

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Inconsistency-tolerant Query Answering for Existential Rules

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Inconsistency-tolerant Query Answering for Existential Rules / Thomas Lukasiewicz, Enrico Malizia, Maria Vanina Martinez, Cristian Molinaro, Andreas Pieris, Gerardo I. Simari. - In: ARTIFICIAL INTELLIGENCE. - ISSN 0004-3702. - STAMPA. - 307:(2022), pp. 103685.1-103685.39. [10.1016/j.artint.2022.103685]

Availability:

This version is available at: <https://hdl.handle.net/11585/879287> since: 2022-03-23

Published:

DOI: <http://doi.org/10.1016/j.artint.2022.103685>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Lukasiewicz, T., Malizia, E., Martinez, M. V., Molinaro, C., Pieris, A., & Simari, G. I. (2022). Inconsistency-tolerant query answering for existential rules. *Artificial Intelligence*, 307

The final published version is available online at <https://dx.doi.org/10.1016/j.artint.2022.103685>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Inconsistency-tolerant Query Answering for Existential Rules*

Thomas Lukasiewicz^a, Enrico Malizia^b, Maria Vanina Martinez^{c,d}, Cristian Molinaro^e,
Andreas Pieris^{f,g}, and Gerardo I. Simari^{h,i}

^aDepartment of Computer Science, University of Oxford, UK

^bDISI, University of Bologna, Italy

^cDepartment of Computer Science, Universidad de Buenos Aires, Argentina

^dInstitute for Computer Science Research (ICC UBA–CONICET), Argentina

^eDIMES, University of Calabria, Italy

^fSchool of Informatics, University of Edinburgh, UK

^gDepartment of Computer Science, University of Cyprus, Cyprus

^hDept. of Computer Sci. and Engineering, Universidad Nacional del Sur (UNS), Argentina

ⁱInstitute for Computer Science and Engineering (UNS–CONICET), Argentina

Abstract

Querying inconsistent knowledge bases is an intriguing problem that gave rise to a flourishing research activity in the knowledge representation and reasoning community during the last years. It has been extensively studied in the context of description logics (DLs), and its computational complexity is rather well-understood. Although DLs are popular formalisms for modeling ontologies, it is generally agreed that rule-based ontologies are well-suited for data-intensive applications, since they allow us to conveniently deal with higher-arity relations, which naturally occur in standard relational databases. The goal of this work is to perform an in-depth complexity analysis of querying inconsistent knowledge bases in the case of the main decidable classes of existential rules, based on the notions of guardedness, linearity, acyclicity, and stickiness, enriched with negative (a.k.a. denial) constraints. Our investigation concentrates on three central inconsistency-tolerant semantics: the ABox repair (AR) semantics, considered as the standard one, and its main sound approximations, the intersection of repairs (IAR) semantics and the intersection of closed repairs (ICR) semantics.

1 Introduction

The purpose of an ontology is to provide an explicit specification via an unambiguous language, typically based on logic, of an abstract model of a domain of interest. A relatively recent, and admittedly quite successful, application of ontologies is ontology-based data access (OBDA) [43], which in turn has emerged as an exciting application of knowledge representation and reasoning technologies in information management systems. The goal of OBDA is to facilitate access to data by separating the user from the raw data sources. This is done using an ontology that provides a unified conceptual view of the data, and makes it accessible via queries solely formulated in the vocabulary of the ontology without any knowledge of the actual structure of the data. In addition, the ontology enriches the possibly incomplete data sources with domain knowledge, enabling more complete answers to queries, typically conjunctive queries.

In real-life OBDA scenarios, involving large amounts of data, it is likely that the raw data is inconsistent with the ontology. Since standard ontology languages adhere to the classical first-order logic semantics, inconsistencies are nothing else than logical contradictions. Therefore, the classical inference semantics fails when faced with an inconsistency, since everything follows from a logical contradiction. This demonstrates the need of developing alternative semantics.

© 2022. This manuscript version is made available under the CC BY-NC-ND 4.0 license. The formal publication of this manuscript is available via the DOI: 10.1016/j.artint.2022.103685.

*This paper is a substantially extended and revised version of the papers [34, 37, 39, 40].

There has been a significant effort on the development of inconsistency-tolerant semantics for query answering purposes. Consistent query answering, first developed for relational databases [1], and then generalized as the *ABox repair* (AR) semantics for several DLs [28], is the most widely accepted semantics for querying inconsistent knowledge bases. The AR semantics is based on the idea that an answer is considered to be valid if it can be inferred from each of the repairs of the extensional data set D , i.e., the \subseteq -maximal consistent subsets of D . However, obtaining the set of consistent answers under the AR semantics is known to be a hard problem, even for lightweight ontology languages [28]. For this reason, several other semantics have been developed with the aim of approximating the set of consistent answers [3, 9, 28, 38]; the list is by no means exhaustive, and we refer the reader to [4] for a comprehensive survey. We also refer the reader to the related work section (Section 9) for further details on inconsistency-tolerant semantics.

The two main approximations of the AR semantics that are of special interest for the present work are the following:

1. The *intersection of ABox repairs* (IAR) semantics: an answer must be inferred from the intersection of the repairs and the ontology [28].
2. The *intersection of closed repairs* (ICR) semantics: an answer must be inferred from the intersection of the *closure* of the repairs and the ontology [3].

Apart from being natural approximations of the AR semantics, the IAR and ICR semantics are coming with additional advantages of practical relevance. Recent work on explanation in the context of inconsistency-tolerant query answering shows that explanations are much easier to define and compute for the IAR and ICR semantics [7]. Moreover, the IAR and ICR semantics are amenable to preprocessing, since the intersection of the (closed) repairs can be computed offline, and then standard query answering algorithms can be employed online. Indeed, the latter approach has been adopted in the implementation of the IAR semantics [30], while for the ICR semantics, it has been already remarked in [5].

The complexity of query answering under different inconsistency-tolerant semantics, including the ones above, has been extensively studied for a wide spectrum of DLs; see, e.g., [3, 6, 8, 30, 44]. But what about rule-based ontology languages? It is agreed that rule-based ontologies are well-suited for data-intensive applications such as OBDA, since they allow us to conveniently deal with higher-arity relations, which naturally occur in standard relational databases. Therefore, analyzing and understanding the complexity of query answering under inconsistency-tolerant semantics in the presence of rule-based ontologies is a highly relevant task that deserves our attention. This is the main concern of the present work.

Towards this direction, we focus on ontologies modeled using existential rules (also called tuple-generating dependencies), i.e., first-order sentences of the form

$$\forall \bar{x} \forall \bar{y} (\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})),$$

with ϕ and ψ being conjunctions of atoms, and negative (a.k.a. denial) constraints, which are first-order sentences of the form

$$\forall \bar{x} (\phi(\bar{x}) \rightarrow \perp),$$

where ϕ is a conjunction of atoms, and \perp denotes the truth constant false. Note that rules of the above form are also known in the literature as Datalog $^\pm$ rules [12]. It is known, however, that query answering under arbitrary existential rules (even without negative constraints) is undecidable (see, e.g., [10]). This has led to an intensive research activity for identifying restrictions on existential rules that lead to decidability. The main decidable paradigms are (i) guardedness [10] (which includes linearity [11]) that is based on the relativization of quantifiers by atomic formulas, (ii) acyclicity [37] that forbids recursion, and (iii) stickiness [13] that forces variables that appear more than once in the left-hand side of an existential rule to be propagated to the right-hand side of the rule. The goal underlying stickiness is to express non-guarded statements without forbidding recursion. The formal definitions of the above paradigms are given in Section 3.

Our main goal is to perform an in-depth complexity analysis of querying inconsistent knowledge bases under the AR, IAR, and ICR semantics, when the ontology is modeled using one of the main decidable classes of existential rules discussed above enriched with negative constraints. Note that we can easily inherit from DLs that our problem in the case of the AR semantics is intractable, even when the ontology and the query are fixed [28]. Thus, another objective of our work, apart from clarifying the complexity landscape, is to understand whether the IAR and ICR semantics reduce the complexity of the problem in question, especially when the ontology and the query are fixed. Our contributions are as follows:

- Before studying inconsistency-tolerant query answering, we first need to understand standard query answering under existential rules without inconsistencies. The latter is well-understood in the case of guardedness, linearity, and stickiness. Surprisingly, query answering under acyclic existential rules has not been explicitly studied before the conference paper [37], which is one of the works on which the present journal paper is based. We show that the problem is NEXPTIME-complete even if we bound the arity, NP-complete if we fix the ontology, and in AC_0 when both the ontology and the query are fixed; these results are summarized in Proposition 3.3.
- The complete picture concerning inconsistency-tolerant query answering in the case of the AR semantics is given by Theorem 5.1. The main outcome is that the problem is intractable, actually coNP-complete, even when the ontology and the query are fixed. This is not surprising in view of the fact that the problem is already coNP-hard for lightweight DLs such as DL-Lite [28].
- The complete picture for the IAR semantics is given by Theorem 6.1. It turned out that the IAR semantics does not reduce the complexity in the case of guardedness; the only difference is a minor decrease from Π_2^P to Θ_2^P when the ontology is fixed. However, in the case of linear, acyclic, and sticky existential rules, we observed some significant differences: the complexity decreases from Π_2^P to NP when the ontology is fixed, and from coNP to AC_0 when both the ontology and the query are fixed. This is due to a central property known as first-order rewritability. Note that all the complexity results, apart from AC_0 , are completeness results.
- The complete picture for the ICR semantics is given by Theorem 7.1. The ICR semantics does not reduce the complexity of our problem, no matter which class of existential rules we consider. We only observed a minor decrease from Π_2^P -complete to Θ_2^P -complete when the ontology is fixed.
- Finally, for the sake of completeness, we consider the central class of full existential rules (i.e., rules without existentially quantified variables), as well as the main extensions of guarded, acyclic, and sticky existential rules that generalize full existential rules; the complete picture for all the inconsistency-tolerant semantics in question is given by Theorem 8.1. It turned out that the analysis performed for the less expressive classes allowed us to easily complete the picture for the more expressive classes of existential rules.

The rest of the paper is organized as follows. The basics on relational databases, conjunctive queries, existential rules, and negative constraints, as well as basic complexity classes, are recalled in Section 2. An overview of query answering under the main decidable classes of existential rules is given in Section 3. The main inconsistency-tolerant semantics for query answering under existential rules are introduced in Section 4. Our complexity results on inconsistency-tolerant query answering w.r.t. the AR, IAR, and ICR semantics are presented in Sections 5, 6, and 7, respectively. In Section 8, we consider the class of full existential rules, as well as the main extensions of guarded, acyclic, and sticky existential rules that generalize full existential rules. A rather comprehensive overview of the main inconsistency-tolerant semantics, which go beyond the AR, IAR, and ICR semantics, that have been proposed and studied during the last decade is given in Section 9. We finally conclude in Section 10 with a brief discussion and directions for future research.

2 Preliminaries

In this section, we recall the basics on relational databases, (unions of) conjunctive queries, ontological query answering (including tuple-generating dependencies¹ and negative constraints), the chase procedure, and the complexity classes encountered in this paper. Throughout the paper, we assume the disjoint countably infinite sets \mathbf{C} , \mathbf{N} , and \mathbf{V} of *constants*, (*labeled*) *nulls*, and *variables*, respectively. We also refer to constants, nulls, and variables as *terms*.

2.1 Relational databases and (unions of) conjunctive queries

A (*relational*) *schema* \mathbf{S} is a finite set of relation symbols (or predicates) with associated arity. We write R/n to denote that the arity of the relation symbol R is $n \geq 0$. A *relational atom* (or simply *atom*) over \mathbf{S} is an expression of the form $R(\bar{t})$, where R is an n -ary relation symbol from \mathbf{S} , and \bar{t} is an n -tuple of

¹Henceforth, as customary in the literature, we adopt the more traditional term tuple-generating dependency instead of existential rule.

terms. An atom is *ground* if it mentions only constants of \mathbf{C} . An *instance* over \mathbf{S} is a (possibly infinite) set of atoms over \mathbf{S} that contain constants and nulls, while a *database* over \mathbf{S} is a finite set of ground atoms over \mathbf{S} , i.e., a finite instance without nulls. For an instance I , we write $\text{dom}(I)$ for the set of all terms occurring in I .

Consider two sets of terms T and S . A *substitution* from T to S is a function $h: T \rightarrow S$. The restriction of h to $T' \subseteq T$, denoted $h|_{T'}$, is the function from T' to S such that, for every $t \in T'$, $h|_{T'}(t) = h(t)$. Consider now two sets of atoms A and B . A *homomorphism* from A to B is a substitution h from the set of terms in A to the set of terms in B , i.e., from $\text{dom}(A)$ to $\text{dom}(B)$, such that (i) $t \in \mathbf{C}$ implies $h(t) = t$, and (ii) $R(t_1, \dots, t_n) \in A$ implies $R(h(t_1), \dots, h(t_n)) \in B$.

A *conjunctive query* (CQ) over a schema \mathbf{S} is a formula of the form

$$q(\bar{x}) := \exists \bar{y} (R_1(\bar{z}_1) \wedge \dots \wedge R_m(\bar{z}_m)),$$

where each $R_i(\bar{z}_i)$, for $i \in \{1, \dots, m\}$, is an atom over \mathbf{S} without nulls, each variable occurring in a tuple \bar{z}_i appears either in \bar{x} or \bar{y} , and \bar{x} contains all the *free variables* of q . If \bar{x} is empty, then q is a *Boolean conjunctive query* (BCQ). The evaluation of CQs is defined in terms of homomorphisms. Given an instance I , the *evaluation of $q(\bar{x})$ over I* , denoted $q(I)$, is the set of all tuples $\bar{t} \in \mathbf{C}^{|\bar{x}|}$ such that there exists a homomorphism h from $q(\bar{x})$ to I with $h(\bar{x}) = \bar{t}$. By abuse of notation, we sometimes treat a tuple of variables as a set of variables, and a conjunction of atoms as a set of atoms. Note that in the case of Boolean CQs, the only possible answer is the empty tuple.

A *union of conjunctive queries* (UCQ) over \mathbf{S} is a formula of the form

$$q(\bar{x}) := q_1(\bar{x}) \vee \dots \vee q_n(\bar{x}),$$

where each $q_i(\bar{x})$ is a CQ over \mathbf{S} . The evaluation of q over an instance I , denoted $q(I)$, is defined as the set of tuples $\bigcup_{i \in \{1, \dots, n\}} q_i(I)$. By abuse of notation, we may treat a UCQ $q(\bar{x})$ as the one above as the set of CQs $\{q_1(\bar{x}), \dots, q_n(\bar{x})\}$.

2.2 Ontological query answering

A *tuple-generating dependency* (TGD) σ is a (constant-free) sentence

$$\forall \bar{x} \forall \bar{y} (\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})),$$

where \bar{x} , \bar{y} , and \bar{z} are tuples of variables of \mathbf{V} , and ϕ and ψ are conjunctions of atoms. For brevity, we write σ as $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$, and use comma for joining atoms. We refer to $\phi(\bar{x}, \bar{y})$ and $\psi(\bar{x}, \bar{z})$ as the *body* and *head* of σ , denoted $\text{body}(\sigma)$ and $\text{head}(\sigma)$, respectively. An instance I satisfies a TGD σ as the one above, written $I \models \sigma$, if the following holds: whenever there exists a homomorphism h such that $h(\phi(\bar{x}, \bar{y})) \subseteq I$, then there exists an *extension* h' of $h|_{\bar{x}}$, i.e., $h' \supseteq h|_{\bar{x}}$, such that $h'(\psi(\bar{x}, \bar{z})) \subseteq I$. The instance I satisfies a set Σ of TGDs, written $I \models \Sigma$, if $I \models \sigma$ for each $\sigma \in \Sigma$. Let TGD be the class of (finite) sets of TGDs.

A *negative constraint* (NC) σ is a sentence of the form

$$\forall \bar{x} (\phi(\bar{x}) \rightarrow \perp),$$

where \bar{x} is a tuple of variables of \mathbf{V} , ϕ is a conjunction of atoms, and \perp denotes the truth constant false. For brevity, we write σ as $\phi(\bar{x}) \rightarrow \perp$, and use comma for joining atoms. We refer to $\phi(\bar{x})$ as the *body* of σ , denoted $\text{body}(\sigma)$. An instance I satisfies an NC σ as the one above, written $I \models \sigma$, if there is *no* homomorphism h such that $h(\phi(\bar{x})) \subseteq I$ (observe the inversion: satisfies NC \rightarrow no homomorphism exists). The instance I satisfies a set Σ of NCs, written $I \models \Sigma$, if $I \models \sigma$ for each $\sigma \in \Sigma$. Let NC be the class of (finite) sets of NCs.

Consider a database D and a set Σ of TGDs and NCs; henceforth, we denote by $\tau(\Sigma)$ and $\nu(\Sigma)$ the set of TGDs and NCs, respectively, occurring in Σ . A *model* of D and Σ is an instance $I \supseteq D$ such that $I \models \tau(\Sigma)$ and $I \models \nu(\Sigma)$. We write $\text{mods}(D, \Sigma)$ for the set of models of D and Σ . The *certain answers* to a CQ q w.r.t. D and Σ is defined as the set of tuples

$$\text{cert}(q, D, \Sigma) = \bigcap_{I \in \text{mods}(D, \Sigma)} q(I).$$

A problem that is central for our work is to compute the certain answers to a CQ w.r.t. a database and a set of TGDs and NCs that falls in $\mathbf{C} \cup \text{NC}$, where $\mathbf{C} \subseteq \text{TGD}$ is a class of TGDs; concrete classes of TGDs are given below. For brevity, given a class \mathbf{C} of TGDs, we write \mathbf{C}_\perp as an abbreviation for $\mathbf{C} \cup \text{NC}$. As is customary when studying the complexity of this problem, we focus on its decision version:

PROBLEM :	QAns(\mathbf{C}_\perp)
INPUT :	A database D , a set $\Sigma \in \mathbf{C}_\perp$, a CQ $q(\bar{x})$, and $\bar{c} \in \text{dom}(D)^{ \bar{x} }$.
QUESTION :	Does $\bar{c} \in \text{cert}(q, D, \Sigma)$?

