompose φ into two parts: a sentence φ_{FO} in the guarded fragment and a sentence $\varphi_{\mathcal{R}}$ in RGF^2, augmented with some additional synchronizing information concerning the set of realizable 1-types $\boldsymbol{\alpha}_0$ and (first-order) 2-types $\boldsymbol{\beta}_0$. Both formulæ express that precisely all 1-types from $\boldsymbol{\alpha}_0$ are realized. However, they differ in how they handle the realization of 2-types from $\boldsymbol{\beta}_0$: $\varphi_{\mathcal{R}}$ asserts that no 2-type outside of $\boldsymbol{\beta}_0$ is realized, whereas φ_{FO} requires the realization of all 2-types in $\boldsymbol{\beta}_0$ (and possibly others as well). We then verify the satisfiability of φ_{FO} and $\varphi_{\mathcal{R}}$ separately and, by carefully combining multiple copies of models of these two sentences, we construct a model of φ. Our approach heavily relies on the fact that regular predicates occur only in guards (otherwise our logic is undecidable).
For a 2-type $\beta \in \boldsymbol{\beta}_\varphi$, we define β^- (called the *regular-free reduction* of β) as the set of formulæ obtained by simultaneously removing from β all $\Sigma_{\mathcal{R}}$-literals involving both variables x_1 and x_2 (preserving those involving only a single variable). For a structure $\mathfrak{A}$, we use $\mathfrak{A}^-$ to denote its substructure obtained by restricting the interpretation of

all regular predicates to equal-element tuples (*i.e.* self-loops).

The formulæ. We first construct the formulæ φ_{FO} and $\varphi_{\mathcal{R}}$, parametrized by (FO-compatible with φ, parametrization in names omitted due to space reasons) sets of types $\boldsymbol{\alpha}_0 \subseteq \boldsymbol{\alpha}_\varphi$ and $\boldsymbol{\beta}_0 \subseteq \boldsymbol{\beta}_\varphi$. Define $\varphi_{\mathrm{FO}} \coloneqq \varphi_{\mathrm{FO},1} \wedge \varphi_{\mathrm{FO},2} \wedge \varphi_{\mathrm{FO},3}$, where

1. Our $\varphi_{\mathrm{FO},1}$ comprises all $\forall^{\mathrm{FO}}$- and $\exists^{\mathrm{FO}}$-conjuncts of φ.

2. The formula $\varphi_{\mathrm{FO},2}$ enforces that precisely the 1-types in $\boldsymbol{\alpha}_0$ are realized in the intended model:

$$\varphi_{\mathrm{FO},2} \coloneqq \textstyle\bigwedge_{\alpha \in \boldsymbol{\alpha}_0} \exists x_1\, \alpha(x_1) \quad \wedge \quad \forall x_1 \bigvee_{\alpha \in \boldsymbol{\alpha}_0} \alpha(x_1).$$

3. The formula $\varphi_{\mathrm{FO},3}$ ensures realizations of regular-free reductions of (at least) 2-types $\beta \in \boldsymbol{\beta}_0$ in the intended model:

$$\varphi_{\mathrm{FO},3} \coloneqq \textstyle\bigwedge_{\beta \in \boldsymbol{\beta}_0} \exists x_1 \exists x_2.\, \mathrm{Aux}(x_1, x_2) \wedge \beta^-(x_1, x_2),$$

where Aux is an auxiliary binary predicate included solely for syntactic correctness. Treating all $\Sigma_{\mathcal{R}}$-predicates in φ_{FO} as standard predicates, we conclude that φ_{FO} is in GF.

Next, we define $\varphi_{\mathcal{R}} \coloneqq \varphi_{\mathcal{R},1} \wedge \varphi_{\mathcal{R},2} \wedge \varphi_{\mathrm{FO},2}$ as follows.

1. Our $\varphi_{\mathcal{R},1}$ comprises all $\forall^{\mathcal{R}}$- and $\exists^{\mathcal{R}}$-conjuncts of φ.

2. Our $\varphi_{\mathcal{R},2}$ ensures that 2-types realized in the intended models belong to $\boldsymbol{\beta}_0$. More specifically, it has the form:

$$\varphi_{\mathcal{R},2} \coloneqq \textstyle\bigwedge_{\mathrm{G}} \bigwedge_{\bar{x}} \forall x_1 \forall x_2\, \mathrm{G}(\bar{x}) \rightarrow \bigvee_{\beta \in \boldsymbol{\beta}_0} \beta^-(x_1 x_2),$$

where G is any guard from φ (a predicate over the signature of φ or program from φ) and $\bar{x}$ is a variable sequence, composed solely of x_1 and x_2, of length matching the arity of G. We stress that $\varphi_{\mathcal{R}}$, as written, is not in RGF due to ocurrences outside guard of literals like $\pm\mathrm{P}(x_1 x_1)$ for regular predicates $\mathrm{P} \in \Sigma_{\mathcal{R}}$ in β^-. This can be remedied by replacing each such literal with (an equivalent guarded) formula $\exists x_2 \mathrm{P}(x_1 x_1) \wedge \top$. The following lemma is the main ingredient of our reduction.

Lemma 4.5. *φ is (finitely) satisfiable iff for some (bounded exponentially in $|\varphi|$) FO-compatible sets of types $\boldsymbol{\alpha}_0$ and $\boldsymbol{\beta}_0$ both formulæ φ_{FO} and $\varphi_{\mathcal{R}}$ are (finitely) satisfiable.* ◀

Lemma 4.5 tells us that to check the (finite) satisfiability of φ, it suffices to iterate over possible sets of 1-types $\boldsymbol{\alpha}_0 \subseteq \boldsymbol{\alpha}_\varphi$ and "sparse" sets $\boldsymbol{\beta}_0 \subseteq \boldsymbol{\beta}_\varphi$—a process that can be performed in doubly-exponential time as done in the previous section—then construct φ_{FO} and $\varphi_{\mathcal{R}}$ and apply the appropriate algorithms to test their satisfiability. Since the constructed formulæ are exponential in $|\varphi|$, and the satisfiability problems of the target logics are 2ExpTime-complete, this naively appears to yield a triply-exponential-time algorithm. However, the second part of Thm. 2.1 shows that the satisfiability of φ_{FO} can be tested in time doubly-exponential in $|\varphi|$. The same holds for $\varphi_{\mathcal{R}}$ by examining the proof of Thm. 4.4 (by incorporating $\boldsymbol{\alpha}_0$ and $\boldsymbol{\beta}_0$ directly into the φ-tables instead of explicitly including $\varphi_{\mathcal{R},2}$ and $\varphi_{\mathrm{FO},2}$ in $\varphi_{\mathcal{R}}$). We conclude:

Theorem 4.6. $\mathsf{Sat}(\mathrm{RGF})$ *is* 2ExpTime*-complete.* ◀

We stress that the finite satisfiability problem for ICPDL is a notorious open question, and the decidability status even for its tiny fragment LoopPDL (Danecki 1984) is open for more than 40 years. Our result however, can be used for conditional lifting the decidability of FSat from ICPDL to RGF.

