QUERYING THE GUARDED FRAGMENT

VINCE BÁRÁNY, GEORG GOTTLOB, AND MARTIN OTTO

Department of Mathematics, Technische Universität Darmstadt
e-mail address: vbarany@mathematik.tu-darmstadt.de

Oxford University Computing Laboratory, Wolfson Building, Parks Rd., OX1 3QD Oxford, UK
e-mail address: georg.gottlob@comlab.ox.ac.uk

Department of Mathematics, Technische Universität Darmstadt
e-mail address: otto@mathematik.tu-darmstadt.de

Abstract. Evaluating a Boolean conjunctive query q against a guarded first-order theory φ is equivalent to checking whether $\varphi \wedge \neg q$ is unsatisfiable. This problem is relevant to the areas of database theory and description logic. Since q may not be guarded, well known results about the decidability, complexity, and finite-model property of the guarded fragment do not obviously carry over to conjunctive query answering over guarded theories, and had been left open in general. By investigating finite guarded bisimilar covers of hypergraphs and relational structures, and by substantially generalising Rosati's finite chase, we prove for guarded theories φ and (unions of) conjunctive queries q that (i) $\varphi \models q$ iff $\varphi \models_{fin} q$, that is, iff q is true in each finite model of φ and (ii) determining whether $\varphi \models q$ is 2EXPTIME-complete. We further show the following results: (iii) the existence of polynomial-size conformal covers of arbitrary hypergraphs; (iv) a new proof of the finite model property of the clique-guarded fragment; (v) the small model property of the guarded fragment with optimal bounds; (vi) a polynomial-time solution to the canonisation problem modulo guarded bisimulation, which yields (vii) a capturing result for guarded bisimulation invariant PTIME.

1998 ACM Subject Classification: F.4.1, F.2.2, I.2.4, H.2.3, H.2.4.

Key words and phrases: guarded fragment, finite model theory, descriptive complexity, hypergraph covers, conjunctive queries, tuple-generating dependencies,description logics.

This is an improved and extended version of [2].

Vince Bárány gratefully acknowledges funding received from ERC Starting Grant Sosna while affiliated with the Institute of Informatics at Warsaw University.

Georg Gottlob's research was funded by the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. 246858. Georg Gottlob also gratefully acknowledges a Royal Society Wolfson Research Merit Award.

Martin Otto's research has been partially supported by DFG grant no. OT 147/5-1.

LOGICAL METHODS
IN COMPUTER SCIENCE

DOI:10.2168/LMCS-???

1. Introduction

The *guarded fragment* of first-order logic (GF), defined through the relativisation of quantifiers by atomic formulas, was introduced by Andréka, van Benthem, and Németi [1], who proved that the satisfiability problem for GF is decidable. Grädel [16] proved that every satisfiable guarded first-order sentence has a finite model, i.e., that GF has the *finite model property* (FMP). In the same paper, Grädel also proved that satisfiability of GF-sentences is complete for 2ExpTime, and ExpTime-complete for sentences involving relations of bounded arity. The guarded fragment has since been intensively studied and extended in various ways. For example, the *clique guarded fragment* (CGF) [17] properly extends GF but still enjoys the finite model property as shown by Hodkinson [23], see also [24] for a simpler proof. Guardedness has emerged as a main new paradigm for decidability and other benign properties such as the FMP, and has applications in various areas of computer science. While the guarded fragment was originally introduced to embed and naturally extend propositional modal logics within first-order logic [1], it has various applications and was more recently shown to be relevant to description logics [15] and to database theory [41, 7]. Fragments of GF were recently studied for query answering in such contexts, see e.g. [7, 9, 8, 41, 10, 6, 37]. The main problems studied in the present paper are motivated by such applications.

1.1. **Main problems studied.** In the present paper we study the problem of *querying guarded theories* using conjunctive queries or unions of conjunctive queries. A Boolean conjunctive query (BCQ) q consists of an existentially closed conjunction of atoms. A union of Boolean conjunctive queries (UCQ) is a disjunction of a finite number of BCQ. If φ is a guarded sentence (or, equivalently, a finite guarded theory), we say that a query q evaluates to true against φ, iff $\varphi \models q$. In this context, we consider the following non-trivial main questions.

Finite controllability. Is it true that for each GF-sentence φ and each UCQ q if all finite models of φ satisfy q then the same is true also for all infinite models of φ, in symbols, that $\varphi \models q \iff \varphi \models_{\mathrm{fin}} q$? Since the query q may not be guarded, the finite model property of the guarded fragment is not sufficient to answer this question positively. Rather, this question amounts to whether for each φ and q as above, whenever $\varphi \wedge \neg q$ is consistent, it also has a finite model. This is equivalent to the finite model property of the extended fragment GF^+ of GF, where universally quantified Boolean combinations of negative atoms can be conjoined to guarded sentences. The concept of finite controllability was introduced by Rosati [41, 42].[1]

Size of finite models. How can we bound the size of finite models? In particular, in case $\varphi \not\models q$, how can we bound the size of the smallest finite models $\mathfrak{M}$ of φ for which $\mathfrak{M} \models \neg q$? Note that any recursive bound on the size of such models $\mathfrak{M}$ immediately yields the decidability of query answering. On the other hand, if φ is consistent and $\varphi \models q$, then the existence of a finite model $\mathfrak{M}$ such that $\mathfrak{M} \models q$ follows trivially from the FMP of GF, because every model $\mathfrak{M}$ of φ is also a model of q. However, little was known about the size of the smallest finite models of a satisfiable guarded sentence φ. Grädel's finite-model construction in [16], in case of unbounded arities, first transforms φ into a doubly exponentially sized structure,

[1] Rosati's definition is slightly stronger (see Proposition 1); the definition given here is, however, better suited for the full guarded fragment.

which is then input to a transformation according to Herwig's theorem [22], requiring a further exponential blow-up in the worst case. This suggests a triple-exponential upper bound. Can we do better?

Hypergraph covers. Approaching the above problems on a slightly more abstract level we construct hypergraph covers satisfying certain acyclicity criteria, which we refer to as weakly N-acyclic covers (see Section 2.2). Very informally, a hypergraph cover is a bisimilar companion $\hat{\mathfrak{H}}$ of $\mathfrak{H}$ with a bisimulation induced by a homomorphic projection $\pi : \hat{\mathfrak{H}} \to \mathfrak{H}$. And weak N-acyclicity implies that every subset of $\hat{\mathfrak{H}}$ of size at most N projects via π into an acyclic (not necessarily induced) sub-hypergraph of $\mathfrak{H}$. This notion is thus intimately related to the query answering problem and the existence of finite covers is a key to finite controllability. This problem has been previously studied by the third author, who in [34] gave a non-elementary construction of weakly N-acyclic hypergraph covers.

Of further interest, in particular in connection with finite controllability of query answering for the more general clique-guarded fragment, is the existence of hypergraph covers that are both conformal and weakly N-acyclic for a suitable N. Existence of conformal covers, with no regard to acycliciy constraints, was established in [24]; their doubly exponential construction being the only known bound.

Is it possible to find better, possibly polynomial constructions of hypergraph covers with the above properties?

Decidability and complexity. Is UCQ-answering over guarded theories decidable, and if so, what is the complexity of deciding whether $\varphi \models q$ for a UCQ q and a guarded sentence φ?

Canonisation and capturing. As a further problem of independent interest, which can be solved on the basis of the methods developed for the above questions, is PTIME canonisation — the problem of providing a unique representative for each guarded bisimulation equivalence class of structures, to be computed in PTIME from any given member of that class. This has implications for capturing the guarded bisimulation invariant fragment of PTIME in the sense of descriptive complexity.

We provide answers to all these questions. Before summarising our results, let us briefly explain how the above questions relate to database theory and description logic.

1.2. **Applications to databases and description logic.** In the database area, *query answering under integrity constraints* plays an important role. In this context a relational database D, consisting of a finite set (conjunction) of ground atoms is given, and a set Σ of *integrity constraints* is specified on D. The database D does not necessarily satisfy Σ, and may thus be "incomplete". The problem of answering a BCQ q on D under Σ consists of determining whether $D \cup \Sigma \models q$, also written as $(D, \Sigma) \models q$.

An important class of integrity constraints in this context are the so-called tuple-generating dependencies [5]. Given a relational schema (i.e., signature) $\mathcal{R}$, a *tuple-generating dependency* (TGD) σ over $\mathcal{R}$ is a first-order formula of the form $\forall\overline{x}\forall\overline{y}\big(\Phi(\overline{x},\overline{y}) \to \exists\overline{z}\,\Psi(\overline{x},\overline{z})\big)$, where $\Phi(\overline{x},\overline{y})$ and $\Psi(\overline{x},\overline{z})$ are conjunctions of atoms over $\mathcal{R}$, called the *body* and the *head* of σ, respectively. It is well known that database query answering under TGD is undecidable, see [4], even for very restricted cases [7]. For the relevant class of *guarded TGD* [7], however, query answering is decidable and actually 2EXPTIME-complete [7]. A TGD σ is *guarded*

(GTGD) if it has an atom in its body that contains all universally quantified variables of σ. For example, the sentence

$$\forall M, N, D \ \big((\mathit{Emp}(M,N,D) \wedge \mathit{Manages}(M,D)) \rightarrow \\ \exists E, N' \ (\mathit{Emp}(E,N',D) \wedge \mathit{Reportsto}(E,M))\big)$$

is a GTGD stating that if M is a manager named N belonging to and managing department D, then there must be at least one employee E having some name N' in department D reporting to M. In general, GTGD are, strictly speaking, not guarded sentences, because their heads may be unguarded. However, by using "harmless" auxiliary predicates and splitting up TGD heads into several rules, each set of GTGD can be rewritten into a guarded sentence that is (for all relevant purposes) equivalent to the original set. For instance, the above TGD can be rewritten into the following three guarded sentences

$$\forall M, N, D \ \big((\mathit{Emp}(M,N,D) \wedge \mathit{Manages}(M,D)) \rightarrow \exists E, N' \ \mathit{aux}(M,D,E,N')\big);$$
$$\forall M, D, E, N' \ \big(\mathit{aux}(M,D,E,N') \ \rightarrow \ \mathit{Emp}(E,N',D)\big);$$
$$\forall M, D, E, N' \ \big(\mathit{aux}(M,D,E,N') \ \rightarrow \ \mathit{Reportsto}(E,M)\big).$$

The class of *inclusion dependencies* (ID) is a simple subclass of the class of GTGD. An ID has the logical form $\forall\,\overline{x},\overline{y}\,(\alpha(\overline{x},\overline{y}) \rightarrow \exists \overline{z}\,\beta(\overline{x},\overline{z}))$, where α and β are single atoms. In [26] it was shown that query answering under ID is decidable and, more precisely, PSpace-complete in the general case and NP-complete for bounded arities. One very important problem was left open in [26]: the finite controllability of query answering in the presence of IDs. Given that in the database world attention is limited to *finite databases*, a Boolean query that would be false in infinite models of $D \cup \Sigma$ only, would still be finitely satisfied by $D \cup \Sigma$ and should be answered positively. Do such queries exist? This problem was solved by Rosati [41], who, by using a finite model generation procedure called *finite chase*, showed that query answering in the presence of IDs is finitely controllable. Rosati's result is actually formulated as follows.

Proposition 1 (Rosati [42])**.** *For every finite set of facts D and set $\mathcal{I}$ of ID and for every N there exists a finite structure $\mathfrak{C}$ extending D and satisfying $\mathcal{I}$ and such that for every Boolean conjunctive query q comprised of at most N atoms $\mathfrak{C} \models q$ iff $D, \mathcal{I} \models q$.*

Description logics are used for ontological reasoning in the Semantic Web and in other contexts. Some description logics such as DL-Lite$_{core}$ and DL-Lite$_{\mathcal{R}}$ [6] are essentially based on IDs, and are thus finitely controllable. The already mentioned class of GTGD and the yet more expressive class of *weakly guarded TGD (WGTGD)* have been introduced and studied in [7, 9] as powerful tools for data integration, data exchange [13], and ontological reasoning. As shown in [9], the class of GTGD augmented to also allow rules with the truth constant $\perp$ (= "false") as their head, generalizes the main DL-Lite description logic families. The WGTGD class is yet more general, and captures, unlike the GTGD class, plain Datalog. The finite controllability of GTD and WGTD theories, however, was left as an open problem. Unfortunately, Rosati's finite chase cannot be directly applied to GTGD or to WGTGD.

Let us briefly sketch how finite controllability of query answering in the presence of GTGD and WGTGD follows from the finite controllability of query answering against GF, which is the main result of the present paper. For GTGD theories this is easy. As explained above, they can be rewritten as guarded sentences and can thus be considered a sub-fragment of GF. Let us now turn our attention to WGTGD theories, and first give some intuition about how they are defined. Roughly, for a TGD set Σ, the set of all argument positions

Π of all atoms of predicates of Σ can be partitioned into sets Π_A and Π_U. Π_A are the so-called *affected positions*, where, when the rules are executed over a database D, Skolem terms (i.e., new instance values of existentially quantified head-variables) may need to be introduced, whereas P_U are those argument positions, which never need to hold Skolem terms (see [7] for a more precise definition). A TGD set Σ is *weakly guarded* (i.e., Σ is a WGTGD set) if each rule body has an atom (a *weak guard*) that covers all those body variables that only occur in affected positions. Note that weakly guarded TGD sets are, in general, unguarded. However, each theory (D, Σ), where D is a database and Σ a WGTDG set can be replaced by an equivalent theory (D, Σ') where Σ' is a GTGD set as follows: Σ' is obtained from Σ by replacing each rule σ of Σ by all possible instantiations of variables in unaffected positions in Σ by constants from D. It is easy to see that for each UCQ q, $(D, \Sigma) \models q$ iff $(D, \Sigma') \models q$ and, moreover, if $(D, \Sigma') \not\models q$, then for each model M of D and Σ' such that $M \not\models q$ it also holds that $M \models \Sigma$. It follows that query answering under WGTGD theories is finitely controllable if query answering under GTGD theories is finitely controllable. Thus, if we can establish that query answering under GF theories is finitely controllable, then query answering under GTGD theories (constituting a sub-fragment of GF) is finitely controllable, and so is query answering under WGTGD theories.

1.3. **Summary of results.**

Finite Controllability. That answering UCQ against guarded sentences is finitely controllable was already implicit in the report [34], although not formulated in this terminology. The finite models constructed in [34] are of non-elementary size and do not yield meaningful complexity results. The following central result of our paper, derived by a completely new proof, yields a much better size bound.

Theorem 2. *For every* GF *sentence* φ *and every* UCQ q, $\varphi \models q \iff \varphi \models_{\mathrm{fin}} q$. *More specifically, if* $\varphi \wedge \neg q$ *is satisfiable then it has a model of size* $2^{|\varphi||q|^{c|q|}}$, *where* c *depends solely on the signature of* φ.

Corollary 3. *Answering* UCQ *against GTGD or WGTGD theories is finitely controllable.*

More refined estimates on the size of finite models are provided in Section 4. To obtain Theorem 2, we establish new results on hypergraph covers, which are of independent interest.

Hypergraph Covers. We relate finite controllability to the concept of hypergraph covers. A *hypergraph cover* for a given hypergraph $\mathfrak{A}$ consists of a hypergraph $\mathfrak{B}$ together with a homomorphism $\pi\colon \mathfrak{B} \xrightarrow{\sim} \mathfrak{A}$ that induces a hypergraph bisimulation between $\mathfrak{B}$ and $\mathfrak{A}$. This notion naturally extends to relational structures $\mathfrak{A}, \mathfrak{B}$ on the basis of homomorphism-induced guarded bisimulations. The following main technical result is used to derive most other results (for definitions of notions mentioned see Section 2).

Theorem 4 (Main Technical Result)**.** *Given* $N \geq 2$ *and a hypergraph* $\mathfrak{A}$ *one can construct an* N*-conformal hypergraph* $\mathfrak{R}_N$ *constituting a weakly* N*-acyclic hypergraph cover of* $\mathfrak{A}$. *In particular,* $\mathfrak{R}_N$ *is conformal whenever* $N > w$, *where* w *is the width of* $\mathfrak{A}$; *moreover,* $|\mathfrak{R}_N| = |\mathfrak{A}|^{w^{\mathcal{O}(N)}}$ *and, for fixed* w *and* N, $\mathfrak{R}_N$ *can be computed in polynomial time.*
The direct analogue for guarded covers of relational structures follows.
We call $\mathfrak{R}_N$ *the* Rosati cover *of* $\mathfrak{A}$.

Let us explain very informally the role of the Rosati covers in establishing Theorem 2. In an easy but key step (Lemma 13 in Section 2) we first reduce a problem instance $\varphi \models q$ for a GF-sentence φ and a UCQ q to the equivalent question of entailment $\varphi \models \chi_q$, where χ_q is a disjunction of *acyclic queries* stemming from the original query q. Crucially, being acyclic, χ_q can be equivalently reformulated as a GF-sentence, ultimately reducing the initial query answering problem to the unsatisfiability of the GF-sentence $\varphi \wedge \neg\chi_q$. It is more difficult to show that this reduction is also valid over finite models, i.e., that $\varphi \models_{\mathrm{fin}} q \iff \varphi \models_{\mathrm{fin}} \chi_q$. In particular, that given a finite $\mathfrak{A} \models \varphi \wedge \neg\chi_q$, a finite model of $\varphi \wedge \neg q$ can also be found. Observe that the "unravelling" of $\mathfrak{A}$ yields a tree-like model $\mathfrak{A}^*$ of $\varphi \wedge \neg\chi_q$ and an acyclic cover of $\mathfrak{A}$. Thus, by virtue of acyclicity, $\mathfrak{A}^* \models \neg q$. However, $\mathfrak{A}^*$ is typically infinite. The challenge is to find a finite cover of $\mathfrak{A}$ retaining a "sufficient degree of acyclicity" so as not to render it a model of q. This is captured by the notion of a weakly N-acyclic cover $\mathfrak{A}^{(N)}$ of $\mathfrak{A}$, which ensures that, similarly to tree unravellings, $\mathfrak{A}^{(N)} \models q$ implies $\mathfrak{A} \models \chi_q$, but with the qualification that $|q| \leq N$. Theorem 4 shows that such covers can be constructed.

Conformal covers. Hodkinson and Otto showed in [24] that all hypergraphs admit guarded bisimilar covers by conformal hypergraphs (for definitions, see Section 2). While the construction in [24] involves a doubly exponential blow-up in size, we here obtain a polynomial construction of conformal covers as a corollary to Theorem 4.

Corollary 5. *Every hypergraph $\mathfrak{H}$ of width w admits a conformal hypergraph cover of size $|\mathfrak{H}|^{w^{\mathcal{O}(w)}}$. For bounded width, we thus obtain polynomial size conformal covers.*

Finite model property of the clique-guarded fragment. As it happens, our construction used for Theorem 2 also yields an extension of Theorem 2 to the clique-guarded fragment, CGF.