This general formulation refers to the *combined complexity* of the problem, that is, the database, the set of TGDs and NCs, the CQ, and the candidate tuple are considered part of the input. It is common, however, to study also refined measures that are more realistic in practice. Here, we consider the *bounded-arity combined complexity*, where the arity of the underlying schema is bounded by an integer, the *fixed-program combined complexity*,² where the set of TGDs and NCs is fixed, and the *data complexity*, where the CQ and the set of TGDs and NCs are fixed. Henceforth, for brevity, we write **c**-, **ba**-, **fp**-, and **d**-complexity for combined, bounded-arity combined, fixed-program combined, and data complexity, respectively.

It should be clear that if the given database D and set Σ of TGDs and NCs are *inconsistent*, i.e., $\text{mods}(D, \Sigma) = \emptyset$,³ then the set of certain answers to a CQ $q(\bar{x})$ consists of all the tuples $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$. This is because the definition of certain answers relies on the classical first-order semantics. That is,

$$\text{cert}(q, D, \Sigma) = \{\bar{c} \in \text{dom}(D)^{|\bar{x}|} \mid D \wedge \Sigma \models_{\text{FO}} q(\bar{c})\},$$

where $D \wedge \Sigma$ denotes the first-order sentence $\bigwedge_{\alpha \in D} \alpha \wedge \bigwedge_{\sigma \in \Sigma} \sigma$, $q(\bar{c})$ is the first-order sentence obtained by instantiating the free variables of q with \bar{c} , and \models_{FO} denotes the standard first-order entailment. Therefore, if D and Σ are inconsistent, this means that the sentence $D \wedge \Sigma$ does not admit a model, i.e., is a logical contradiction, and thus, everything is entailed from it:

$$\text{mods}(D, \Sigma) = \emptyset \implies \text{cert}(q, D, \Sigma) = \{\bar{c} \mid \bar{c} \in \text{dom}(D)^{|\bar{x}|}\}.$$

Another crucial observation is that, if $\text{mods}(D, \Sigma) \neq \emptyset$ (i.e., D and Σ are consistent), then for computing the certain answers to q w.r.t. D and Σ , it suffices to focus on the TGDs of Σ . That is:

$$\text{mods}(D, \Sigma) \neq \emptyset \implies \text{cert}(q, D, \Sigma) = \text{cert}(q, D, \tau(\Sigma)).$$

By exploiting the above two implications, we can easily show that:

$$\bar{c} \in \text{cert}(q, D, \Sigma) \iff \text{mods}(D, \Sigma) = \emptyset \quad \text{or} \quad \bar{c} \in \text{cert}(q, D, \tau(\Sigma)). \quad (1)$$

It is not difficult to see that the problem of checking whether $\text{mods}(D, \Sigma) = \emptyset$ boils down to the problem of checking whether the database and the set of TGDs entail at least one NC. Given an NC σ of the form $\phi(\bar{x}) \rightarrow \perp$, we write q_σ for the Boolean CQ $\exists \bar{x} \phi(\bar{x})$. Then:

$$\text{mods}(D, \Sigma) = \emptyset \iff \text{there is } \sigma \in \nu(\Sigma) \text{ s.t. } \text{cert}(q_\sigma, D, \tau(\Sigma)) \neq \emptyset. \quad (2)$$

From (1) and (2), we immediately get the following folklore result:

Proposition 2.1. *Consider a database D , a set Σ of TGDs and NCs, a CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$. The following are equivalent:*

1. $\bar{c} \in \text{cert}(q, D, \Sigma)$.
2. There is $\sigma \in \nu(\Sigma)$ such that $\text{cert}(q_\sigma, D, \tau(\Sigma)) \neq \emptyset$, or $\bar{c} \in \text{cert}(q, D, \tau(\Sigma))$.

2.3 The chase procedure

The above proposition shows that for checking whether a tuple is a certain answer to a CQ w.r.t. a database and a set Σ of TGDs and NCs, we actually have to reason with the TGD component of Σ . The *chase procedure* is a useful algorithmic tool when reasoning with TGDs that takes as an input a database D and a set Σ of TGDs, and constructs a (possibly infinite) instance I such that $I \supseteq D$ and $I \models \Sigma$; see, e.g., [10, 21]. We start by defining the notion of chase step.

²A set of TGDs and NCs is sometimes called a program, and hence the term fixed-program.

³If Σ consists only of TGDs, then D and Σ are always consistent, i.e., $\text{mods}(D, \Sigma) \neq \emptyset$. In this case, a model of D and Σ can be constructed via the chase procedure introduced in Section 2.3.

⁴Notice that q_σ is Boolean, and by $\text{cert}(q_\sigma, D, \tau(\Sigma)) \neq \emptyset$, we simply mean that the only possible answer (i.e., the empty tuple) is a certain answer.

Consider an instance I and a TGD σ of the form $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$. We say that σ is *applicable* w.r.t. I if there exists a homomorphism h from $\phi(\bar{x}, \bar{y})$ to I . In this case, *the result of applying σ over I with h* is the instance $J = I \cup h'(\psi(\bar{x}, \bar{z}))$, where h' is an extension of $h|_{\bar{x}}$ that maps each variable of \bar{z} to a distinct null not in I , and for each pair (z, w) of distinct variables of \bar{z} , $h'(z) \neq h'(w)$. For such a single chase step, we write $I \langle \sigma, h \rangle J$.

The main idea of the chase is, starting from a database D , to exhaustively apply TGDs over the instance constructed so far. This simple idea is formalized via the notion of chase sequence. We distinguish the two cases where a chase sequence is finite or infinite. Consider a set Σ of TGDs:

- A finite sequence I_0, I_1, \dots, I_n of instances, where $n \geq 0$, is a *chase sequence for I_0 under Σ* if: (i) for each $0 \leq i < n$, $I_i \langle \sigma, h \rangle I_{i+1}$ for some $\sigma \in \Sigma$ and homomorphism h from $\text{body}(\sigma)$ to I_i , (ii) for each $0 \leq i < j < n$, assuming that $I_i \langle \sigma_i, h_i \rangle I_{i+1}$ and $I_j \langle \sigma_j, h_j \rangle I_{j+1}$, $\sigma_i = \sigma_j$ implies $h_i \neq h_j$, i.e., h_i and h_j are different homomorphisms, and (iii) there is no TGD of Σ that is applicable w.r.t. I_n . The result of the chase is the (finite) instance I_n .
- An infinite sequence I_0, I_1, \dots of instances is a *chase sequence for I_0 under Σ* if: (i) for each $i \geq 0$, $I_i \langle \sigma, h \rangle I_{i+1}$ for some $\sigma \in \Sigma$ and homomorphism h from $\text{body}(\sigma)$ to I_i , (ii) for each $i, j > 0$ such that $i \neq j$, assuming that $I_i \langle \sigma_i, h_i \rangle I_{i+1}$ and $I_j \langle \sigma_j, h_j \rangle I_{j+1}$, $\sigma_i = \sigma_j$ implies $h_i \neq h_j$, and (iii) for each $i \geq 0$, and for every $\sigma \in \Sigma$ that is applicable w.r.t. I_i due to a homomorphism h , there exists $j \geq i$ such that $I_j \langle \sigma, h \rangle I_{j+1}$. The latter is known as the *fairness condition*, and guarantees that all the applicable TGDs eventually will be applied. The result of the chase is $\bigcup_{i \geq 0} I_i$.

Since we consider the *oblivious* version of the chase, i.e., a TGD is applied whenever its body is satisfied no matter whether its head is satisfied, every chase sequence for I under Σ leads to the same result (up to isomorphism). Thus, we can refer to *the* result of the chase for I under Σ , denoted $\text{chase}(I, \Sigma)$.

The following is a well-known result, which exposes the usefulness of the chase in relation with ontological query answering. The key reason why this result holds is because, given a database D and a set Σ of TGDs, the instance $\text{chase}(D, \Sigma)$ is not only a model, but is a *universal model* of D and Σ , which means that $\text{chase}(D, \Sigma)$ can be homomorphically mapped to every instance $I \in \text{mods}(D, \Sigma)$.

Proposition 2.2 (see, e.g., [19, 21]). *Consider a database D , a set Σ of TGDs, and a CQ q . It holds that $\text{cert}(q, D, \Sigma) = q(\text{chase}(D, \Sigma))$.*

2.4 Complexity classes

We assume that the reader has some background in computational complexity theory, including the notions of Turing machine, and hardness and completeness of a problem for a complexity class, as can be found in standard textbooks, e.g., in [27, 42]. In what follows, we briefly recall the complexity classes that we encounter in our complexity results. The complexity class PSPACE (resp., PTIME, EXPTIME, and 2EXPTIME) contains all the decision problems that can be solved in polynomial space (resp., polynomial, exponential, and double exponential time) via a deterministic Turing machine. The complexity classes NP and NEXPTIME contain all the decision problems that can be solved in polynomial and exponential time via a non-deterministic Turing machine, respectively, while coNP and coNEXPTIME are their complementary classes, where “Yes” and “No” instances are interchanged. The class Θ_2^P is the class of all decision problems that can be decided in polynomial time by a deterministic Turing machine using a logarithmic number of calls to an NP-oracle. The class Σ_2^P is the class of problems that can be solved in non-deterministic polynomial time using an NP-oracle, and Π_2^P is the complement of Σ_2^P . The complexity class AC₀ is the class of all languages that are decidable via uniform families of Boolean circuits of polynomial size and constant depth. The above complexity classes and their inclusion relationships (which are all currently believed to be strict) are shown below:

$$\begin{aligned} \text{AC}_0 \subseteq \text{PTIME} \subseteq \text{NP}, \text{coNP} \subseteq \Theta_2^P \subseteq \Sigma_2^P, \Pi_2^P \subseteq \text{PSPACE} \subseteq \text{EXPTIME} \\ \subseteq \text{NEXPTIME}, \text{coNEXPTIME} \subseteq \text{P}^{\text{NEXPTIME}} \subseteq 2\text{EXPTIME}. \end{aligned}$$

3 Ontological query answering: Overview and new results

It is well known that $\text{QAns}(\text{TGD}_\perp)$ is undecidable; this is implicit in [2], which studies the implication problem for database dependencies. A stronger result of this kind can be found in [10], where it is shown that $\text{QAns}(\text{TGD}_\perp)$ is undecidable even in data complexity. Actually, the above negative results hold

	c-complexity	ba-complexity	fp-complexity	d-complexity
G_{\perp}	2EXPTIME	EXPTIME	NP	P TIME
L_{\perp}	PSPACE	NP	NP	in AC_0
A_{\perp}	NEXPTIME	NEXPTIME	NP	in AC_0
S_{\perp}	EXPTIME	NP	NP	in AC_0

Table 1: Complexity of $QAns(C_{\perp})$, where $C \in \{G, L, A, S\}$. Apart from the AC_0 upper bounds, the rest are completeness results.

even without considering NCs. Let us stress that Propositions 2.1 and 2.2 do not provide a chase-based decision procedure for query answering, since the chase is (in general) infinite. This has led to an intensive research activity for identifying syntactic restrictions on sets of TGDs that lead to decidability. Such restrictions can be classified into three main syntactic paradigms: guardedness (which includes linearity), acyclicity, and stickiness. We proceed to recall each of those paradigms, and discuss the complexity of query answering (summarized in Table 1).

3.1 Guardedness

A TGD σ is called *guarded* if $\text{body}(\sigma)$ has an atom, called *guard*, that contains all the variables occurring in $\text{body}(\sigma)$. Although the chase under a set of guarded TGDs does not necessarily terminate, query answering is decidable. This follows from the fact that the result of the chase procedure is “treelike”, or, in other words, has finite treewidth [10]. Let G be the class of sets of guarded TGDs. Then:

Proposition 3.1 ([10]). $QAns(G_{\perp})$ is 2EXPTIME-complete in c-complexity, EXP-TIME-complete in ba-complexity, NP-complete in fp-complexity, and PTIME-complete in d-complexity.

A key subclass of guarded TGDs is the class of *linear* TGDs, i.e., TGDs whose body consists of a single atom [11]. Let L be the class of sets of linear TGDs.

Proposition 3.2 ([10]). $QAns(L_{\perp})$ is PSPACE-complete in c-complexity, NP-complete in ba-complexity and fp-complexity, and in AC_0 in d-complexity.

3.2 Acyclicity

The *predicate graph* of a set Σ of TGDs is defined as follows: its nodes are the predicates occurring in Σ , and there is an edge from P to R iff there is a TGD $\sigma \in \Sigma$ such that P occurs in $\text{body}(\sigma)$ and R occurs in $\text{head}(\sigma)$. We call Σ *acyclic* (a.k.a. *non-recursive*) if its *predicate graph* contains no directed cycles. It is easy to see that acyclicity ensures the termination of the chase. Thus, by Propositions 2.1 and 2.2, we immediately get the decidability of $QAns(A_{\perp})$, where A denotes the class of acyclic sets of TGDs. However, the exact complexity of the problem has not been studied before the conference paper [37], which is one of the works on which the present journal paper is based. We proceed to show that:

Proposition 3.3. $QAns(A_{\perp})$ is NEXPTIME-complete in c- and ba-complexity, NP-complete in fp-complexity, and in AC_0 in d-complexity.

To establish the upper bounds for $QAns(A_{\perp})$, Proposition 2.1 suggests that it suffices to establish the same upper bounds for $QAns(A)$. At this point, one may be tempted to think that this can be done by simply constructing the chase instance I , and then checking whether the given tuple belongs to the evaluation of the CQ over I . Although this simple algorithm shows that indeed $QAns(A)$ is in NP in fp-complexity, it does not lead to optimal upper bounds in the case of c-complexity, ba-complexity, and d-complexity. In particular, it yields a 2EXPTIME upper bound in c-complexity and ba-complexity, and a PTIME upper bound in d-complexity. The next example shows that indeed this is the best that we can achieve via the naive procedure that explicitly constructs the chase instance.

Example 3.1. Consider the family of acyclic sets of TGDs

$$\{\Sigma_n = \{R_i(x), R_i(y) \rightarrow \exists z P_{i+1}(x, y, z), R_{i+1}(z)\}_{i \in \{0, \dots, n\}}\}_{n \geq 0}.$$

Consider also the family of databases

$$\{D_m = \{R_0(c_1), \dots, R_0(c_m)\}\}_{m \geq 0}.$$

The i -th stratum of Σ_n computes all the pairs of terms using the $m^{(2^i)}$ terms stored in the predicate R_i , and for each such pair generates a fresh null that is stored in the predicate R_{i+1} . It is easy to verify that the predicate R_n in $\text{chase}(D_m, \Sigma_n)$ contains $m^{(2^n)}$ nulls. Since each chase step generates exactly one null, the chase procedure starting from D_m and applying TGDs of Σ_n terminates after double-exponentially many steps in n , and polynomially many steps in m .⁵ ■

We proceed to provide more refined procedures that lead to optimal complexity upper bounds for $\text{QAns}(\mathbf{A})$. In Section 3.2.1 we establish the desired upper bounds, and we show that they are indeed worst-case optimal in Section 3.2.2.

3.2.1 Upper bounds

Our main technical result, which in turn allows us to obtain the desired upper bounds for $\text{QAns}(\mathbf{A})$, essentially shows that, for query answering purposes under acyclic sets of TGDs, it suffices to apply exponentially many chase steps, while this exponential bound does not depend on the input database. To formalize this statement, we first need to recall the notion of stratification.

Consider a set Σ of TGDs. Let $\text{sch}(\Sigma)$ be the set of predicates occurring in Σ . A *stratification* of Σ is a partition $\{\Sigma_1, \dots, \Sigma_n\}$, where $n \geq 1$, of Σ such that, for some function $f : \text{sch}(\Sigma) \rightarrow \{1, \dots, n\}$, the following holds:

- For each predicate $R \in \text{sch}(\Sigma)$, all the TGDs with R in their head belong to $\Sigma_{f(R)}$, i.e., they belong to the same set of the partition.
- If there exists a TGD in Σ such that the predicate R appears in its body, while the predicate P appears in its head, then $f(R) < f(P)$.

The *depth* of a set Σ of TGDs that admits a stratification, denoted $\text{depth}(\Sigma)$, is defined as the cardinality of its smallest stratification.

It is easy to verify that if a set of TGDs is acyclic, then it admits a stratification, and thus we can refer to its depth. Consider now a CQ q and a set $\Sigma \in \mathbf{A}$. Let $|q|$ be the number of atoms occurring in q , and $\text{width}(\Sigma)$ be the maximum number of atoms occurring in the body of a TGD of Σ . We define the function

$$g(q, \Sigma) = \begin{cases} |q| \cdot \left\lfloor \frac{\text{width}(\Sigma)^{\text{depth}(\Sigma)+1} - 1}{\text{width}(\Sigma) - 1} \right\rfloor & \text{if } \text{width}(\Sigma) > 1 \\ |q| \cdot \text{depth}(\Sigma) & \text{if } \text{width}(\Sigma) = 1. \end{cases}$$

Roughly speaking, $g(q, \Sigma)$ provides an upper bound on the number of chase steps that we need to apply in order to be able to safely conclude whether the query q is entailed by the instance $\text{chase}(D, \Sigma)$ for any database D .