Corollary 4.7. *If the finite satisfiability for* ICPDL *is decidable and* 2ExpTime*-complete, then so is* $\mathsf{FSat}(\mathrm{RGF})$. ◀

In the remainder of this section, we prove Lemma 4.5. The left-to-right implication is easy. Let $\mathfrak{A} \models \varphi$, and define $\boldsymbol{\alpha}_1$ and $\boldsymbol{\beta}_1$ as the sets of 1-types and 2-types realised in $\mathfrak{A}$. Since $\boldsymbol{\beta}_1$ may be doubly exponential in $|\varphi|$, we seek a smaller subset $\boldsymbol{\beta}_0 \subseteq \boldsymbol{\beta}_1$ to complete the proof. First, construct an auxiliary $\varphi_{\mathcal{R}}$ based on $\boldsymbol{\alpha}_1$ and $\boldsymbol{\beta}_1$. Note that $\mathfrak{A} \models \varphi_{\mathcal{R}}$. Second, applying the construction from Lemma 4.3 we obtain a model $\mathfrak{B} \models \varphi_{\mathcal{R}}$ that realize only exponentially (w.r.t. $|\varphi|$) many 2-types. Take $\boldsymbol{\alpha}_0$ and $\boldsymbol{\beta}_0$ to be the sets of 1- and 2-types realised in $\mathfrak{B}$. Observe that $\boldsymbol{\alpha}_0 = \boldsymbol{\alpha}_1$ and $\boldsymbol{\beta}_0 \subseteq \boldsymbol{\beta}_1$. Now construct φ_{FO} and $\varphi_{\mathcal{R}}$ for $\boldsymbol{\alpha}_0$ and $\boldsymbol{\beta}_0$. It can be readily verified that $\mathfrak{B} \models \varphi_{\mathcal{R}}$, and that the expansion of $\mathfrak{A}$ (interpreting Aux as the full relation) satisfies φ_{FO}. For the other direction, fix $\boldsymbol{\alpha}_0$ and $\boldsymbol{\beta}_0$ as in the lemma's statement, and consider models $\mathfrak{A}_{\mathrm{FO}} \models \varphi_{\mathrm{FO}}$ and $\mathfrak{A}_{\mathcal{R}} \models \varphi_{\mathcal{R}}$. W.l.o.g. we impose some extra assumptions on $\mathfrak{A}_{\mathrm{FO}}$ and $\mathfrak{A}_{\mathcal{R}}$. Indeed:

(C1) As $\varphi_{\mathcal{R}}$ is a two-variable sentence, no relation in $\mathfrak{A}_{\mathcal{R}}$ contains tuples composed of more than two distinct elements.

(C2) $\mathfrak{A}_{\mathrm{FO}}$ contains no pair of distinct elements a, b with $\mathfrak{A}_{\mathrm{FO}} \models \mathrm{T}[\mathrm{a}, \mathrm{b}]$ for a regular $\mathrm{T} \in \Sigma_{\mathcal{R}}$ (all atoms involving regular predicates from φ_{FO} employ only a sole variable).

(C3) For convenience, all conjuncts of $\varphi_{\mathrm{FO},3}$ are witnessed in $\mathfrak{A}_{\mathrm{FO}}$ by pairs of distinct elements (*i.e.* $\mathfrak{A}_{\mathrm{FO}} \models \beta^-[\mathrm{ab}]$ for some $\mathrm{a} \not\approx \mathrm{b}$). It is easy to rebuild $\mathfrak{A}_{\mathrm{FO}}$ if this is not the case.

(C4) The domains of $\mathfrak{A}_{\mathrm{FO}}$ and $\mathfrak{A}_{\mathcal{R}}$ are of equal cardinality (finite or $\aleph_0$). Each 1-type is realized in $\mathfrak{A}_{\mathrm{FO}}$ and $\mathfrak{A}_{\mathcal{R}}$ an equal number of times. In particular, if $\mathfrak{A}_{\mathcal{R}}$ is finite, then the model of φ produced by our construction will be finite as well. This follows from $\varphi_{\mathrm{FO},2}$, the FMP for GF (Grädel 1999), downward Löwenheim-Skolem property for $\mathcal{L}_{\omega_1\omega}$, and some standard "model surgeries" used in the context of guarded logics.

We now combine $\mathfrak{A}_{\mathrm{FO}}$ and $\mathfrak{A}_{\mathcal{R}}$ into a model of φ by constructing a *fusion structure* $\mathfrak{F}$. From a bird's eye view, $\mathfrak{F}$ is a (potentially infinite) two-dimensional square grid in which each row is an isomorphic copy of $\mathfrak{A}_{\mathrm{FO}}$, each column is an isomorphic copy of $\mathfrak{A}_{\mathcal{R}}$, and the interpretation of regular predicates is confined to the columns of $\mathfrak{F}$. This structure is nearly a model of φ; the only remaining issue is that some "vertical" pairs of elements may lack witnesses for certain $\exists^{\mathrm{FO}}$-conjuncts in φ. We solve this with a "circular witnessing scheme" akin to the one by Grädel et al. (1997).

Step I: Domain of $\mathfrak{F}$ and 1-types. Let K denote the cardinality of the domains of $\mathfrak{A}_{\mathrm{FO}}$ and $\mathfrak{A}_{\mathcal{R}}$ (equal due to (C4)). If $\mathrm{K} = \aleph_0$ then, after a bijective renaming, we may assume that the domains of $\mathfrak{A}_{\mathrm{FO}}$ and $\mathfrak{A}_{\mathcal{R}}$ are both equal to $\mathbb{Z}$. We then let $\mathrm{F} \coloneqq \mathbb{Z} \times \mathbb{Z}$ and we interpret predicates from Σ minimally to fulfil $\mathrm{tp}_1^{\mathfrak{F}}(k, l) = \mathrm{tp}_1^{\mathfrak{A}_{\mathrm{FO}}}(k{+}l)$ for all $k, l \in \mathbb{Z}$. Otherwise, after a bijective renaming, the domains of $\mathfrak{A}_{\mathrm{FO}}$ and $\mathfrak{A}_{\mathcal{R}}$ are both equal to $\mathbb{Z}_{\mathrm{K}} \coloneqq \{0, 1, \ldots, \mathrm{K}{-}1\}$. We let $\mathrm{F} \coloneqq \mathbb{Z}_{\mathrm{K}} \times \mathbb{Z}_{\mathrm{K}}$ and interpret predicates from Σ minimally to fulfil $\mathrm{tp}_1^{\mathfrak{F}}(k, l) = \mathrm{tp}_1^{\mathfrak{A}_{\mathrm{FO}}}(k{+}l \bmod \mathrm{K})$ for all $k, l \in \mathbb{Z}_{\mathrm{K}}$.