Theorem 6. *For every $\varphi \in$ CGF and every $q \in$ UCQ we have $\varphi \models q \iff \varphi \models_{\mathrm{fin}} q$. More specifically, if $\varphi \wedge \neg q$ is satisfiable then it has a finite model of size $2^{(|\varphi|+|\tau|^{\mathcal{O}(h)})(hw)^{\mathcal{O}(hw)}}$, where h is the height of q, τ is the signature of φ, and $w = \max\{\mathrm{width}(\varphi), \mathrm{width}(\tau)\}$.*

In particular, we obtain finite models of any satisfiable clique-guarded formula, and thereby a new proof of the Finite Model Property of the clique-guarded fragment. In fact, our construction yields more compact finite models than hitherto known.

Small model property. Through our new method of finite-model construction, we are able to improve the bounds implicit in [16] for GF and the overhead for CGF implicit in [23, 24] on the size of the smallest finite model of a satisfiable (clique-)guarded sentence.

Theorem 7. *Every satisfiable formula of* CGF (*and thus of* GF) *has a finite model of size exponential in the length and doubly exponential in the width of the formula. Moreover, for every $k \geq 2$, the k-variable fragment of* CGF (GF) *has finite models of exponential size in the length of the formula.*

Another important fragment of first-order logic that has the finite model property is 2-variable first-order logic, denoted by FO^2. It was shown in [19] that if an FO^2 formula ϕ is satisfiable, then it has a model of cardinality singly exponential in the size of ϕ, improving an earlier doubly exponential bound by Mortimer [31]. As a consequence, [19] also proved NExpTime-completeness of the satisfiability problem for FO^2. A more powerful fragment that embeds a number of key features of description logics is the extension of the 2-variable fragment with counting quantifiers, denoted C^2. In contrast to FO^2, C^2 does not have the finite model property, but computationally it is no more difficult than FO^2: both satisfiability and finite satisfiability of C^2-formulas are decidable and NExpTime-complete [20, 40, 38].

In [15] Grädel proposed the guarded fragment of FO^2 as a testbed for simple description logics. A more suitable fragment for this purpose is obtained by imposing on the one hand a restriction to guarded quantification while allowing on the other hand the use of counting quantifiers. The two-variable guarded fragment with counting quantifiers, GC^2, properly subsumes the description logic $\mathcal{ALCQI}$, cf. e.g. [11], for which finite satisfiability was shown decidable in ExpTime by Lutz et al. [30]. In [27] Kazakov gave a polynomial, satisfiability-preserving translation from GC^2 to the 3-variable guarded fragment GF^3 thus establishing ExpTime-completenes of satisfiability of GC^2-formulas. Finite satisfiability for GC^2 was also shown to be ExpTime-complete by Pratt-Hartmann [39]. The latter decision method is based on a reduction to integer programming. It is interesting to note that the optimal lower bound on the size of smallest finite models of finitely satisfiable GC^2-formulas is doubly exponential [39].

Complexity of query answering. In [16] Grädel proved that satisfiability of GF-sentences is complete for 2ExpTime, and ExpTime-complete in case of bounded arity. We show that, more generally, answering UCQ on the class of models of a GF-sentence can also be performed in 2ExpTime, which solves the initially posed complexity question about query answering over guarded theories. It follows from the work of Lutz [28], however, that query answering remains 2ExpTime-complete for BCQ even in the bounded arity case. Considering unions of *acyclic* conjunctive queries we derive an ExpTime solution to the query answering problem, and prove ExpTime-completeness already for a particular fixed GF sentence.

Our algorithms are built around Grädel's solution of the satisfiability problem for GF. The first step consists in reducing $\varphi \models q$ to $\varphi \models \chi_q$, where χ_q is a disjunction of acyclic queries stemming from the original query. The formula χ_q may, however, be of exponential size in terms of the length of q, demanding a closer inspection of the contribution of different dimensions of χ_q to the overall complexity of checking (un)satisfiability of the guarded sentence $\varphi \wedge \neg\chi_q$.

We also investigate the problem of query answering over models of a fixed guarded sentence, and provide a number of useful bounds. Our bounds for fixed sentences φ are not all tight and leave room for future research.

Canonisation and capturing. As a further consequence of the proof method employed for Theorem 4, we find a polynomial solution to the canonisation problem for guarded bisimulation equivalence $\sim_{\mathrm{g}}$. This allows us to capture the $\sim_{\mathrm{g}}$-invariant fragment of Ptime in the sense of descriptive complexity, i.e., to provide effective syntax with Ptime model checking for the Ptime queries closed under guarded bisimulation equivalence. Canonisation is achieved through inversion of the natural game invariant $\mathbb{I}(\mathfrak{A})$ that uniquely characterises the guarded bisimulation class, or the complete GF-theory, of a given finite structure $\mathfrak{A}$. A Ptime reconstruction of a model from the abstract specification of its equivalence class yields Ptime canonisation.

Theorem 8. *For every relational signature τ there exists a* Ptime *algorithm computing from a given invariant $\mathbb{I}(\mathfrak{A})$ of an unspecified τ-structure $\mathfrak{A}$ a finite τ-structure* $\mathrm{can}(\mathfrak{A})$ *such that* $\mathbb{I}(\mathrm{can}(\mathfrak{A})) = \mathbb{I}(\mathfrak{A})$*; hence* $\mathrm{can}(\mathfrak{A}) \sim_{\mathrm{g}} \mathfrak{A}$ *and* $\mathrm{can}(\mathfrak{A}') = \mathrm{can}(\mathfrak{A})$ *whenever* $\mathfrak{A} \sim_{\mathrm{g}} \mathfrak{A}'$.

Corollary 9. *The class of all those* PTIME *Boolean queries that are invariant under guarded bisimulation,* PTIME$/\!\sim_{\mathrm{g}}$*, can be captured in the sense of descriptive complexity.*

Organisation. Section 2 defines the main concepts and introduces guarded bisimilar hypergraph covers as a main tool. It also states the above-mentioned Lemma 13. Section 3 presents the construction of the Rosati cover. From this and Lemma 13, the finite controllability of GF is proven in Section 4. Section 5 establishes our new complexity results. Section 6 deals with canonisation and capturing.

2. Hypergraphs and guarded fragments

We work with finite relational signatures. Let us fix such a signature τ and let width(τ) denote the the maximal arity of any of the relation symbols in τ.

2.1. Guarded fragments.

The guarded fragment of first-order logic, GF, as introduced by Andréka et al. [1], is the collection of first-order formulas with certain syntactic restrictions in the quantification pattern, which is analogous to the relativised nature of modal logic. The set of $\mathsf{GF}(\tau)$ formulas is the smallest set

- containing all atomic formulas of signature τ and equalities between variables;
- closed under Boolean connectives: $\neg$, $\wedge$, $\vee$, $\rightarrow$, $\leftrightarrow$;
- and such that whenever $\psi(\overline{x},\overline{y})$ is a $\mathsf{GF}(\tau)$ formula with all free variables indicated and $\alpha(\overline{x},\overline{y})$ is a τ-atom (or an equality) involving all free variables of ψ, then the following are in $\mathsf{GF}(\tau)$ as well:

$$(\forall\,\overline{x}.\,\alpha)\,\psi := \forall\overline{x}(\alpha \rightarrow \psi) \quad \text{and} \quad (\exists\,\overline{x}.\,\alpha)\,\psi := \exists\overline{x}(\alpha \wedge \psi).$$

In a τ-structure $\mathfrak{A}$ a non-empty set X of elements is said to be *guarded* if it is a singleton or there is an atom $R^{\mathfrak{A}}(\overline{a})$ such that every member of X occurs in $\overline{a}$. A *maximal guarded set* is one not properly included in any other guarded set. A tuple $\overline{b}$ of elements is guarded if the set of its components is guarded.

While GF provides an important extension of the modal fragment, guarded quantification is too restrictive to express some basic temporal operators. To remedy this shortcoming various relaxations of the notion of guardedness and corresponding fragments have been introduced, chief among them the clique-guarded fragment.

The *clique guarded fragment*, CGF, relaxes the constraints on guards α in GF to allow existentially quantified conjunctions of atoms as guards that guarantee that the tuple of free variables is clique-guarded. A set X of elements of a structure $\mathfrak{A}$ is *clique-guarded* if every pair of elements of X is guarded, equivalently, if X induces a clique in the Gaifman graph of $\mathfrak{A}$. A tuple $\overline{a}$ is clique-guarded whenever the set of its components is. Observe that while guarded sets are bounded in size by the width of the signature, there can be arbitrarily large clique-guarded sets whenever the width is at least 2. Recall that the *width* of a formula φ, width(φ) is the maximal number of free variables in any of its subformulas. In a clique-guarded formula φ the maximal size of a clique-guarded set quantified over is bounded by width(φ).

Observe that guardedness and clique-guardedness (of tuples of any fixed arity) are definable in the corresponding logic. That is, there are formulas $\mathtt{guarded}_{\mathcal{L}}(\overline{x})$ expressing that the tuple $\overline{x}$ is guarded in the sense appropriate for the fragment. E.g., $\mathtt{guarded}_{\mathsf{GF}}(\overline{x}) = \bigvee_{\alpha} \exists \overline{y} \alpha(\overline{x}, \overline{y})$ where α ranges over all τ-atoms in tuples of variables comprising at least $\overline{x}$, as indicated. A formula $\mathtt{guarded}_{\mathsf{CGF}}(\overline{x})$ can be similarly defined.

An *atomic τ-type* $t(x_1, \dots, x_n)$ is a maximal consistent set of τ-literals (atoms or negated atoms, including (in)equalities) whose constituent terms are among the variables $x_1, \dots, x_n$ and the constants from τ. An atomic type $t(x_1, \dots, x_n)$ determines, for every choice of indices $\bar{\imath} = (i_1, \dots, i_k)$, its *restriction* to components $\bar{\imath}$, which is an atomic type in k variables $(x_{i_1}, \dots, x_{i_k})$ denoted $t|_{\bar{\imath}}$; conversely we say that t is an *extension* of $t|_{\bar{\imath}}$. In a τ-structure $\mathfrak{A}$ the atomic type $\mathrm{atp}_{\mathfrak{A}}(\overline{a})$ of a tuple $\overline{a}$ is the unique atomic type $t(\overline{x})$ such that $\mathfrak{A} \models t(\overline{a})$. One says that t is *realised* by $\overline{a}$ in $\mathfrak{A}$. Each atomic type can be identified with the isomorphism type of the sub-structure induced by any tuple realising it. Over a signature of r many relational symbols of maximal arity w and k constants there are $2^{\mathcal{O}(r(n+k)^w)}$ many atomic types in n variables. We identify each atomic type with the conjunction of its literals.

Guarded bisimulation. The notion of guarded bisimulation [16], denoted $\sim_{\mathrm{g}}$, can be defined either in terms of the guarded bisimulation game, a variant of the Ehrenfeucht-Fraïssé style pebble game in which the set of pebbles must at any given time be guarded, or as a back-and-forth system of partial isomorphisms whose domain and image are both guarded. GF is preserved under guarded bisimulation [1], see also [16, 18]:

$$\mathfrak{A} \sim_{\mathrm{g}} \mathfrak{B} \quad \Longrightarrow \quad \text{for all } \varphi \in \mathsf{GF}: \ \mathfrak{A} \models \varphi \ \Leftrightarrow \ \mathfrak{B} \models \varphi.$$

Given a relational structure $\mathfrak{A}$, its *guarded bisimulation game graph*, denoted $\mathbb{G}(\mathfrak{A})$, has as its vertices the set $G(\mathfrak{A})$ of all *maximal guarded tuples* of $\mathfrak{A}$, each labeled by its atomic type, equivalently, by the isomorphism type of the substructure induced by the tuple. Two such tuples $\overline{a}$ and $\overline{b}$ are linked by an edge labeled by a partial bijection $\rho \subseteq \{1, \dots, |\overline{a}|\} \times \{1, \dots, |\overline{b}|\}$ whenever $a_i = b_j$ for all $(i,j) \in \rho$. Note that structures $\mathfrak{A}$ and $\mathfrak{B}$ are guarded bisimilar iff $\mathbb{G}(\mathfrak{A})$ and $\mathbb{G}(\mathfrak{B})$ are bisimilar in the modal sense [18].

The *guarded bisimulation invariant* $\mathbb{I}(\mathfrak{A})$ of $\mathfrak{A}$ is defined as the bisimulation quotient of $\mathbb{G}(\mathfrak{A})$. Vertices of $\mathbb{I}(\mathfrak{A})$ correspond to $\sim_{\mathrm{g}}$-classes of maximal guarded tuples of $\mathfrak{A}$, labeled by their atomic types (induced isomorphism types). A ρ-labeled edge links vertices v and w if there are guarded tuples $\overline{a}$ and $\overline{b}$ in $\mathfrak{A}$ realising the $\sim_{\mathrm{g}}$-classes represented by v and by w, respectively, and such that $a_i = b_j$ for all $(i,j) \in \rho$.

Scott normal form and satisfiability criterion. Grädel's analysis of decidability for GF [16] uses the following Scott normal form corresponding to a relational Skolemisation.

Lemma 10 ([16, Lemma 3.1])**.** *To every (clique-)guarded τ-sentence φ one can associate a companion (clique-)guarded $\tau \cup \sigma$-sentence*

$$\psi = \bigwedge_j (\forall \overline{x}.\alpha_j)\, \vartheta_j(\overline{x}) \ \wedge\ \bigwedge_i (\forall \overline{x}.\beta_i)\, (\exists \overline{y}.\gamma_i)\ \psi_i(\overline{x}, \overline{y}) \tag{2.1}$$

such that $\psi \models \varphi$ and every $\mathfrak{A} \models \varphi$ has a $\tau \cup \sigma$-expansion $\mathfrak{B} \models \psi$. Here $|\sigma| \leq |\varphi|$, $\mathrm{width}(\psi) = \mathrm{width}(\varphi)$ *and the ϑ_j, ψ_i are quantifier-free.*

A guarded bisimulation game graph G or, similarly, a guarded-bisimulation invariant I is said to satisfy the formula ψ in Scott normal form (2.1) if

- its vertices are labeled by atomic types in the signature of ψ that are guarded and that satisfy the universal conjuncts of ψ; and
- for each vertex v with label $t(\overline{xz})$ and for each conjunct $(\forall\,\overline{x}.\beta_i)\,(\exists\,\overline{y}.\gamma_i)\;\psi_i(\overline{x},\overline{y})$ of ψ such that $t(\overline{xz}) \models \beta_i(\overline{x})$ there exists a vertex w labeled with some type $s(\overline{x}'\overline{y}) \models \psi_i(\overline{x}',\overline{y})$ such that $s|_{\overline{x}'} = t|_{\overline{x}}$ and v and w are linked by an edge labeled with the mapping $\rho : \overline{x} \to \overline{x}'$.

Proposition 11 (cf. [16, Lemma 3.4])**.** *Let ψ be the normal form of φ as in* (2.1)*. Then φ is satisfiable if, and only if, there exists a guarded bisimulation invariant $\mathfrak{I}$ satisfying ψ and such that vertices of $\mathfrak{I}$ are labeled by distinct guarded atomic types.*

2.2. **Hypergraphs, acyclicity and covers.** A *hypergraph* is a pair $H = (V, S)$ with V its set of elements and $S \subseteq \mathcal{P}(V)$ a set of subsets of V, which are called hyperedges. For a set of hyperedges S, let $S\Downarrow$ stand for the closure of S under subsets. A set X of elements of H is guarded if $X \in S\Downarrow$. The Gaifman graph $\Gamma(H)$ of H is the undirected graph having vertex set V and, as edges, all non-degenerate guarded pairs of H. The maximal size of any hyperedge is referred to as the *width* of H. To every τ-structure $\mathfrak{A}$ one associates in a natural way a hypergraph $H[\mathfrak{A}]$ with V the universe of $\mathfrak{A}$ and S the collection of maximal guarded subsets of $\mathfrak{A}$. The width of $H[\mathfrak{A}]$ is then bounded by $\mathrm{width}(\tau)$. The Gaifman graph of $\mathfrak{A}$ is $\Gamma[\mathfrak{A}] := \Gamma(H[\mathfrak{A}])$.

A *homomorphism* $h\colon H \to H'$ between hypergraphs $H = (V, S)$ and $H' = (V', S')$ is a map from V to V' such that $h(s) \in S'$ for all $s \in S$. A hypergraph homomorphism h is *rigid* if $|h(s)| = |s|$ for every hyperedge s. Every homomorphism $h\colon \mathfrak{A} \to \mathfrak{A}'$ between relational structures induces a hypergraph homomorphism from $H[\mathfrak{A}]$ to $H[\mathfrak{A}']$.

Game graphs $G(H)$ and invariants $I(H)$ are defined similarly for hypergraphs H as for relation structures, where instead of guarded bisimulation we use the natural notion of *hypergraph bisimulation*. It is safe to think of hypergraph bisimulation as of guarded bisimulation stripped of all atomic relational information. Vertices in the game graph are maximal hyperedges each labeled with the isomorphism type of the sub-hypergraph induced by it, in other words, the label carries the information about all hyperedges lying inside a maximal hyperedge. Edges of the game graph connect overlapping maximal hyperedges and are in bijection with and are labeled by partial bijections compatible with the actual overlap.

A hypergraph H is (N-)*conformal* if every clique in $\Gamma(H)$ (of size at most N) is covered by a hyperedge of H. A structure $\mathfrak{A}$ is (N-)conformal whenever $H[\mathfrak{A}]$ is, i.e., if every k-clique ($k \leq N$) in its Gaifman graph is covered by a ground atom. Over conformal structures guarded quantification is as powerful as clique-guarded quantification.

A hypergraph H is (N-)*chordal* if all cycles in $\Gamma(H)$ of length greater than 3 (and at most N) have a chord in $\Gamma(H)$. An analogous notion for relational structures $\mathfrak{A}$ is similarly defined in terms of the Gaifman graph $\Gamma(H(\mathfrak{A}))$.

A hypergraph is (N-)*acyclic* if it is both (N-)chordal and (N-)conformal. For finite hypergraphs acyclicity is equivalent to tree decomposability. A finite hypergraph is *tree decomposable* if it can be reduced to the empty hypergraph by iteratively deleting some non-maximal hyperedge or some vertex contained in at most one hyperedge (Graham's algorithm) cf. [3]. We say that a relational structure $\mathfrak{A}$ is *guarded tree decomposable* if $\mathfrak{A}$ allows a tree decomposition in the sense of Robertson–Seymour with guarded bags. This is equivalent to $H[\mathfrak{A}]$ being tree decomposable, i.e. acyclic.