Lemma 3.1. *Consider a database D , a set $\Sigma \in \mathbf{A}$ of TGDs, a CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$. If $\bar{c} \in \text{cert}(q, D, \Sigma)$, then there exists a sequence of instances $(J_i)_{0 \leq i \leq g(q, \Sigma)}$ with $J_0 = D$ and $J_i \langle \sigma, \mu \rangle J_{i+1}$ for some TGD $\sigma \in \Sigma$ and homomorphism μ from $\text{body}(\sigma)$ to $\text{dom}(J_i)$ such that $\bar{c} \in q(J_{g(q, \Sigma)})$.*


Proof. By hypothesis, there is a chase sequence $s = (I_i)_{0 \leq i \leq n}$ with $I_i \langle \sigma_i, h_i \rangle I_{i+1}$ for D under Σ such that $\bar{c} \in q(I_n)$ with $h(\bar{x}) = \bar{c}$. Consider an arbitrary atom α of q . It is clear that $h(\alpha) \in I_n$. We are going to establish an upper bound on the number of chase steps of s that are really needed to generate $h(\alpha)$. To this end, we first need to recall the so-called *chase relation* of s , denoted \prec_s , which is a binary relation over the atoms of I_n such that $(\beta, \gamma) \in \prec_s$ iff there exists $i \in \{0, \dots, n-1\}$ with $\beta \in h_i(\text{body}(\sigma_i))$ and $\gamma \in I_{i+1} \setminus I_i$. In simple words, \prec_s encodes which atoms generate some other atom via a single chase step. For convenience, we write $\beta \prec_s \gamma$ for the fact that (β, γ) belongs to \prec_s . We also write \prec_s^* for the transitive closure of \prec_s . Let $\prec_{s, h(\alpha)}$ be the subrelation of \prec_s defined as

$$\{(\beta, \gamma) \mid \beta \prec_s \gamma \text{ and } \gamma = h(\alpha) \text{ or } \gamma \prec_s^* h(\alpha)\}.$$

Roughly, $\prec_{s, h(\alpha)}$ collects only the sequences of atoms that lead to $h(\alpha)$. Observe that the inverse relation of $\prec_{s, h(\alpha)}$, denoted $\prec_{s, h(\alpha)}^-$, is a directed acyclic graph, where its nodes are atoms of I_n , with $h(\alpha)$ being the root, i.e., is the only node without an incoming edge. We can now transform $\prec_{s, h(\alpha)}^-$ into the binary

⁵This example provides a lower bound even if we consider the restricted version of the chase, where a TGD σ is applicable w.r.t. an instance I only if it is really violated, i.e., the homomorphism that maps $\text{body}(\sigma)$ to I cannot be extended to a homomorphism that maps $\text{head}(\sigma)$ to I .

relation $\prec_{s,h(\alpha)}^{-,\Delta}$ that represents a rooted tree, where $h(\alpha)$ is the root node, and each node has at most $\text{width}(\Sigma)$ children. This is done in the obvious way by duplicating some nodes of $\prec_{s,h(\alpha)}^-$ as shown in the figure below; the formal definition is omitted:



It should be clear that by providing an upper bound on the number of nodes occurring in $\prec_{s,h(\alpha)}^{-,\Delta}$, we immediately get an upper bound on the number of nodes in $\prec_{s,h(\alpha)}^-$. Due the fact that Σ is acyclic, the depth of the rooted tree $\prec_{s,h(\alpha)}^{-,\Delta}$ is at most $\text{depth}(\Sigma)$. Thus, the number of nodes in $\prec_{s,h(\alpha)}^{-,\Delta}$ is at most

$$\hat{g}(\Sigma) = \begin{cases} \left\lfloor \frac{\text{width}(\Sigma)^{\text{depth}(\Sigma)+1} - 1}{\text{width}(\Sigma) - 1} \right\rfloor & \text{if } \text{width}(\Sigma) > 1 \\ \text{depth}(\Sigma) & \text{if } \text{width}(\Sigma) = 1. \end{cases}$$

This allows us to conclude that $\prec_{s,h(\alpha)}$ has at most $\hat{g}(\Sigma)$ nodes, which means that the number of atoms in I_n that are needed to generate $h(\alpha)$ is at most $\hat{g}(\Sigma)$. This implies that the number of chase steps of s that are needed to generate $h(\alpha)$ is at most $\hat{g}(\Sigma)$. Hence, we can construct from s a sequence of instances $(J_i)_{0 \leq i \leq |q| \cdot \hat{g}(\Sigma)}$ with $J_0 = D$ and $J_i(\sigma, h)J_{i+1}$ for some $\sigma \in \Sigma$ and homomorphism μ from $\text{body}(\sigma)$ to $\text{dom}(J_i)$ such that $\bar{c} \in q(J_{|q| \cdot \hat{g}(\Sigma)})$. By definition, $|q| \cdot \hat{g}(\Sigma) = g(q, \Sigma)$, and the claim follows. \square

Having the above lemma in place, it is now easy to devise a non-deterministic algorithm for $\text{QAns}(\mathbf{A})$ that runs in exponential time in general, and in polynomial time whenever the set of TGDs is fixed: guess a sequence of instances $(J_i)_{0 \leq i \leq g(q, \Sigma)}$ with $J_0 = D$, a sequence of pairs $(\sigma_i, h_i)_{0 \leq i \leq g(q, \Sigma) - 1}$, where $\sigma_i \in \Sigma$ and h_i is a substitution from the set of variables occurring in $\text{body}(\sigma_i)$ to $\text{dom}(J_i)$, and a substitution h from the variables occurring in $q(\bar{x})$ to $\text{dom}(J_{g(q, \Sigma)})$ with $h(\bar{x}) = \bar{c}$, and then check that $J_i(\sigma_i, h_i)J_{i+1}$ for each $i \in \{0, \dots, g(q, \Sigma) - 1\}$, and h maps q to $J_{g(q, \Sigma)}$. Therefore, we obtain that $\text{QAns}(\mathbf{A})$ is in NEXPTIME in \mathbf{c} - and \mathbf{ba} -complexity, and in NP in \mathbf{fp} -complexity. However, the above algorithm provides only a PTIME upper bound for $\text{QAns}(\mathbf{A})$ in \mathbf{d} -complexity. For the latter type of complexity we need to argue a bit more.

It is implicit in [11] that Lemma 3.1 above implies that the class of acyclic sets of TGDs is *UCQ-rewritable*: given a set $\Sigma \in \mathbf{A}$ of TGDs and a CQ $q(\bar{x})$, we can construct a (finite) UCQ $Q_{q, \Sigma}(\bar{x})$ such that, for every database D and tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$, $\bar{c} \in \text{cert}(q, D, \Sigma)$ iff $\bar{c} \in Q_{q, \Sigma}(D)$. Since evaluating a *fixed* UCQ over a database is in AC_0 [47], we get that $\text{QAns}(\mathbf{A})$ is in AC_0 in \mathbf{d} -complexity.

3.2.2 Lower bounds

We now proceed to establish the complexity lower bounds stated in Proposition 3.3. Actually, the NP -hardness of $\text{QAns}(\mathbf{A})$ in \mathbf{fp} -complexity is inherited from the well-known fact that deciding whether a tuple of constants belongs to the evaluation of a CQ over a database is NP -hard, even if the underlying schema is fixed. It remains to show the following:

Lemma 3.2. *$\text{QAns}(\mathbf{A})$ is NEXPTIME -hard in \mathbf{ba} -complexity.*

The above result is shown via a reduction from the standard exponential tiling problem. A *tiling system* is a tuple $\mathcal{T} = (n, m, H, V, s)$, where n and m are numbers in unary, H and V are subsets of $\{1, \dots, m\} \times \{1, \dots, m\}$, and s is a sequence of numbers of $\{1, \dots, m\}$; let $s[i]$ be the i -th element of s . An *exponential tiling* for \mathcal{T} is a function $f : \{0, \dots, 2^n - 1\} \times \{0, \dots, 2^n - 1\} \rightarrow \{1, \dots, m\}$ such that:

- $f(i, 0) = s[i]$, for each $0 \leq i \leq (|s| - 1)$,
- $(f(i, j), f(i + 1, j)) \in H$, for each $0 \leq i \leq 2^n - 2$ and $0 \leq j \leq 2^n - 1$, and
- $(f(i, j), f(i, j + 1)) \in V$, for each $0 \leq i \leq 2^n - 1$ and $0 \leq j \leq 2^n - 2$.

The exponential tiling problem is defined as follows:


PROBLEM :	ExpTiling
INPUT :	A tiling system \mathcal{T} .
QUESTION :	Is there an exponential tiling for \mathcal{T} ?

The goal is to provide a polynomial time reduction from ExpTiling to QAns(A). Given a tiling system $\mathcal{T} = (n, m, H, V, s)$, we are going to construct in polynomial time a database $D_{\mathcal{T}}$, and a set $\Sigma_{\mathcal{T}} \in \mathbf{A}$ of TGDs that mentions only predicates of fixed arity such that \mathcal{T} has an exponential tiling iff $\text{cert}(\text{Yes}(), D_{\mathcal{T}}, \Sigma_{\mathcal{T}}) \neq \emptyset$

The database $D_{\mathcal{T}}$. It simply stores the horizontal and vertical compatibility relations H and V , respectively, together with the sequence of numbers s :

$$D_{\mathcal{T}} = \{H(i, j) \mid (i, j) \in H\} \cup \{V(i, j) \mid (i, j) \in V\} \cup \{S_i(s[i])\}_{i \in \{0, \dots, |s|-1\}}.$$

The set of TGDs $\Sigma_{\mathcal{T}}$. The idea underlying $\Sigma_{\mathcal{T}}$ is, during the chase, to inductively construct tilings of size $2^i \times 2^i$ from tilings of size $2^{i-1} \times 2^{i-1}$. This exploits the following simple fact, which has been already observed in [18] where Datalog with complex values is studied: the left square in the figure below of size $2^i \times 2^i$, for $i > 1$, with each x_i, y_i, z_i, w_i being of size $2^{i-2} \times 2^{i-2}$, satisfies the horizontal and vertical compatibility relations iff the nine subsquares of size $2^{i-1} \times 2^{i-1}$ depicted on the right in the following figure satisfy the compatibility relations. Let us clarify that the origin of a grid is considered to be the upper-left cell.



To achieve this construction, we encode $2^i \times 2^i$ squares, for $i > 0$, of the form

ul	ur
ll	lr

where ul, ur, ll, lr are squares of size $2^{i-1} \times 2^{i-1}$, as relational atoms of the form

$$T_i(id, ul, ur, ll, lr),$$

where id is the identity of the encoded square, and ul, ur, ll, lr are the identities of its four subsquares as shown above.

We first construct squares of size 2×2 that satisfy the compatibility relations directly from H and V stored in the database $D_{\mathcal{T}}$. This is done via the TGD

$$H(x_1, x_2), H(x_3, x_4), V(x_1, x_3), V(x_2, x_4) \rightarrow \exists z T_1(z, x_1, x_2, x_3, x_4).$$

The inductive construction of squares of size $2^i \times 2^i$, for $i \in \{2, \dots, n\}$, which satisfy the compatibility relations, from squares of size $2^{i-1} \times 2^{i-1}$, is done via the following TGDs. For each $i \in \{1, \dots, n-1\}$, we have the TGD

$$\begin{aligned} &T_i(u_1, x_1, x_2, x_3, x_4), T_i(u_2, x_2, y_1, x_4, y_3), T_i(u_3, y_1, y_2, y_3, y_4), \\ &T_i(u_4, x_3, x_4, z_1, z_2), T_i(u_5, x_4, y_3, z_2, w_1), T_i(u_6, y_3, y_4, w_1, w_2), \\ &T_i(u_7, z_1, z_2, z_3, z_4), T_i(u_8, z_2, w_1, z_4, w_3), T_i(u_9, w_1, w_2, w_3, w_4) \rightarrow \exists u T_{i+1}(u, u_1, u_3, u_7, u_9). \end{aligned}$$

We now need to verify the initial condition. To this end, we need to extract from the squares of size $2^n \times 2^n$ the tiles at positions $(0, 0), (1, 0), \dots, (|s| - 1, 0)$. This is done by defining relational atoms of the form

$$\text{Top}_i^j(x, y),$$

where $1 \leq i \leq n$ and $0 \leq j \leq |s| - 1$, to express that in the $2^i \times 2^i$ square x , at position $(j, 0)$, we have the tile $y \in \{1, \dots, m\}$. We then need to add the TGDs

$$T_1(x, x_1, x_2, x_3, x_4) \rightarrow \text{Top}_1^0(x, x_1), \text{Top}_1^1(x, x_2),$$

for each $i \in \{2, \dots, \lceil \log |s| \rceil\}$,

$$T_i(x, x_1, x_2, x_3, x_4), \text{Top}_{i-1}^0(x_1, y_0), \dots, \text{Top}_{i-1}^{2^{i-1}-1}(x_1, y_{2^{i-1}-1}) \rightarrow \text{Top}_i^0(x, y_0), \dots, \text{Top}_i^{2^{i-1}-1}(x, y_{2^{i-1}-1}),$$

$$T_i(x, x_1, x_2, x_3, x_4), \text{Top}_{i-1}^0(x_2, y_0), \dots, \text{Top}_{i-1}^{2^{i-1}-1}(x_2, y_{2^{i-1}-1}) \rightarrow \text{Top}_i^{2^{i-1}}(x, y_0), \dots, \text{Top}_i^{2^i-1}(x, y_{2^i-1}),$$

and, for each $i \in \{\lceil \log |s| \rceil + 1, \dots, n\}$,

$$T_i(x, x_1, x_2, x_3, x_4), \text{Top}_{i-1}^0(x_1, y_0), \dots, \text{Top}_{i-1}^{|s|-1}(x_1, y_{|s|-1}) \rightarrow \text{Top}_i^0(x, y_0), \dots, \text{Top}_i^{|s|-1}(x, y_{|s|-1}).$$

We finally add the TGD


$$\text{Top}_n^0(x, y_0), S_0(y_0), \dots, \text{Top}_n^{|s|-1}(x, y_{|s|-1}), S_{|s|-1}(y_{|s|-1}) \rightarrow \text{Yes}(),$$

which simply checks whether there exists a square of size $2^n \times 2^n$ that complies with H and V , and, in addition, the tile at position $(i, 0)$, for $i \in \{0, \dots, |s| - 1\}$, is $s[i]$, i.e., it checks whether an exponential tiling for \mathcal{T} has been found.


It should be clear that \mathcal{T} has an exponential tiling iff the atom $\text{Yes}()$ occurs in $\text{chase}(D_{\mathcal{T}}, \Sigma_{\mathcal{T}})$, which implies that the above construction is correct. It should be also clear that $D_{\mathcal{T}}$ and $\Sigma_{\mathcal{T}}$ can be constructed in polynomial time, while $\Sigma_{\mathcal{T}}$ mentions only predicates of arity at most five. This completes the proof of Lemma 3.2.

3.3 Stickiness

This condition, introduced in [13], is inherently different from guardedness and acyclicity. It ensures neither finite treewidth nor termination of the chase. Instead, the decidability of query answering is obtained via backward-chaining techniques. The goal of stickiness is to capture joins among variables that are not expressible via guarded TGDs, but without forcing the chase to terminate. The key property underlying this condition is that, during the chase, terms that unify with variables that appear more than once in the body of a TGD (i.e., join variables) are always propagated (or “stick”) to the inferred atoms; this is graphically illustrated as



where the first set of TGDs is sticky, while the second is not. The formal definition is based on an inductive procedure that marks the variables that may violate the semantic property described above. Roughly, during the base step of this procedure, a variable that appears in the body of a TGD σ but not in every head atom of σ is marked. Then, the marking is inductively propagated from head to body as follows



Stickiness requires every marked variable to appear only once in the body of a TGD. Let us now give the formal definition.

Consider a set Σ of TGDs; we can always assume that the TGDs in Σ do not share variables. For brevity, given an atom $R(\bar{t})$ and a variable $x \in \bar{t}$, $\text{pos}(R(\bar{t}), x)$ is the set of positions in $R(\bar{t})$ at which x occurs; a position $R[i]$ identifies the i -th attribute of the predicate R . Let $\sigma \in \Sigma$ and x a variable in the body of σ . We inductively define when x is marked in Σ as follows:

- If there is an atom $R(\bar{t}) \in \text{head}(\sigma)$ such that $x \notin \bar{t}$, then x is marked in Σ .
- Assuming that there exists an atom $R(\bar{t}) \in \text{head}(\sigma)$ such that $x \in \bar{t}$, if there is $\sigma' \in \Sigma$ that has in its body an atom of the form $R(\bar{t}')$, and each variable in $R(\bar{t}')$ at a position of $\text{pos}(R(\bar{t}), x)$ is marked in Σ , then x is marked in Σ .

The set Σ is *sticky* if there is no TGD whose body contains two occurrences of a variable that is marked in Σ . Let \mathbf{S} be the class of sticky sets of TGDs. Then:

Proposition 3.4 ([13, 23]). $\text{QAns}(\mathbf{S}_\perp)$ is *EXPTIME-complete in c-complexity, NP-complete in ba-complexity and fp-complexity, and in AC₀ in d-complexity*.

Let us clarify that in [13], where stickiness has been introduced, only the *c-complexity*, *fp-complexity*, and *d-complexity* have been considered. Moreover, the query answering algorithm devised in [13] does not provide an NP upper bound in the case of *ba-complexity*. This result is implicit in [23], where a result analogous to Lemma 3.1 is shown. In particular, there is a function $g(x, y)$, which is polynomial in x and exponential in y , such that, for every database D , set $\Sigma \in \mathbf{S}$, CQ $q(\bar{x})$, and tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$, $\bar{c} \in \text{cert}(q, D, \Sigma)$ implies there exists a sequence of instances $(I_i)_{0 \leq i \leq g(|q|+n, m)}$ with $J_0 = D$ and $I_i \langle \sigma, h \rangle I_{i+1}$ for some TGD $\sigma \in \Sigma$ and homomorphism h , where n is the number of predicates in Σ , and m the maximum arity over all those predicates, such that $\bar{c} \in q(I_{g(|q|+n, m)})$. Thus, if the arity of the underlying schema is bounded by a constant, $g(|q| + n, m)$ is a polynomial, which in turn leads to an easy guess and check algorithm for $\text{QAns}(\mathbf{S})$ that runs in polynomial time. Therefore, $\text{QAns}(\mathbf{S})$ is in NP in *ba-complexity*.

4 Consistent ontological query answering

As already discussed in Section 2, the set of certain answers to a CQ $q(\bar{x})$ w.r.t. a database D and a set Σ of TGDs and NCs that are inconsistent consists of all the tuples $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$. This exposes the main weakness of the standard certain answers semantics as defined above: the answers that we obtain from databases that are inconsistent with the given set of TGDs and NCs are not meaningful in practice. For this reason, several inconsistency-tolerant semantics have been proposed in the literature. All these inconsistency-tolerant semantics are based on the key notion of repair, which is essentially a \subseteq -maximal consistent subset of the given database.

Definition 4.1 (Repairs). Consider a database D and a set Σ of TGDs and NCs. A *repair* of D and Σ is a database $D' \subseteq D$ such that

1. $\text{mods}(D', \Sigma) \neq \emptyset$, and
2. there is no $\alpha \in D \setminus D'$ such that $\text{mods}(D' \cup \{\alpha\}, \Sigma) \neq \emptyset$.

We write $\text{reps}(D, \Sigma)$ for the set of repairs of D and Σ . ■

A simple example that illustrates the notion of repair follows. Note that this example will also serve as a running example in the rest of the section for illustrating the different inconsistency-tolerant semantics that we consider in our work.