Step II: Isomorphic copies. As already said, we naturally view $\mathfrak{F}$ as a two-dimensional square table and speak about *columns* (resp. *rows*) of $\mathfrak{F}$, meaning the sets of all elements from $\mathfrak{F}$ sharing the same 1st (resp. 2nd) coordinate. We also

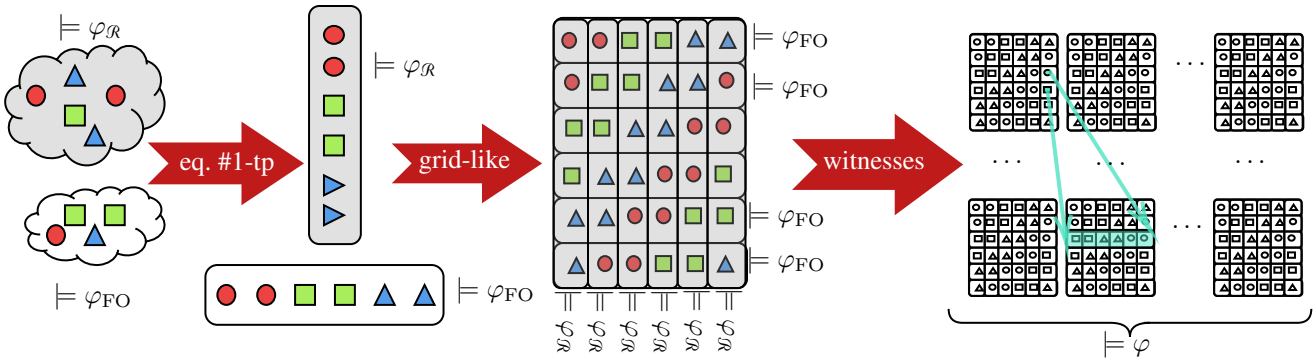

Figure 2: A visualisation of the "fusion" procedure from Section 4.2.

speak about *vertical* (resp. *horizontal*) tuples from $\mathfrak{F}$, meaning the (guarded, *i.e.* contained in some relation) tuples having all its elements in the same column (resp. row) of $\mathfrak{F}$. By construction, each row and column of $\mathfrak{F}$ contains the same number of realizations of each 1-type α. This implies that from every row there exists a 1-type-preserving bijection to $\mathfrak{A}_{\mathrm{FO}}$, and similarly for columns and $\mathfrak{A}_{\mathcal{R}}$. Thus, we may define the interpretation of relations in $\mathfrak{F}$ minimally to ensure that every row of $\mathfrak{F}$ becomes an isomorphic copy of $\mathfrak{A}_{\mathrm{FO}}$ and that every column of $\mathfrak{F}$ becomes an isomorphic copy of $\mathfrak{A}_{\mathcal{R}}$. A few important observations about $\mathfrak{F}$. First, note that by condition (C2), no program connect elements from different columns of $\mathfrak{F}$. Second, all guarded tuples from $\mathfrak{F}$ consisting of a single element are both vertical and horizontal. Lastly, each guarded tuple in $\mathfrak{F}$ is either vertical or horizontal.

Step III: "Almost" satisfaction. We show that $\mathfrak{F}$ is "almost" a model of φ, *i.e.* it satisfies all the conjuncts (with their usual components from Def. 4.1) from φ except for $\exists^{\mathrm{FO}}$-ones. We deal with each conjunct separately, proving its satisfaction.

- $\mathfrak{F} \models \forall x_1\, \lambda(x_1)$ and $\mathfrak{F} \models \forall \bar{x}\, \eta_i^{\forall}(\bar{x}) \to \psi_i^{\forall}(\bar{x})$ (for all i).

It suffices to show that every 1-type and 2-type realized in $\mathfrak{F}$ lies in $\boldsymbol{\alpha}_0 \cup \boldsymbol{\beta}_0$, which consists of types FO-compatible types with φ only. For 1-types, those realized in $\mathfrak{F}$ also appear in $\mathfrak{A}_{\mathcal{R}}$ and $\mathfrak{A}_{\mathrm{FO}}$ (by construction), and thus belong to $\boldsymbol{\alpha}_0$ by $\mathfrak{A}_{\mathcal{R}}, \mathfrak{A}_{\mathrm{FO}} \models \varphi_{\mathrm{FO},2}$. For 2-types, any guarded tuple $\bar{\mathrm{a}}$ is either horizontal or vertical, and so appears in $\mathfrak{A}_{\mathrm{FO}}$ or $\mathfrak{A}_{\mathcal{R}}$, respectively. In the first case, we invoke $\mathfrak{A}_{\mathrm{FO}} \models \varphi_{\mathrm{FO},1}$; Otherwise, as $\bar{\mathrm{a}}$ uses at most 2 elements, we apply $\mathfrak{A}_{\mathcal{R}} \models \varphi_{\mathcal{R},2}$.

- $\mathfrak{F} \models \forall x_1 x_2\, \pi_i^{\forall}(x_1 x_2) \to \phi_i^{\forall}(x_1 x_2)$ (for all i).

Let $\mathfrak{F} \models \pi_i^{\forall}[\mathrm{a}, \mathrm{b}]$ for some a, b. By (C2), we infer that (a, b) is vertical (otherwise $\mathrm{a} = \mathrm{b}$, so it is vertical as well), and that the path witnessing $(\mathrm{a}, \mathrm{b}) \in (\pi_i^{\forall})^{\mathfrak{F}}$ is entirely contained within the column of $\mathfrak{F}$ determined by a (or b). Since each column satisfies $\varphi_{\mathcal{R},1}$, we are done.

- $\mathfrak{F} \models \forall x_1\, \gamma_i^{\exists}(x_1) \to \exists x_2\, \pi_i^{\exists}(x_1 x_2) \wedge \phi_i^{\exists}(x_1 x_2)$ (for all i).

Suppose $\mathfrak{F} \models \gamma_i^{\exists}[\mathrm{a}]$. By construction, the column of a is isomorphic to $\mathfrak{A}_{\mathcal{R}}$ and it satisfies $\varphi_{\mathcal{R},1}$. Hence, the element a finds its witness for the satisfaction of $\pi_i^{\exists}(x_1 x_2) \wedge \phi_i^{\exists}(x_1 x_2)$.