Definition 1 (Cover)**.** *A* guarded bisimilar cover *(or just* cover *for short)* $\pi\colon \mathfrak{B} \xrightarrow{\sim} \mathfrak{A}$ *is an onto homomorphism* $\pi\colon \mathfrak{B} \to \mathfrak{A}$ *inducing a guarded bisimulation* $\{(\bar{b}, \pi(\bar{b})) \mid \bar{b}$ *maximal guarded tuple in* $\mathfrak{B}\}$ *between relational structures* $\mathfrak{B}$ *and* $\mathfrak{A}$ *of the same vocabulary. A cover* $\pi\colon \mathfrak{B} \xrightarrow{\sim} \mathfrak{A}$ *is* weakly N-conformal *if the image under* π *of any clique of size up to* N *in the Gaifman graph of* $\mathfrak{B}$ *is guarded in* $\mathfrak{A}$*; similarly it is* weakly N-chordal *if the image under* π *of every cycle of length greater than* 3 *and up to* N *in the Gaifman graph of* $\mathfrak{B}$ *has a chord in the Gaifman graph of* $\mathfrak{A}$*; and it is* weakly N-acyclic *if it is both weakly* N*-conformal and weakly* N*-chordal.*

Analogous notions of hypergraph covers are defined mutatis mutandis with the additional stipulation that the restriction of a cover homomorphism to every hyperedge is expressly required to be injective, i.e. that the cover homomorphism is to be rigid. (In the case of guarded bisimilar covers among relations structures the analogous condition is implied by the above definition. That is, every guarded bisimular cover $\pi\colon \mathfrak{B} \xrightarrow{\sim} \mathfrak{A}$ induces a hypergraph cover $\hat{\pi}\colon H[\mathfrak{B}] \xrightarrow{\sim} H[\mathfrak{A}]$ such that $\hat{\pi}$ is a rigid homomorphism of hypergraphs).

A homomorphism $h\colon H \to H'$ into the hypergraph $H' = (V', S')$ is called *tree decomposable* if there is some $S'' \subseteq S'{\downarrow}$ such that $h\colon H \to H''$ is a homomorphism into $H'' = (V', S'')$ and H'' is tree decomposable. This extends to guarded tree-decomposable relational structures in the usual manner. Every homomorphism into a (guarded) tree-decomposable hypergraph (structure) is trivially (guarded) tree decomposable. More generally observe the following.

Remark 12. *A hypergraph cover* $\pi\colon \mathfrak{B} \xrightarrow{\sim} \mathfrak{A}$ *is weakly* N*-acyclic iff for every homomorphism* $h\colon \mathcal{Q} \to \mathfrak{B}$ *from a hypergraph* $\mathcal{Q}$ *of at most* N *elements* $\pi \circ h$ *is tree decomposable. The situation for guarded bisimilar covers is strictly analogous.*

2.3. **Conjunctive queries.** *Conjunctive queries* (CQ) are formulas of the form $\exists \overline{x} \bigwedge_i \alpha_i$, where the α_i are positive literals. A Boolean conjunctive query (BCQ) is one with no free variables. A union of (Boolean) conjunctive queries (UCQ) is a disjunction of BCQ. The size $|q|$ of a UCQ q is its length as a formula, and its *height* is the maximal size of its disjuncts (constituent CQ).

To every BCQ $Q = \exists \overline{x} \bigwedge_i \alpha_i$ of signature τ one can associate the τ-structure $\mathcal{Q}$ having as its universe the set of variables in $\overline{x}$ and atoms as prescribed by the α_i of Q. Then $\mathfrak{A} \models Q$ iff there exists a homomorphism $h : \mathcal{Q} \to \mathfrak{A}$, [12]. We say that $Q \in \mathsf{CQ}$ is *acyclic* if the associated structure $\mathcal{Q}$ is acyclic. Note that this is equivalent to the existence of a *guarded* conjunctive query equivalent to Q, i.e., one that is both in GF and in CQ [14].

For each BCQ Q we define its *treeification in signature* τ, denoted χ^τ_Q, as the disjunction of all *acyclic* BCQ T in the signature τ comprised of *at most three times as many atoms* as Q and such that $T \models Q$. Further, for $q = \bigvee_i Q_i$ a UCQ we set $\chi^\tau_q = \bigvee_i \chi_{Q_i}$. It is obvious that $\chi^\tau_q \models q$ for every q. In the following τ will always be an expansion of the signature of q, and will be omitted whenever clear from the context or of no import.

Lemma 13. *Let* τ *be a signature consisting of* r *relation symbols of maximal arity* w *(the width of* τ*). Consider a* UCQ $q = \bigvee_i Q_i$ *over* τ *and let* $h = \max_i |Q_i|$ *(the height of* q*).*

(i) For a τ*-structure* $\mathfrak{A}$ *we have* $\mathfrak{A} \models \chi^\tau_q$ *if there is a guarded tree decomposable homomorphism* $\eta : \mathcal{Q}_i \to \mathfrak{A}$ *for some* Q_i*.*

(ii) In particular, for all $\varphi \in \mathsf{GF}[\tau]$*:* $\varphi \models q$ *iff* $\varphi \models \chi^\tau_q$*.*

(iii) *The size of the treeification* χ_q^τ *is at most* $r^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}$*;*
moreover, χ_q^τ *can be constructed in time* $|q|r^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}$.

Proof. *(i)* Let $\eta : \mathcal{Q}_i \to \mathfrak{A}$ be a guarded tree decomposable homomorphism. This means that $\eta : \mathcal{Q}_i \to \mathfrak{B} \subseteq \mathfrak{A}$ for some guarded tree decomposable $\mathfrak{B}$, a (not-necessarily induced) substructure of $\mathfrak{A}$.

Consider a fixed guarded tree decomposition of $\mathfrak{B}$ represented as $(V, \preceq, \gamma)$ with $(V, \preceq)$ a forest with transitive edges and γ assigning to each node an atom of $\mathfrak{B}$, the guard of the corresponding bag of the tree decomposition.

As η maps $\mathcal{Q}_i$ homomorphically into $\mathfrak{B}$, for each atom α of $\mathcal{Q}_i$ we can pick a node $v_\alpha \in V$ with $\gamma(v_\alpha)$ guarding the image of α under η. Let W be the closure of the set of all these v_α under greatest lower bounds w.r.t. $\preceq$. Then $(W, \preceq, \gamma|_W)$ represents a guarded tree decomposition of the structure $\mathcal{T}$ consisting of those atoms in the image of η together with atoms of the form $\gamma(w)$ for $w \in W$. Note that at least half of the nodes in W are of the form v_α, therefore $\mathcal{T}$ has at most three times as many atoms as $\mathcal{Q}_i$. And we have $\eta : \mathcal{Q}_i \to \mathcal{T} \subseteq \mathfrak{B}$.

To $\mathcal{T}$ corresponds an acyclic BCQ T, whose models are precisely those structures containing a homomorphic image of $\mathcal{T}$. Then $T \models Q_i$ and hence T is one of the disjuncts in χ_q^τ. Therefore $\mathfrak{A} \models \chi_q^\tau$.

(ii) Since $\chi_q^\tau \models q$, trivially $\varphi \models \chi_q^\tau$ implies $\varphi \models q$. To prove the converse implication assume indirectly that $\varphi \models q$ but $\varphi \wedge \neg\chi_q^\tau$ were satisfiable. Note that the latter is equivalent to a guarded formula. Then, by the tree model property of GF [16], there is a guarded tree decomposable model $\mathfrak{T} \models \varphi \wedge \neg\chi_q^\tau$. By our assumption $\mathfrak{T} \models q$, i.e. $\mathfrak{T} \models Q_i$ for some BCQ Q_i in q, which means that there is a homomorphism $\eta : \mathcal{Q}_i \to \mathfrak{T}$. Given that $\mathfrak{T}$ is guarded tree decomposable, so is η. By (i) therefore $\mathfrak{T} \models \chi_q^\tau$, contradicting our assumption.

(iii) Recall that, for a BCQ Q, the formula χ_Q^τ is a disjunction of several (acyclic) BCQ T, each of which has at most $3|Q|$ many atoms and therefore requires no more than $3|Q|w$ many variables; the overall number of constituent BCQ of these dimensions is bounded by $(r(3|Q|w)^w)^{3|Q|}$, and each such T has length $\mathcal{O}(|Q|w)$. All in all, $|\chi_Q^\tau| \leq (r(3|Q|w)^w)^{3|Q|}\mathcal{O}(|Q|w) = r^{\mathcal{O}(|Q|)}(|Q|w)^{\mathcal{O}(|Q|w)}$.

For a UCQ $q = \bigvee_i Q_i$ we have, by definition, $\chi_q^\tau = \bigvee_i \chi_{Q_i}^\tau$, and, if the height of q is h, then $|\chi_q^\tau| = r^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}$ by the previous estimate. One way to compute χ_q^τ is to exhaustively enumerate all acyclic CQ of the right dimensions and to check each one for entailment of some Q_i (verifiable in time $(hw)^{\mathcal{O}(hw)}$ for each i). Such a procedure can be carried out in time $|q|r^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}$. □

Concerning the size of treeifications, note that for a fixed signature the figure from (iii) simplifies to $|\chi_q^\tau| = h^{\mathcal{O}(h)}$ and that a $2^{\Omega(h)}$ lower bound can be established even if we require treeifications to be free of redundant disjuncts. Indeed, in the signature $\tau = \{E, T\}$, where E is binary and T is ternary, it is easy to see that the BCQ Q_n for which $\mathcal{Q}_n$ is a simple E-cycle of length n, the number of triangulations of $\mathcal{Q}_n$ and hence the number of disjuncts in $\chi_{Q_n}^\tau$ is $2^{\Omega(n)}$.

The next key fact is a direct consequence of Lemma 13 (i) and Remark 12 that highlights the role of query treeification and motivates our interest in weakly N-acyclic covers.

Fact 1. For every weakly N-acylic cover $\pi\colon \mathfrak{B} \xrightarrow{\sim} \mathfrak{A}$ of τ-structures, for every $\varphi \in \mathsf{GF}[\tau]$ and every $q \in \mathsf{UCQ}[\tau]$ of height at most N:

$$\mathfrak{B} \models q \implies \mathfrak{A} \models \chi_q^\tau \qquad \text{and hence} \qquad \mathfrak{A} \models \varphi \wedge \neg\, \chi_q^\tau \implies \mathfrak{B} \models \varphi \wedge \neg\, q\ . \tag{2.2}$$

Using this fact and the finite model property of the guarded fragment, Theorem 2 will follow, once it is established that every finite relational structure admits weakly N-acyclic covers for all N. This is precisely the content of Theorem 4. Before engaging in the proof of this main technical result let us point out a noteworthy consequence of item (ii) of Lemma 13.

An interpolation property. Consider the fragments GF and UCQ (equivalently, the positive existential fragment) of first-order logic in a relational signature. They are incomparable with respect to expressive power and, as mentioned above, the intersection of the two fragments comprises (up to semantic equivalence) precisely the unions of guarded conjunctive queries, or unions of acyclic conjunctive queries (ACQ). In one reading, Lemma 13 (ii) states that the fragments GF and UCQ have a strong form of interpolation with ACQ interpolants.[2] Indeed, consider some $\varphi \in \mathsf{GF}$ and $q \in \mathsf{UCQ}$ in signature τ. Then

$$\varphi \models q \qquad \implies \qquad \varphi \models \chi_q^\tau \ \text{ and } \ \chi_q^\tau \models q \tag{2.3}$$

and it is interesting to note that the treeification χ_q^τ of the query q is a uniform interpolant for all $\varphi \in \mathsf{GF}[\tau]$ that entail q. As a consequence of our Theorem 2 we will see that the interpolation property (2.3) remains intact when the semantics is restricted to finite models.

3. The Rosati cover

Rosati proved Proposition 1 using a "finite chase" procedure [41, 42] that safely reuses variables and results in very compact finite models. However, his proof of correctness of the finite chase with respect to conjunctive query answering is very intricate. We adapt the core idea of his model construction to give a more general guarded bisimilar cover construction for finite models, and a conceptually cleaner and simpler proof of faithfulness with respect to conjunctive queries of bounded size.

Theorem 14. *Given $N \geq 2$ and a bisimulation invariant $\mathfrak{I} = \mathbb{I}(\mathfrak{A})$ of an unspecified hypergraph $\mathfrak{A}$, one can construct hypergraphs $\mathfrak{R}_N$ and $\mathfrak{R}$ such that $I(\mathfrak{R}_N) = I(\mathfrak{R}) = \mathfrak{I}$ and $\mathfrak{R}_N$ is N-conformal and a weakly N-acyclic cover of $\mathfrak{R}$. We have $|\mathfrak{R}_N| = |\mathfrak{I}|^{w^{\mathcal{O}(N)}}$, where w is the width [3] of $\mathfrak{A}$ and, for fixed w and N, $\mathfrak{R}_N$ can be computed in polynomial time.*

The analogous claim for the guarded bisimulation invariant $\mathfrak{I} = \mathbb{I}(\mathfrak{A})$ of a finite relational structure $\mathfrak{A}$, and concerning guarded bisimilar covers, follows.

It is not hard to see how this formulation entails the statement of Theorem 4 as given in the introduction. Observe that $\mathbb{I}(\mathfrak{A}) = \mathbb{G}(\mathfrak{A})$, where $\mathbb{G}(\mathfrak{A})$ is the bisimulation game graph of the given $\mathfrak{A}$ (without passage to a non-trivial quotient), can be enforced by introducing new predicates to distinguish each individual guarded tuple of $\mathfrak{A}$. Then $\mathfrak{R}_N$ is a weakly N-acyclic (guarded) bisimilar cover of $\mathfrak{A}$ itself. Moreover, since N-conformality implies conformality at large assuming $N > w = \mathrm{width}(\mathfrak{A})$, $\mathfrak{R}_{w+1}$ is a conformal cover of $\mathfrak{A}$ and we also obtain Corollary 5.

[2] We thank Damian Niwinski for this observation.

[3] Note that this width is determined by $\mathfrak{I}$.

We first define the Rosati cover of a given finite hypergraph (or relational structure), for fixed N. After preliminary observations much resembling some of Rosati's key lemmas [41, 42] we prove its two crucial properties: weak N-chordality and N-conformality of $\mathfrak{R}_N$ over $\mathfrak{A}$. In fact, $\mathfrak{R}_N$ will be the top layer of a chain of covers of increasing degrees of weak chordality, similar to the construction of [34].

3.1. **The definition of $\mathfrak{R}_N$.** Let w be the width of $\mathfrak{A}$ (apparent from $\mathfrak{I} = \mathbb{I}(\mathfrak{A})$), i.e. the maximal size of any of its hyperedges (guarded sets). We assume throughout that $w > 1$, since width 1 is trivial. For the rest of this section we also fix $m \geq N \geq 2$.

Consider a relational structure $\mathfrak{A}$ and its guarded-bisimulation invariant $\mathfrak{I} = \mathbb{I}(\mathfrak{A})$. Recall that the vertices of $\mathbb{I}(\mathfrak{A})$ represent complete GF-types of maximal guarded tuples $\overline{a}$ such that $a_i \neq a_j$ for all $i \neq j$ and are labeled by the isomorphism type of the sub-structure induced by any (and all) corresponding tuple. We denote vertices of $\mathfrak{I}$ by symbols $d, e, \ldots$ and for each $e = [\overline{a}]_{\sim_{\mathrm{g}}}$ we let $[e] = \{1, \ldots, |\overline{a}|\}$. Edges of $\mathfrak{I}$ are triples $\rho = (d, [\rho], e)$, where $d = [\overline{a}]_{\sim_{\mathrm{g}}}$ and $e = [\overline{b}]_{\sim_{\mathrm{g}}}$ and the label $[\rho]$ is a non-empty partial injection $[d] \to [e]$ such that $a_i = b_j$ for all $(i, j) \in [\rho]$. (In which case ρ induces a $\sim_{\mathrm{g}}$-preserving partial isomorphism $\overline{a}' \restriction \mathrm{dom}[\rho] \to \overline{b}' \restriction \mathrm{img}[\rho]$ for any $\overline{a}' \sim_{\mathrm{g}} \overline{a}$ and any $\overline{b}' \sim_{\mathrm{g}} \overline{b}$.) Symbols ρ, σ, etc. will refer to edges of $\mathfrak{I}$, their respective labels will be denoted $[\rho], [\sigma]$, etc.

We adapt the same notation in the case of a hypergraph and its hypergraph bisimulation invariant. Recall that the latter is the bisimulation quotient of the hypergraph bisimulation game graph as described earlier. In the following we often blur the distinction in phrasing and notation between the cases of hypergraphs and of relational structures, opting to treat these perfectly analogous cases as one.

We associate to the invariant $\mathfrak{I}$ a set of constant and function symbols as follows.

- To every vertex e, every $i \in [e]$ and $0 \leq j < w^{m+2}$ we associate a constant symbol $c^j_{e,i}$.
- To every edge $\rho = (d, [\rho], e)$ and every $i \in [e] \setminus \mathrm{img}[\rho]$ and $0 \leq j < w^{m+2}$ we associate a function symbol $f^j_{\rho,i}$ of arity $|\mathrm{dom}[\rho]|$.

We work with well-formed terms in the above signature. As shorthand we write $\mathbf{c}^j_e$ for $(c^j_{e,i})_{1 \leq i \leq k}$, and $\mathbf{f}^j_\rho(\bar{t})$ for $(f^j_{\rho,i}(\bar{t}))_{i \notin \mathrm{img}[\rho]}$ and for every tuple $\bar{t} = (t_1, \ldots, t_l)$ we let $\{\bar{t}\}$ stand for $\{t_1, \ldots, t_l\}$. For each term t let $J(t)$ denote the set of "j-values" occurring in the superscript of a function symbol at any depth within t. This notion extends naturally to tuples of terms. Thus $J(c^j_{e,i}) = \{j\}$ and $J(f^j_{\rho,i}(\bar{t})) = \{j\} \cup J(\bar{t})$. The *truncation* of a term t at depth κ, denoted $t/_\kappa$, is defined by the following recursive rules and is extended to tuples of terms and to sets of terms in the obvious way.

$$c^j_{e,i}/_\kappa = c^j_{e,i} \qquad \begin{array}{lcll} f^j_{\rho,i}(\bar{t})/_0 & = & c^j_{e,i} & (\rho = (d, [\rho], e)) \\ f^j_{\rho,i}(\bar{t})/_{\kappa+1} & = & f^j_{\rho,i}(\bar{t}/_\kappa) & \end{array} \tag{3.1}$$

The N-th Rosati cover $\mathfrak{R}_N$ is made up of terms of height at most N and is built to realise all guarded bisimulation types in $\mathfrak{I}$. To that end we first define the sets $\mathcal{K}^r_N(e)$ of "instances of e at height r" for each $e \in \mathfrak{I}$ and $r \geq 0$ by simultaneous recursion.

$$\begin{array}{rcll} \mathcal{K}^0_N(e) & = & \{\, \mathbf{c}^j_e \mid j < w^{m+2} \} & \\ \mathcal{K}^{r+1}_N(e) & = & \{\, \boldsymbol{\rho}^j(\overline{s} \restriction \mathrm{dom}[\rho]) \;\mid & \overline{s} \in \mathcal{K}^r_N(d),\, \rho = (d, [\rho], e), \\ & & & j < w^{m+2},\, j \notin J(\overline{s} \restriction \mathrm{dom}[\rho]) \,\} \end{array} \tag{3.2}$$

where for each edge $\rho = (d, [\rho], e)$ in $\mathfrak{I}$ and terms $\overline{s}$ and j as appropriate $\boldsymbol{\rho}^j(\overline{s} \restriction \mathrm{dom}[\rho])$ denotes the tuple $(u_1, \ldots, u_{|[e]|})$ such that

$$u_i = \begin{cases} s_l & (\,(l,i) \in [\rho]\,) \\ f^j_{\rho,i}(\overline{s}_{/N-1}) & (\,i \notin \mathrm{img}[\rho]\,) \end{cases} \quad \text{for each } i \in [e]\ .$$

Obviously $\boldsymbol{\rho}^j(\overline{s} \restriction \mathrm{dom}[\rho])$ depends solely on $\overline{t} = \overline{s} \restriction \mathrm{dom}[\rho]$, wherefore more often than not we shall simply write $\boldsymbol{\rho}^j(\overline{t})$ so that, in particular, $\{\boldsymbol{\rho}^j(\overline{t})\} = \{\mathbf{f}^j_\rho(\overline{t}_{/N-1})\} \cup \{\overline{t}\}$.