Example 4.2. Consider the database

$$D = \{\text{Professor}(p), \text{Postdoc}(p), \text{Group}(g), \text{LeaderOf}(p, g)\},$$

asserting that p is a professor, a postdoc, and the leader of the research group g . Consider also the set Σ of TGDs and NCs consisting of

$$\begin{aligned} \text{Professor}(x) &\rightarrow \exists y \text{Researcher}(x), \text{WorksOn}(x, y), \text{Project}(y) \\ \text{Postdoc}(x) &\rightarrow \exists y \text{Researcher}(x), \text{WorksOn}(x, y), \text{Project}(y) \\ \text{LeaderOf}(x, y) &\rightarrow \text{Professor}(x), \text{Group}(y) \\ \text{Professor}(x), \text{Postdoc}(x) &\rightarrow \perp, \end{aligned}$$

expressing that professors and postdocs are researchers who work on some project, the domain (range) of the relation $\text{LeaderOf}(\cdot, \cdot)$ consists of professors (research groups), and professors and postdocs form disjoint sets. Clearly, $\text{mods}(D, \Sigma) = \emptyset$ since p violates the disjointness assertion. The repairs of D and Σ are

$$\begin{aligned} D_1 &= \{\text{Professor}(p), \text{Group}(g), \text{LeaderOf}(p, g)\} \\ D_2 &= \{\text{Postdoc}(p), \text{Group}(g)\}. \end{aligned}$$

To obtain D_1 it suffices to remove the atom $\text{Postdoc}(p)$ from D . To obtain D_2 , apart from removing $\text{Professor}(p)$, we also need to remove $\text{LeaderOf}(p, g)$, which, together with the third TGD above, implies $\text{Professor}(p)$. ■

4.1 ABox repair semantics

Having the notion of repair in place, we can now recall the main inconsistency-tolerant semantics, i.e., the ABox repair (AR) semantics [28]. Notice that this semantics has been proposed in the context of description logics, where the database is called assertional box (ABox); hence the name ABox repair. The underlying idea is very simple: a tuple is a certain answer if it is entailed by every repair.

Definition 4.3 (AR Semantics). Consider a database D , and a set Σ of TGDs and NCs with $\text{reps}(D, \Sigma) = \{D_1, \dots, D_n\}$. The *AR-certain answers* to q w.r.t. D and Σ is defined as $\text{cert}_{\text{AR}}(q, D, \Sigma) = \bigcap_{i \in \{1, \dots, n\}} \text{cert}(q, D_i, \Sigma)$. ■

A simple example that illustrates the AR semantics follows:

Example 4.4. Consider the database D and the set Σ of TGDs and NCs in Example 4.2. Consider also the Boolean CQs

$$\begin{aligned} q_1 &= \exists x \text{Group}(x) & q_2 &= \text{Researcher}(p) \\ q_3 &= \exists x \text{Project}(x) & q_4 &= \text{Professor}(p). \end{aligned}$$

The query q_1 asks whether a group exists, q_2 whether p is a researcher, q_3 whether a project exists, and q_4 whether p is a professor. Recall that $\text{reps}(D, \Sigma) = \{D_1, D_2\}$ as in Example 4.2. Observe that, for each $i \in \{1, 2\}$ and $j \in \{1, 2, 3\}$, it holds that $\text{cert}(q_j, D_i, \Sigma) \neq \emptyset$; hence, $\text{cert}_{\text{AR}}(q_j, D, \Sigma) \neq \emptyset$. On the other hand, even if $\text{cert}(q_4, D_2, \Sigma) \neq \emptyset$, $\text{cert}(q_2, D_2, \Sigma) = \emptyset$, and thus $\text{cert}_{\text{AR}}(q_4, D, \Sigma) = \emptyset$. ■

Although the AR semantics provides an elegant way to deal with inconsistency in ontological query answering, finding AR-certain answers is most commonly in coNP, and very often coNP-hard in d-complexity [28]. The reason is the fact that we need to consider many repairs, in general, exponentially many in the size of the database. This led to a large body of work on defining sound approximations of the AR semantics with the aim of reducing the complexity of query answering. Two of the most prominent such approximations are the intersection of repairs (IAR) semantics [28] and the intersection of closed repairs (ICR) semantics [3].

4.2 Intersection of ABox repairs semantics

The key idea is, instead of considering all the possible repairs, to focus on one repair that approximates all the others in a sound way. The obvious repair with this property is the one obtained by computing the intersection of all the repairs.

Definition 4.5 (IAR Semantics). Consider a database D , and a set Σ of TGDs and NCs with $\text{reps}(D, \Sigma) = \{D_1, \dots, D_n\}$. The *IAR-certain answers* to a CQ q w.r.t. D and Σ is defined as $\text{cert}_{\text{IAR}}(q, D, \Sigma) = \text{cert}(q, \bigcap_{i \in \{1, \dots, n\}} D_i, \Sigma)$. ■

Clearly, $\text{cert}_{\text{IAR}}(q, D, \Sigma) \subseteq \text{cert}_{\text{AR}}(q, D, \Sigma)$, i.e., the IAR semantics is indeed a sound approximation of the AR semantics. Here is a simple example:

Example 4.6. Consider the database D and the set Σ of TGDs and NCs in Example 4.2, and the CQs in Example 4.4. Recall that $\text{reps}(D, \Sigma) = \{D_1, D_2\}$ given in Example 4.2, and thus $D_1 \cap D_2 = \{\text{Group}(g)\}$. Hence, $\text{cert}_{\text{IAR}}(q_1, D, \Sigma) = \text{cert}(q_1, D_1 \cap D_2, \Sigma) \neq \emptyset$, while $\text{cert}_{\text{IAR}}(q_i, D, \Sigma) = \emptyset$ for each $i \in \{2, 3, 4\}$. ■

	c-complexity	ba-complexity	fp-complexity	d-complexity
G_{\perp}	2EXPTIME	EXPTIME	Π_2^P	coNP
L_{\perp}	PSPACE	Π_2^P	Π_2^P	coNP
A_{\perp}	$P^{NEXPTIME}$	$P^{NEXPTIME}$	Π_2^P	coNP
S_{\perp}	EXPTIME	Π_2^P	Π_2^P	coNP

Table 2: Complexity of $\text{QAns}_{\text{AR}}(\mathcal{C}_{\perp})$, where $\mathcal{C} \in \{\text{G}, \text{L}, \text{A}, \text{S}\}$; these are completeness results.

4.3 Intersection of closed repairs semantics

The key idea is, as in the case of the IAR semantics, to focus on one repair that approximates all the others in a sound way. This can be done in a more refined way than by considering the intersection of all repairs, which corresponds to closing the repairs w.r.t. the given set of TGDs before intersecting them. Given a database D and a set Σ of TGDs, we denote by $\text{cl}(D, \Sigma)$ the set of ground atoms that can be entailed by D and Σ , i.e., the set of atoms $\bigcap \text{mods}(D, \Sigma)$.

Definition 4.7 (ICR Semantics). Consider a database D , and a set Σ of TGDs and NCs with $\text{reps}(D, \Sigma) = \{D_1, \dots, D_n\}$. The *ICR-certain answers* to a CQ q w.r.t. D and Σ is $\text{cert}_{\text{ICR}}(q, D, \Sigma) = \text{cert}(q, \bigcap_{i \in \{1, \dots, n\}} \text{cl}(D_i, \tau(\Sigma)), \Sigma)$. ■

Here is a simple example based on our running example:

Example 4.8. Consider the database D and the set Σ of TGDs and NCs in Example 4.2, and the CQs in Example 4.4. Recall that $\text{reps}(D, \Sigma) = \{D_1, D_2\}$ as in Example 4.2; thus, $\text{cl}(D_1, \tau(\Sigma)) \cap \text{cl}(D_2, \tau(\Sigma)) = \{\text{Research}(p), \text{Group}(g)\}$. Hence, $\text{cert}_{\text{ICR}}(q_i, D, \Sigma) \neq \emptyset$ for $i \in \{1, 2\}$, and $\text{cert}_{\text{ICR}}(q_i, D, \Sigma) = \emptyset$ for $i \in \{3, 4\}$. ■

Observe that in the scenario adopted in the above simple examples, where the repairs of D and Σ are the databases D_1 and D_2 given in Example 4.2, it holds that $D_1 \cap D_2 \subset \text{cl}(D_1, \tau(\Sigma)) \cap \text{cl}(D_2, \tau(\Sigma))$, which explains the fact that more queries are entailed in the case of the ICR semantics. In general, we can show that

$$\text{cert}_{\text{IAR}}(q, D, \Sigma) \subseteq \text{cert}_{\text{ICR}}(q, D, \Sigma) \subseteq \text{cert}_{\text{AR}}(q, D, \Sigma),$$

which justifies the statement that the ICR semantics is a finer approximation of the AR semantics than the IAR semantics.

4.4 Inconsistency-tolerant ontological query answering

Having the above semantics in place, we are now ready to revisit ontological query answering in order to ensure conceptually meaningful answers. The intention is not to compute the certain answers, but the s -certain answers, where s is one of the inconsistency-tolerant semantics described above, i.e., $s \in \{\text{AR}, \text{IAR}, \text{ICR}\}$. This gives rise to the following problems; as usual, \mathcal{C} denotes a class of TGDs:

PROBLEM :	$\text{QAns}_s(\mathcal{C}_{\perp})$
INPUT :	A database D , a set $\Sigma \in \mathcal{C}_{\perp}$, a CQ $q(\bar{x})$, and $\bar{c} \in \text{dom}(D)^{ \bar{x} }$.
QUESTION :	Does $\bar{c} \in \text{cert}_s(q, D, \Sigma)$?

We may also write $\text{QAns}_s(\text{NC})$ for the problem that takes as input a database D , a set Σ of NCs (i.e., Σ does not contain any TGDs), a CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$, and asks whether $\bar{c} \in \text{cert}_s(q, D, \Sigma)$. Pinpointing the exact complexity of the above problems is the main goal of the present work. We are going to consider different complexity measures, that is, the combined complexity, the bounded-arity and fixed-program combined complexity, and the data complexity of $\text{QAns}_s(\mathcal{C}_{\perp})$, which are defined in the obvious way. Each one of the next three sections focuses on one of the semantics in question.

5 ABox repair semantics

We first focus on $\text{QAns}_{\text{AR}}(\mathcal{C}_{\perp})$, where \mathcal{C} is one of the classes of TGDs discussed above. The main result of this section follows:

Input: database D , set Σ of TGDs and NCs, CQ $q(\bar{x})$, and tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$
Output: accept, if $\bar{c} \notin \text{cert}_{\text{AR}}(q, D, \Sigma)$; otherwise, reject

```

guess a database  $D' \subseteq D$ 
if there exists  $\sigma \in \nu(\Sigma)$  such that  $\text{cert}(q_\sigma, D', \tau(\Sigma)) \neq \emptyset$  then
  | return reject
foreach  $\alpha \in D \setminus D'$  do
  | if there is no  $\sigma \in \nu(\Sigma)$  such that  $\text{cert}(q_\sigma, D' \cup \{\alpha\}, \tau(\Sigma)) \neq \emptyset$  then
  | | return reject
if  $\bar{c} \in \text{cert}(q, D', \tau(\Sigma))$  then
  | return reject
else
  | return accept

```

Algorithm 1: AlgorithmAR

Theorem 5.1. *The t -complexity of $\text{QAns}_{\text{AR}}(\mathcal{C}_\perp)$, where $t \in \{\text{c}, \text{ba}, \text{fp}, \text{d}\}$ and $\mathcal{C} \in \{\text{G}, \text{L}, \text{A}, \text{S}\}$, is as shown in Table 2.*

The rest of the section is devoted to establishing the above result. We first show, in Section 5.1, the upper bounds, and then, in Section 5.2, the lower bounds.

5.1 Upper bounds

Interestingly, all the upper bounds in Table 2 are obtained via the simple algorithm that checks whether there exists a repair that does not entail the given tuple of constants. The formal definition of this algorithm, called **AlgorithmAR**, is given in Algorithm 1. It is clear that this non-deterministic algorithm is correct. It is also easy to see that **AlgorithmAR** (D, Σ, q, \bar{c}) runs in polynomial time, assuming access to an oracle that is powerful enough for solving the problem $\text{QAns}(\mathcal{C})$, where \mathcal{C} is the class of TGDs from which $\tau(\Sigma)$ is coming from. Therefore:

Lemma 5.1. *For a class \mathcal{C} of TGDs, $\text{QAns}_{\text{AR}}(\mathcal{C}_\perp)$ is in $\text{coNP}^{\mathcal{C}}$ in t -complexity, where $t \in \{\text{c}, \text{ba}, \text{fp}, \text{d}\}$, assuming that $\text{QAns}(\mathcal{C})$ is in \mathcal{C} in t -complexity.*

The desired upper bounds given in Table 2 are obtained from Propositions 3.1, 3.2, 3.3, and 3.4, which provide the t -complexity of $\text{QAns}(\mathcal{C})$, for $t \in \{\text{c}, \text{ba}, \text{fp}, \text{d}\}$, Lemma 5.1, and the following complexity facts:

$$\text{coNP}^{\mathcal{C}} = \begin{cases} \text{coNP} & \text{if } \mathcal{C} \subseteq \text{PTIME} \\ \Pi_2^P & \text{if } \mathcal{C} = \text{coNP} \\ \mathcal{C} & \text{if } \mathcal{C} \in \{\text{PSPACE}, \text{EXPTIME}, 2\text{EXPTIME}\} \\ \text{P}^{\text{NEXPTIME}} & \text{if } \mathcal{C} = \text{NEXPTIME}. \end{cases}$$

The first three facts are actually well-known. We only need to argue why the fourth one holds. The complexity class $\text{NP}^{\text{NEXPTIME}}$ lies at a higher level of the so-called strong exponential hierarchy. We know that the strong exponential hierarchy collapses to its Δ_2 level, which implies that $\text{NP}^{\text{NEXPTIME}} = \text{P}^{\text{NEXPTIME}}$ [26]. Observe that the class $\text{P}^{\text{NEXPTIME}}$ is a deterministic one, since the oracle machines in terms of which it is defined are deterministic, and therefore $\text{coP}^{\text{NEXPTIME}} = \text{P}^{\text{NEXPTIME}}$.

5.2 Lower bounds

We now proceed to establish the complexity lower bounds claimed in Table 2. In fact, the \mathcal{C} -hardness results, where $\mathcal{C} \in \{\text{PSPACE}, \text{EXPTIME}, 2\text{EXPTIME}\}$, are coming for free, since already $\text{QAns}(\mathcal{C})$ is \mathcal{C} -hard. Therefore, to complete the picture, it suffices to establish the following results:

1. $\text{QAns}_{\text{AR}}(\text{A}_\perp)$ is $\text{P}^{\text{NEXPTIME}}$ -hard in **ba**-complexity.
2. $\text{QAns}_{\text{AR}}(\text{NC})$ is Π_2^P -hard in **fp**-complexity.

3. $\text{QAns}_{\text{AR}}(\text{NC})$ is coNP -hard in \mathbf{d} -complexity.

Notice that the last two statements refer only to negative constraints. The rest of the section is devoted to establishing the above three lower bounds.

Theorem 5.2. $\text{QAns}_{\text{AR}}(\mathbf{A}_{\perp})$ is P^{NEXPTIME} -hard in \mathbf{ba} -complexity.

Actually, the above result has been recently shown in [20] by relying on the construction, given in the proof of Lemma 3.2, for showing that $\text{QAns}(\mathbf{A})$ is NEXPTIME -hard. As we will need it later, we recall the proof of this result, which relies on a P^{NEXPTIME} -hard variation of the exponential tiling problem, introduced in [20]. An *extended tiling system* is a tuple $\mathcal{E} = (k, n, m, H_1, V_1, H_2, V_2)$, where k, n , and m are numbers in unary, and H_1, V_1, H_2 , and V_2 are subsets of $\{1, \dots, m\} \times \{1, \dots, m\}$. We say that \mathcal{E} is *valid* if the following holds: for every sequence s of length k of numbers from $\{1, \dots, m\}$, there is no exponential tiling for the tiling system $\mathcal{T}_1 = (n, m, H_1, V_1, s)$, or there is an exponential tiling for the tiling system $\mathcal{T}_2 = (n, m, H_2, V_2, s)$. The *extended exponential tiling problem* follows:

PROBLEM : `ExtendedExpTiling`
INPUT : `An extended tiling system \mathcal{E} .`
QUESTION : `Is \mathcal{E} valid?`

We are now ready to recall the proof of Theorem 5.2 given in [20].

Proof of Theorem 5.2. Let $\mathcal{E} = (k, n, m, H_1, V_1, H_2, V_2)$ be an extended tiling system. We construct a database $D_{\mathcal{E}}$, and a set $\Sigma_{\mathcal{E}} \in \mathbf{A}_{\perp}$ that mentions only predicates of bounded arity, such that \mathcal{E} is valid iff $\text{cert}_{\text{AR}}(\text{Yes}(), D_{\mathcal{E}}, \Sigma_{\mathcal{E}}) \neq \emptyset$.

The database $D_{\mathcal{E}}$. As expected, it stores the horizontal and vertical compatibility relations H_i and V_i , for $i \in \{1, 2\}$, respectively. In addition, it stores all the possible sequences of length k of numbers from $\{1, \dots, m\}$. More precisely,

$$D_{\mathcal{E}} = \{H_{\ell}(i, j) \mid (i, j) \in H_{\ell}\}_{\ell \in \{1, 2\}} \cup \{V_{\ell}(i, j) \mid (i, j) \in V_{\ell}\}_{\ell \in \{1, 2\}} \cup \{S_i^1(j), S_i^2(j)\}_{i \in \{0, \dots, k-1\}, j \in \{1, \dots, m\}} \cup \{\text{No}_1()\}.$$

The atom $\text{No}_1()$ is an auxiliary atom that will help us to check whether, for every sequence s of length k , there is no exponential tiling for (n, m, H_1, V_1, s) .