We now address the satisfaction of $\exists^{\mathrm{FO}}$-conjuncts of the form $\forall \bar{x}, \eta_i^{\exists}(\bar{x}) \to \exists \bar{y}, \vartheta_i^{\exists}(\bar{x}\bar{y}) \wedge \psi_i^{\exists}(\bar{x}\bar{y})$ in $\mathfrak{F}$. Observe that all horizontal guarded tuples $\bar{\mathrm{a}}$ with $\mathfrak{F} \models \eta_i^{\exists}[\bar{\mathrm{a}}]$ have the required witnesses within their rows—this follows directly from the satisfaction of $\varphi_{\mathrm{FO},1}$ by every row of $\mathfrak{F}$. Thus, the only obstruction to $\mathfrak{F}$ being a model of φ is the potential lack of witnesses for some vertical guarded tuples. By (C1), such vertical tuples consist of exactly two elements (single-element tuples have witnesses in their rows). To resolve this, we take multiple copies of $\mathfrak{F}$ and, for each vertical guarded pair lacking witnesses, embed it into a row of a different copy of $\mathfrak{F}$, thereby ensuring the presence of $\exists^{\mathrm{FO}}$-witnesses.

Step IV: Providing missing $\exists^{\mathrm{FO}}$-witnesses. Since the case of $\mathrm{K} = \aleph_0$ can be handled analogously, let us focus only on the case when K is a natural number. We define a structure $\mathfrak{G}$ as the disjoint union of 3K isomorphic copies of $\mathfrak{F}$, *i.e.*:

- The domain G of $\mathfrak{G}$ is equal to $\{0,1,2\} \times \mathbb{Z}_{\mathrm{K}} \times F$.
- The interpretation of all predicates is minimal to make $\mathfrak{G}$ restricted to $\{i\} \times \{j\} \times F$, for each i and j, isomorphic to $\mathfrak{F}$.

For convenience, let $G_i := \{i\} \times \mathbb{Z}_{\mathrm{K}} \times F$ for each $i \in \{0,1,2\}$, and let $\mathfrak{G}_i$ be the restriction of $\mathfrak{G}$ to G_i. When referring to rows of $\mathfrak{G}$ or $\mathfrak{G}_i$ we mean rows in their $\mathfrak{F}$-components, that is sets of the form $\{i\} \times \{j\} \times E$, for some i, j, where E forms a row in $\mathfrak{F}$. Analogously for columns and vertical/horizontal tuples. Consider a vertical guarded pair $(\mathrm{b}_1, \mathrm{b}_2)$ of distinct elements in $\mathfrak{G}$. If the pair (b_1, b_2) lacks some $\exists^{\mathrm{FO}}$-witnesses, that is, there is at least one i such that $\mathfrak{G} \models \eta_i^{\exists}[\bar{\mathrm{c}}]$ for some tuple $\bar{\mathrm{c}}$ consisting solely of b_1 and b_2 (each occurring at least once), then we provide a witnesses by "connecting" $(\mathrm{b}_1, \mathrm{b}_2)$ to (to be selected) row $\mathfrak{E}$ as follows.

As $\mathfrak{G} \models \varphi_{\mathcal{R},2}$, we infer $\mathfrak{G} \models \beta^-[\mathrm{b}_1, \mathrm{b}_2]$ for some $\beta \in \boldsymbol{\beta}_0$. By $\mathfrak{E} \models \varphi_{\mathrm{FO},3}$, there are $\mathrm{a}_1, \mathrm{a}_2 \in E$ with $\mathfrak{E} \models \beta^-[\mathrm{a}_1, \mathrm{a}_2]$. By condition (C3), we can assume that $\mathrm{a}_1 \neq \mathrm{a}_2$. For any tuple $\bar{\mathrm{a}}$ containing at least one of $\mathrm{b}_1, \mathrm{b}_2$, some elements from $E \setminus \{\mathrm{a}_1, \mathrm{a}_2\}$, and a relation $\mathrm{R} \in \Sigma_{\mathrm{FO}}$ of arity $|\bar{\mathrm{a}}|$, we add $\mathrm{R}(\bar{\mathrm{a}})$ to $\mathfrak{G}$ if and only if $\mathfrak{E} \models \mathrm{R}[\mathfrak{h}(\bar{\mathrm{a}})]$, where $\mathfrak{h}$ maps $\mathrm{b}_s \mapsto \mathrm{a}_s$ (for all $s \in \{1,2\}$) and acts as the identity on E. Thus, $\mathfrak{h}$ defines an isomorphism between $(E \cup \{\mathrm{b}_1, \mathrm{b}_2\}) \setminus \{\mathrm{a}_1, \mathrm{a}_2\}$ and E, if regular predicates are ignored. No other facts are added—specifically, no fresh regular atoms.

Note that since $\mathfrak{E} \models \varphi_{\mathrm{FO},1}$, all tuples built out of elements of E have in $\mathfrak{E}$ all the $\exists^{\mathrm{FO}}$-witnesses required by φ. After the above described connection of $\mathrm{b}_1, \mathrm{b}_2$ to $\mathfrak{E}$ this becomes true also for all tuples $\bar{\mathrm{a}}$ built out of elements from $\{\mathrm{b}_1, \mathrm{b}_2\} \cup E \setminus \{\mathrm{a}_1, \mathrm{a}_2\}$ (because $\bar{\mathrm{a}}$ starts to mimic the behaviour of the tuple $\mathfrak{h}(\mathrm{a})$ from E, if regular predicates are ignored).

We repeat the process for all vertical guarded $(\mathrm{b}_1, \mathrm{b}_2)$. To avoid conflicts, if $\mathrm{b}_1, \mathrm{b}_2$ are in G_i, we select a row $\mathfrak{E}$ from the $G_{(i+1) \bmod 3}$ that is not yet used by any other pair from their column. Since each G_i has K copies of $\mathfrak{F}$, each with K rows, and the number of such pairs is less than K^2, there are enough rows to complete the process. This circular strategy, using three disjoint sets G_i, guarantees that if an element b is connected to a row $\mathfrak{E}$, no element of $\mathfrak{E}$ is ever connected back to b's row. One can check that this scheme preserves all $\forall^{\mathrm{FO}}$-, $\forall^{\mathcal{R}}$-, and $\exists^{\mathcal{R}}$-conjuncts, yielding a model of φ.

4.3 A Slightly Extended Setting

In description logics and ontology-based data access, the focus often shifts from satisfiability to the more general problem of query entailment over knowledge bases (see the survey by Ortiz and Šimkus for extra background). We show how this general entailment problem reduces to plain satisfiability, assuming that the queries are acyclic. Without this restriction, the problem becomes undecidable (see Sec. 5).
An RGF-knowledge base (KB) is a pair $\mathcal{K} \coloneqq (\mathcal{A}, \mathcal{T})$, where $\mathcal{A}$ (ABox) is a finite set of ground facts and $\mathcal{T}$ (TBox) a finite set of RGF-sentences. A structure $\mathfrak{A}$ satisfies $\mathcal{K}$ if it satisfies all elements of $\mathcal{A} \cup \mathcal{T}$. A query q is entailed by $\mathcal{K}$ (written: $\mathcal{K} \models \mathrm{q}$) if it holds in all models of $\mathcal{K}$. A CQ q is *tree-shaped* if it contains only atoms of arity at most two, and its variables can be bijectively mapped to the nodes of some tree in a way that each binary atom $\mathrm{R}(x, y)$ satisfies one of: (i) $x = y$, (ii) x is a child of y, or (iii) x is the parent of y. In the (tree-shaped) CQ entailment problem over RGF-KBs we ask if an input (tree-shaped) CQ is entailed by an input RGF-KB.