The sets $\mathcal{H}^r_N(e)$ of *hyperedges above* e *(at height* r*)* are obtained from the above by simply forgetting the tuple ordering: $\mathcal{H}^r_N(e) = \{\,\{\overline{t}\} \mid \overline{t} \in \mathcal{K}^r_N(e)\,\}$; further set $\mathcal{H}_N(e) = \bigcup_r \mathcal{H}^r_N(e)$ and $\mathcal{H}_N = \bigcup_{e \in \mathfrak{I}} \mathcal{H}_N(e)$. All terms t appearing in some hyperedge in $\mathcal{H}_N$ have height at most N and, due to the stipulation $j \notin J(\overline{t})$ in (3.2), no function symbol at the root of a subterm of t occurs again within that subterm.

Observe that every $h \in \mathcal{H}_N$ is either of the form $\{\boldsymbol{\rho}^j(\overline{t})\} = \{\mathbf{f}^j_\rho(\overline{t}_{/N-1})\} \cup \{\overline{t}\} \in \mathcal{H}^{r+1}_N(e)$ for some r and e the target of ρ or is equal to some $\{\mathbf{c}^j_e\} \in \mathcal{H}^0_N(e)$. Crucially, under the assumption $N \geq 2$ the constraint $j \notin J(\overline{t})$ of (3.2) ensures that the former partitioning of h is unique and we say that $h \in \mathcal{H}^{r+1}_N(e)$ is obtained by ρ*-extension* of some (not necessarily unique) hyperedge $h' \in \mathcal{H}^r_N(d)$, with $\rho = (d, [\rho], e)$, and denote this using the shorthand $h' \xrightarrow{\rho} h$. Note, in particular, that the sets $\mathcal{H}(e)$ partition $\mathcal{H}$. Henceforth we often omit the subscript N writing $\mathcal{H}$, $\mathcal{H}(e)$, etc.

A hyperedge h will be called a *primary guard of* X if it is a guard of X, viz. $X \subseteq h$, and is not the ρ-extension of some h' also guarding X.

Lemma 15. *Assume* $m \geq N \geq 2$*. Then for every guarded set* X *of terms there is an* $e_X \in \mathfrak{I}$ *such that all primary guards of* X *belong to* $\mathcal{H}(e_X)$*.*

Proof. Consider a hyperedge h that is a primary guard of X. If $h = \{\mathbf{c}^j_e\}$ for appropriate e and j, then h is the only primary guard of X, and we can set $e_X = e$. Otherwise we have $h = \{\boldsymbol{\rho}^j(\overline{t})\} = \{\mathbf{f}^j_\rho(\overline{t}_{/N-1})\} \cup \{\overline{t}\} \in \mathcal{H}^{r+1}(e)$ for an appropriate edge $\rho : d \to e$ in $\mathfrak{I}$, some superscript j, and terms $\overline{t}$. Because h is by choice a primary guard of X, it cannot be that $X \subseteq \{\overline{t}\}$. So there is some $f^j_{\rho,i}(\overline{t}/_{N-1})$ in X, and we set $e_X = e$ to be the target of ρ.

Suppose indirectly that our choice of e_X was not unique, i.e. that there is some $e' \neq e$ and a primary guard of X of the form $h' = \{\mathbf{f}^{j'}_\sigma(\overline{s}_{/N-1})\} \cup \{\overline{s}\} \in \mathcal{H}^{r+1}(e')$. Then, by the previous argument, some $f^{j'}_{\sigma,i'}(\overline{s}_{/N-1})$ would have to be in X. This, however, would imply that $f^j_{\rho,i}(\overline{t}/_{N-2})$ had to be among $\overline{s}/_{N-1}$ and vice versa $f^{j'}_{\sigma,i'}(\overline{s}_{/N-2})$ among $\overline{t}/_{N-1}$. Given that $N \geq 2$ this would contradict the requirement that j and j' each have but one occurrence in these terms. $\square$

Let $\overline{\mathcal{H}}_N$ be comprised of the hyperedges in $\mathcal{H}_N$ together with sub-hyperedges $h' \subseteq h$ for each $h \in \mathcal{H}(e)$ precisely as specified by the type τ_e labeling $e \in \mathfrak{I}$. It follows from the above that whether some such h' is included in $\overline{\mathcal{H}}_N$ does not depend on the choice of h. Indeed, by Lemma 15, we may assume that h is a primary guard of h' since for every $\rho = (d, [\rho], e)$ the types $\tau_d|_{\mathrm{dom}[\rho]}$ and $\tau_e|_{\mathrm{img}[\rho]}$ are identical.

Definition 2 (Rosati cover)**.**
We define $\mathfrak{R}^m_N$ *as having universe* $\bigcup \mathcal{H}_N$ *and hyperedges* $\overline{\mathcal{H}}_N$*, and set* $\mathfrak{R}_N = \mathfrak{R}^N_N$*.*

Using similar reasoning as in Lemma 15 one can verify that $\mathfrak{I}$ is indeed the guarded bisimulation invariant of $\mathfrak{R}^m_N$, i.e., that $\mathfrak{R}^m_N \sim_{\mathrm{g}} \mathfrak{A}$ for any $\mathfrak{A}$ with $\mathbb{I}(\mathfrak{A}) = \mathfrak{I}$.

Lemma 16. *For all $m \geq N \geq 2$ it holds that $\mathbb{I}(\mathfrak{R}^m_N) = \mathfrak{I}$. In particular, for each $e \in \mathfrak{I}$ all hyperedges in $\mathcal{H}_N(e)$ realise the guarded bisimulation type represented by $e \in \mathfrak{I}$.*

Proof. Consider $h_0 \in \mathcal{H}(e_0)$ and $g_0 \in \mathcal{H}(d_0)$ such that $X = h_0 \cap g_0 \neq \emptyset$. As in Lemma 15 we can find primary guards $h_r, g_s \in \mathcal{H}(e_X)$ of X by tracing backward from h_0 and from g_0, respectively, through extension sequences

$$\begin{array}{l} h_0 \xleftarrow{\rho_1} h_1 \xleftarrow{\rho_2} h_2 \cdots \xleftarrow{\rho_r} h_r \in \mathcal{H}(e_X) \quad \text{and} \\ g_0 \xleftarrow{\sigma_1} g_1 \xleftarrow{\sigma_2} g_2 \cdots \xleftarrow{\sigma_s} g_s \in \mathcal{H}(e_X)\ . \end{array}$$

Let $h_i \in \mathcal{H}(e_i)$ for all $0 \leq i < r$ and $g_l \in \mathcal{H}(d_l)$ for all $0 \leq l < s$. Then in $\mathfrak{I}$ we have the following paths.

$$e_0 \xleftarrow{[\rho_1]} e_1 \ \cdots \ e_{r-1} \xleftarrow{[\rho_r]} e_X \xrightarrow{[\sigma_s]} d_{s-1} \cdots \ d_1 \xrightarrow{[\sigma_1]} d_0$$

Given the nature of edges in a (guarded) bisimulation invariant as representing partial isomorphisms they are invertible and compositional in the sense that for each $v \xrightarrow{[\rho]} w$ there is also $w \xrightarrow{[\rho]^{-1}} v$ and then for every $w \xrightarrow{[\sigma]} u$ there is also $v \xrightarrow{[\sigma]\circ[\rho]} u$ as long as $[\sigma] \circ [\rho] \neq \emptyset$. This means that for any non-empty $[\pi] \subseteq [\sigma_1] \circ \cdots \circ [\sigma_s] \circ [\rho_r]^{-1} \circ \cdots \circ [\rho_1]^{-1}$ there is an edge $\pi = (e_0, [\pi], d_0)$ in $\mathfrak{I}$ and now there is one such $[\pi]$ that maps the projection of X in e_0 to the projection of X in d_0.

It follows that all moves made from any $h_0 \in \mathcal{H}_N(e_0)$ to any $g_0 \in \mathcal{H}_N(d_0)$ in the guarded bisimulation game on $\mathfrak{R}^m_N$ have corresponding edges from e_0 to d_0 in $\mathfrak{I}$. The converse of this being enforced by the very definition of $\mathfrak{R}^m_N$, we can establish that the (guarded) bisimulation invariant of $\mathfrak{R}^m_N$ is no other than $\mathfrak{I}$. $\square$

Lemma 17. *$\mathcal{H}^r_N(e)/_k = \mathcal{H}^r_k(e)$ for all $m \geq N > k \geq 2$, all r, and all $e \in \mathfrak{I}$. Truncation of terms at depth k thus acts as a homomorphic projection from $\mathfrak{R}^m_N$ onto $\mathfrak{R}^m_k$ inducing a guarded bisimulation. In other words, we have the following chain of covers.*

$$\mathfrak{R}^N_N \xrightarrow{\sim} \mathfrak{R}^N_{N-1} \xrightarrow{\sim} \cdots \mathfrak{R}^N_3 \xrightarrow{\sim} \mathfrak{R}^N_2$$

Proof. For $\sigma \in \mathrm{Sym}([w^{m+2}])$ a permutation of j-values and t a term let t^σ denote the term obtained by translating all superscripts j in t according to σ.

$$\begin{array}{rcl} (c^j_{e,i})^\sigma & = & c^{\sigma(j)}_{e,i} \\ (f^j_{\rho,i}(\overline{t}))^\sigma & = & f^{\sigma(j)}_{\rho,i}(\overline{t}^\sigma) \end{array}$$

Based on definitions (3.1) and (3.2) it is straightforward to verify by induction on N and on r that $\mathcal{H}^r_N(e)/_{N-1} \subseteq \mathcal{H}^r_{N-1}(e)$ and that $\mathcal{H}^r_N(e)$ is closed under translations $\cdot^\sigma$ for all e.

Using the latter one can in fact show by induction that $\mathcal{H}^r_N(e)/_{N-1} = \mathcal{H}^r_{N-1}(e)$ for all r, $e \in \mathfrak{I}$ and $m \geq N$. This amounts to proving that all hyperedges $h = \{\boldsymbol{\rho}^j(\overline{t}|_{\mathrm{dom}[\rho]})\} \in \mathcal{H}^{r+1}_{N-1}(e)$ obtained by ρ-extension of some $g = \{\overline{t}\} \in \mathcal{H}^r_{N-1}(d)$ can also be obtained as truncations of hyperedges in $\mathcal{H}^{r+1}_N(e)$, assuming, by the induction hypothesis, that $\mathcal{H}^r_{N-1}(d) = \mathcal{H}^r_N(d)/_{N-1}$, i.e., that there is a $\hat{g} = \{\overline{u}\} \in \mathcal{H}^r_N(d)$ such that $g = \hat{g}/_{N-1}$. While $j \notin J(\overline{t}|_{\mathrm{dom}[\rho]})$ in this case, cf. (3.2), it is conceivable that j does occur in $\overline{u}|_{\mathrm{dom}[\rho]}$ at depth N. If so, then take a permutation $\sigma \in \mathrm{Sym}([w^{m+2}])$ that fixes $J(\overline{u}|_{\mathrm{dom}[\rho]}) \setminus \{j\}$ pointwise but does not fix j (that such a permutation exists follows from $N \leq m$), otherwise let $\sigma = \mathrm{id}$. Then, by closure under translations, $\hat{g}^\sigma = \{\overline{u}^\sigma\} \in \mathcal{H}^r_N(d)$ and by the choice of σ we have $\overline{u}^\sigma/_{N-1} = \overline{u}/_{N-1}$ and $j \notin J(\overline{u}^\sigma|_{\mathrm{dom}[\rho]})$. Consequently $\hat{h} = \{\boldsymbol{\rho}^j(\overline{u}^\sigma|_{\mathrm{dom}[\rho]})\}$ is a hyperedge in $\mathcal{H}^{r+1}_N(e)$ and $\hat{h}/_{N-1} = \{\mathbf{f}^j_\rho(\overline{u}^\sigma|_{\mathrm{dom}[\rho]}/_{N-2})\} \cup \{\overline{u}^\sigma|_{\mathrm{dom}[\rho]}/_{N-1}\} = h$ as needed.

It follows that $\mathcal{H}^r_N(e)/_\kappa = \mathcal{H}^r_\kappa(e)$ for all $m \geq N > \kappa$. Truncation of terms at depth $N-1$ is therefore a homomorphism from $\mathfrak{R}^m_N$ to $\mathfrak{R}^m_{N-1}$ that is onto. By Lemma 16 it also induces a (guarded) bisimulation $\mathfrak{R}^m_N \xrightarrow{\sim} \mathfrak{R}^m_{N-1}$ yielding a chain of covers as claimed. □

3.2. **Size of the Rosati cover.** The size of $\mathfrak{R}^m_N$ can be bounded as follows. Let w be the width of $\mathfrak{I}$, assume that $w \geq 2$ and let $J = w^{m+2}$. Then there are $J|\mathfrak{A}|^{\mathcal{O}(w)}$ many constants $c^j_{e,i}$ and function symbols $f^j_{\rho,i}$ altogether, and each term of height up to N contains at most w^{N+1} many such symbols. For $m = N$, therefore, the total number of terms in $\mathfrak{R}_N$ is at most $(J|\mathfrak{I}|^{\mathcal{O}(w)})^{w^{N+1}} = |\mathfrak{I}|^{w^{\mathcal{O}(N)}}$ as stated in Theorem 14.

3.3. **Auxiliary notions.** Consider a hyperedge $h = \{\mathbf{f}^j_\rho(\overline{t}_{/N-1})\} \cup \{\overline{t}\} \in \mathcal{H}^{r+1}_N(e)$. The elements of $\{\mathbf{f}^j_\rho(\overline{t}_{/N-1})\}$ will be referred to as *siblings*; we denote the sibling relation as $f^j_{\rho,i}(\overline{t}_{/N-1}) \equiv f^j_{\rho,l}(\overline{t}_{/N-1})$. We also say that these terms are *introduced in the hyperedge* h and that h *is a* ρ*-extension*. Furthermore, elements of $\{\overline{t}\}$ are said to be *predecessors* of those in $\{\mathbf{f}^j_\rho(\overline{t}_{/N-1})\}$, and we denote this by writing $t_l \prec f^j_{e,i}(\overline{t}_{/N-1})$, for l and i as appropriate. Constants covered by a hyperedge $\{\mathbf{c}^j_e\} \in \mathcal{H}^0_N(e)$ are also regarded as siblings introduced in that hyperedge. Compare Lemmas 4–9 of [42] for some of the following properties.

Lemma 18. *Let $m \geq N \geq 2$ as before.*

(i) The relations $\equiv$, $\prec$, and its inverse $\succ$ partition the set of all guarded pairs of $\mathfrak{R}^m_N$.
(ii) $\equiv$ is an equivalence relation having guarded equivalence classes.
(iii) Whenever $t^0 \prec t^1 \equiv t^2$ then $\{t^0, t^1, t^2\}$ is guarded and $t^0 \prec t^2$.
(iv) $\mathfrak{R}^m_N$ has no directed $\prec$-cycles of length $\leq N$.
(v) If h is a primary guard of X then some $\prec$-maximal element of h must be in X.
(vi) Assuming $m \geq N \geq 3$, the relation $\prec$ is transitive on every guarded set of terms.
(vii) If $m \geq N \geq 3$ and $h \in \mathfrak{R}^m_N$ is a (primary) guard of $X \subseteq \mathfrak{R}^m_N$ then $h/_{N-1}$ is a (primary) guard of $X/_{N-1} \subseteq \mathfrak{R}^m_{N-1}$. In particular, $e_X = e_{X/_{N-1}}$ for every guarded set $X \subseteq \mathfrak{R}^m_N$.

Proof. As to item (i), observe that the sibling and predecessor relationships are reflected in the terms themselves. Siblings are identical terms for all but the indices in the subscript of their respective root symbol, and the $(N-1)$-truncation of each predecessor of a term occurs in it as an immediate subterm of the root symbol. Given that $N \geq 2$ and that the j superscripts are by definition unique within each term, it is impossible for some $t \in \mathfrak{R}^m_N$ to have $t/_{N-2}$ as a subterm at depth 2, and hence it is impossible to have some $t \prec t' \prec t$.

Item (ii) can be equivalently stated in a form similar to that of (iii), asserting that whenever $t^0 \equiv t^1 \equiv t^2$ then $\{t^0, t^1, t^2\}$ is guarded and also $t^0 \equiv t^2$ holds. Let us first verify that $\{t^0, t^1, t^2\}$ is guarded in both these cases. For item (ii) this is obviously the case if t^0, t^1, t^2 are sibling constants belonging to some $\mathbf{c}^j_e$. Otherwise let $h = \{\mathbf{f}^j_\rho(\overline{u}_{/N-1})\} \cup \{\overline{u}\} \in \mathcal{H}^{r+1}(e_{\{t^0,t^1\}})$ be any primary guard of the pair $\{t^0, t^1\}$. Then, whether $t^0 \equiv t^1$ (ii) or $t^0 \prec t^1$ (iii) the term t^1 must have been introduced in the hyperedge h and must therefore take the form $t^1 = f^j_{\rho,i}(\overline{u}/_{N-1})$. Being a sibling of t^1, $t^2 = f^j_{\rho,l}(\overline{u}/_{N-1})$ and as such is contained in h, which therefore guards $\{t^0, t^1, t^2\}$. Now it is obvious from the definition of the sibling and predecessor relations that $t^0 \equiv t^2$ or $t^0 \prec t^2$, according to whether $t^0 \equiv t^1$ or $t^0 \prec t^1$.

Item (iv) is a trivial consequence of the requirement that superscripts j must not occur twice in any term in $\mathfrak{R}^m_N$. Indeed, if $t^{k-1} \prec \ldots \prec t^2 \prec t^1 \prec t^0$ is a predecessor chain of length $k \leq N$ then $t^r/_{N-r}$ is, for each $r < k$, a subterm of t^0 at depth r. This implies that $t^0, t^1, \ldots, t^{k-1}$ are pairwise distinct.