The set of TGDs and NCs $\Sigma_{\mathcal{E}}$. We first add the following NCs. For each $i \in \{0, \dots, k-1\}$, $j \in \{1, 2\}$, and $\ell, \ell' \in \{1, \dots, m\}$ such that $\ell \neq \ell'$, we have

$$S_i^j(\ell), S_i^j(\ell') \rightarrow \perp.$$

For each $i \in \{0, \dots, k-1\}$, and $\ell, \ell' \in \{1, \dots, m\}$ such that $\ell \neq \ell'$, we also have

$$S_i^1(\ell), S_i^2(\ell') \rightarrow \perp.$$

The above set of NCs guarantees that in each repair exactly one atom of the form $S_i^j(\ell)$ occurs, for each $j \in \{1, 2\}$, where $\ell \in \{1, \dots, m\}$. That is, in each repair, we keep a proper sequence s of length k of numbers from $\{1, \dots, m\}$ such that $\mathcal{T}_i = (n, m, H_i, V_i, s)$, for each $i \in \{1, 2\}$, are proper tiling systems.

We then add the TGDs $\Sigma_{\mathcal{T}_1}$ and $\Sigma_{\mathcal{T}_2}$ that encode the tiling systems \mathcal{T}_1 and \mathcal{T}_2 as defined in the proof of Lemma 3.2; $\Sigma_{\mathcal{T}_1}$ and $\Sigma_{\mathcal{T}_2}$ use different predicates. Assuming that the final TGD of $\Sigma_{\mathcal{T}_i}$, which checks for the existence of an exponential tiling for \mathcal{T}_i , has in its head the atom $\text{Yes}_i()$, we finally add the following NC and TGDs:

$$\text{Yes}_1(), \text{No}_1() \rightarrow \perp \quad \text{No}_1() \rightarrow \text{Yes}() \quad \text{Yes}_2() \rightarrow \text{Yes}().$$

The above NC ensures that $\text{No}_1()$ appears in *every* repair iff the atom $\text{Yes}_1()$ is not entailed, or, equivalently, for every sequence s of length k , the tiling system (n, m, H_1, V_1, s) does not admit an exponential tiling. It should then be clear that the following are equivalent:

1. For every $D \in \text{reps}(D_{\mathcal{E}}, \Sigma_{\mathcal{E}})$, the atom $\text{Yes}()$ occurs in $\text{chase}(D, \tau(\Sigma_{\mathcal{E}}))$.
2. For every sequence s of length k of numbers from $\{1, \dots, m\}$, there is no exponential tiling for $\mathcal{T}_1 = (n, m, H_1, V_1, s)$, or there is an exponential tiling for $\mathcal{T}_2 = (n, m, H_2, V_2, s)$, i.e., \mathcal{E} is valid.

Hence, the above reduction is correct. It is also easy to verify that $D_{\mathcal{E}}$ and $\Sigma_{\mathcal{E}}$ can be constructed in polynomial time, and $\Sigma_{\mathcal{E}}$ mentions only predicates of bounded arity. This completes the proof of Theorem 5.2. \square

Theorem 5.3. $\text{QAns}_{\text{AR}}(\text{NC})$ is Π_2^P -hard in fp-complexity.

The proof of the above result exploits the satisfiability problem for quantified Boolean formulas with two alternations of quantifiers starting with universal quantifiers (2QBF $_{\forall}$). We actually consider formulas of the form

$$\varphi = \forall x_1 \cdots \forall x_n \exists y_1 \cdots \exists y_m \psi,$$

where $\psi = C_1 \wedge \cdots \wedge C_k$ is a 3CNF formula with C_i being a clause of the form $(\ell_i^1 \vee \ell_i^2 \vee \ell_i^3)$, while each literal ℓ_i^j is either a variable or the negation of a variable. The formula φ is *satisfiable* if, for every assignment of truth values to variables x_1, \dots, x_n , there is an assignment of truth values to variables y_1, \dots, y_m such that ψ evaluates to true. The Π_2^P -hard problem of interest follows:

PROBLEM : 2QBF $_{\forall}$ -SAT
INPUT : A 2QBF $_{\forall}$ formula φ .
QUESTION : Is φ satisfiable?

We are now ready to give the proof of Theorem 5.3.

Proof of Theorem 5.3. Given a 2QBF $_{\forall}$ formula φ as above, we construct a database D_{φ} , and a BCQ q_{φ} , such that φ is satisfiable iff $\text{cert}_{\text{AR}}(q_{\varphi}, D_{\varphi}, \Sigma) \neq \emptyset$ with

$$\Sigma = \{S(x, _, z), S(_, x, z) \rightarrow \perp\},$$

where “ $_$ ” denotes a “don’t care” variable that occurs only once.

The database D_{φ} . Roughly, in D_{φ} we store, for each clause C , all the valuations that make C true. A valuation for C is a function f from the variables in C to $\{0, 1\}$. For a literal $\ell = x$ (resp., $\ell = \neg x$), $f(\ell) = f(x)$ (resp., $f(\ell) = \neg f(x)$). A valuation f satisfies a clause $C = (\ell_1 \vee \ell_2 \vee \ell_3)$ if $(f(\ell_1) \vee f(\ell_2) \vee f(\ell_3))$ evaluates to true. For a clause C , let $T(C)$ be the set of valuations for C that make C true. For a literal ℓ , we write $\text{var}(\ell)$ for its variable. The database D_{φ} is defined as

$$\bigcup_{1 \leq i \leq k} \bigcup_{f \in T(C_i)} \bigcup_{1 \leq j \leq 3} \left\{ P_i^j \left(c_i^f, f(\text{var}(\ell_i^j)) \right) \right\} \cup \bigcup_{1 \leq i \leq n} \{S(0, 1, d_i), S(1, 0, d_i)\}.$$

The atom $P_i^j(c_i^f, f(\text{var}(\ell_i^j)))$ simply states the following: according to the valuation f , the variable of the literal ℓ_i^j of C_i takes the value $f(\text{var}(\ell_i^j))$. The S -atoms are auxiliary atoms, and their purpose is explained below.

The query q_{φ} . Observe that the set $\text{reps}(D, \Sigma)$ consists of all the subsets of D that can be formed by keeping either the atom $S(0, 1, d_i)$ or the atom $S(1, 0, d_i)$, for each $i \in \{1, \dots, n\}$. In fact, each repair D' corresponds to a possible assignment $\mu_{D'}$ of truth values to the universally quantified variables of φ . More precisely, the atom $S(0, 1, d_i)$ (resp., $S(1, 0, d_i)$) states that the universally quantified variable x_i is assigned the value 0 (resp., 1). Therefore, it suffices to check whether, for every $D' \in \text{reps}(D, \Sigma)$, there exists a valuation for φ , which is compatible with $\mu_{D'}$, that makes φ true. This is achieved via the Boolean CQ

$$q_{\varphi} = \bigwedge_{1 \leq i \leq k} \bigwedge_{1 \leq j \leq 3} P_i^j(z_i, \text{var}(\ell_i^j)) \wedge \bigwedge_{1 \leq i \leq n} S(x_i, _, d_i),$$

where all the variables occurring in q_{φ} are existentially quantified. This completes the proof of Theorem 5.3. \square

Theorem 5.4. $\text{QAns}_{\text{AR}}(\text{NC})$ is coNP-hard in d-complexity.

The above result relies on a variation of the unsatisfiability problem for Boolean formulas. A $(2+2)$ Boolean formula is a CNF formula where each clause has two positive and two negative literals. The coNP-hard problem of interest follows:

	c-complexity	ba-complexity	fp-complexity	d-complexity
G_{\perp}	2EXPTIME	EXPTIME	Θ_2^P	coNP
L_{\perp}	PSPACE	Π_2^P	NP	in AC ₀
A_{\perp}	P ^{NEXPTIME}	P ^{NEXPTIME}	NP	in AC ₀
S_{\perp}	EXPTIME	Π_2^P	NP	in AC ₀

Table 3: Complexity of $\text{QAns}_{\text{IAR}}(C_{\perp})$, where $C \in \{G, L, A, S\}$. Recall that $\Theta_2^P = \text{P}^{\text{NP}[O(\log n)]}$, i.e., it collects all the problems that are solvable in polynomial time with logarithmically many calls to an NP-oracle. Apart from the AC₀ upper bounds, the rest are completeness results.

PROBLEM :	2+2UNSAT
INPUT :	A (2 + 2) Boolean formula φ .
QUESTION :	Is φ unsatisfiable?

We are now ready to give the proof of Theorem 5.4.

Proof of Theorem 5.4. Given a (2 + 2) formula $\varphi = C_1 \wedge \dots \wedge C_n$ over the variables x_1, \dots, x_m , we define the database D_{φ} as follows:

$$\{\text{Pos}_1(i, x_i^1), \text{Pos}_2(i, x_i^2), \text{Neg}_1(i, x_i^3), \text{Neg}_2(i, x_i^4) \mid C_i = x_i^1 \vee x_i^2 \vee x_i^3 \vee x_i^4\} \\ \cup \{\text{True}(x_i), \text{False}(x_i) \mid 1 \leq i \leq m\},$$

which essentially stores the formula φ , and also assigns to each variable in φ both the value 1 and the value 0. It is not difficult to verify that φ is unsatisfiable iff $\text{cert}_{\text{AR}}(q, D_{\varphi}, \Sigma) \neq \emptyset$, where

$$\Sigma = \{\text{True}(x), \text{False}(x) \rightarrow \perp\}$$

and q is the Boolean CQ

$$\exists x \exists y_1 \dots \exists y_4 (\text{Pos}_1(x, y_1) \wedge \text{False}(y_1) \wedge \text{Pos}_2(x, y_2) \wedge \text{False}(y_2) \wedge \\ \text{Neg}_1(x, y_3) \wedge \text{True}(y_3) \wedge \text{Neg}_2(x, y_4) \wedge \text{True}(y_4)).$$

It is clear that, for each variable x of φ , a repair of $\text{reps}(D, \Sigma)$ keeps either the atom $\text{True}(x)$ or the atom $\text{False}(x)$, i.e., each $D' \in \text{reps}(D, \Sigma)$ corresponds to a possible assignment $\mu_{D'}$ of truth values to the variables of φ . The query q checks that each such assignment evaluates φ to false, which is actually done by checking that at least one clause of φ evaluates to false. \square

6 Intersection of repairs semantics

We now concentrate on $\text{QAns}_{\text{IAR}}(C_{\perp})$, where C is one of the classes of TGDs discussed above. The main result of this section follows:

Theorem 6.1. *The t-complexity of $\text{QAns}_{\text{IAR}}(C_{\perp})$, where $t \in \{c, \text{ba}, \text{fp}, \text{d}\}$ and $C \in \{G, L, A, S\}$, is as shown in Table 3.*

The rest of the section is devoted to establishing the above result. We first show, in Section 6.1, the upper bounds, and then, in Section 6.2, the lower bounds.

6.1 Upper bounds

Although for the ABox repair semantics we were able to establish all the upper bounds (see Table 2) in a uniform way via the algorithm `AlgorithmAR`, this is not the case for the intersection of repairs semantics. We are able, however, to partition the cells of Table 3 into four groups, and establish the upper bounds claimed in the cells of each such group in a uniform way. The four groups are as follows:

1. The c-complexity and the ba-complexity for C_{\perp} , where $C \in \{G, L, A, S\}$, as well as the d-complexity for G_{\perp} .

Input: database D , set Σ of TGDs and NCs, CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$
Output: accept, if $\bar{c} \notin \text{cert}_{\text{IAR}}(q, D, \Sigma)$; otherwise, reject

```

guess a database  $D^* \subseteq D$ 
foreach  $\alpha \in D \setminus D^*$  do
  guess a database  $D_\alpha \subseteq D$ 
  if  $\alpha \in D_\alpha$  then
    | return reject
  else
    | foreach  $\beta \in D \setminus D_\alpha$  do
      | | if there is no  $\sigma \in \nu(\Sigma)$  s.t.  $\text{cert}(q_\sigma, D_\alpha \cup \{\beta\}, \tau(\Sigma)) \neq \emptyset$  then
        | | | return reject
if  $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$  then
  | return reject
else
  | return accept

```

Algorithm 2: AlgorithmIAR1

2. The fp-complexity for G_\perp .
3. The fp-complexity for C_\perp , where $C \in \{L, A, S\}$.
4. The d-complexity for C_\perp , where $C \in \{L, A, S\}$.

We proceed to give more details for each of the above groups.

6.1.1 The c-complexity and the ba-complexity for C_\perp , where $C \in \{G, L, A, S\}$, as well as the d-complexity for G_\perp

The upper bounds are obtained via the simple procedure AlgorithmIAR1, depicted in Algorithm 2, that checks whether there exists a superset of the intersection of repairs that does not entail the given tuple \bar{c} of constants. More precisely, the algorithm guesses a subset D^* of the input database D , and then checks that for every atom $\alpha \in D \setminus D^*$, there exists $D_\alpha \in \text{reps}(D, \Sigma)$, where Σ is the input set of TGDs and NCs, such that $\alpha \notin D_\alpha$, and thus, α is not in the intersection of repairs. This implies that D^* is a superset of the intersection of repairs. Finally, the algorithm rejects if $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$; otherwise, it accepts. Indeed, this algorithm is correct due to the following lemma:

Lemma 6.1. *Consider a database D , a set Σ of TGDs and NCs, a CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$. The following are equivalent:*

1. $\bar{c} \notin \text{cert}_{\text{IAR}}(q, D, \Sigma)$.
2. There exists $D^* \supseteq \bigcap_{D' \in \text{reps}(D, \Sigma)} D'$ such that $\bar{c} \notin \text{cert}(q, D^*, \tau(\Sigma))$.

Proof. The fact that (1) implies (2) holds trivially with D^* being exactly the intersection of repairs. The other direction follows by the monotonicity of the CQ q , i.e., for every two instances I_1 and I_2 , $I_1 \subseteq I_2$ implies $q(I_1) \subseteq q(I_2)$. \square

It is also easy to see that the non-deterministic algorithm AlgorithmIAR1 runs in polynomial time, assuming access to an oracle that can solve $\text{QAns}(C)$, where C is the class of TGDs from which the input set of TGDs is coming from. Therefore:

Lemma 6.2. *For a class C of TGDs, $\text{QAns}_{\text{IAR}}(C_\perp)$ is in coNP^C in t -complexity, where $t \in \{c, \text{ba}, d\}$, assuming that $\text{QAns}(C)$ is in C in t -complexity.*

Since, by Lemma 6.1, AlgorithmIAR1 is correct, the desired upper bounds for Group 1 given in Table 3 are obtained from Propositions 3.1, 3.2, 3.3 and 3.4, Lemma 6.2, and the following complexity facts that

Input: database D , set Σ of TGDs and NCs, CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$
Output: accept if $\bar{c} \in \text{cert}_{\text{IAR}}(q, D, \Sigma)$; otherwise, reject

```

 $D^* := D$ 
foreach  $\alpha \in D$  do
  | if there exists  $D_\alpha \in \text{reps}(D, \Sigma)$  such that  $\alpha \notin D_\alpha$  then
  | |  $D^* := D^* \setminus \{\alpha\}$ 
if  $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$  then
  | return accept
else
  | return reject

```

Algorithm 3: AlgorithmIAR2

have been already discussed in the previous section and we recall again for the sake of readability:

$$\text{coNP}^{\mathcal{C}} = \begin{cases} \text{coNP} & \text{if } \mathcal{C} \subseteq \text{PTIME} \\ \Pi_2^P & \text{if } \mathcal{C} = \text{coNP} \\ \mathcal{C} & \text{if } \mathcal{C} \in \{\text{PSPACE}, \text{EXPTIME}, 2\text{EXPTIME}\} \\ \text{P}^{\text{NEXPTIME}} & \text{if } \mathcal{C} = \text{NEXPTIME}. \end{cases}$$

6.1.2 The fp-complexity for G_\perp

We need to establish a $\Theta_2^P = \text{P}^{\text{NP}[O(\log n)]}$ upper bound. To this end, we exploit the procedure AlgorithmIAR2, depicted in Algorithm 3, that constructs the intersection of repairs D^* , and accepts if the given tuple \bar{c} belongs to $\text{cert}(q, D^*, \tau(\Sigma))$; otherwise, it rejects. The intersection of repairs D^* is constructed by starting from the input database D , and removing all the atoms α for which there exists at least one repair $D_\alpha \in \text{reps}(D, \Sigma)$ such that $\alpha \notin D_\alpha$. More precisely, D^* is constructed via polynomially many parallel calls to an NP-oracle. In fact, for each atom $\alpha \in D$, we call in parallel an NP-oracle that does the following:

1. Guess a database $D_\alpha \subseteq D$.
2. If $\alpha \in D_\alpha$, then reject.
3. For each atom $\beta \in D \setminus D_\alpha$, if there is no $\sigma \in \nu(\Sigma)$ such that $\text{cert}(q_\sigma, D_\alpha \cup \{\beta\}, \tau(\Sigma)) \neq \emptyset$, then return reject; otherwise; return accept.

It is crucial to clarify that the check whether $\text{cert}(q_\sigma, D_\alpha \cup \{\beta\}, \tau(\Sigma)) \neq \emptyset$ is feasible in polynomial time since q_σ and $\tau(\Sigma)$ are fixed – recall that we are studying the fp-complexity of $\text{QAns}_{\text{IAR}}(G_\perp)$, while, by Proposition 3.1, $\text{QAns}(G)$ is in PTIME in d-complexity. Therefore, the above oracle is indeed an NP-oracle. It is clear that, for an atom $\alpha \in D$, if the above oracle returns **accept**, then there exists a repair D_α such that $\alpha \notin D_\alpha$, and thus, α does not belong to the intersection of repairs. Consequently, the intersection of repairs D^* is constructed by simply removing from D all the atoms α for which the oracle returns **accept**. Since D^* can be constructed in polynomial time via parallel NP-oracle calls, we can conclude that it can also be constructed in polynomial time via logarithmically many NP-oracle calls; see, e.g., [42]. Once we have D^* in place, we need one more call to an NP-oracle for checking whether $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$; the latter is indeed in NP since, by Proposition 3.1, $\text{QAns}(G)$ is in NP in fp-complexity. Hence, we get the desired $\Theta_2^P = \text{P}^{\text{NP}[O(\log n)]}$ upper bound.