Theorem 4.8. *The entailment problem of tree-shaped CQ over* RGF*-KBs is* 2ExpTime*-complete.* ◀

Proof. We reduce the problem to formula satisfiability in polynomial time. Given a tree-shaped CQ q, the well-known "rolling-up" technique (Horrocks and Tessaris 2000) yields an RGF-sentence φ_{q} such that, for all models $\mathfrak{A}$ of $\mathcal{K}$, we have that $\mathfrak{A} \models \mathrm{q}$ iff $\mathfrak{A} \models \varphi_{\mathrm{q}}$ (cf. Section 3.2 of Bednarczyk's PhD thesis). It remains to encode the KB as a single formula. Let N be the number of individual names in $\mathcal{A}$. Introduce a fresh N-ary predicate aux, and let $\mathcal{A}^\dagger$ be obtained from $\mathcal{A}$ by bijectively renaming these names to $x_1, \dots, x_{\mathrm{N}}$. Then

$$\Big(\exists x_1 \dots x_{\mathrm{N}}\ \mathrm{aux}(x_1, \dots, x_{\mathrm{N}}) \wedge \bigwedge \mathcal{A}^\dagger\Big) \ \wedge \ \bigwedge \mathcal{T} \ \wedge \ \neg\varphi_{\mathrm{q}}.$$

is unsatisfiable iff $\mathcal{K} \not\models \mathrm{q}$, concluding the proof. □

5 Undecidability of Querying

We prove undecidability of (both general and finite) CQ-entailment problems for RGF, already for its very restricted fragment: fluted GF^2 with a *single* transitive guard. This significantly tightens the undecidability of the unrestricted UCQ-entailment over GF^2+TG by Gottlob et al. (2013, Th. 2) in three ways: (i) our logic is stricter [fluting], (ii) we utilize only a single CQ [rather than a UCQ], and (iii) our undecidability applies also to the finite case [which was open].
We first introduce "triangular grid-like" σ-structures dubbed *snakes*, where $\sigma \coloneqq \{\mathrm{B}, \mathrm{T}_0, \mathrm{T}_1, \mathrm{T}_2, \mathrm{E}, \mathrm{H}, \mathrm{V}, \mathrm{P}\}$ and all but the last three predicates in σ (which are binary) are unary. In snakes, the predicates H and V denote the *horizontal* and the *vertical* successors in the intended grid. By "triangular" we mean that the i-th column of a snake contains $i{+}1$ (resp. $i{+}2$) elements if i is even (resp. odd). Additionally, E labels *even* columns, B labels the *bottom* of the grid, T_is label the elements exactly i steps from the *top* of the grid, and (a transitive) P relates the elements reachable via a directed *path* in the underlying graph. Quite unusually, the direction of V alternates between columns. Consult Figure 3 for intuitions.

Definition 5.1. *Fix* $\mathrm{N} \in \mathbb{N} \cup \{\infty\}$. *The set* $\mathbb{S}_{\mathrm{N}} \subseteq \mathbb{N}^2$ *consists of all pairs* (n, m) *with* $m \leq n \leq \mathrm{N}$ *and* $m{-}1 \leq n \leq \mathrm{N}$ *if* n *is even and odd, respectively. The* snake ordering $<_{\mathbb{S}}$ *(also known as boustrophedon) on* $\mathbb{S}_{\mathrm{N}}$ *is defined by cases:* $(n, m) <_{\mathbb{S}} (n', m')$ *if* $n < n'$ *or if* $n = n'$ *and either (i)* n *is odd and* $m < m'$*, or (ii)* n *is even and* $m' < m$.
The N-snake $\mathfrak{S} \coloneqq (\mathbb{S}_{\mathrm{N}}, \cdot^{\mathfrak{S}})$ *is a structure defined as:*

- $\mathrm{E}^{\mathfrak{S}} \coloneqq \{(n, m) \in \mathbb{S}_{\mathrm{N}} \mid n \textit{ even}\}$; $\mathrm{B}^{\mathfrak{S}} \coloneqq \{(0, m) \in \mathbb{S}_{\mathrm{N}}\}$;
- $\mathrm{T}_i^{\mathfrak{S}} \coloneqq \{(n, m) \in \mathbb{S}_{\mathrm{N}} \mid (n, m{+}i{+}1) \notin \mathbb{S}_{\mathrm{N}}, (n, m{+}i) \in \mathbb{S}_{\mathrm{N}}\}$;
- $\mathrm{V}^{\mathfrak{S}} \coloneqq \mathrm{V}^{\mathfrak{S}}_{\textit{even}} \cup \mathrm{V}^{\mathfrak{S}}_{\textit{odd}}$, *where:*
  - $\mathrm{V}^{\mathfrak{S}}_{\textit{even}} \coloneqq \{((n, m), (n, m{-}1)) \in \mathbb{S}^2_{\mathrm{N}} \mid n \textit{ is even}\}$,
  - $\mathrm{V}^{\mathfrak{S}}_{\textit{odd}} \coloneqq \{((n, m), (n, m{+}1)) \in \mathbb{S}^2_{\mathrm{N}} \mid n \textit{ is odd}\}$;
- $\mathrm{H}^{\mathfrak{S}} \coloneqq \{((n, m), (n{+}1, m)) \in \mathbb{S}^2_{\mathrm{N}}\}$; *and* $\mathrm{P}^{\mathfrak{S}}$ *is* $<_{\mathbb{S}}$ *on* $\mathbb{S}_{\mathrm{N}}$.