Property (v) is a straightforward consequence of the definitions. Consider h a primary guard of X. Either $h = \{\mathbf{c}^j_e\}$ and each $c^j_{e,i}$ is $\prec$-maximal within h, or $h = \{\mathbf{f}^j_\rho(\bar{t}/_{N-1})\} \cup \{\bar{t}\}$ introduces some $f^j_{\rho,i}(\bar{t}/_{N-1}) \in X$, which is then $\prec$-maximal within h.

Assuming $m \geq N \geq 3$, property (vi) follows from the prior ones. Indeed, let $t^0 \prec t^1 \prec t^2$ such that $\{t^0, t^1, t^2\}$ is guarded. Then, according to (i) either $t^2 \equiv t^0$ or $t^2 \prec t^0$ or $t^0 \prec t^2$. In the first case we have $t^1 \prec t^0$ by (iii) and thus a two-cycle $t^0 \prec t^1 \prec t^0$, in the second case we have a three-cycle $t^0 \prec t^1 \prec t^2 \prec t^0$, both contradicting (iv). Therefore $t^0 \prec t^2$.

Finally, towards (vii) consider a hyperedge h that is a guard (i.e. superset) of $X \subset \mathfrak{R}^m_N$. By Lemma 17, also $h/_{N-1}$ is a hyperedge in $\mathfrak{R}^m_{N-1}$, and it guards $X/_{N-1}$. If $h/_{N-1}$ is *not* a primary guard of $X/_{N-1}$ then $h/_{N-1}$ is of the form $\{\mathbf{f}^j_\rho(\bar{t}/_{N-2})\} \cup \{\bar{t}\} \in \mathcal{H}^{r+1}_{N-1}(e)$ and $X/_{N-1} \subseteq \{\bar{t}\}$. By Lemma 17 again, $h = \{\mathbf{f}^j_\rho(\bar{u}/_{N-1})\} \cup \{\bar{u}\} \in \mathcal{H}^{r+1}_N(e)$ for some terms $\bar{u}$ such that $\bar{u}/_{N-1} = \bar{t}$. Suppose now that h is a primary guard of X and thus there is some term $f^j_{\rho,l}(\bar{u}/_{N-1})$ belonging to X. Then $f^j_{\rho,l}(\bar{u}/_{N-1})/_{N-1} = f^j_{\rho,l}(\overline{u/}_{N-2}) = f^j_{\rho,l}(\bar{t}/_{N-2})$ belongs to $X/_{N-1}$. Given that $N-1 \geq 2$, this contradicts the assumption $X/_{N-1} \subseteq \{\bar{t}\}$, i.e. that $h/_{N-1}$ is not a primary guard of $X/_{N-1}$. □

3.4. ***N*-conformality of $\mathfrak{R}_N$.** Consider $3 \leq l \leq N$ and an l-clique $\{t^0, \ldots, t^{l-1}\}$ in $\mathfrak{R}_N$, i.e., such that all pairs $\{t^i, t^j\}$ are guarded. By Lemma 18 there are no predecessor-cycles in $\{t^0, \ldots, t^{l-1}\}$ but there is a term, wlog. t^0, such that every one of $t^1, \ldots, t^{l-1}$ is either a predecessor or a sibling of t^0.

Observe that the projection of any primary guard of t^0 to $\mathfrak{R}^m_{N-1}$ guards $\{t^0_{/N-1}, \ldots, t^{l-1}_{/N-1}\}$. This would already be sufficient to establish a weaker form of Theorem 14 still yielding Theorem 2 for GF. However, we can show that the entire l-clique is guarded already in $\mathfrak{R}_N$.

Proposition 19. *Assume that for some $2 \leq l \leq N$ there are $t^0, \ldots, t^{l-1}$ in $\mathfrak{R}_N$ such that all pairs $\{t^i, t^j\}$ are guarded. Then the entire clique $\{t^0, \ldots, t^{l-1}\}$ is guarded in $\mathfrak{R}_N$.*

Proof. From the trivial base case for $l = 2$ we proceed by induction on l. By the preceding observation we may assume wlog. that each of $t^1, \ldots, t^{l-1}$ is either a predecessor or a sibling of t^0. By the induction hypothesis $X = \{t^1, \ldots, t^{l-1}\}$ is guarded.

Consider first the case when $t^0 \equiv t^i$ for some $i \neq 0$. Then t^i is a $\prec$-maximal element of X and as such is necessarily introduced in any primary guard h of X. But then t^0, being a sibling of t^i, is also introduced in h, which therefore guards the entire clique.

Otherwise we know that $t^i \prec t^0$ for all $0 < i < l$. Also, X being guarded it contains a $\prec$-maximal element, wlog. t^1. Let $h^{(0)}$ be a primary guard of the pair $\{t^0, t^1\}$. Given that $t^1 \prec t^0$, then t^0 is introduced in $h^{(0)}$, cf. property (v). In this case t^0 takes the form $f^{j_0}_{\rho_0,i_0}(\bar{u}/_{N-1})$ for some $\rho_0 : e_1 \to e_0$ and appropriate i_0 and $h^{(0)} = \{\boldsymbol{\rho}^{j_0}_0(\bar{u})\} = \{\mathbf{f}^{j_0}_{\rho_0}(\bar{u}/_{N-1})\} \cup \{\bar{u}\} \in \mathcal{H}(e_0)$ where $\{\bar{u}\} = h^{(1)}|_{\mathrm{dom}\rho_0}$ for some $h^{(1)} \in \mathcal{H}(e_1)$.

Note that each $t^i/_{N-1}$ is a subterm of t^0 at depth one, i.e., is among those in $\overline{u}/_{N-1}$. For each i, let u^i denote the component of $\overline{u}$ such that $u^i/_{N-1} = t^i/_{N-1}$ and let $Y = \{u^1, \ldots, u^{l-1}\}$. Thus $X/_{N-1} = Y/_{N-1}$.

Crucially $t^1 = u^1$, for it is included in $h^{(1)}$. Also observe that, because t^1 is $\prec$-maximal in X, for each $1 < i < l$ either $t^i \equiv t^1$ or $t^i/_{N-1} = u^i/_{N-1}$ also occurs as a subterm of t^1.

Tracing backward from $h^{(0)}$ and $h^{(1)}$ as above we can find a chain of expansions

$$h^{(0)} \xleftarrow{\rho_0^{j_0}} h^{(1)} \xleftarrow{\rho_1^{j_1}} \cdots \xleftarrow{\rho_r^{j_r}} h^{(r+1)} \tag{3.3}$$

with $h^{(\lambda)}$ a guard of Y for each $\lambda \leq r$ until we reach some $h^{(r+1)} \in \mathcal{H}(e_Y)$ a primary guard of Y. Let $\{\overline{v}\} = h^{(r+1)}$ and consider some $\{\overline{w}\} = g^{(r+1)} \in \mathcal{H}(e_X)$, a primary guard of X.

Given that $N \geq l \geq 3$ and $X/_{N-1} = Y/_{N-1}$, according to Lemma 18 (vii) we have $e_X = e_{X/_{N-1}} = e_{Y/_{N-1}} = e_Y$, thereby $g^{(r+1)}, h^{(r+1)} \in \mathcal{H}_N(e_X)$. From this it follows that the extension sequence $\langle \rho_r, \ldots, \rho_1, \rho_0 \rangle$ is applicable to $g^{(r+1)}$. However, our aim is to mimic the exact same derivation sequence with the same j_λ-values as in (3.3) starting from $g^{(r+1)}$.

Notice that t^1 being $\prec$-maximal among X, it is also $\prec$-maximal in $g^{(r+1)}$, in accordance with Lemma 18(v). Therefore, every $w_k \in g^{(r+1)}$ is either a sibling or a predecessor of t^1. Similarly, every $v_k \in h^{(r+1)}$ is a sibling or a predecessor, respectively, of u^1. In other words, the exact relationships within $g^{(r+1)}$ are mirrored in $h^{(r+1)}$. Given that $t^1 = u^1$ this implies $h^{(r+1)}/_{N-1} = g^{(r+1)}/_{N-1}$.

Having ascertained $h^{(r+1)}/_{N-1} = g^{(r+1)}/_{N-1}$, it now follows that the same extension sequence $\langle \rho_r^{j_r}, \ldots, \rho_1^{j_1}, \rho_0^{j_0} \rangle$ as in (3.3) is applicable to $g^{(r+1)}$ – with the very same j_λ-values – producing an analogous derivation to that of $h^{(0)}$ from $h^{(r+1)}$:

$$g^{(0)} \xleftarrow{\rho_0^{j_0}} g^{(1)} \xleftarrow{\rho_1^{j_1}} \cdots \xleftarrow{\rho_r^{j_r}} g^{(r+1)}$$

ending in some $g^{(0)} \in \mathcal{H}(e_0)$. Note that, because each $h^{(\lambda)}$ is a guard of Y, also each $g^{(\lambda)}$ is a guard of X. Moreover, a simple induction shows that $h^{(\lambda)}/_{N-1} = g^{(\lambda)}/_{N-1}$ for all $\lambda \leq r+1$. In particular, $g^{(0)} = \{\boldsymbol{\rho}_0^{j_0}(\overline{v})\} = \{\mathbf{f}_{\rho_o}^{j_0}(\overline{v}|_{N-1})\} \cup \{\overline{v}\}$ where $\{\overline{v}\} = g^{(1)}|_{\mathrm{dom}\rho_0}$ and $\overline{v}/_{N-1} = \overline{u}|_{N-1}$. As such, $g^{(0)}$ introduces

$$f_{\rho_0,i_0}^{j_0}(\overline{v}/_{N-1}) = f_{\rho_0,i_0}^{j_0}(\overline{u}/_{N-1}) = t^0$$

and thus guards the entire clique. □

3.5. Weak N-chordality of $\mathfrak{R}_N$ over $\mathfrak{A}$.

Proposition 20. *Consider an l-cycle $C = \{\{t^0, t^1\}, \{t^1, t^2\}, \ldots, \{t^i, t^{i+1}\}, \ldots, \{t^{l-1}, t^0\}\}$ in the Gaifman graph of $\mathfrak{R}_N^m$, with $3 \leq l \leq N \leq m$. Then the projection $C/_{N-l+3}$ of C into $\mathfrak{R}_{N-l+3}^m$ admits a guarded triangulation (in particular a chordal decomposition) in $\mathfrak{R}_{N-l+3}^m$.*

Proof. By Proposition 19 all 3-cycles are guarded in $\mathfrak{R}_N^m$, proving the case of $N = 3$. We proceed by induction on N.

Given an l-cycle in $\mathfrak{R}_N^m$ as above, by Lemma 18 (i) we know that for every i either $t^i \equiv t^{i+1}$ or $t^i \prec t^{i+1}$ or $t^{i+1} \prec t^i$. According to Lemma 18 (iv) there are no predecessor cycles of length $\leq N$ in $\mathfrak{R}_N^m$, hence it cannot be the case that $t^i \prec t^{i+1}$ for all i, nor that $t^{i+1} \prec t^i$ for all i.

Then for some i one of the following cases must hold:

- $t^{i-1} \equiv t^i \equiv t^{i+1}$: then, by Lemma 18 (ii), $\{t^{i-1}, t^i, t^{i+1}\}$ is guarded and $t^{i-1} \equiv t^{i+1}$;

- $t^{i-1} \prec t^i \equiv t^{i+1}$: then, by Lemma 18 (iii), $\{t^{i-1}, t^i, t^{i+1}\}$ is guarded and $t^{i-1} \prec t^{i+1}$;
- $t^{i+1} \prec t^i \equiv t^{i-1}$: then, similarly, $\{t^{i-1}, t^i, t^{i+1}\}$ is guarded and $t^{i+1} \prec t^{i-1}$;
- $t^{i-1} \prec t^i \succ t^{i+1}$: then both $t^{i-1}_{/N-1}$ and $t^{i+1}_{/N-1}$ are maximal proper subterms of t^i; therefore, the projection $h/_{N-1}$ of any hyperedge h of $\mathfrak{R}^m_N$ in which t^i was introduced guards $\{t^{i-1}, t^i, t^{i+1}\}/_{N-1}$ in $\mathfrak{R}^m_{N-1}$.

In each case we have found some i such that $\{t^0, \dots, t^{i-1}, t^{i+1}, \dots, t^{l-1}\}_{/N-1}$ constitutes in $\mathfrak{R}^m_{N-1}$ a cycle of length $l-1$ and $\{t^{i-1}, t^i, t^{i+1}\}/_{N-1}$ a guarded triangle. The claim follows by the induction hypothesis. □

4. Finite controllability and small models

Relying on Theorem 4 one can show that UCQ answering against GF and even against CGF sentences is finitely controllable. The sharper Theorem 14 also yields optimal upper bounds on the minimal size of finite models for each of these fragments as expressed in Theorems 2 & 7 below. Matching lower bounds are implicit in [16].

Theorem 2. *For every $\varphi \in \mathsf{GF}$ and every $q \in \mathsf{UCQ}$:*

$$\varphi \models q \iff \varphi \models_{\text{fin}} q.$$

More specifically, if $\varphi \wedge \neg\, q$ is satisfiable then it has a finite model of size $2^{(|\varphi|+|\tau|^{\mathcal{O}(h)})(hw)^{\mathcal{O}(hw)}}$, where h is the height of q, τ is the signature of φ, and w the width of τ.

Proof. Recall the properties of χ_q from Lemma 13. We establish the claim by proving the following equivalences.

$$\varphi \models q \text{ iff } \varphi \models \chi_q \text{ iff } \varphi \models_{\text{fin}} \chi_q \text{ iff } \varphi \models_{\text{fin}} q$$

The first equivalence was proved in Lemma 13 (ii) and the second equivalence follows from the finite model property of the guarded fragment. Also $\varphi \models_{\text{fin}} \chi_q \Rightarrow \varphi \models_{\text{fin}} q$ is a trivial consequence of $\chi_q \models q$. It remains to be seen that $\varphi \not\models_{\text{fin}} \chi_q$ implies $\varphi \not\models_{\text{fin}} q$. Note that $\varphi \not\models_{\text{fin}} \chi_q$ is the same as $\varphi \not\models \chi_q$ thanks to the finite model property of GF.

So assume that $\varphi \wedge \neg\chi_q$ is satisfiable. Then, by Proposition 11, there is some invariant $\mathfrak{I}$ satisfying the Scott normal form ψ of φ as in Lemma 10. Let h be the height of q, viz. the maximal size of its consituent CQ. Applying Theorem 14 on input $\mathfrak{I}$ with $N = h$ we obtain finite models $\mathfrak{R}^N_2$ and $\mathfrak{R}^N_N$ of $\varphi \wedge \neg\chi_q$, with $\mathfrak{R}^N_N$ a weakly N-acyclic cover of $\mathfrak{R}^N_2$. From Fact 1 it then follows that $\mathfrak{R}^N_N \models \varphi \wedge \neg q$. This concludes the proof of finite controllability.

According to Theorem 14, $|\mathfrak{R}_h| = |\mathfrak{I}|^{w^{\mathcal{O}(h)}}$, where w is the width of the signature τ. From Proposition 11 it follows that $|\mathfrak{I}|$ is bounded by the number of atomic types in the signature of ψ, which is of the order $2^{\mathcal{O}((|\tau|+|\varphi|+|\chi_q|)w^w)}$. Finally, Lemma 13 (iii) gives $|\chi^\tau_q| = |\tau|^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}$. Putting it all together we obtain the estimate $|\mathfrak{R}_h| = 2^{(|\tau|+|\varphi|+|\chi_q|)w^{\mathcal{O}(w+h)}} = 2^{|\varphi|w^{\mathcal{O}(w+h)}+|\tau|^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}} = 2^{(|\varphi|+|\tau|^{\mathcal{O}(h)})(hw)^{\mathcal{O}(hw)}}$ as claimed. □

Naturally, both the size and the width of the signature of a GF-formula are bounded by its length. It is thus easy to see how the above statement of Theorem 2 implies that given in the introduction. In particular, we observe the following corollaries.

Corollary 21. *For every k, every satisfiable sentence of the k-variable guarded fragment has finite models of exponential size in the length of the formula.*

Corollary 22. *For a finite set F of τ-structures let $\mathcal{C}_F$ denote the class of those τ-structures not allowing a homomorphic image of any member of F. If a guarded sentence φ has a model in $\mathcal{C}_F$ then it also has one of size $2^{\mathcal{O}(|\varphi|)}$.*

Another corollary is the validity of the uniform interpolation property (2.3) for GF and the positive existential fragment also in the finite model semantics.

Corollary 23. *Consider some $\varphi \in \mathsf{GF}$ and $q \in \mathsf{UCQ}$ in signature τ. Then*

$$\varphi \models_{\text{fin}} q \qquad \Longrightarrow \qquad \varphi \models_{\text{fin}} \chi_q^\tau \quad \textit{and} \quad \chi_q^\tau \models_{\text{fin}} q \tag{4.1}$$

4.1. **Finite controllability for the clique-guarded fragment.** The above results easily carry over to the clique-guarded fragment with essentially the same bounds. In [24, Section 3.3] a reduction of the (finite) satisfiability problem for CGF to the (finite) satisfiability problem for GF is presented. We borrow their idea with some adaptations to keep the blow-up in formula size to a minimum.

Our reduction maps a given clique-guarded sentence $\varphi \in \mathsf{CGF}[\tau]$ to a guarded sentence $\varphi^* \in \mathsf{GF}[\tau, G]$, where G is a fresh relation symbols of arity $w = \max\{\mathrm{width}(\tau), \mathrm{width}(\varphi)\}$. First, we translate φ to φ' by replacing each clique-guarded quantifier occurring in φ according to the pattern[4]

$$\begin{array}{rcl} [\,(\exists\overline{y}.\alpha(\overline{xy}))\,\psi\,]' & = & (\exists\overline{y}.G(\overline{xy}))\,\big(\alpha(\overline{xy}) \wedge \psi'\big) \\ [\,(\forall\overline{y}.\alpha(\overline{xy}))\,\psi\,]' & = & (\forall\overline{y}.G(\overline{xy}))\,\big(\alpha(\overline{xy}) \rightarrow \psi'\big) \end{array} \tag{4.2}$$

otherwise trivially commuting with Boolean connectives. The intended role of $G(\overline{z})$ is to reflect guardedness of $\overline{z}$ in the expanded signature $\tau \cup \{G\}$. Accordingly, φ^* is defined as the conjunction of φ' and

$$\bigwedge_R \bigwedge_{\{\overline{u}\}\subseteq\{\overline{z}\}} (\forall\overline{z}.\, R(\overline{z}))\ G(\overline{u}) \tag{4.3}$$

where R ranges over $\tau \cup \{G\}$ and $\overline{z}$ and $\overline{u}$ are of the appropriate arity such that all variables in $\overline{u}$ also occur in $\overline{z}$. The following properties of this translation are readily verified.