6.1.3 The fp-complexity for C_\perp , where $C \in \{L, A, S\}$

We need to establish an NP upper bound. A crucial notion that we are going to exploit is that of culprit, which is essentially a minimal inconsistent subset of a database D w.r.t. a set Σ of TGDs and NCs. Formally, a *culprit* of D w.r.t. Σ is a subset D' of D such that the following conditions hold:

1. $\text{mods}(D', \Sigma) = \emptyset$, and

Input: database D , set Σ of TGDs and NCs, CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$
Output: accept if $\bar{c} \in \text{cert}_{\text{IAR}}(q, D, \Sigma)$; otherwise, reject

$D^* := D \setminus \bigcup_{D' \in \text{culprit}(D, \Sigma)} D'$
if $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$ **then**
 | **return** accept
else
 | **return** reject

Algorithm 4: AlgorithmIAR3

2. there is no $D'' \subsetneq D'$ such that $\text{mods}(D'', \Sigma) = \emptyset$.

Intuitively, a culprit is a minimal subset of D that, together with $\tau(\Sigma)$, entails some NC $\sigma \in \nu(\Sigma)$; a culprit for σ is a “minimal explanation” [14] of q_σ . We write $\text{culprit}(D, \Sigma)$ for the set of all culprits of D w.r.t. Σ . By deleting from D a *minimal* set of facts S intersecting *all* culprits from $\text{culprit}(D, \Sigma)$,⁶ we obtain a repair $R = D \setminus S$. By this, it is an easy exercise to show that the intersection of the repairs of D w.r.t. Σ is precisely the subset of D obtained after eliminating its culprits (w.r.t. Σ):

Lemma 6.3. *Consider a database D , and a set Σ of TGDs and NCs. It holds that*

$$\bigcap_{D' \in \text{reps}(D, \Sigma)} D' = D \setminus \bigcup_{D' \in \text{culprit}(D, \Sigma)} D'.$$

From the above lemma, we immediately get the algorithm AlgorithmIAR3, depicted in Algorithm 4, that explicitly constructs the intersection of repairs D^* , and accepts if the given tuple \bar{c} belongs to $\text{cert}(q, D^*, \tau(\Sigma))$; otherwise, it rejects. It remains to argue how we get the desired NP upper bounds. To this end, it suffices to show that $\text{culprit}(D, \Sigma)$ can be computed in polynomial time when the set Σ is fixed and falls in \mathbf{C}_\perp , where $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$. In such a case, D^* can be computed in polynomial time, and the claim follows by Propositions 3.2, 3.3 and 3.4, which state that $\text{QAns}(\mathbf{C})$ for $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$ is in NP in fp-complexity.

The key property of the classes in question that we are going to exploit is *UCQ-rewritability* shown in [24]: given a set of TGDs $\Sigma \in \mathbf{C}$ and a CQ $q(\bar{x})$, we can construct a UCQ $Q_{q, \Sigma}(\bar{x})$ such that, for *every* database D and tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$,

$$\bar{c} \in \text{cert}(q, D, \Sigma) \iff \bar{c} \in Q_{q, \Sigma}(D).$$

Consider now a set $\Sigma \in \mathbf{C}_\perp$, where $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$. By exploiting UCQ-rewritability, we can first “embed” the TGDs of $\tau(\Sigma)$ into the NCs of $\nu(\Sigma)$, which leads to a set of NCs denoted $\langle \Sigma \rangle$, and then compute the culprits of a database D w.r.t. Σ by collecting all the images of the so-called minimal specializations of the NCs of $\langle \Sigma \rangle$ in D via injective mappings. Let us formalize the above discussion.

- Let Q_Σ be the UCQ $\bigcup_{\sigma \in \nu(\Sigma)} Q_{q_\sigma, \tau(\Sigma)}$; recall that q_σ denotes the Boolean CQ that corresponds to σ . The *embedding* of $\tau(\Sigma)$ into $\nu(\Sigma)$ is the set

$$\langle \Sigma \rangle = \{\phi(\bar{x}) \rightarrow \perp \mid \exists \bar{x} \phi(\bar{x}) \in Q_\Sigma\}.$$

- A *specialization* of a NC σ is a NC obtained from σ by identifying some of the variables occurring in $\text{body}(\sigma)$. For example, $R(x, y, x), S(x) \rightarrow \perp$ is a specialization of $R(x, y, z), S(z) \rightarrow \perp$ obtained by identifying x and z . Notice that a NC is trivially a specialization of itself. We write $\text{sp}(\sigma)$ for the set of all specializations of a NC σ , and for a set of NCs Σ' we define $\text{sp}(\Sigma')$ as the set $\bigcup_{\sigma \in \Sigma'} \text{sp}(\sigma)$. Moreover, we define $\text{msp}(\Sigma')$ as the largest subset of $\text{sp}(\Sigma')$ such that, for every $\sigma \in \text{msp}(\Sigma')$, there is no $\sigma' \in \text{msp}(\Sigma')$ with $\text{body}(\sigma') \subsetneq \text{body}(\sigma)$ (up to variable renaming). Clearly, $\text{msp}(\Sigma')$ is unique (up to variable renaming). Finally, for a database D , let

$$I_{D, \Sigma} = \{h(\text{body}(\sigma)) \mid \sigma \in \text{msp}(\langle \Sigma \rangle) \text{ and } h \text{ is an injective}$$

homomorphism from $\text{body}(\sigma)$ to $D\}$.

⁶Such a deletion S is a (hypergraph) transversal of $\text{culprit}(D, \Sigma)$, and hence S is minimal iff, for each fact $f \in S$, there is a culprit $C \in \text{culprit}(D, \Sigma)$ such that $C \cap S = \{f\}$ [16, 22].

We proceed to show the following:

Lemma 6.4. *For a database D and a set $\Sigma \in \mathbf{C}_\perp$, where $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$, it holds that $\text{culprit}(D, \Sigma) = I_{D, \Sigma}$.*

Proof. (\subseteq) Consider an arbitrary $C \in \text{culprit}(D, \Sigma)$. There is $\sigma \in \langle \Sigma \rangle$ and a homomorphism h from $\text{body}(\sigma)$ to C , and for every $\sigma' \in \langle \Sigma \rangle$, there is no homomorphism from $\text{body}(\sigma')$ to a strict subset of C . Let σ_h be the NC with $\text{body}(\sigma_h)$ being the conjunction of atoms obtained from C by converting each constant c into a variable x_c . We claim that $\sigma_h \in \text{msp}(\langle \Sigma \rangle)$ (up to variable renaming), which in turn implies that $C \in I_{D, \Sigma}$. By contradiction, assume that this is not the case. It is clear that σ_h is a specialization of σ (up to variable renaming), and thus $\sigma_h \in \text{sp}(\langle \Sigma \rangle)$ (up to variable renaming). This implies that there exists $\sigma' \in \text{sp}(\langle \Sigma \rangle)$ such that $\text{body}(\sigma') \subsetneq \text{body}(\sigma)$ (up to variable renaming). Therefore the NC $\sigma'' \in \langle \Sigma \rangle$ from which σ' is obtained via specialization can be mapped via a homomorphism to a strict subset of C , which is a contradiction.

(\supseteq) Consider now an arbitrary $C \in I_{D, \Sigma}$. By definition, there exists $\hat{\sigma} \in \text{msp}(\langle \Sigma \rangle)$ and an injective mapping \hat{h} that maps $\text{body}(\hat{\sigma})$ to C . By contradiction, assume that $C \notin \text{culprit}(D, \Sigma)$. Suppose that $\hat{\sigma}$ is the specialized version of $\sigma \in \langle \Sigma \rangle$. Clearly, there exists a homomorphism from $\text{body}(\sigma)$ to C . Hence, the fact that $C \notin \text{culprit}(D, \Sigma)$ allows us to conclude that there exists $\sigma' \in \langle \Sigma \rangle$ and a homomorphism h' from $\text{body}(\sigma')$ to a strict subset C' of C . Consider the NC σ'' with $\text{body}(\sigma'')$ being the conjunction of atoms obtained from C' by converting each constant c into a variable x_c . Clearly, $\sigma'' \in \text{sp}(\langle \Sigma \rangle)$ (up to variable renaming). But since $\hat{h}(\text{body}(\hat{\sigma})) = C$ with \hat{h} being an injective homomorphism, we can conclude that $\text{body}(\sigma'') \subsetneq \text{body}(\hat{\sigma})$ (up to variable renaming). But this contradicts the fact $\hat{\sigma} \in \text{msp}(\langle \Sigma \rangle)$. Therefore, $C \in \text{culprit}(D, \Sigma)$, and the claim follows. \square

The fact that for a fixed set $\Sigma \in \mathbf{C}_\perp$, where $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$, the culprits of a database D w.r.t. Σ can be computed in polynomial time is an easy consequence of Lemma 6.4. Indeed, we can compute in constant time the set of NCs $\text{msp}(\langle \Sigma \rangle)$. This in turn allows us to compute the set of databases $I_{D, \Sigma}$ in polynomial time in the size of D , which, by Lemma 6.4, coincides with $\text{culprit}(D, \Sigma)$.

6.1.4 The d-complexity for \mathbf{C}_\perp , where $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$

Our goal is to establish an AC_0 upper bound. To this end, we are going to show the following technical result, which essentially states that computing the IAR-certain answers to a CQ in the case of linear, acyclic and sticky sets of TGDs boils down to evaluating a first-order query over the input database.

Lemma 6.5. *Consider a set of TGDs and NCs $\Sigma \in \mathbf{C}_\perp$, where $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$, and a CQ $q(\bar{x})$. We can construct a first-order query $\Phi_{q, \Sigma}(\bar{x})$ such that, for every database D , $\text{cert}_{\text{IAR}}(q, D, \Sigma) = \Phi_{q, \Sigma}(D)$.*

Having the above lemma in place, it is straightforward to obtain the desired AC_0 upper bound. The key fact is that the first-order query $\Phi_{q, \Sigma}$ does not depend on the input database. Therefore, for a fixed set Σ of TGDs and NCs, and a fixed CQ $q(\bar{x})$, we can construct the first-order query $\Phi_{q, \Sigma}(\bar{x})$ in constant time. Then, given a database D and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$, we simply need to check whether $\bar{c} \in \Phi_{q, \Sigma}(D)$. The latter is in AC_0 since first-order query evaluation is in AC_0 in data complexity. It remains to establish Lemma 6.5.

We first establish the following useful auxiliary result:

Lemma 6.6. *Consider a set of TGDs and NCs $\Sigma \in \mathbf{C}_\perp$, where $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$, and a CQ $q(\bar{x})$. We can construct a first-order query $\Psi_{q, \Sigma}(\bar{x})$ such that, for every database D , $q\left(D \setminus \bigcup_{D' \in \text{culprit}(D, \Sigma)} D'\right) = \Psi_{q, \Sigma}(D)$.*

Proof. We assume that $q(\bar{x}) = \exists \bar{y} \phi(\bar{x}, \bar{y})$. We also assume, w.l.o.g., that the NCs of $\langle \Sigma \rangle$ and q do not share variables. We define $\Psi_{q, \Sigma}(\bar{x})$ as

$$\exists \bar{y} \left(\phi(\bar{x}, \bar{y}) \wedge \bigwedge_{\psi(\bar{z}) \rightarrow \perp \in \text{msp}(\langle \Sigma \rangle)} \forall \bar{z} \left(\psi(\bar{z}) \wedge \bigwedge_{v, w \in \bar{z}, v \neq w} v \neq w \rightarrow \bigwedge_{\alpha \in \phi(\bar{x}, \bar{y}), \beta \in \psi(\bar{z})} \alpha \neq \beta \right) \right).$$

Given a database D , $\Psi_{q, \Sigma}(D)$ is the set of tuples $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$ such that $q(\bar{c})$ is mapped to D via a homomorphism h , while $h(D)$ does not contain an atom of $I_{D, \Sigma}$, and thus, by Lemma 6.4, $h(D)$ does not contain an atom of $\text{culprit}(D, \Sigma)$. In other words, $\Psi_{q, \Sigma}(D) = q\left(D \setminus \bigcup_{D' \in \text{culprit}(D, \Sigma)} D'\right)$, and the claim follows. \square

We are now ready to give the proof of Lemma 6.5. We claim that the desired first-order query $\Phi_{q,\Sigma}(\bar{x})$ is the query

$$\bigvee_{q' \in Q_{q,\tau(\Sigma)}} \Psi_{q',\Sigma}(\bar{x}),$$

where $\Psi_{q',\Sigma}(\bar{x})$ is the first-order query provided by Lemma 6.6. Given a database D , by Lemma 6.3 and UCQ-rewritability, we get that

$$\text{cert}_{\text{IAR}}(q, D, \Sigma) = \bigcup_{q' \in Q_{q,\tau(\Sigma)}} q' \left(D \setminus \bigcup_{D' \in \text{culprit}(D,\Sigma)} D' \right).$$

Since $\Psi_{q',\Sigma}(D)$ is precisely $q' \left(D \setminus \bigcup_{D' \in \text{culprit}(D,\Sigma)} D' \right)$, the claim follows.

6.2 Lower bounds

We now proceed to establish the complexity lower bounds claimed in Table 3. The \mathcal{C} -hardness results, where $\mathcal{C} \in \{\text{NP}, \text{PSPACE}, \text{EXPTIME}, 2\text{EXPTIME}\}$, are coming for free since already $\text{QAns}(\mathcal{C})$ is \mathcal{C} -hard. Therefore, to complete the picture, it suffices to establish the following hardness results:

1. $\text{QAns}_{\text{IAR}}(\mathbf{A}_{\perp})$ is $\text{P}^{\text{NEXP TIME}}$ -hard in **ba**-complexity.
2. $\text{QAns}_{\text{IAR}}(\text{NC})$ is Π_2^P -hard in **ba**-complexity.
3. $\text{QAns}_{\text{IAR}}(\mathbf{G}_{\perp})$ is Θ_2^P -hard in **fp**-complexity.
4. $\text{QAns}_{\text{IAR}}(\mathbf{G}_{\perp})$ is **coNP**-hard in **d**-complexity.

Notice that the second statement refers only to negative constraints. The rest of the section is devoted to establishing the above four lower bounds.

Theorem 6.2. $\text{QAns}_{\text{IAR}}(\mathbf{A}_{\perp})$ is $\text{P}^{\text{NEXP TIME}}$ -hard in **ba**-complexity.

Proof. The proof is via a reduction from the extended exponential tiling problem, which has been used for showing Theorem 5.2, that is, the $\text{P}^{\text{NEXP TIME}}$ -hardness in **ba**-complexity of $\text{QAns}_{\text{AR}}(\mathbf{A}_{\perp})$. In fact, the proof is an adaptation of the proof of Theorem 5.2. Recall that in the proof of Theorem 5.2, given an exponential tiling system $\mathcal{E} = (k, n, m, H_1, V_1, H_2, V_2)$, we construct the database

$$D_{\mathcal{E}} = \{H_{\ell}(i, j) \mid (i, j) \in H_{\ell}\}_{\ell \in \{1,2\}} \cup \{V_{\ell}(i, j) \mid (i, j) \in V_{\ell}\}_{\ell \in \{1,2\}} \cup \{S_i^1(j), S_i^2(j)\}_{i \in \{0, \dots, k-1\}, j \in \{1, \dots, m\}} \cup \{\text{No}_1()\}$$

and a set $\Sigma_{\mathcal{E}} \in \mathbf{A}_{\perp}$, which mentions only predicates of bounded arity, such that

$$\mathcal{E} \text{ is valid} \iff \text{cert}_{\text{AR}}(\text{Yes}(), D_{\mathcal{E}}, \Sigma_{\mathcal{E}}) \neq \emptyset. \quad (3)$$

For showing Theorem 6.2, we adapt $D_{\mathcal{E}}$ and $\Sigma_{\mathcal{E}}$ as follows:

$$\begin{aligned} \hat{D}_{\mathcal{E}} &= D_{\mathcal{E}} \cup \{\text{Yes}(), \text{No}_2()\} \\ \hat{\Sigma}_{\mathcal{E}} &= \Sigma_{\mathcal{E}} \cup \\ &\quad \{\phi(\bar{x}), \text{Yes}(), \text{No}() \rightarrow \perp, \\ &\quad \text{Yes}_2(), \text{No}_2() \rightarrow \perp, \\ &\quad \text{Yes}_1(), \text{No}_2() \rightarrow \text{No}()\}, \end{aligned}$$

where $\phi(\bar{x})$ is the conjunction of atoms

$$\bigwedge_{\ell \in \{1,2\}} \left(\bigwedge_{(i,j) \in H_{\ell}} H_{\ell}(i, j) \wedge \bigwedge_{(i,j) \in V_{\ell}} V_{\ell}(i, j) \wedge \bigwedge_{i \in \{0, \dots, k-1\}} S_i^{\ell}(x_i) \right).$$

We first show that

$$\text{cert}_{\text{AR}}(\text{Yes}(), D_{\mathcal{E}}, \Sigma_{\mathcal{E}}) \neq \emptyset \iff \text{cert}_{\text{AR}}(\text{Yes}(), \hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset. \quad (4)$$

For the direction (\Rightarrow), assume that $\text{cert}_{\text{AR}}(\text{Yes}(), \hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}}) = \emptyset$. This implies that there exists $D' \in \text{reps}(\hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}})$ such that $\text{cert}(\text{Yes}(), D', \hat{\Sigma}_{\mathcal{E}}) = \emptyset$. It is clear that $\text{Yes}() \notin D'$. Observe that $\text{cert}(q, D', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$, where q is the Boolean CQ $\exists \bar{x} \phi(\bar{x}), \text{No}()$, since otherwise $\text{mods}(D' \cup \text{Yes}(), \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$, which contradicts the fact that $D' \in \text{reps}(\hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}})$. This in turn implies $\text{cert}(\text{Yes}_1(), D', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$. Let $D'' = D' \setminus \{\text{No}_2()\}$. We proceed to show that $D'' \in \text{reps}(D_{\mathcal{E}}, \Sigma_{\mathcal{E}})$ and $\text{cert}(\text{Yes}(), D'', \Sigma_{\mathcal{E}}) = \emptyset$, which in turn implies that $\text{cert}_{\text{AR}}(\text{Yes}(), D_{\mathcal{E}}, \Sigma_{\mathcal{E}}) = \emptyset$, as needed. It is clear that $D'' \subseteq D_{\mathcal{E}}$. Moreover, since $D'' \subseteq D'$ and $\Sigma_{\mathcal{E}} \subseteq \hat{\Sigma}_{\mathcal{E}}$, we immediately get that $\text{cert}(\text{Yes}(), D'', \Sigma_{\mathcal{E}}) = \emptyset$ (since $\text{cert}(\text{Yes}(), D', \hat{\Sigma}_{\mathcal{E}}) = \emptyset$), and $\text{mods}(D'', \Sigma_{\mathcal{E}}) \neq \emptyset$ (since $\text{mods}(D', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$). To show that $D'' \in \text{reps}(D_{\mathcal{E}}, \Sigma_{\mathcal{E}})$, it remains to show that there is no $\alpha \in D_{\mathcal{E}} \setminus D''$ such that $\text{mods}(D'' \cup \{\alpha\}, \Sigma_{\mathcal{E}}) \neq \emptyset$. We first observe that in $D_{\mathcal{E}} \setminus D''$ we only have atoms of the form $S_i^j(\ell)$ for $i \in \{0, \dots, k-1\}$, $j \in \{1, 2\}$ and $\ell \in \{1, \dots, m\}$, and the atom $\text{No}_1()$. It is clear that adding to D'' an atom α of the form $S_i^j(\ell)$ leads to inconsistency, i.e., $\text{mods}(D'' \cup \{\alpha\}, \Sigma_{\mathcal{E}}) = \emptyset$. Observe now that $\text{cert}(\text{Yes}_1(), D', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$ implies $\text{cert}(\text{Yes}_1(), D'', \Sigma_{\mathcal{E}}) \neq \emptyset$ since the facts (resp., the TGDs and NCs) in $D' \setminus D''$ (resp., $\hat{\Sigma}_{\mathcal{E}} \setminus \Sigma_{\mathcal{E}}$) do not affect the entailment of $\text{Yes}_1()$. Therefore, $\text{cert}(\text{Yes}_1(), D'' \cup \{\text{No}_1()\}, \Sigma_{\mathcal{E}}) \neq \emptyset$, which in turn implies that $\text{mods}(D'' \cup \{\text{No}_1()\}, \Sigma_{\mathcal{E}}) = \emptyset$. We conclude that $D'' \in \text{reps}(D_{\mathcal{E}}, \Sigma_{\mathcal{E}})$, and the claim follows.