Given $(a, b) \in \mathbb{N}^2$, *let* $\mathfrak{S}_{(a,b)}$ *denote the restriction of the* ∞*-snake to* $\{(n, m) \in \mathbb{S}_\infty \mid (n, m) \leq_{\mathbb{S}} (a, b)\}$. *The pairs* $(\mathrm{N}, 0)$ *and* $(\mathrm{N}, \mathrm{N}{+}1)$ *for an even and odd* N*, respectively, are called* final. *Let* $\mathfrak{S}_{\mathrm{N}}$ *denote* $\mathfrak{S}_{(\mathrm{N},\mathrm{M})}$ *for the final* (N, M). ◀

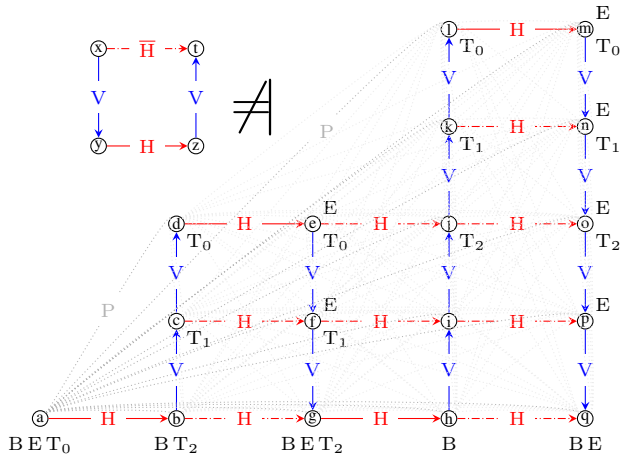

Figure 3: Above depicts the 4-snake $\mathfrak{S}_4$ violating the conjunctive query q_{clo}. The solid arrows are the ones enforced by the formula φ_{grd}, while the dashed arrows are the ones implied by (the violation of) q_{clo}. The dotted arrows represent the predicate P. The snake ordering $<_{\mathbb{S}}$ coincide with the alphabetic order of elements.

We employ snakes in our reduction from the (finite) query non-entailment problem to the undecidable tiling problem of a (finite) octant (Bresolin et al. 2010, p. 699). This is achieved through $\varphi_{\mathfrak{S}} \coloneqq \varphi_{\text{pro}} \wedge \varphi_{\text{con}} \wedge \varphi_{\text{ini}} \wedge \varphi_{\text{grd}} \wedge \neg\mathrm{q}_{\text{clo}}$ describing natural properties of snakes. Here, φ_{pro} *pro*pagates unary

atoms via the grid relations; φ_{con} enforces basic *con*sistency conditions (*e.g.* disjointness of all T_is); φ_{ini} captures the *ini*tial element of $\mathfrak{S}$; φ_{grd} builds a skeleton of a *grid* (*cf.* solid arrows from Fig. 3); and q_{clo} is a CQ that "*clo*ses" the grid (*cf.* dashed arrows appearing in Fig. 3). To manufacture $\varphi_{\mathfrak{S}}$ we rely on two "helper" predicates: a unary $\underline{\mathrm{L}}$ (denoting the *last* column of a finite snake), and a binary $\overline{\mathrm{H}}$ (denoting the relative *complement of* H w.r.t. P). The second predicate is vital for the query design as CQs do not contain negation. It is not difficult to provide the full definitions of the first four conjuncts of $\varphi_{\mathfrak{S}}$. As an example, the formula $\forall x_1 \neg \mathrm{L}(x_1) \rightarrow \exists x_2 [\mathrm{P}(x_1 x_2) \wedge \mathrm{H}(x_1 x_2)]$ captures existence of H-successors in grids (V-successors are treated analogously). In our construction the grid relations always follow P. Such a "P-relativised" quantification used for constructing grids makes P, by its transitivity, resemble the "universal relation". Importantly, with the formula $\forall x_1 \forall x_2\, \mathrm{P}(x_1 x_2) \rightarrow \big(\mathrm{H}(x_1 x_2) \leftrightarrow \neg \overline{\mathrm{H}}(x_1 x_2)\big]\big)$ one can correctly axiomatize that $\overline{\mathrm{H}}$ is interpreted as a (relative to P) complement of H. Finally, the role of (negated) query q_{clo} is to identify the cases of missing H-connections between two consecutive grid cells. The conjunctive query $\mathrm{q}_{\text{clo}} \coloneqq \exists x \exists y \exists z \exists t\, \mathrm{V}(xy) \wedge \mathrm{H}(yz) \wedge \mathrm{V}(zt) \wedge \overline{\mathrm{H}}(xt)$ then does the job. Our main technical (inductive) lemma is stated below.

Lemma 5.2. *Consider* $\mathfrak{A} \models \varphi_{\mathfrak{S}}$*, a pair* $(\mathrm{N}, \mathrm{M}) \in \mathbb{S}_\infty$*, and a semi-strong homomorphism* $\mathfrak{f}$ *witnessing* $\mathfrak{S}_{(\mathrm{N},\mathrm{M})} \hookrightarrow\!\!\!\!\rightarrow_{\mathfrak{f}} \mathfrak{A}$*. Assume that either* (N, M) *is not final or* $\mathfrak{f}(\mathrm{N}, \mathrm{M}) \notin \mathrm{L}^{\mathfrak{A}}$*. Then* $\mathfrak{f}$ *extends to a homomorphism* $\mathfrak{g}$ *witnessing* $\mathfrak{S}_{(\mathrm{N}',\mathrm{M}')} \hookrightarrow\!\!\!\!\rightarrow_{\mathfrak{g}} \mathfrak{A}$*, where* $(\mathrm{N}', \mathrm{M}')$ *is the successor of* (N, M) *w.r.t.* $<_{\mathbb{S}}$*.* ◀

Based on Lemma 5.2 we show correctness of our encoding, *i.e.* that models of $\varphi_{\mathfrak{S}}$ indeed "encode" snakes. Formally:

Lemma 5.3. *(A) All (finite) snakes expand to (finite) models of* $\varphi_{\mathfrak{S}}$*. (B) If* $\mathfrak{A}$ *satisfies* $\varphi_{\mathfrak{S}} \wedge \forall x_1 \neg \mathrm{L}(x_1)$ *then* $\mathfrak{S}_\infty \hookrightarrow\!\!\!\!\rightarrow \mathfrak{A}$*. (C) For any finite* $\mathfrak{A} \models \varphi_{\mathfrak{S}}$ *we have* $\mathfrak{S}_{\mathrm{N}} \hookrightarrow\!\!\!\!\rightarrow_{\mathfrak{f}} \mathfrak{A}$ *for some* N *and a map* $\mathfrak{f}$ *with* $\mathfrak{f}(\mathrm{N}, \mathrm{M}) \in \mathrm{L}^{\mathfrak{A}} \cap \mathrm{T}_0^{\mathfrak{A}}$ *for the final* (N, M)*.* ◀