(1) $|\varphi^*| = \mathcal{O}(|\varphi|) + |\tau| w^{\mathcal{O}(w)}$ where w is as above.
(2) Every model of φ can be expanded to a model of φ^* by interpreting G as the universal relation of arity w.
(3) Every *conformal* model of φ^* is also a model of φ: in conformal models every clique-guarded tuple is also guarded and hence, by (4.3), guarded by a G-atom. All conformal models of (4.3) thus satisfy $\forall\overline{z}\,(\,(G(\overline{z}) \wedge \alpha(\overline{z})) \leftrightarrow \alpha(\overline{z})\,)$ for every clique-guard α and, therefore, also $\varphi \leftrightarrow \varphi'$.

These properties enable us to extend the scope of the reduction from mere satisfiability (in the finite) to the more general query entailment problem (in the finite).

[4]where in each individual case $G(\overline{xy})$ is to be understood as referring to the padding of $\overline{xy}$ to a w-tuple, say, by repeated occurrences of the last variable of the tuple $\overline{y}$. This is merely to render $G(\overline{xy})$ a well-formed atom; in the context of (4.3) the actual choice of padding has no import.

Lemma 24. *Let $\varphi \in \mathsf{CGF}$ and $q \in \mathsf{UCQ}$ be arbitrary, and let $\varphi^* \in \mathsf{GF}$ be the translation of φ as explained above. Then*

$$\varphi \models_{(\mathrm{fin})} q \quad \iff \quad \varphi^* \models_{(\mathrm{fin})} q$$

Proof. Assume first that $\varphi^* \models_{(\mathrm{fin})} q$ and let $\mathfrak{A}$ be a (finite) model of φ. Then, by property (2) of the translation, $\mathfrak{A}$ has an expansion $\mathfrak{A}^* \models \varphi^*$. Obviously, $\mathfrak{A}^*$ is finite whenever $\mathfrak{A}$ is finite. So by assumption, $\mathfrak{A}^* \models q$, which trivially implies $\mathfrak{A} \models q$, given that the interpretation of G has no bearing on q.

Assume now that $\varphi \models_{(\mathrm{fin})} q$ and take any (finite) model $\mathfrak{B} \models \varphi^*$. Let $w = \mathrm{width}(\mathfrak{B}) \leq \mathrm{width}(\varphi^*)$, and let $\mathfrak{B}^{(N)}$ be the N-th Rosati cover of $\mathfrak{B}$ as in Theorem 4 for $N = w+1$. Then $\mathfrak{B}^{(N)} \models \varphi^*$. Furthermore, $\mathfrak{B}^{(N)}$ is N-conformal and hence also conformal, since $N > w$. Therefore, by property (3) of the translation, we have $\mathfrak{B}^{(N)} \models \varphi$, and so, by assumption, $\mathfrak{B}^{(N)} \models q$, since $\mathfrak{B}^{(N)}$ is finite whenever $\mathfrak{B}$ is finite. Then, according to Fact 1, $\mathfrak{B} \models \chi_q$, and via Lemma 13 item (ii) we conclude that $\mathfrak{B} \models q$. $\square$

Combining the above with Theorem 2 yields its generalisation to CGF as follows.

Theorem 6. *For every $\varphi \in \mathsf{CGF}$ and every $q \in \mathsf{UCQ}$ we have $\varphi \models q \iff \varphi \models_{\mathrm{fin}} q$. More specifically, if $\varphi \wedge \neg q$ is satisfiable then it has a finite model of size $2^{(|\varphi|+|\tau|^{\mathcal{O}(h)})(hw)^{\mathcal{O}(hw)}}$, where h is the height of q, τ is the signature of φ, and $w = \max\{\mathrm{width}(\varphi), \mathrm{width}(\tau)\}$.*

Proof. Theorem 2 together with Lemma 24 provide the following chain of equivalences

$$\varphi \models q \iff \varphi^* \models q \iff \varphi^* \models_{\mathrm{fin}} q \iff \varphi \models_{\mathrm{fin}} q$$

proving the first assertion. Towards the size bound, if $\varphi \wedge \neg q$ is satisfiable then, by Lemma 24, so is $\varphi^* \wedge \neg q$. Recall that $|\varphi^*| = \mathcal{O}(|\varphi|) + |\tau| w^{\mathcal{O}(w)}$ by property (1) of the translation. According to Theorem 2, there exists a model $\mathfrak{B}$ of $\varphi^* \wedge \neg q$ of size

$$2^{(|\varphi^*|+|\tau|^{\mathcal{O}(h)})(hw)^{\mathcal{O}(hw)}} = 2^{(|\varphi|+|\tau|w^{\mathcal{O}(w)}+|\tau|^{\mathcal{O}(h)})(hw)^{\mathcal{O}(hw)}} = 2^{(|\varphi|+|\tau|^{\mathcal{O}(h)})(hw)^{\mathcal{O}(hw)}}$$

Finally, as in the proof of Lemma 24 we construct the model $\mathfrak{B}^{(N)}$ of $\varphi \wedge \neg q$ by taking the N-th Rosati cover of $\mathfrak{B}$ with $N = w+1$. By Theorem 4, $|\mathfrak{B}^{(N)}| = |\mathfrak{B}|^{w^{\mathcal{O}(w)}}$, which is still of the same order of magnitude $2^{(|\varphi|+|\tau|^{\mathcal{O}(h)})(hw)^{\mathcal{O}(hw)}}$ as $|\mathfrak{B}|$, as claimed. $\square$

Theorem 7 as announced in the introduction is a straightforward corollary of the above.

5. Complexity of query answering

In this paper *query answering* is the problem of deciding $\varphi \models q$ for a given $\varphi \in \mathsf{GF}$ and q a UCQ. By Lemma 13 (ii) this amounts to testing unsatisfiability of the guarded sentence $\varphi \wedge \neg\chi_q$, known to be 2ExpTime-complete and in $\mathrm{DTime}(2^{\mathcal{O}((r+|\varphi|+|\chi_q|)w^w)})$, where r is the size and w the width of τ [16]. With these parameters for τ recall from Lemma 13 that $|\chi_q^\tau| = r^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}$ and that χ_q^τ is computable in time $|q|r^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}$ for any UCQ q of height h. Query answering is thus 2ExpTime-complete, even for a fixed query, and in $\mathrm{DTime}(|q|r^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)} + 2^{(r+|\varphi|)w^w + r^{\mathcal{O}(h)}(hw)^{\mathcal{O}(hw)}})$. Notice that there is a double-exponential dependence only in terms of the height h of queries and the width w of the signature. Under increasing constraints on the variability of signatures we can break down and simplify the time complexity as follows:

- $2^{(|q||\varphi|)^{\mathcal{O}(|q||\varphi|)}}$ under no restrictions on q nor on φ nor on τ;

- $|q|2^{(h|\varphi|)^{\mathcal{O}(h|\varphi|)}}$ without restrictions but highlighting the influence of query height;
- $|q|(h|\varphi|)^{\mathcal{O}(h)} + 2^{(h|\varphi|)^{\mathcal{O}(h)}} \leq |q|2^{(h|\varphi|)^{\mathcal{O}(h)}}$ when the width of τ is bounded (a matching double-exponential lower bound in this case follows from the work of Lutz [28]);
- $|q|h^{O(h)} + 2^{\mathcal{O}(|\varphi|)+h^{\mathcal{O}(h)}} \leq |q|2^{\mathcal{O}(|\varphi|)+h^{\mathcal{O}(h)}}$ for any fixed signature τ (for this case we provide a single-exponential lower bound as stated in Theorem 25 and proved in Proposition 27 below);
- $|q||\varphi|^{O(1)} + 2^{|\varphi|^{\mathcal{O}(1)}} \leq |q|2^{|\varphi|^{\mathcal{O}(1)}}$ for queries of bounded height and over signatures of bounded width;

In [28, 29] Lutz considered the query answering problem against specifications in various description logics, among them a certain $\mathcal{ALCI}$, which can be naturally seen as a fragment of GF. Lutz proved that answering BCQ against $\mathcal{ALCI}$ specifications is 2EXPTIME-complete. Given that description logics are interpreted over relational structures involving unary and binary predicates only, this implies that query answering against GF is 2EXPTIME-complete already for BCQ and on signatures of width two.

A further important particular case is that of *acyclic queries*. Below ACQ are unions of acylcic Boolean conjunctive queries. Observe that for q an ACQ the exponential blow-up in passing from q to χ_q^τ can be avoided by rewriting q as a guarded existential sentence q^* of essentially the same length as q. Query answering for ACQ reduces in polynomial time to GF-satisfiability. Regarding query answering against a fixed $\varphi \in \mathsf{GF}$ we thus find that for ACQ the complexity reduces to EXPTIME. In fact, it can also be shown to be EXPTIME-complete for certain φ, cf. Proposition 27 below.

For a fixed $\psi \in \mathsf{GF}[\tau \cup \sigma]$ *target query answering* is the problem of deciding $D \wedge \psi \models q$ on input q a UCQ and D a τ-structure (given as a conjunction of ground atoms with elements of D as individual constants). The next theorem summarises our observations on query answering and some results on subproblems of target query answering.

Theorem 25.

(1) Deciding $\varphi \models q$, on input $\varphi \in \mathsf{GF}$ and q a UCQ, is 2EXPTIME-complete already for a fixed query q [16]*, or with the width of φ bounded and q a BCQ* [28].

(2) For each $\varphi \in \mathsf{GF}$, deciding $\varphi \models q$ on input q an ACQ is in EXPTIME*; and it is* EXPTIME*-complete for certain φ.*

(3) There is a GF-sentence ψ such that deciding $D \wedge \psi \models Q$, on input Q a BCQ and D a conjunction of atoms of bounded width*, is* PSPACE*-hard.*

(4) For all universal *$\psi \in \mathsf{GF}$, deciding $D \wedge \psi \models q$, on input q a UCQ and D a conjunction of atoms, is in Π_2^P; and for certain* universal *ψ it is Π_2^P-complete already for CQ q.*

(5) For all $\psi \in \mathsf{GF}$ and q a UCQ, deciding $D \wedge \psi \models q$ on input D, is in co-NP *and* co-NP*-complete already for $q = \bot$ and certain* universal *ψ. Hence, satisfiability of $D \wedge \psi$ on input D is in* NP*, and is* NP*-complete for certain* universal *$\psi \in \mathsf{GF}$.*

Observe that item (2) implies that satisfiability for GF can be EXPTIME-complete already for a fixed signature, see Proposition 27 below. This strengthens a result of [16], where EXPTIME-completeness of satisfiability was shown for GF formulas over bounded arity but variable signatures.

Corollary 26. *For some relational signature τ satisfiability for $\mathsf{GF}[\tau]$ is* EXPTIME*-complete.*

Note that item (1) is but a restatement of results of Grädel [16] and Lutz [28], listed here for the sake of completeness. Similarly, the upper bound in item (2) is a consequence of [16] as remarked at the beginning of this section. We prove each of the remaining claims separately in the propositions to follow.

Proposition 27. *There is a* GF*-sentence* φ *such that deciding* $\varphi \models Q$ *for* ACQ Q *is* ExpTime*-hard.*

Proof. It is well known that deterministic exponential time equals alternating polynomial space. We show how to encode the behaviour of polynomial-space alternating Turing machines into the query answering problem over a fixed formula. The formula φ in question will merely provide the means of alternation blindly generating all trees potentially suitable for encoding any strategy of the existential player in any alternating run of any ATM on any input. Then, for any given ATM M and input w we craft an appropriate UCQ $q_{M,w}$ comprising as disjuncts various conjunctive queries, each matching a different source of error in the encoding of the behaviour or acceptance of M on input w. Ultimately the goal is to have $\varphi \models q_{M,w}$ if, and only if, M has no accepting run on w, i.e., if the existential player has no winning strategy in the game corresponding to the computation of M on input w.

By standard arguments we may restrict attention to normalised ATM in which universal and existential states alternate in any run, which have a single universal initial state, disjoint sets of accepting and rejecting states, and every configuration of which has precisely two successor configurations (including accepting and rejecting configurations, which have only accepting or rejecting successor configurations, respectively). Moreover we may assume that on each input of length n the run of the normalised ATM uses precisely $p(n)$ amount of space for a polynomial p.
Let φ be the conjunction of the following guarded formulas, where $\oplus$ denotes exclusive or.

$$\begin{array}{rl}
(\ \exists x.\, R(x)\ ) & B(x) \wedge A(x) \\
(\forall x.\, B(x)\ ) & E(x) \oplus A(x) \\
(\forall x.\, B(x)\ ) & \top(x) \oplus \bot(x) \\
(\forall x.\, B(x)\ ) & T(x)\ \oplus\ \exists y.S(x,y) \\
(\ \forall xy.\, S(x,y)\ ) & B(y)\ \wedge\ (\,E(x) \leftrightarrow E(y)\,) \\
(\ \forall x.\, T(x)\ ) & E(x) \rightarrow \exists y.F(x,y) \\
(\ \forall x.\, T(x)\ ) & A(x) \rightarrow \exists y_1.A_1(x,y_1) \\
(\ \forall x.\, T(x)\ ) & A(x) \rightarrow \exists y_2.A_2(x,y_2) \\
(\ \forall xy.\, F(x,y)\ ) & B(y) \wedge A(y) \\
(\ \forall xy.\, A_1(x,y)\ ) & B(y) \wedge E(y) \\
(\ \forall xy.\, A_2(x,y)\ ) & B(y) \wedge E(y)
\end{array} \tag{5.1}$$

Intuitively speaking, every model of φ (or rather its guarded unravelling) represents a tree whose vertices correspond to instances of the variables x and y in the above formulation. There are three kinds of vertices: plain B-vertices, the root R-vertex and T-vertices. Each vertex represents a bit, hence the letter B, whose value is either $\top$ or $\bot$ as witnessed by the predicates of the same name. There are four kinds of successor edges in every such tree corresponding to a model of φ: S,F,A_1 and A_2-successors. Every vertex is either a T-vertex or it has an S-successor. The intention is that maximal S-successor chains connecting T-vertices encode individual configurations. In this sense a T-vertex is a terminal vertex of the configuration it belongs to and encodes its last bit. The predicates E and A mark whether the state of a given configuration is existential or universal, respectively. All vertices belonging to the same configuration (in between consecutive T-vertices) carry the same E

or A marking, which is passed down along S-edges. A T-vertex terminating a configuration in an existential state has an F-successor vertex beginning a new configuration. A T-vertex belonging to a configuration with a universal state has both an A_1- and an A_2-successor vertex, each starting a successor configurations.

Note that for query answering one can always restrict attention to *minimal models* of φ, which have only one R-vertex and at any vertex have at most one successor of whatever kind required and no other kind of successor vertices, and every vertex of which is reachable from the R-vertex via a sequence of overlapping S-, F-, A_1- and A_2-atoms. Minimal models of φ are well suited to encode strategies of the existential player in any game determined by a normalised ATM and an input word.

Using the framework provided by (minimal) models of φ the power of unions of conjunctive queries suffices to filter out those models that do *not* represent a winning strategy for the existential player in the game defined by a given ATM M on a given input w. To demonstrate this we must first choose an appropriate encoding of Turing machine configurations. As is customary we write $\alpha q\beta$ for the configuration with tape contents $\alpha\beta$ when the machine is in state q and its head is positioned on the first letter of β. Wlog. the tape alphabet is binary, i.e. $\alpha, \beta \in \{0,1\}^*$.

A configuration $\alpha q\beta$ will be encoded as a bit string $\tilde{\alpha}\tilde{q}\tilde{\beta}$, where $\tilde{\alpha}[2i] = \alpha[i]$ and $\tilde{\alpha}[2i-1] = 0$ for all $1 \leq i \leq |\alpha|$, and similarly for β and $\tilde{\beta}$, and where $\tilde{q} = (11)^j(10)^{r-j}$ if q is the j-th of the r many states of M in some fixed enumeration. In brief: the state is encoded in unary interleaved with 1 digits at the point of the head position and around it the tape contents are interleaved with 0 digits to clearly identify the position where the state is encoded.

Let M have r states and use precisely $n = p(|w|)$-space on inputs of size $|w|$. The query $q_{M,w}$ will then consist of a disjunction of acyclic conjunctive queries (each of size $\mathcal{O}(p(|w|))$ exhausting the reasons a model of φ could fail to encode a winning strategy for the existential player in the game of M on w:

- CQ asserting the existence of an S-successor chain connecting T-vertices too short to represent a configuration, or the existence of too long an S-successor chain:

  $$\exists x_0, \ldots, x_m\ F(x_0, x_1) \wedge \bigwedge_{i<m} S(x_i, x_{i+1}) \ \wedge\ T(x_m)$$

  for $m < 2r + 2n$ and similarly with $A_1(x_0, x_1)$ and $A_2(x_0, x_1)$ or $R(x_1)$ in place of $F(x_0, x_1)$, and

  $$\exists x_1, \ldots, x_{2r+2n+1} \bigwedge_{i<2r+2n+1} S(x_i, x_{i+1})$$

- CQ asserting that on odd positions of a maximal S-successor chain the bit values do not constitute a word in $0^*1^r0^*$:

  $$\exists x_1, \ldots, x_{2r+2n} \bigwedge_{i<2r+2n} S(x_i, x_{i+1}) \ \wedge\ \top(x_{2i-1}) \wedge \bot(x_{2j-1}) \wedge \top(x_{2l-1})$$

  for $1 \leq i < j < l \leq r + n$, along with

  $$\exists x_1, \ldots, x_{2r+1} \bigwedge_{i<2r+2n} S(x_i, x_{i+1}) \ \wedge\ \top(x_1) \wedge \top(x_{2r+1})$$

  etc.

- CQ asserting that the configuration beginning with the root vertex is not the initial configuration for the given input word: one CQ for each bit in the encoding of the initial configuration looking for a wrong bit value, such as

$$\exists x_1, \ldots, x_m \; R(x_1) \wedge \bigwedge_{i<m} S(x_i, x_{i+1}) \wedge \; \bot(x_m)$$

  for $m < 2r$ odd or

$$\exists x_1, \ldots, x_m \; R(x_1) \wedge \bigwedge_{i<m} S(x_i, x_{i+1}) \wedge \; \top(x_m)$$

  for $m > 2r$ odd, and (assuming for simplicity that the initial state is the 1st one) for $2 < m \leq 2r$ even; further

$$\exists x_1, \ldots, x_{2r+2k} \; R(x_1) \wedge \bigwedge_{i<2r+2k} S(x_i, x_{i+1}) \wedge \; \bot(x_{2r+2k})$$

  for $w[k] = 1$ and similarly for $w[k] = 0$;
- CQ checking that consecutive S-successor chains do not represent successor configurations, e.g., by asserting that in a given tape position not under the head of the ATM the bit values in the two configurations are not identical: for instance as in

$$\exists x_{-1}, x_0, x_1, \ldots, x_{2r+2n}, x_{2r+2n+1}, x_{2r+2n+2} \bigwedge_{-1 \leq i < 2r+2n, i \neq 2l} S(x_i, x_{i+1}) \wedge$$
$$F(x_{2l}, x_{2l+1}) \wedge \bot(x_{-1}) \wedge \bot(x_1) \wedge \bot(x_3) \wedge \top(x_2) \wedge \bot(x_{2r+2n+2})$$

  and, similarly, with $A_1(x_{2l}, x_{2l+1})$ or $A_2(x_{2l}, x_{2l+1})$ in place of $F(x_{2l}, x_{2l+1})$ and with the bit values at corresponding positions x_2 and $x_{2r+2n+2}$ swapped, and all this for each $l \leq r + n$;
- CQ asserting that the A_i-successor configuration of a universal configuration was not derived by the i-th of the two applicable transitions;
- CQ asserting the existence of a configuration in reject state:

$$\exists x_1, \ldots, x_{2r} \bigwedge_{1 \leq i < r} S(x_{2i-1}, x_{2i}) \wedge S(x_{2i}, x_{2i+1}) \; \wedge \top(x_{2i-1}) \wedge$$
$$\bigwedge_{i \leq q} \top(x_{2i}) \wedge \bigwedge_{q < i \leq r} \bot(x_{2i})$$

  for each rejecting state q.