For the direction (\Leftarrow), assume that $\text{cert}_{\text{AR}}(\text{Yes}(), D_{\mathcal{E}}, \Sigma_{\mathcal{E}}) = \emptyset$, which, by equivalence (3), implies that \mathcal{E} is not valid. This means that there exists $D' \in \text{reps}(D_{\mathcal{E}}, \Sigma_{\mathcal{E}})$ such that $\text{cert}(\text{Yes}_1(), D', \Sigma_{\mathcal{E}}) \neq \emptyset$ and $\text{cert}(\text{Yes}_2(), D', \Sigma_{\mathcal{E}}) = \emptyset$. Observe that $\text{cert}(\text{No}_1(), D', \Sigma_{\mathcal{E}}) = \emptyset$ because otherwise $\text{mods}(D', \Sigma_{\mathcal{E}}) = \emptyset$, which contradicts the fact that $D' \in \text{reps}(D_{\mathcal{E}}, \Sigma_{\mathcal{E}})$. Let $D'' = D' \cup \{\text{No}_2()\}$. We proceed to show that $D'' \in \text{reps}(\hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}})$ and $\text{cert}(\text{Yes}(), D'', \hat{\Sigma}_{\mathcal{E}}) = \emptyset$, which in turn implies that $\text{cert}_{\text{AR}}(\text{Yes}(), \hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}}) = \emptyset$, as needed. It is clear that $D'' \subseteq \hat{D}_{\mathcal{E}}$. Moreover, it can be easily verified that adding $\text{No}_2()$ to D' , and adding $\hat{\Sigma}_{\mathcal{E}} \setminus \Sigma_{\mathcal{E}}$ to $\Sigma_{\mathcal{E}}$, do not affect the non-entailment of $\text{No}_1()$ and $\text{Yes}_2()$, and thus, $\text{cert}(\text{No}_1(), D'', \hat{\Sigma}_{\mathcal{E}}) = \emptyset$ and $\text{cert}(\text{Yes}_2(), D'', \hat{\Sigma}_{\mathcal{E}}) = \emptyset$. As a consequence, we have that $\text{cert}(\text{Yes}(), D'', \hat{\Sigma}_{\mathcal{E}}) = \emptyset$. The last three statements allow us also to conclude that $\text{mods}(D'', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$. To show that $D'' \in \text{reps}(\hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}})$, it remains to show that there is no $\alpha \in \hat{D}_{\mathcal{E}} \setminus D''$ such that $\text{mods}(D'' \cup \{\alpha\}, \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$. Notice that in $\hat{D}_{\mathcal{E}} \setminus D''$ we only have atoms of the form $S_i^j(\ell)$ for $i \in \{0, \dots, k-1\}$, $j \in \{1, 2\}$ and $\ell \in \{1, \dots, m\}$, and the atoms $\text{No}_1()$ and $\text{Yes}()$. It is clear that adding to D'' an atom α of the form $S_i^j(\ell)$ leads to inconsistency, i.e., $\text{mods}(D'' \cup \{\alpha\}, \hat{\Sigma}_{\mathcal{E}}) = \emptyset$. Since $\text{cert}(\text{Yes}_1(), D', \Sigma_{\mathcal{E}}) \neq \emptyset$, $D' \subseteq D''$, and $\Sigma_{\mathcal{E}} \subseteq \hat{\Sigma}_{\mathcal{E}}$, we get that $\text{cert}(\text{Yes}_1(), D'', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$, which implies that $\text{mods}(D'' \cup \{\text{No}_1()\}, \hat{\Sigma}_{\mathcal{E}}) = \emptyset$. Finally, we show that $\text{cert}(\text{No}(), D'', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$ and $\text{cert}(q, D'', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$, where q is the Boolean CQ $\exists \bar{x} \phi(\bar{x})$, which imply that $\text{mods}(D'' \cup \{\text{Yes}()\}, \hat{\Sigma}_{\mathcal{E}}) = \emptyset$, as needed. The fact that $\text{cert}(\text{No}(), D'', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$ follows from $\text{cert}(\text{No}_2(), D'', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$ since $\text{No}_2() \in D''$, and $\text{cert}(\text{Yes}_1(), D'', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$ since $\text{cert}(\text{Yes}_1(), D', \Sigma_{\mathcal{E}}) \neq \emptyset$, $D' \subseteq D''$ and $\Sigma_{\mathcal{E}} \subseteq \hat{\Sigma}_{\mathcal{E}}$. The fact that $\text{cert}(q, D'', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$ follows from $\text{cert}(q, D', \Sigma_{\mathcal{E}}) \neq \emptyset$, $D' \subseteq D''$ and $\Sigma_{\mathcal{E}} \subseteq \hat{\Sigma}_{\mathcal{E}}$.

It is also not difficult to verify that

$$\text{cert}_{\text{AR}}(\text{Yes}(), \hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset \iff \text{cert}_{\text{IAR}}(\text{Yes}(), \hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset. \quad (5)$$

The direction (\Leftarrow) holds trivially, since the IAR semantics is an approximation of the AR semantics. For the direction (\Rightarrow), we observe that, for each repair $D' \in \text{reps}(\hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}})$ such that $\text{cert}(\text{Yes}(), D', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$, it holds that $\text{Yes}() \in D'$. Since, by hypothesis, for each $D' \in \text{reps}(\hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}})$, we have that $\text{cert}(\text{Yes}(), D', \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$, we conclude that each repair of $\text{reps}(\hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}})$ contains the atom $\text{Yes}()$. Therefore, $\bigcap_{D' \in \text{reps}(\hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}})} D'$ contains $\text{Yes}()$, and $\text{cert}_{\text{IAR}}(\text{Yes}(), \hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset$ follows.

By putting together the equivalences (3), (4), and (5), we conclude that

$$\mathcal{E} \text{ is valid} \iff \text{cert}_{\text{IAR}}(\text{Yes}(), \hat{D}_{\mathcal{E}}, \hat{\Sigma}_{\mathcal{E}}) \neq \emptyset,$$

which shows the correctness of our reduction, and Theorem 6.2 follows. \square

Theorem 6.3. $\text{QAns}_{\text{IAR}}(\text{NC})$ is Π_2^P -hard in ba-complexity.

The proof of the above result exploits a variant of 2QBF $_{\forall}$ -SAT. A 2NQBFF $_{\forall}$ formula is a 2QBF $_{\forall}$ formula

$$\varphi = \forall x_1 \dots \forall x_n \exists y_1 \dots \exists y_m \psi,$$

where ψ is a 3CNF formula of the form

$$\bigwedge_{1 \leq i \leq k} C_i \wedge \bigwedge_{1 \leq i \leq n} (C_i^+ \wedge C_i^-)$$

with $C_i = (\ell_i^1 \vee \ell_i^2 \vee \ell_i^3)$, $C_i^+ = (x_i \vee \neg y_{t(i)})$ and $C_i^- = (\neg x_i \vee y_{t(i)})$ for some $t(i) \in \{1, \dots, m\}$, and for each $i \in \{1, \dots, n\}$, the universally quantified variable x_i occurs only in the clauses C_i^+ and C_i^- . In other words, the latter says that, for each $i \in \{1, \dots, n\}$, the truth value of the variable $y_{t(i)}$ is determined by the value of x_i . The Π_2^P -hard problem of interest follows [25, 45]:

PROBLEM : 2NQBF $_{\forall}$ -SAT
INPUT : A 2NQBF $_{\forall}$ formula φ .
QUESTION : Is φ satisfiable?

We are now ready to give the proof of Theorem 6.3.

Proof of Theorem 6.3. For a 2NQBF $_{\forall}$ formula φ as defined above, our goal is to construct a database D_φ and a set $\Sigma_\varphi \in \text{NC}$, which mention only predicates of bounded arity, such that φ is satisfiable iff $\text{cert}_{\text{IAR}}(\text{Sat}(), D_\varphi, \Sigma_\varphi) \neq \emptyset$.

The database D_φ . Our intention is to store (i) the values that a universally quantified variable can take, (ii) an auxiliary atom $\text{Sat}()$, which indicates that φ is satisfiable, (iii) all the satisfying assignments for each clause C_i (not C_i^+ or C_i^-) of φ , (iv) auxiliary atoms that would allow us to force the existentially quantified variable $y_{t(i)}$, for $i \in \{1, \dots, n\}$, to take the same value as x_i , and (v) ‘‘consistency’’ atoms that would allow us to ensure that an assignment to the existentially quantified variables of φ is consistent among its clauses, i.e., a variable is assigned the same value in every clause that it appears. The formal definition of D_φ follows:

$$\begin{aligned} & \bigcup_{1 \leq i \leq n} \{\text{Value}(x_i, 0), \text{Value}(x_i, 1)\} \cup \{\text{Sat}()\} \\ & \bigcup_{1 \leq i \leq k} \bigcup_{\substack{b_1, b_2, b_3 \in \{0,1\}, \\ b_1 \vee b_2 \vee b_3 = 1}} \{\text{Clause}(c_i, \ell_i^1, b_1, \ell_i^2, b_2, \ell_i^3, b_3)\} \cup \\ & \bigcup_{1 \leq i \leq n} \left(\bigcup_{b \in \{0,1\}} \{\text{Force}(x_i, b, y_{t(i)}, b)\} \cup \bigcup_{\substack{b_1, b_2 \in \{0,1\} \\ b_1 \oplus b_2 = 1}} \{\text{Force}(x_i, b_1, \neg y_{t(i)}, b_2)\} \right) \cup \\ & \bigcup_{1 \leq i \leq m} \left(\bigcup_{b \in \{0,1\}} \{\text{Cons}(y_i, b, y_i, b), \text{Cons}(\neg y_i, b, \neg y_i, b)\} \cup \right. \\ & \left. \bigcup_{\substack{b_1, b_2 \in \{0,1\} \\ b_1 \oplus b_2 = 1}} \{\text{Cons}(y_i, b_1, \neg y_i, b_2), \text{Cons}(\neg y_i, b_1, y_i, b_2)\} \right). \end{aligned}$$

The set Σ_φ of NCs. This set consists of three NCs. Note that, in what follows, we use x_i and y_j for the actual constants used in the database D_φ in order to represent the variables of φ . To avoid notational clutter, for variables we will use only the symbols z and w (possibly with subscripts and superscripts). The first NC of Σ_φ simply states that a universally quantified variable can take only one value:

$$\text{Value}(z, 0), \text{Value}(z, 1) \rightarrow \perp.$$

The second NC encodes the satisfiability of φ once an assignment to the universally quantified variables has been fixed (which is provided by a repair due to the NC above). Before defining this NC, let us introduce some auxiliary conjunctions of atoms, which will eventually give rise to the desired NC. The first one is

$$\text{Config} = \bigwedge_{\alpha \in \text{Struct}} \alpha,$$

where $\text{Struct} = D_\varphi \setminus (\{\text{Sat}()\} \cup \{\text{Value}(x_i, 0), \text{Value}(x_i, 1)\}_{1 \leq i \leq n})$. The second conjunction is defined as

$$\forall \text{Assign} = \bigwedge_{1 \leq i \leq n} \text{Value}(x_i, w_i),$$

which “reads” the assignment to the universally quantified variables. The third conjunction aims at “copying” the assignment for the universally quantified variables to the associated (according to C_i^+ and C_i^-) existentially quantified variables:

$$\text{Copy} = \bigwedge_{1 \leq i \leq n} \bigwedge_{\substack{1 \leq j \leq k, 1 \leq r \leq 3, \\ \text{var}(\ell_j^r) = y_{t(i)}}} \text{Force}(x_i, w_i, z_j^r, w_j^r).$$

The fourth conjunction is defined as

$$\exists\text{Consistency} = \bigwedge_{\substack{1 \leq j_1, j_2 \leq k, \\ 1 \leq r_1, r_2 \leq 3, \\ \text{var}(\ell_{j_1}^{r_1}) = \text{var}(\ell_{j_2}^{r_2})}} \text{Cons}(z_{j_1}^{r_1}, w_{j_1}^{r_1}, z_{j_2}^{r_2}, w_{j_2}^{r_2}),$$

which states that an assignment to the existentially quantified variables is consistent among the clauses of φ . Finally, the fifth conjunction is defined as

$$\text{Satisfied} = \bigwedge_{1 \leq i \leq k} \text{Clause}(c_i, z_i^1, w_i^1, z_i^2, w_i^2, z_i^3, w_i^3),$$

which simply encodes the fact the φ is satisfiable. Having the above conjunctions in place, the desired NC is defined as

$$\text{Config}, \forall\text{Assign}, \text{Copy}, \exists\text{Consistency}, \text{Satisfied} \rightarrow \perp.$$

We also add to Σ_φ the NC

$$\text{Config}, \forall\text{Assign}, \text{Sat}() \rightarrow \perp.$$

This completes the construction of Σ_φ . We proceed to show that the above reduction is correct, i.e., φ is satisfiable iff $\text{cert}_{\text{IAR}}(\text{Sat}(), D_\varphi, \Sigma_\varphi) \neq \emptyset$.

(\Rightarrow) Consider an arbitrary repair $D \in \text{reps}(D_\varphi, \Sigma_\varphi)$. It suffices to show that $\text{Sat}() \in D$. We proceed by considering the following two cases on the shape of D :

1. There exists a universally quantified variable x_i such that none of the atoms $\text{Value}(x_i, b)$, for $b \in \{0, 1\}$, occurs in D , which means that D does not assign a value to x_i . In this case, $\text{Sat}()$ necessarily belongs to D ; otherwise, D is not a repair since adding $\text{Sat}()$ would not violate any of the NCs.
2. Assume now that D assigns a value to every universally quantified variable of φ . By hypothesis, φ is satisfiable, which allows us (by definition of Σ_φ) to conclude that $\text{Struct} \setminus D \neq \emptyset$; otherwise, the body of the second NC is satisfied, which cannot be the case since D is a repair. But then $\text{Sat}()$ necessarily belongs to D because otherwise D is not a repair since adding $\text{Sat}()$ would not violate any of the NCs; in particular, the third NC would not be violated as $\text{Struct} \setminus D \neq \emptyset$.

(\Leftarrow) Assume now that φ is not satisfiable. Thus, there exists an assignment μ to the universally quantified variables such that, for every (valid) assignment to the existentially quantified variables, φ is not satisfied. Let D be the subset of D_φ that keeps only one Value-atom for each universally quantified variable as dictated by μ , and all the other atoms of D_φ apart from $\text{Sat}()$. It should be clear that D satisfies Σ_φ . In particular, the second NC is satisfied since φ is not satisfiable, while the third NC is satisfied since $\text{Sat}()$ is not in D . Moreover, D is maximal since by adding either $\text{Sat}()$ or a Value-atom, one of the NCs will be violated. Therefore, $D \in \text{reps}(D_\varphi, \Sigma_\varphi)$, which in turn implies that $\text{cert}_{\text{IAR}}(\text{Sat}(), D_\varphi, \Sigma_\varphi) = \emptyset$. \square

Remark. Interestingly, the above proof applies even if we consider the AR semantics. Thus, we get an alternative proof for the fact that $\text{QAns}_{\text{AR}}(\text{NC})$ is Π_2^P -hard in ba-complexity. However, the proof given in Section 5 shows that the Π_2^P -hardness holds even for fp-complexity, but it exploits a more complex CQ.

Theorem 6.4. $\text{QAns}_{\text{IAR}}(\text{G}_\perp)$ is Θ_2^P -hard in fp-complexity.

The proof of the above result relies on a Θ_2^P -hard variant of 3SAT that involves counting of satisfiable formulas [32, 33]. For a set A of 3CNF formulas, we write $\#A$ for the cardinality of $\{\varphi \in A \mid \varphi \text{ is satisfiable}\}$. The problem follows:

PROBLEM : Comp3SAT
INPUT : Two sets A and B of 3CNF formulas.
QUESTION : Is $\#A > \#B$?

The above problem remains Θ_2^P -hard even if we pose several simplifying assumptions on A and B . In particular, we can assume that $|A| = |B|$, all formulas in A and B are over the same set of variables and have the same number of clauses, and $A = \{\varphi_1, \dots, \varphi_m\}$, $B = \{\psi_1, \dots, \psi_m\}$ are such that φ_{i+1} (resp., ψ_{i+1}) is satisfiable implies φ_i (resp., ψ_i) is satisfiable, for $i \in \{1, \dots, m-1\}$. It should be clear, due to the last assumption, that $\#A > \#B$ iff there exists $i \in \{1, \dots, m\}$ such that φ_i is satisfiable and ψ_i is unsatisfiable.