Given an $\mathrm{N} \in \mathbb{N} \cup \{\infty\}$, by the N-octant (a.k.a. the lower triangle of $\mathbb{N} \times \mathbb{N}$) we mean the set $\mathbb{O}_{\mathrm{N}}$ of pairs $(n, m) \in \mathbb{N}^2$ satisfying $m \leq n \leq \mathrm{N}$. A *tiling system* $\mathcal{D} \coloneqq (\mathcal{T}, \mathcal{V}, \mathcal{H})$ is composed of a finite set $\mathcal{T}$ of *tiles*, and binary relations $\mathcal{V}, \mathcal{H}$ on $\mathcal{T}$. A system $\mathcal{D}$ *covers* $\mathbb{O}$ if there is a map $\xi \colon \mathbb{O}^2 \to \mathcal{T}$ satisfying $(\xi(n, m), \xi(n{+}1, m)) \in \mathcal{H}$ and $(\xi(n, m), \xi(n, m{+}1)) \in \mathcal{V}$ for all pairs $(n, m), (n, m{+}1) \in \mathbb{O}^2$. The (finite) octant tiling problem asks if $\mathcal{D}$ covers some (finite) octant. By Lemma 5.3, which encodes grids in models of $\varphi_{\mathfrak{S}}$, the reduction is standard. For all $t \in \mathcal{T}$ we introduce fresh unary T_t, and construct a formula $\varphi_{\mathcal{D}}$ asserting that all elements from snakes (except for the top odd ones, which are irrelevant to the reduction) are labelled by exactly one T_t, and that H- and V-connected elements respect $\mathcal{H}$ and $\mathcal{V}$. We show that $\varphi_{\mathfrak{S}} \wedge \varphi_{\mathcal{D}}$ has a (finite) model iff $\mathcal{D}$ covers some (finite) octant. Thus:

Theorem 5.4. *Both finite and general CQ entailment problems are undecidable for* RGF*, already for its fluted two-variable fragment with a single transitive guard.* ◀

We conclude by translating our results into the setting of description logics (see the standard textbook for any missing definitions). Let $\mathcal{ALC}^{\cap_\top}$ extend $\mathcal{ALC}$ with concept expressions of the form $\exists(\mathrm{r} \cap \mathrm{s}).\top$, denoting elements that have at least one r-filler also serving as an s-filler. We also consider the axiom $\mathtt{tr}(\mathrm{r})$, asserting that r is transitive, and inclusion axioms of the form $\mathrm{r} \subseteq \mathrm{s} \cup \mathrm{t}$, interpreted by $\mathcal{I}$ as $\mathrm{r}^{\mathcal{I}} \subseteq \mathrm{s}^{\mathcal{I}} \cup \mathrm{t}^{\mathcal{I}}$. With a routine rewriting of first-order formulæ into a (slightly extended) syntax of $\mathcal{ALC}$, we prove the following theorem.

Corollary 5.5. *Fix role names* $\mathrm{p}, \mathrm{h}, \bar{\mathrm{h}}$*. Both finite and general CQ entailment problems are undecidable for TBoxes of the form* $\mathcal{T} \cup \{\mathtt{tr}(p), \mathrm{p} \subseteq \mathrm{h} \cup \bar{\mathrm{h}}\}$*, where* $\mathcal{T}$ *is an* $\mathcal{ALC}^{\cap_\top}$*-TBox, already for queries without the (transitive) role* p*.* ◀

6 Decreasing Computational Complexity

We conclude by examining whether RGF can be restricted to a reasonable subfragment with lower computational complexity (modulo standard complexity-theoretic assumptions). Our guiding requirements are twofold: first we aim to retain some form of recursion–specifically, by preserving operators such as $\cdot^*$ or $\cdot^+$; second, we require the ability to express non-trivial properties of relations with unbounded arity. There are three main sources of 2ExpTime-hardness of RGF:

- The first one is the full GF itself (Grädel 1999). Fortunately, $\mathsf{Sat}(\mathrm{GF}^k)$ is only ExpTime-complete each $k > 1$.
- The second one is ICPDL (Lange and Lutz 2005). Fortunately, when the input formulæ are of bounded iwidth, the satisfiability problem for ICPDL is ExpTime-complete.
- The last one is GF^2+TG (Kieroński 2006), namely the two-variable GF with transitive guards, a strict sublogic of $\mathrm{RGF}^2[\cdot^+]$. Fortunately, GF^2 with *one-way transitive guards* (*i.e.* imposing that all atoms involving a transitive T have the form $\mathrm{T}(x_1 x_1)$, $\mathrm{T}(x_2 x_2)$, or $\mathrm{T}(x_1 x_2)$) is ExpSpace-complete with hardness holding already for its fluted sublogic (Pratt-Hartmann and Tendera 2023).

The references make it clear that we cannot hope for a nice fragment of RGF with complexity below ExpSpace. To identify a suitable ExpSpace-complete logic, we must: (i) restrict ourselves to a proper subfragment of GF, (ii) limit the interaction between operators in ICPDL-programs, and (iii) enforce a one-way behavior for regular guards, in particular by disallowing the converse operator. A natural base logic satisfying all three criteria is the forward GF (denoted FGF), which has an ExpTime-complete satisfiability problem (Bednarczyk 2021) and fits well with the notion of one-way guards introduced by Kieroński (2006). The main result of this section (available only in the extended version of this paper) establishes ExpSpace-completeness of $\mathrm{FRGF}^2[\cdot^+, \cdot^*, ?]$. We also show that adding any further operators to just Kleene's plus results in a higher complexity.

Theorem 6.1. *The satisfiability problem for* $\mathrm{FRGF}[\cdot^+, \cdot^*, ?]$ *is* ExpSpace*-complete.* ◀

Theorem 6.2. *Both finite and general satisfiability problems for* $\mathrm{FRGF}^2[\cdot^+, \circ]$*,* $\mathrm{FRGF}^2[\cdot^+, \cup]$*, and* $\mathrm{FRGF}^2[\cdot^+, \cap]$ *are already* 2ExpTime*-hard.* ◀

Acknowledgements

Work supported by NCN grants 2024/53/N/ST6/01390 (B. Bednarczyk) and 2021/41/B/ST6/00996 (E. Kieroński).

References

Andréka, H.; Németi, I.; and van Benthem, J. 1998. Modal Languages and Bounded Fragments of Predicate Logic. *Journal of Philosophical Logic* 27(3):217–274.

Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press.

Bárány, V.; Gottlob, G.; and Otto, M. 2014. Querying the Guarded Fragment. *Logical Methods in Computer Science* 10(2).

Bárány, V.; ten Cate, B.; and Segoufin, L. 2015. Guarded Negation. *Journal of the ACM* 62(3):22:1–22:26.

Bednarczyk, B., and Kieroński, E. 2025. Guarded Fragments Meet Dynamic Logic: The Story of Regular Guards (Extended Version).

Bednarczyk, B. 2021. Exploiting Forwardness: Satisfiability and Query-Entailment in Forward Guarded Fragment. In *Proceedings of JELIA*, 179–193.

Bednarczyk, B. 2024. *Database-Inspired Reasoning Problems in Description Logics With Path Expressions*. Ph.D. Dissertation, Dresden University of Technology, Germany.