Much as those examples illustrated above, all of the flaws in the encoding of an accepting run can likewise be expressed using acyclic conjunctive queries, polynomially many in total, and each of size $\mathcal{O}(|M| + p(|w|))$. □

Proposition 28. *There is a* GF*-sentence ψ (using constants) such that deciding $D \wedge \psi \models Q$ on input Q a* BCQ *and D a conjunction of atoms (of bounded width) is* PSPACE*-hard.*

Proof. The proof is by reduction from QBF. Wlog. we may assume that the QBF instances are sentences in prenex normal form

$$\exists X_0 \, \forall X_1 \, \exists X_2 \ldots \forall X_{2m-1} \, \exists X_{2m} \; \vartheta \tag{5.2}$$

with ϑ a 3CNF-formula with free variables among $\{X_0, X_1, \ldots, X_{2m}\}$. To represent "valuation strategies" for the existentially quantified variables we use a variant of the formula (5.1).

On the one hand, the encoding is greatly simplified to single bit "configurations" representing Boolean values. On the other hand, Boolean values need to be encoded as elements of the domain of models to allow for a stronger form of pattern matching. To that end we rely on constants 0 and 1 and encode bit values using a binary predicate $V(x, b)$ where b is either 0 or 1. Let ψ be the conjunction of the following (where, again, $\oplus$ stands for exclusive or).

$$\begin{array}{rl} & R(r,1,0) \\ (\ \forall xtf.\, R(x,t,f)\ ) & B(x,t,f) \wedge A(x) \\ (\forall xtf.\, B(x,t,f)\ ) & E(x) \oplus A(x) \\ (\forall xtf.\, B(x,t,f)\ ) & V(x,t) \oplus V(x,f) \\ (\ \forall xtf.\, B(x,t,f)\ ) & E(x) \rightarrow \exists y.S(x,y,t,f) \\ (\ \forall xtf.\, B(x,t,f)\ ) & A(x) \rightarrow (\ \exists y_1.S(x,y_1,t,f)\ )V(y_1,t) \\ (\ \forall xtf.\, B(x,t,f)\ ) & A(x) \rightarrow (\ \exists y_2.S(x,y_2,t,f)\ )V(y_2,f) \\ (\ \forall xytf.\, S(x,y,t,f)\ ) & B(y,t,f) \wedge E(x) \oplus E(y) \end{array} \tag{5.3}$$

We may think of models of ψ (more precisely their guarded unravelling) as representing infinite Boolean valuation trees with all possible assignments encoded at every odd level and encoding a choice of a Boolean value at each node at an even distance from the root. In minimal models of ψ nodes at even levels (those marked with A) have precisely two successors representing the two Boolean values, whereas all nodes at odd levels (marked with E) have precisely one successor representing an existential choice of a Boolean value. For the encoding of Boolean valuations of the variables of (5.2) only the first $2m$ levels of these trees will play a role.

For each Q3CNF formula (5.2) with matrix ϑ consisting of k clauses we assign an input data structure D_ϑ defined as the conjunction

$$X(v_0, v_1) \wedge X(v_1, v_2) \wedge \ldots \wedge X(v_{2m-1}, v_{2m})$$

together with the conjunction of atoms $C(c_i, v_j, b)$ for all $1 \leq i \leq k$ and $0 \leq j \leq 2m$ and $b \in \{0,1\}$ such that setting X_j to b does *not* satisfy the i-th clause. The elements $v_0, \ldots, v_{2m}$ and $c_1, \ldots, c_k$ of the "database" D_ϑ are perceived as constants, however, unlike 0 and 1, in ψ they are not accessible by (constant) names.

As in the proof of Proposition 27 we design queries Q_ϑ to hold true in a model precisely when it contains a branch corresponding to an assignment of the variables falsifying (one of the clauses in) ϑ. In fact, we use Q_ϑ merely as a general clause checking pattern, hence our reduction will use the same query $Q_\vartheta = Q_m$ for every Q3CNF formula with $2m$ alternately quantified variables:

$$\exists \overline{x}, \overline{y}, \overline{z}, t, f, c\ R(y_0, t, f)\ \wedge \bigwedge_{i<2m} (\ X(x_i, x_{i+1})\ \wedge\ S(y_i, y_{i+1}, t, f)\ \wedge\ V(y_i, z_i)\ \wedge C(c, x_i, z_i)\ )\ .$$

In other words, the query Q_m matches the variables $\overline{x}$ to the corresponding constants $\overline{v}$ of D and the variables $\overline{y}$ to a single branch in any given model of $D \wedge \psi$ and $\overline{z}$ to the sequence of Boolean values assigned to the variables $X_0, \ldots, X_{2m}$ on that branch. Then Q_m is satisfied in a given model (which encodes a particular choice of Skolem functions for the existentially quantified variables $X_0, X_2, \ldots$) iff there is a clause c that is falsified by an assignment to the universally quantified variables corresponding to some path within the model. Therefore, Q_m is true in all models of $D_\vartheta \wedge \psi$ iff (5.2) is false. $\square$

Proposition 29. *The problem of deciding $D \wedge \psi \models Q$ for a fixed universal* GF *sentence ψ on input consisting of a conjunction of atoms D and a* UCQ *Q is in Π_2^P for each ψ, and for some ψ is Π_2^P-complete already for* BCQ *Q.*

Proof. As in the previous propositions we may restrict attention to minimal models of $D \wedge \psi$ as regards query answering. Note that now, due to the universality of ψ, all minimal models of $D \wedge \psi$ have the same universe as D. Hence, to check $D \wedge \psi \models Q$, one merely has to universally choose a model of ψ on the universe of D and existentially guess elements realising Q in the model chosen. This involves first universally choosing then existentially guessing a polynomial number of bits in terms of $|D|$, whence membership in Π_2^P.

We show Π_2^P-hardness by reduction from the validity problem for quantified propositional formulas of the form

$$\forall X_1, \ldots, X_n \, \exists X_{n+1}, \ldots, X_{n+m} \, \vartheta \,,$$

with ϑ a 3CNF formula with free variables $X_1, \ldots, X_{n+m}$.

In the input structure $D = D_n$ the universally quantified variables $X_1, \ldots, X_n$ are encoded as a successor chain

$$S(v_1, v_2) \wedge S(v_2, v_3) \wedge \ldots \wedge S(v_{n-1}, v_n) \;\wedge\; X(v_1, 1, 0) \wedge \ldots \wedge X(v_n, 1, 0) \,,$$

along with the entire table of satisfying assignments of each of the three-literal clauses

$$\bigwedge_{(x,y,z) \in \{0,1\}^3} \; \bigwedge_{(i,j,k) \in \{0,1\}^3 \setminus \{(x,y,z)\}} R_{(x,y,z)}(i, j, k) \;.$$

The formula ψ is devised so that (minimal) models of $D \wedge \psi$ will correspond to all possible assignments of the variables $X_1, \ldots, X_n$, for whatever n. We set

$$\psi \;=\; (\,\forall v.\, X(v, t, f)\,)\; V(v, t) \oplus V(v, f)\,.$$

Finally, we use the query Q_ϑ to guess Boolean values for the existentially quantified variables $X_{n+1}, \ldots, X_{n+m}$ and to test satisfaction of ϑ by checking all triplets of Boolean values for variables occurring together in a clause against the truth table of the type of that clause. Let

$$Q_\vartheta = \exists z_1, \ldots, z_n, x_1, \ldots, x_{n+m} \bigwedge_{1 \le i < n} (S(z_i, z_{i+1}) \wedge V(z_i, x_i)) \;\wedge\; V(z_n, x_n) \;\wedge \bigwedge_{\substack{(-1)^p X_i \vee (-1)^q X_j \vee (-1)^r X_k \\ \text{a clause in } \vartheta}} R_{(p,q,r)}(x_i, x_j, x_k) \;.$$

It is now easy to verify that $D_n \wedge \psi \models Q_\vartheta$ iff $\forall X_1, \ldots, X_n \, \exists X_{n+1}, \ldots, X_{n+m} \, \vartheta$ is valid. □

Proposition 30. *For any fixed $\psi \in$* GF *and input D a conjunction of atoms, satisfiability of $D \wedge \psi$ is in NP and is NP-complete for certain universal ψ. Deciding $D \wedge \psi \models Q$ on input D for any fixed ψ as before and fixed* UCQ *Q is in co-NP; and it is co-NP-complete already for $Q = \bot$ and certain universal ψ.*

Proof. To test satisfiability of $D \wedge \psi$ for a fixed ψ one can pre-compute the Scott normal form Ψ of ψ as in (2.1), as well as the set of admissible guarded atomic types in its signature, i.e. those atomic types that can be realised in some model of Ψ. Having done that, for each given input D it remains to be verified that its atoms can be assigned admissible atomic types wrt. Ψ such that (i) the type assigned to each atom actually contains that atom; (ii) overlapping atoms are assigned consistent types restricted to their overlap; and (iii) the resulting structure satisfies Ψ. Such an assignment can be guessed and verified in NP.

To show NP-hardness of satisfiability of $D \wedge \psi$ for an appropriate ψ observe that 3-colourability of graphs can be directly formalised within this problem. Every simple graph $G = (V, E)$ can be identified with the conjunction $\bigwedge_{(u,v) \in E} E(u, v)$ with vertices $v \in V$ perceived as distinct constants. Then G is 3-colourable iff $G \wedge \psi$ is satisfiable, where

$$\psi = \forall v \bigvee_{i<3} \big(C_i(v) \wedge \bigwedge_{i \neq j < 3} \neg C_j(v) \big) \ \wedge \ (\forall uv.E(u,v)) \bigwedge_i \neg \big(C_i(u) \wedge C_i(v) \big) \ . \tag{5.4}$$

Turning to the problem $D \wedge \psi \models Q$ for fixed ψ and Q and input D, note that this is equivalent to the unsatisfiability of $D \wedge \psi \wedge \neg \chi_Q$, where χ_Q is the "treeification" of Q, that can now be precomputed since Q is fixed. Thus, by the above argument, $D \wedge \psi \models Q$ on input D can be verified in co-NP for any fixed ψ and Q; and it is co-NP-complete for ψ of (5.4) and $Q = \bot$. □

6. Canonisation and capturing

The abstract version of the *capturing Ptime problem* asks for an effective (recursive, syntactic) representation of all polynomial-time computable Boolean queries over finite relational structures. The core problem is not so much the effective representation of the class of all polynomial-time algorithms, which is easy via polynomially clocked Turing machines, say. Rather it lies in the requirement that these machines or algorithms must represent *queries* on finite structures, i.e., they need to respect isomorphism in the sense of producing the same answer on isomorphic structures (or on inputs that represent isomorphic structures). In notation to be used below, we indicate this constraint explicitly in writing PTIME/$\simeq$ for the set of those PTIME algorithms that respect isomorphism.

This undecidable semantic constraint is almost trivially enforcible and therefore mostly goes unnoticed when dealing with queries on *linearly ordered* finite structures. This is because there is an obvious canonisation procedure on the class of all linearly ordered $\tau_<$-structures (we let $< \in \tau_<$ be the distinguished binary relation that is interpreted as a linear ordering in the class $\mathcal{C}^<$ of all linearly ordered finite $\tau_<$-structures).

By *canonisation* (w.r.t. isomorphism over $\mathcal{C}$) we here mean a map can: $\mathcal{C} \to \mathcal{C}$ such that $\mathrm{can}(\mathfrak{A}) \simeq \mathfrak{A}$ and $\mathrm{can}(\mathfrak{A}) = \mathrm{can}(\mathfrak{A}')$ whenever $\mathfrak{A} \simeq \mathfrak{A}'$. Indeed, for $\mathfrak{A} \in \mathcal{C}^<$ we may just identify the linearly ordered universe of $\mathfrak{A}$, $(A, <^{\mathfrak{A}})$, with an initial segment of $(\mathbb{N}, <)$ to obtain such a canonical representative of the isomorphism type of $\mathfrak{A}$. Then the application of arbitrary polynomial time decision procedures to $\mathrm{can}(\mathfrak{A})$ for $\mathfrak{A} \in \mathcal{C}^<$ – i.e., the application of semantically unconstrained algorithms after pre-processing with can – provides an effective representation of the class of all polynomial time computable Boolean queries on $\mathcal{C}$. It is well known from the fundamental results of Immerman [25] and Vardi [43] that this abstract capturing result finds a concrete logical counterpart in the logics LFP (least fixpoint logic) and IFP (inductive fixpoint logic). The open question whether PTIME may also be captured, abstractly or by some suitable logic, over all not necessarily ordered finite structures has driven much of the development of descriptive complexity in finite model theory. Interesting variations of this question concern

(a) restricted classes of finite structures other than $\mathcal{C}^<$; and
(b) rougher equivalence relations than $\simeq$.

We point to the work of Grohe and his survey [21] for successes with larger and larger natural classes of structures in the sense of (a), and to [32] for a very simple but interesting capturing result in the sense of (b) concerning bisimulation-invariant Ptime (Ptime/$\sim$). The class Ptime/$\sim$ consists of those Ptime Boolean queries that respect bisimulation equivalence; it can be regarded as the class of Ptime queries in the modal world. In that case, canonisation is obtained through passage to a definably ordered version of the bisimulation quotient of the given structure, $\mathbb{I}^<(\mathfrak{A}) := (\mathbb{I}(\mathfrak{A}), <)$ where $\mathbb{I}(\mathfrak{A}) = \mathfrak{A}/{\sim}$ such that $\mathbb{I}(\mathfrak{A}) \sim \mathfrak{A}$ is trivially satisfied. For reasons indicated above, the distinction between (canonical) standard representations and definably linearly ordered versions of structures is often blurred, and in fact immaterial for our concerns. Hence we may avoid explicit passage to a standard representation of a linearly ordered structure like $\mathbb{I}^<(\mathfrak{A})$ and seemingly weaken the requirement that $\mathbb{I}^<(\mathfrak{A}) = \mathbb{I}^<(\mathfrak{A}')$ for $\mathfrak{A} \sim \mathfrak{A}'$ to $\mathbb{I}^<(\mathfrak{A}) \simeq \mathbb{I}^<(\mathfrak{A}')$.

A polynomial time canonisation procedure $\mathrm{can}\colon \mathcal{C} \to \mathcal{C}$ w.r.t. some equivalence $\approx$ (on $\mathcal{C}$) will always yield an abstract capturing result for the class of all Ptime computable Boolean queries on finite structures from $\mathcal{C}$ that respect $\approx$, which we denote by Ptime/$\approx$ (over $\mathcal{C}$). Since pre-processing with the canonisation procedure can be performed in Ptime and enforces $\approx$-invariance, it can be coupled with any effective representation of otherwise unconstrained Ptime decision algorithms to capture Ptime/$\approx$:

$$\textsc{Ptime}/{\approx} \;\equiv\; \textsc{Ptime} \circ \mathrm{can},$$

in a notation that suggests how canonisation acts as a filter to guarantee the required semantic invariance. If canonisation produces linearly ordered output structures, which we indicate notationally as in $\mathrm{can}^<(\mathfrak{A}) = (\mathrm{can}(\mathfrak{A}), <)$, then Ptime/$\approx$ is in fact captured by LFP or IFP over the canonisation results by the Immerman–Vardi Theorem:

$$\textsc{Ptime}/{\approx} \;\equiv\; \mathsf{LFP} \circ \mathrm{can}^< \;\equiv\; \mathsf{IFP} \circ \mathrm{can}^<.$$

In the case of Ptime/$\sim$, [32] correspondingly translates the abstract capturing result into capturing by a suitable extension of the modal μ-calculus, which exactly matches the expressive power of LFP over the (internally interpretable) linearly ordered canonisations $\mathrm{can}^<(\mathfrak{A}) := \mathbb{I}^<(\mathfrak{A})$ indicated above.

Here we primarily want to provide an abstract capturing result for $\sim_{\mathrm{g}}$-invariant Ptime, Ptime/$\sim_{\mathrm{g}}$, which corresponds to Ptime in the guarded world. This is achieved through a polynomial canonisation w.r.t. guarded bisimulation equivalence which produces (definably) linearly ordered representatives of the complete guarded bisimulation types of given finite relational structures. This canonisation may be of interest beyond our application to the capturing issue. The proposed canonisation produces linearly ordered output structures $\mathrm{can}^<(\mathfrak{A})$ that are uniformly interpretable over powers of the original structures $\mathfrak{A}$ in IFP and LFP in a $\sim_{\mathrm{g}}$-invariant manner. An adaptation of the approach of [32] therefore also entails a concrete logical capturing result by means of some higher-dimensional guarded fixpoint logic, but we do not pursue this here.

We fix some terminology, similar to the one discussed e.g. in [33], which makes sense for arbitrary equivalences $\approx$ between structures; we are going to use these notions solely with reference to guarded bisimulation equivalence.

Definition 3.

(1) A complete invariant $\mathbb{I}$ *for $\approx$ on $\mathcal{C}$ (with values in some set $\mathcal{D}$) is a map*

$$\begin{array}{rcl} \mathbb{I}\colon \mathcal{C} & \longrightarrow & \mathcal{D} \\ \mathfrak{A} & \longmapsto & \mathbb{I}(\mathfrak{A}) \end{array}$$

such that $\mathbb{I}(\mathfrak{A}) = \mathbb{I}(\mathfrak{A}')$ *for all* $\mathfrak{A} \approx \mathfrak{A}' \in \mathcal{C}$.[5]

(2) An inversion of the invariant $\mathbb{I}$ *is then a map* $F \colon \mathcal{D} \to \mathcal{C}$ *that acts as a right inverse to* $\mathbb{I}$*:* $\mathbb{I} \circ F = \mathrm{id}$*, or, equivalently,* $F(\mathbb{I}(\mathfrak{A})) \approx \mathfrak{A}$ *for all* $\mathfrak{A} \in \mathcal{C}$*.*

(3) Canonisation *w.r.t.* $\approx$ *over* $\mathcal{C}$ *is a map*

$$\begin{array}{rcl} \mathrm{can} \colon \mathcal{C} & \longrightarrow & \mathcal{C} \\ \mathfrak{A} & \longmapsto & \mathrm{can}(\mathfrak{A}) \end{array}$$

such that $\mathrm{can}(\mathfrak{A}) \approx \mathfrak{A}$ *and* $\mathrm{can}(\mathfrak{A}) = \mathrm{can}(\mathfrak{A}')$ *for all* $\mathfrak{A} \approx \mathfrak{A}' \in \mathcal{C}$.[5]

Note that canonisations are complete invariants whose values are structures of the original kind while an invariant in general may produce values of a different format. Clearly an inversion of an invariant always yields a canonisation of the form $\mathrm{can} := F \circ \mathbb{I}$. Note also that (1) says that $\approx$ is induced by equality of $\mathbb{I}$-images.