Benedikt, M.; Bourhis, P.; and Vanden Boom, M. 2016. A Step Up in Expressiveness of Decidable Fixpoint Logics. In *Proceedings of LICS*, 817–826.

Bourhis, P.; Morak, M.; and Pieris, A. 2017. Making Cross Products and Guarded Ontology Languages Compatible. In Sierra, C., ed., *Proceedings of IJCAI*, 880–886.

Bresolin, D.; Della Monica, D.; Goranko, V.; Montanari, A.; and Sciavicco, G. 2010. Metric Propositional Neighborhood Logics: Expressiveness, Decidability, and Undecidability. In *Proceedings of ECAI*, 695–700.

Calvanese, D.; Eiter, T.; and Ortiz, M. 2009. Regular Path Queries in Expressive Description Logics with Nominals. In *Proceedings of IJCAI*, 714–720.

Calvanese, D.; Giacomo, G. D.; and Lenzerini, M. 2008. Conjunctive Query Containment and Answering Under Description Logic Constraints. *ACM Transactions on Computational Logic* 9(3):22:1–22:31.

Christoph Jung, J.; Lutz, C.; Martel, M.; and Schneider, T. 2018. Querying the Unary Negation Fragment with Regular Path Expressions. In *Proceedings of ICDT*, 15:1–15:18.

Christoph Jung, J.; Lutz, C.; Pulcini, H.; and Wolter, F. 2021. Separating Data Examples by Description Logic Concepts with Restricted Signatures. In *Proceedings of KR*, 390–399.

Danecki, R. 1984. Propositional Dynamic Logic with Strong Loop Predicate. In *Proceedings of MFCS*, 573–581.

Figueira, D., and Figueira, S. 2025. A Common Ancestor of PDL, Conjunctive Queries, and Unary Negation First-order. *ArXiv (version 1)*.

Figueira, D.; Figueira, S.; and Baque, E. P. 2020. Finite Controllability for Ontology-Mediated Query Answering of CRPQ. In *Proceedings of KR*, 381–391.

Ganzinger, H.; Meyer, C.; and Veanes, M. 1999. The Two-Variable Guarded Fragment with Transitive Relations. In *Proceedings of LICS*, 24–34.

Göller, S.; Lohrey, M.; and Lutz, C. 2009. PDL with Intersection and Converse: Satisfiability and Infinite-State Model Checking. *Journal of Symbolic Logic* 74(1):279–314.

Göller, S. 2008. *Computational Complexity of Propositional Dynamic Logics*. Ph.D. Dissertation, University of Leipzig.

Gottlob, G.; Pieris, A.; and Tendera, L. 2013. Querying the Guarded Fragment with Transitivity. In *Proceedings of ICALP*, 287–298.

Grädel, E., and Otto, M. 2014. The Freedoms of (Guarded) Bisimulation. In *Johan van Benthem on Logic and Information Dynamics*. Springer. 3–31.

Grädel, E., and Walukiewicz, I. 1999. Guarded Fixed Point Logic. In *Proceedings of LICS*, 45–54.

Grädel, E.; Kolaitis, P.; and Vardi, M. Y. 1997. On the Decision Problem for Two-Variable First-Order Logic. *The Bulletin of Symbolic Logic* 3(1):53–69.

Grädel, E. 1998. Description Logics and Guarded Fragments of First Order Logic. In *Proceedings of DL*.

Grädel, E. 1999. On The Restraining Power of Guards. *Journal of Symbolic Logic* 64(4):1719–1742.

Harel, D.; Pnueli, A.; and Vardi, M. 1982. Two Dimensional Temporal Logic and PDL with Intersection. (Unpublished).

Horrocks, I., and Tessaris, S. 2000. Answering Conjunctive Queries over DL ABoxes: A Preliminary Report. In Baader, F., and Sattler, U., eds., *Proceedings of DL*, 173–182.

Horrocks, I.; Kutz, O.; and Sattler, U. 2006. The Even More Irresistible SROIQ. In *Proceedings of KR*, 57–67.

Kazakov, Y. 2006. *Saturation-Based Decision Procedures For Extensions Of The Guarded Fragment*. Ph.D. Dissertation, Universität des Saarlandes.

Kieroński, E., and Rudolph, S. 2021. Finite Model Theory of the Triguarded Fragment and Related Logics. In *Proceedings of LICS*, 1–13.

Kieroński, E.; Pratt-Hartmann, I.; and Tendera, L. 2017. Equivalence Closure in the Two-Variable Guarded Fragment. *Journal of Logic and Computation* 27(4):999–1021.

Kieroński, E. 2005. Results on the Guarded Fragment with Equivalence or Transitive Relations. In *Proceedings of CSL*, 309–324.

Kieroński, E. 2006. On the Complexity of the Two-Variable Guarded Fragment with Transitive Guards. *Information and Computation* 204(11):1663–1703.

Lange, M., and Lutz, C. 2005. 2-ExpTime Lower Bounds for Propositional Dynamic Logics with Intersection. *The Journal of Symbolic Logic* 70(4):1072–1086.

Lutz, C., and Manière, Q. 2024. Adding Circumscription to Decidable Fragments of First-Order Logic: A Complexity Rollercoaster. In *PProceedings of KR*.

Michaliszyn, J. 2008. Konstrukcje PDL w Logice ze Strażnikami. Master's thesis, University of Wrocław. Available (only in polish) at https://ii.uni.wroc.pl/~jmi/papers/magisterka.PDF.

Michaliszyn, J. 2009. Decidability of the Guarded Fragment with the Transitive Closure. In *Proceedings of ICALP*, 261–272.

Nakamura, Y. 2025. Guarded Negation Transitive Closure Logic is 2-EXPTIME-complete. *ArXiv (version 2)*.

Ortiz, M., and Šimkus, M. 2012. Reasoning and query answering in description logics. In *Proceedings of Reasoning Web*, 1–53.

Otto, M. 2012. Highly Acyclic Groups, Hypergraph Covers, and the Guarded Fragment. *Journal of the ACM* 59(1):5:1–5:40.

Pratt-Hartmann, I., and Tendera, L. 2023. Adding Transitivity and Counting to the Fluted Fragment. In *Proceedings of CSL*, 32:1–32:22.

Segoufin, L., and ten Cate, B. 2013. Unary Negation. *Logical Methods in Computer Science* 9(3).

Szwast, W., and Tendera, L. 2004. The Guarded Fragment with Transitive Guards. *Annals of Pure and Applied Logic* 128(1-3):227–276.

Tendera, L. 2017. Finite Model Reasoning in Expressive Fragments of First-Order Logic. In *Proceedings of M4M*, 43–57.

Zheng, S., and Schmidt, R. A. 2020. Deciding the Loosely Guarded Fragment and Querying Its Horn Fragment Using Resolution. In *Proceedings of AAAI*, 3080–3087.