Theorem (Canonisation) 8. *Guarded bisimulation equivalence on finite relational structures admits* PTIME *canonisation. More specifically, definably linearly ordered versions of the guarded bisimulation game invariants* $\mathbb{I}^<(\mathfrak{A}) := (\mathbb{I}(\mathfrak{A}), <)$ *discussed above are* PTIME *computable complete invariants w.r.t.* $\sim_{\mathrm{g}}$ *and admit* PTIME *inversions* F *such that* $\mathrm{can}^< := F^< \circ \mathbb{I}^<$ *produces linearly ordered representatives from the* $\sim_{\mathrm{g}}$*-class of every finite relational structure* $\mathfrak{A}$*. The values of both maps,* $\mathbb{I}^<(\mathfrak{A})$ *and* $\mathrm{can}^<(\mathfrak{A})$ *are uniformly* IFP*- and* LFP*-interpretable as ordered quotients over* $\mathfrak{A}^w$ *in a* $\sim_{\mathrm{g}}$*-invariant manner, for fixed relational vocabulary* τ *of width* w*.*

Corollary 31 (Capturing)**.** *Guarded-bisimulation-invariant* PTIME *can be captured (admits an effective, syntactic representation) in the form*

$$\begin{array}{rcl} \mathrm{PTIME}/{\sim_{\mathrm{g}}} & \equiv & \mathrm{PTIME} \circ \mathrm{can}, \ \ \textit{or} \\ \mathrm{PTIME}/{\sim_{\mathrm{g}}} & \equiv & \mathsf{LFP} \circ \mathrm{can}^< \ \equiv \ \mathsf{IFP} \circ \mathrm{can}^<. \end{array}$$

We fix a finite relational vocabulary τ of width w. For a finite τ-structure $\mathfrak{A}$, we let $G(\mathfrak{A}) \subseteq A^w$ be the set of all maximal guarded tuples of $\mathfrak{A}$ (with repetitions of components where appropriate, to uniformly pad tuples to arity w). The guarded bisimulation game graph $\mathbb{G}(\mathfrak{A})$ introduced in Section 2 has $G(\mathfrak{A})$ as its set of vertices, unary predicates for atomic types, and edge relations $(E_\rho)_{\rho \in \Sigma}$ for the set Σ of all partial bijections $\rho \subseteq \{1, \dots, w\} \times \{1, \dots, w\}$, where

$$(\bar{a}, \bar{b}) \in E_\rho \quad \text{iff} \quad a_i = b_j \text{ for all } (i,j) \in \rho.$$

The guarded bisimulation invariant $\mathbb{I}(\mathfrak{A}) := \mathbb{G}(\mathfrak{A})/{\sim}$, as discussed in Section 2, is obtained as the quotient of this game graph w.r.t. *modal* bisimulation equivalence.[6] Passage from $\mathfrak{A}$ to this quotient $\mathbb{I}(\mathfrak{A})$ almost provides a complete invariant for $\sim_{\mathrm{g}}$ in the sense of Definition 3, but not quite: clearly $\mathfrak{A} \sim_{\mathrm{g}} \mathfrak{A}'$ implies $\mathbb{G}(\mathfrak{A}) \sim \mathbb{G}(\mathfrak{A}')$ and hence $\mathbb{I}(\mathfrak{A}) = \mathbb{G}(\mathfrak{A})/{\sim} \simeq \mathbb{G}(\mathfrak{A}')/{\sim} = \mathbb{I}(\mathfrak{A}')$, but not $\mathbb{I}(\mathfrak{A}) = \mathbb{I}(\mathfrak{A}')$ as required. As discussed above, this defect is overcome as soon as we provide a definable linear ordering of the universes $G(\mathfrak{A})/{\sim_{\mathrm{g}}}$ of $\mathbb{I}(\mathfrak{A})$ and thus turn them into linearly ordered complete invariants $\mathbb{I}^<(\mathfrak{A}) := (\mathbb{I}(\mathfrak{A}), <)$.

Such a definable linear ordering can be obtained in an inductive refinement process, which produces a sequence of pre-orderings $\preceq^i$ on $G(\mathfrak{A})$. This refinement process starts

[5]As above, we use notation $\mathbb{I}^<(\mathfrak{A}) := (\mathbb{I}(\mathfrak{A}), <)$ or $\mathrm{can}^<(\mathfrak{A}) := (\mathrm{can}(\mathfrak{A}), <)$ to indicate an invariant or canonisation that produces as output some relational structure, which is (definably) linearly ordered. In this context we relax e.g. the condition $\mathbb{I}(\mathfrak{A}) = \mathbb{I}(\mathfrak{A}')$ to $\mathbb{I}^<(\mathfrak{A}) \simeq \mathbb{I}^<(\mathfrak{A}')$ without any loss.

[6]Note that the $\sim$-equivalence classes of vertices in $\mathbb{G}(\mathfrak{A})$ are precisely the $\sim_{\mathrm{g}}$-equivalence classes of the maximal guarded tuples in $G(\mathfrak{A})$.

from an arbitrary but fixed ordering of the finite set of atomic types of w-tuples, which is uniformly imposed as a global pre-order so that

$$(\bar{a} \preceq^0 \bar{b} \text{ and } \bar{a} \preceq^0 \bar{b}) \;\Leftrightarrow\; \mathfrak{A}, \bar{a} \;\sim_{\mathrm{g}}^0\; \mathfrak{A}, \bar{b},$$

i.e., the equivalence relation induced by $\preceq^0$ is atomic equivalence, which is $\sim_{\mathrm{g}}^0$-equivalence, of maximal guarded tuples. In other words, $\prec^0$ uniformly defines a linear ordering of the quotient $G(\mathfrak{A})/\sim_{\mathrm{g}}^0$ (or, equivalently, of $\mathbb{G}(\mathfrak{A})/\sim^0$ in terms of modal bisimulation). The refinement proceeds in such a manner that each level $\preceq^i$ induces a linear ordering of the quotient $G(\mathfrak{A})/\sim_{\mathrm{g}}^i$ (or $\mathbb{G}(\mathfrak{A})/\sim^i$). Then the inductive fixpoint of this refinement sequence produces a pre-ordering $\preceq$ that linearly orders the quotient $G(\mathfrak{A})/\sim_{\mathrm{g}}$ (or $\mathbb{G}(\mathfrak{A})/\sim$) and thus yields the desired $\mathbb{I}^<(\mathfrak{A})$.

Besides $\preceq^0$ we fix an arbitrary ordering on the set Σ of edge labels in the game graphs $\mathbb{G}(\mathfrak{A})$. For $\bar{a} \in G(\mathfrak{A})$ we define Boolean incidence functions $\iota_{\rho,\beta}$ for $\rho \in \Sigma$ and $\beta \in G(\mathfrak{A})/\sim_{\mathrm{g}}^i$ according to

$$\iota_{\rho,\beta}(\bar{a}) := 1 \quad \text{iff} \quad \left\{\bar{b} \in G(\mathfrak{A}) \colon (\bar{a}, \bar{b}) \in \rho\right\} \cap \beta \neq \emptyset.$$

Note that the $\iota_{\rho,\beta}$-value precisely describes the existence or non-existence of a move along a ρ-edge to a position in the $\sim_{\mathrm{g}}^i$-class β. A simple analysis of one round in the guarded bisimulation game shows that

$$\mathfrak{A}, \bar{a} \;\sim_{\mathrm{g}}^{i+1}\; \mathfrak{A}, \bar{a}' \quad \text{iff} \quad \iota_{\rho,\beta}(\bar{a}) = \iota_{\rho,\beta}(\bar{a}') \text{ for all } \rho \in \Sigma \text{ and } \beta \in G(\mathfrak{A})/\sim_{\mathrm{g}}^i.$$

It follows that the pre-ordering $\preceq^{i+1}$ defined by a lexicographic ordering of tuples w.r.t. $\iota_{\rho,\beta}$-values is as desired.

Definition 4. *Let $\mathbb{I}^< \colon \mathfrak{A} \mapsto (\mathbb{I}(\mathfrak{A}), <) := (\mathbb{G}(\mathfrak{A}), \prec)/\sim$ be the linearly ordered version of the quotient of the guarded bisimulation game graph $\mathbb{G}(\mathfrak{A})$ described above. We now refer to this linearly ordered structure as the* ordered guarded-bisimulation invariant of $\mathfrak{A}$.

The following is then immediate from the preceding discussion and the fact that the inductive refinement process outlined above is naturally captured as an inductive fixpoint (in the sense of inductive fixpoint logic IFP), and hence, buy the Gurevich–Shelah Theorem also by a least fixpoint process (in the sense of least fixpoint logic LFP).

Lemma 32. *$\mathbb{I}^<$ provides a complete invariant w.r.t. guarded bisimulation equivalence on the class of all finite τ-structures. Moreover, $\mathbb{I}^<(\mathfrak{A})$ is* Ptime *computable from $\mathfrak{A}$ and uniformly interpretable in a $\sim_{\mathrm{g}}$-invariant manner as a quotient over $\mathfrak{A}^w$, where w is the width of τ, in* IFP *and* LFP*.*

In order to prove Theorem 8 and, as our main goal, the abstract capturing result of Corollary 31, it therefore suffices to provide a Ptime computable (and hence also IFP- and LFP-interpretable) inversion for the complete invariant $\mathbb{I}^<$. This, in combination with $\mathbb{I}^<$, produces a Ptime computable (and automatically IFP- and LFP-interpretable) canonisation as follows:

$$\begin{tikzcd} \mathfrak{A} \arrow[r, mapsto, "\mathbb{I}^<"] \arrow[rr, mapsto, bend right, "\mathrm{can}^<"'] & \mathbb{I}^<(\mathfrak{A}) \arrow[r, mapsto, "F^<"] & \mathrm{can}^<(\mathfrak{A}) = F^<(\mathbb{I}^<(\mathfrak{A})) = (\mathrm{can}(\mathfrak{A}), <) \end{tikzcd}$$

yields a linearly ordered representative of the $\sim_{\mathrm{g}}$-equivalence class of $\mathfrak{A}$.

In fact, we obtain $\mathrm{can}^<(\mathfrak{A})$ as an ordered version of the Rosati cover $\mathfrak{R}_2(\mathbb{I}^<(\mathfrak{A}))$ obtained from an ordered version of $\mathbb{I}(\mathfrak{A})$. Indeed, from Theorem 14 we already know that $\mathfrak{R}_2(\mathfrak{I})$ is a suitable candidate for inverting an ordered guarded bisimulation invariant $(\mathfrak{I}, <)$, because

$\mathbb{I}^{<}(\mathfrak{R}_2(\mathfrak{I})) = (\mathfrak{I}, <)$. It remains to define a canonical linear ordering of the Rosati cover. This is a trivial exercise given the term structure of the elements of $\mathfrak{R}_2$. First, using the linear order of a given invariant $(\mathfrak{I}, <)$ and the standard ordering of natural numbers we define a linear ordering of constants $c^j_{e,i}$, say, in a lexicographic manner applied to the corresponding tuples (e, i, j). We also fix a similarly defined ordering of all function symbols $f^j_{\rho,i}$ arising from $\mathfrak{I}$ as introduced in Section 3.1. It is then straighforward to extend these to a linear ordering of all terms (of depth 2 in the case of $\mathfrak{R}_2(\mathfrak{I})$) by stipulating that constants, i.e. terms of height zero precede all terms of height one, which in turn precede all terms of height two in the ordering; and that terms of the same height are ordered first according to their root function symbols, then according to their sets of subterms inductively. It is apparent that such an ordering of $\mathfrak{R}_2(\mathfrak{I})$ can be computed in polynomial time given the ordered invariant $(\mathfrak{I}, <)$. Hence $\mathrm{can}^{<}(\mathfrak{A}) = (\mathfrak{R}_2(\mathbb{I}^{<}(\mathfrak{A})), <)$ is well defined, polynomial-time computable and fulfills the claims of Theorem 8.

References

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

[2] V. Bárány, G. Gottlob, and M. Otto. Querying the guarded fragment. Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, Edinburgh, UK, pp. 1–10, 2010.

[3] C. Beeri, R. Fagin, D. Maier and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30(3):479–513, 1983.

[4] C. Beeri and M. Y. Vardi. The implication problem for data dependencies. ICALP'81, pp. 73–85, 1981.

[5] C. Beeri, M. Y. Vardi. A proof procedure for data dependencies. *JACM*, 31(4):218–741, 1984.

[6] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, R. Rosati. Tractable reasoning and efficient query answering in description logics: the DL-Lite family. *J. Autom. Reasoning*, 39(3):385–429, 2007.

[7] A. Calí, G. Gottlob, and M. Kifer. Taming the infinite chase: query answering under expressive relational constraints. Proc. KR'08, pp. 70–80, 2008.

[8] A. Calí, G. Gottlob, T. Lukasiewicz. A general datalog-based framework for tractable query answering over ontologies. Proc. PODS'09, 2009.

[9] A. Calí, G. Gottlob, T. Lukasiewicz. Datalog$^{\pm}$: a unified approach to ontologies and integrity constraints. Proc. ICDT'09 pp. 14–30, 2009.

[10] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Data complexity of query answering in description logics. Proc. KR'06, 2006.

[11] D. Calvanese. Finite model reasoning in description logics. In Principles of Knowledge Representation and Reasoning (KR'96), pp. 292–303, 1996.

[12] A.K. Chandra and P.M. Merlin. Optimal implementation of conjunctive queries in relational databases. In Proc. 9th ACM Symp. on Theory of Computing, pp. 77–90, 1977.

[13] R. Fagin, Ph. Kolaitis, R. J. Miller and L. Popa. Data exchange: semantics and query answering. *Theor. Comp. Sci.*, 336(1):89-124, 2005.

[14] G. Gottlob, N. Leone and F. Scarcello. Robbers, marshals, and guards: game theoretic and logical characterizations of hypertree width *Journal of Computer and System Sciences*, 66:775?808, 2003.

[15] E. Grädel. Description logics and guarded fragments of first-order logic. In *Proc. DL'98*, CEUR Electronic Workshop Proc., 1998.

[16] E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64(4):1719–1742, 1999.

[17] E. Grädel. Decision procedures for guarded logics. In *Proc. CADE'99*, pp. 31-51, 1999.

[18] E. Grädel and C. Hirsch and M. Otto. Back and forth between guarded and modal logics. *ACM Transactions on Computational Logics*, 3(3):418–463, 2002.

[19] E. Grädel, Ph. Kolaitis, and M.Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symb. Logic*, 3:53–69, 1997.

[20] E. Grädel, M. Otto, and E. Rosen. Two-variable logic with counting is decidable. In Proceedings of the 12th IEEE Symposium on Logic in Computer Science, LICS 1997, pp. 306–317, 1997.

[21] M. Grohe. The quest for a logic capturing PTIME. In Proceedings of the 23rd Annual IEEE Symposium on Logic in Computer Science, LICS 2008, Pittsburgh, PA, USA, pp. 267–271, 2008.

[22] B. Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 155:365–371, 2005.

[23] I. Hodkinson. Loosely guarded fragment of first-order logic has the finite model property. *Studia Logica*, 70:205–240, 2002.

[24] I. Hodkinson and M. Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *Bulletin of Symbolic Logic*, 9:387–405, 2003.

[25] N. Immerman. Relational queries computable in polynomial time. *Inform. and Control*, 68:86–104, 1986.

[26] D. S. Johnson and A. C. Klug. Testing containment of conjunctive queries under functional and inclusion dependencies. JCSS 28(1):167-189, 1984.

[27] Y. Kazakov. A polynomial translation from the two-variable guarded fragment with number restrictions to the guarded fragment. In Logics in Artificial Intelligence: 9th European Conference, JELIA 2004, J. J. Alferes and J. Leite, eds, vol. 3229 of Lecture Notes in Artificial Intelligence, pp. 372–384. 2004.

[28] C. Lutz. Inverse roles make conjunctive queries hard. Proceedings of the 2007 International Workshop on Description Logics (DL2007), CEUR-WS, 2007. Long version available from authors website: http://www.informatik.uni-bremen.de/ clu/papers/archive/dl07a.pdf.

[29] C. Lutz. The complexity of conjunctive query answering in expressive description logics. In Automated Reasoning, 4th International Joint Conference, IJCAR 2008, LNCS 5195, pp. 179–193.

[30] C. Lutz, U. Sattler, and L. Tendera. The complexity of finite model reasoning in description logics. Information and Computation, 199:132171, 2005.

[31] M. Mortimer. On languages with two variables. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:134–140, 1975.

[32] M. Otto. Bisimulation-invariant Ptime and higher-dimensional μ-calculus. *TCS*, 224:237-265, 1999.

[33] M. Otto. Bounded variable logics and counting – A study in finite models. *Lecture Notes in Logic*, vol. 9, Springer-Verlag, 1997.

[34] M. Otto. Avoiding incidental homomorphisms into guarded covers. Techn. rep., TU Darmstadt, 2009.

[35] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, Edinburgh, UK, pp. 11–20, 2010.

[36] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *J. of the ACM*, 59(1), Article No. 5, 2011.

[37] H. Pérez-Urbina, B. Motik, and I. Horrocks. Tractable query answering and rewriting under description logic constraints. *J. Applied Logic*, 8(2):186–209, 2010.

[38] I. Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *J. of Logic, Language and Information*, 14:369–395, 2005.

[39] I. Pratt-Hartmann. Complexity of the guarded two-variable fragment with counting quantifiers. *J. of Logic and Computation*, 17(1):133–155, 2007.

[40] L. Pacholski, W. Szwast, and L. Tendera. Complexity results for first-order two- variable logic with counting, S1AM Journal on Computing, 29(4), 1083-1117, 1999.

[41] R. Rosati. On the decidability and finite controllability of query processing in databases with incomplete information. In *Proc. PODS'06*, pp. 356–365, 2006.

[42] R. Rosati. On the finite controllability of conjunctive query answering in databases under open-world assumption. *Journal of Computer and System Sciences*, 77(3):572–594, 2011.

[43] M. Y. Vardi. The complexity of relational query languages. In *Proc. STOC'82*, pp. 137–146, 1982.