We are now ready to proceed with the proof of Theorem 6.4.

Proof of Theorem 6.4. Given $A = \{\varphi_1, \dots, \varphi_m\}$ and $B = \{\psi_1, \dots, \psi_m\}$, we are going to construct a database $D_{A,B}$ and a Boolean CQ $q_{A,B}$ such that the following are equivalent:

1. There exists $i \in \{1, \dots, m\}$ such that φ_i is satisfiable and ψ_i is unsatisfiable.
2. $\text{cert}_{\text{IAR}}(q_{A,B}, D_{A,B}, \Sigma) \neq \emptyset$ for some fixed set $\Sigma \in \mathbf{G}_\perp$.

Our goal is to devise $D_{A,B}$ and $q_{A,B}$ in a modular way, where the parts that are coming from A are independent from those that are coming from B . To this end, we are first going to construct a database D_A and a CQ $q_A(x)$ such that, for each $i \in \{1, \dots, m\}$, φ_i is satisfiable iff $\langle f_i \rangle \in q_A(D_A)$; the constant f_i should be understood as the identifier for the formula φ_i . Moreover, we are going to construct a database D_B and a CQ $q_B(x)$ such that, for each $i \in \{1, \dots, m\}$, ψ_i is unsatisfiable iff $\langle f_i \rangle \in \text{cert}_{\text{IAR}}(q_B, D_B, \Sigma)$ for some fixed set $\Sigma \in \mathbf{G}_\perp$; here, f_i acts as the identifier for the formula ψ_i . Once we have the above databases and CQs in place, it will be easy to construct $D_{A,B}$ and $q_{A,B}$. In the sequel, we assume that all the formulas in A and B are over the variables x_1, \dots, x_n and have k clauses.

The database D_A and the CQ q_A . Given a 3CNF formula $\varphi_i = C_{i,1} \wedge \dots \wedge C_{i,k}$ from A with $C_{i,j} = (\ell_{i,j}^1 \vee \ell_{i,j}^2 \vee \ell_{i,j}^3)$, the database D_A stores all the truth assignments that make a certain clause of φ_i true. To this end, we use an 8-ary predicate AClause. For example, given the clause $C_{i,j} = x_{i_1} \vee x_{i_2} \vee \neg x_{i_3}$,

$$\text{AClause}(f_i, c_j, x_{i_1}, 1, x_{i_2}, 0, \neg x_{i_3}, 1)$$

encodes the clause itself (recall that f_i is the identifier of φ_i , while c_j is the identifier of $C_{i,j}$), and at the same time encodes the truth assignment that sets x_{i_1} to true, x_{i_2} to false, and x_{i_3} to false (i.e., $\neg x_{i_3}$ to true). Formally, let $D_{A,1}$ be the database

$$\bigcup_{1 \leq i \leq m} \bigcup_{1 \leq j \leq k} \bigcup_{\substack{b_1, b_2, b_3 \in \{0,1\}, \\ b_1 \vee b_2 \vee b_3 = 1}} \{ \text{AClause}(f_i, c_j, \ell_{i,j}^1, b_1, \ell_{i,j}^2, b_2, \ell_{i,j}^3, b_3) \}.$$

To check whether a formula φ_i is satisfiable, we need to check whether each of its clauses is satisfiable; this is the purpose of the CQ q_A given below. To this end, we need a mechanism that allows us to ensure that an assignment for φ_i is consistent among its clauses, i.e., a variable is assigned the same value in every clause that it appears. This can be done via the following ‘‘consistency’’ atoms that form the database $D_{A,2}$:

$$\bigcup_{1 \leq i \leq n} \bigcup_{b \in \{0,1\}} \{ \text{Cons}(x_i, b, x_i, b), \text{Cons}(\neg x_i, b, \neg x_i, b) \} \cup \\ \bigcup_{1 \leq i \leq n} \bigcup_{\substack{b_1, b_2 \in \{0,1\}, \\ b_1 \oplus b_2 = 1}} \{ \text{Cons}(x_i, b_1, \neg x_i, b_2), \text{Cons}(\neg x_i, b_1, x_i, b_2) \}.$$

The database D_A is defined as the union $D_{A,1} \cup D_{A,2}$.

Let us now define the CQ $q_A(x)$. Its purpose is essentially to check whether there exists $\varphi_i \in A$ that is satisfiable, which can be done by checking that each clause of φ_i is satisfiable. This can be easily achieved via the CQ $q_A(x)$ defined below, where all the involved variables, apart from x , are existentially quantified; recall that $\text{var}(\ell)$ is the variable of the literal ℓ :

$$\bigwedge_{1 \leq j \leq k} \text{AClause}(x, c_j, y_j^1, z_j^1, y_j^2, z_j^2, y_j^3, z_j^3) \wedge \bigwedge_{1 \leq i \leq m} \bigwedge_{\substack{1 \leq j_1, j_2 \leq k, \\ 1 \leq r_1, r_2 \leq 3, \\ \text{var}(\ell_{i,j_1}^{r_1}) = \text{var}(\ell_{i,j_2}^{r_2})}} \text{Cons}(y_{j_1}^{r_1}, z_{j_1}^{r_1}, y_{j_2}^{r_2}, z_{j_2}^{r_2}).$$

This completes the definition of q_A . By construction, we get that:

Lemma 6.7. *For each $i \in \{1, \dots, m\}$, φ_i is satisfiable iff $\langle f_i \rangle \in q_A(D_A)$.*

The database D_B and the CQ q_B . Given a 3CNF formula $\psi_i = C_{i,1} \wedge \dots \wedge C_{i,k}$ from B with $C_{i,j} = (\ell_{i,j}^1 \vee \ell_{i,j}^2 \vee \ell_{i,j}^3)$, the database D_B assigns to the variables of ψ_i both the values true and false, and it also stores all the clauses of ψ_i . The latter is achieved via a 5-ary predicate $\text{BClause}^{s_1 s_2 s_3}$, where $s_1, s_2, s_3 \in \{\mathbf{p}, \mathbf{n}\}$. For example, the clause $C_{i,j} = x_{i_1} \vee x_{i_2} \vee \neg x_{i_3}$ is encoded via the atom

$$\text{BClause}^{\text{ppn}}(f_i, c_j, x_{i_1}, x_{i_2}, x_{i_3}),$$

with the superscript ppn indicating that the variable of the first (resp., second, third) literal appears positively (resp., positively, negatively) in $C_{i,j}$. Moreover, D_B stores some auxiliary atoms that would allow us to check (via a fixed set $\Sigma \in \mathbf{G}_\perp$) whether ψ_i is unsatisfiable. We proceed to formally define D_B .

For a literal ℓ , let $\text{sign}(\ell) = \mathbf{p}$ (resp., $\text{sign}(\ell) = \mathbf{n}$) if $\ell = x$ (resp., $\ell = \neg x$) for a variable x . For a clause $C_{i,j}$ of ψ_i , we write $s_{i,j}^r$ for $\text{sign}(\ell_{i,j}^r)$, where $r \in \{1, 2, 3\}$. Recall that $\text{var}(\ell)$ is the variable of the literal ℓ . The database D_B is defined as

$$\begin{aligned} & \bigcup_{1 \leq i \leq m} \bigcup_{1 \leq j \leq n} \{\text{True}(f_i, x_j), \text{False}(f_i, x_j)\} \cup \\ & \bigcup_{1 \leq i \leq m} \bigcup_{1 \leq j \leq k} \left\{ \text{BClause}^{s_{i,j}^1, s_{i,j}^2, s_{i,j}^3}(f_i, c_j, \text{var}(\ell_{i,j}^1), \text{var}(\ell_{i,j}^2), \text{var}(\ell_{i,j}^3)) \right\} \cup \\ & \bigcup_{1 \leq i \leq m} \bigcup_{0 \leq j \leq k-1} \{\text{SuccCl}(f_i, c_j, c_{j+1})\} \cup \\ & \bigcup_{1 \leq i \leq m} \{\text{MinCl}(f_i, c_0), \text{MaxCl}(f_i, c_k), \text{Unsat}(f_i)\}. \end{aligned}$$

The sequence of atoms $(\text{SuccCl}(f_i, c_j, c_{j+1}))_{0 \leq j \leq k-1}$ essentially tells us that in ψ_i the clause $C_{i,j}$ comes immediately after the clause $C_{i,j-1}$. The fact that the first atom of the sequence refers to the clause $C_{i,0}$, which does not exist, is a technicality that will become clear below. The remaining atoms give us access to the (virtually) first clause $C_{i,0}$ and the last clause $C_{i,k}$ of ψ_i , and also state that ψ_i is unsatisfiable.

The CQ $q_B(x)$, is defined as the atomic query $\text{Unsat}(x)$, which simply asks whether there exists a formula in B that is unsatisfiable.

We claim that there exists a set $\Sigma \in \mathbf{G}_\perp$ such that, for each $i \in \{1, \dots, m\}$, ψ_i is unsatisfiable iff $\langle f_i \rangle \in \text{cert}_{\text{IAR}}(q_B, D_B, \Sigma)$. In particular, $\Sigma = \Sigma_{\text{cons}} \cup \Sigma_{\text{sat}}$ with Σ_{cons} being a set of NCs that perform a consistency check (i.e., a variable is either true or false, and a formula is either satisfiable or unsatisfiable), and Σ_{sat} being a set of guarded TGDs that evaluates each formula $\psi_i \in B$ and derives the atom $\text{Sat}(f_i)$ if ψ_i is satisfiable. More precisely, Σ_{cons} consists of the NCs:

$$\begin{aligned} \text{True}(x, y), \text{False}(x, y) & \rightarrow \perp \\ \text{Sat}(x), \text{Unsat}(x) & \rightarrow \perp. \end{aligned}$$

The set Σ_{sat} consists of the following TGDs; a \star symbol is a placeholder for \mathbf{p} or \mathbf{n} , while, as usual, $_$ is a “don’t care” variable that occurs only once:

$$\begin{aligned} \text{BClause}^{\mathbf{p}\star\star}(x, y, z, _, _) , \text{True}(x, z) & \rightarrow \text{SatCl}(x, y) \\ \text{BClause}^{\mathbf{n}\star\star}(x, y, z, _, _) , \text{False}(x, z) & \rightarrow \text{SatCl}(x, y) \\ \text{BClause}^{\star\mathbf{p}\star}(x, y, _, z, _) , \text{True}(x, z) & \rightarrow \text{SatCl}(x, y) \\ \text{BClause}^{\star\mathbf{n}\star}(x, y, _, z, _) , \text{False}(x, z) & \rightarrow \text{SatCl}(x, y) \\ \text{BClause}^{\star\star\mathbf{p}}(x, y, _, _, z) , \text{True}(x, z) & \rightarrow \text{SatCl}(x, y) \\ \text{BClause}^{\star\star\mathbf{n}}(x, y, _, _, z) , \text{False}(x, z) & \rightarrow \text{SatCl}(x, y) \\ \text{MinCl}(x, y) & \rightarrow \text{SatChain}(x, y) \\ \text{SatChain}(x, y), \text{SuccCl}(x, y, z), \text{SatCl}(x, z) & \rightarrow \text{SatChain}(x, z) \\ \text{MaxCl}(x, y), \text{SatChain}(x, y) & \rightarrow \text{Sat}(x). \end{aligned}$$

This completes the construction of Σ . By construction, we get that:

Lemma 6.8. For $i \in \{1, \dots, m\}$, ψ_i is unsatisfiable iff $\langle f_i \rangle \in \text{cert}_{\text{IAR}}(q_B, D_B, \Sigma)$.

The database $D_{A,B}$ and the CQ $q_{A,B}$. We can now easily construct the database $D_{A,B}$ and the Boolean CQ $q_{A,B}$ with the desired property:

$$D_{A,B} = D_A \cup D_B \quad \text{and} \quad q_{A,B} = \exists x(q_A(x) \wedge q_B(x)).$$

Indeed, we can show the following, where $\Sigma \in \mathbf{G}_\perp$ is the set devised above:

Lemma 6.9. *The following are equivalent:*

1. There exists $i \in \{1, \dots, m\}$ such that φ_i is satisfiable and ψ_i is unsatisfiable.
2. $\text{cert}_{\text{IAR}}(q_{A,B}, D_{A,B}, \Sigma) \neq \emptyset$.

Proof. We rely on the following key observation, which is easy to verify since D_A is always consistent with Σ :

$$\bigcap_{D \in \text{reps}(D_{A,B}, \Sigma)} D = D_A \cup \left(\bigcap_{D \in \text{reps}(D_B, \Sigma)} D \right). \quad (6)$$

We can now proceed with the proof of the claim.

(1) \Rightarrow (2). By Lemma 6.7 and 6.8, we get that $\langle f_i \rangle \in q_A(D_A)$ and $\langle f_i \rangle \in \text{cert}_{\text{IAR}}(q_B, D_B, \Sigma)$. Therefore, by (6), $\langle f_i \rangle \in \text{cert}_{\text{IAR}}(q_{A,B}, D_{A,B}, \Sigma)$, which in turn implies that $\text{cert}_{\text{IAR}}(q_{A,B}, D_{A,B}, \Sigma) \neq \emptyset$.

(2) \Rightarrow (1). By hypothesis, there exists $i \in \{1, \dots, m\}$ such that $\langle f_i \rangle \in \text{cert}_{\text{IAR}}(q_{A,B}, D_{A,B}, \Sigma)$. By the equality (6), we can conclude that $\langle f_i \rangle \in q_A(D_A)$ and $\langle f_i \rangle \in \text{cert}_{\text{IAR}}(q_B, D_B, \Sigma)$. Therefore, by Lemma 6.7 and 6.8, we get that φ_i is satisfiable and ψ_i is unsatisfiable, and the claim follows. \square

With Lemma 6.9 in place, we can conclude that

$$\#A > \#B \iff \text{cert}_{\text{IAR}}(q_{A,B}, D_{A,B}, \Sigma) \neq \emptyset$$

for a fixed $\Sigma \in \mathbf{G}_\perp$, and thus, $\text{QAns}_{\text{IAR}}(\mathbf{G}_\perp)$ is Θ_2^P -hard in fp-complexity. \square

Theorem 6.5. $\text{QAns}_{\text{IAR}}(\mathbf{G}_\perp)$ is coNP-hard in d-complexity.

The proof of the above result exploits the unsatisfiability problem of Boolean formulas in negation normal form. A Boolean formula is in *negation normal form* (NNF) if it uses only \neg , \wedge and \vee , and \neg is only applied to variables. The coNP-hard problem of interest follows:

PROBLEM : NNF-UNSAT
INPUT : A Boolean formula φ in NNF.
QUESTION : Is φ unsatisfiable?

We can now proceed with the proof of Theorem 6.5.

Proof of Theorem 6.5. Given a formula φ in NNF over the variables x_1, \dots, x_m , we define the database D_φ as follows:

$$\begin{aligned} & \{\text{And}(\psi, \psi_1, \psi_2) \mid \psi = \psi_1 \wedge \psi_2 \text{ is a subformula of } \varphi\} \\ \cup & \{\text{Or}(\psi, \psi_1, \psi_2) \mid \psi = \psi_1 \vee \psi_2 \text{ is a subformula of } \varphi\} \\ \cup & \{\text{Not}(\psi, \psi') \mid \psi = \neg\psi' \text{ is a subformula of } \varphi\} \\ \cup & \{\text{True}(x_i), \text{False}(x_i) \mid 1 \leq i \leq m\} \cup \{\text{Unsat}(\varphi)\}, \end{aligned}$$

which essentially stores the formula φ , it assigns to each variable in φ both the value 1 and the value 0, and it states that φ is unsatisfiable.

	c-complexity	ba-complexity	fp-complexity	d-complexity
G_{\perp}	2EXPTIME	EXPTIME	Θ_2^P	coNP
L_{\perp}	PSPACE	Π_2^P	Θ_2^P	coNP
A_{\perp}	$P^{NEXPTIME}$	$P^{NEXPTIME}$	Θ_2^P	coNP
S_{\perp}	EXPTIME	Π_2^P	Θ_2^P	coNP

Table 4: Complexity of $QAns_{ICR}(C_{\perp})$, where $C \in \{G, L, A, S\}$. These are completeness results.

It is not difficult to show that φ is unsatisfiable iff $\text{cert}_{IAR}(\text{Unsat}(\varphi), D_{\varphi}, \Sigma) \neq \emptyset$, where Σ consists of the guarded TGDs

$$\begin{aligned}
\text{And}(x, y, z), \text{True}(y), \text{True}(z) &\rightarrow \text{True}(x) \\
\text{Or}(x, y, z), \text{True}(y) &\rightarrow \text{True}(x) \\
\text{Or}(x, y, z), \text{True}(z) &\rightarrow \text{True}(x) \\
\text{Not}(x, y), \text{False}(y) &\rightarrow \text{True}(x),
\end{aligned}$$

which are responsible for evaluating the formula φ , and the NCs

$$\begin{aligned}
\text{True}(x), \text{False}(x) &\rightarrow \perp \\
\text{True}(x), \text{Unsat}(x) &\rightarrow \perp
\end{aligned}$$

with the obvious meaning. We proceed to show that the above is a reduction.

(\Rightarrow) Assume that $\text{cert}_{IAR}(\text{Unsat}(\varphi), D_{\varphi}, \Sigma) = \emptyset$. Hence, there exists $D' \in \text{reps}(D, \Sigma)$ such that $\text{Unsat}(\varphi) \notin D'$. By definition of repairs, $\text{True}(\varphi) \in D'$. Therefore, D' encodes a satisfying assignment for φ ; simply set x_i to 1 (resp., 0) if $\text{True}(x_i) \in D'$ (resp., $\text{False}(x_i) \in D'$). Thus, φ is satisfiable, as needed.

(\Leftarrow) Conversely, assume that φ is satisfiable. Thus, there is $D' \in \text{reps}(D, \Sigma)$ such that $\text{Unsat}(\varphi) \notin D'$; otherwise, the NC $\text{True}(x), \text{Unsat}(x) \rightarrow \perp$ would be violated. Hence, $\text{Unsat}(\varphi) \notin \bigcap_{D' \in \text{reps}(D, \Sigma)} D'$, which in turn implies that $\text{cert}_{IAR}(\text{Unsat}(\varphi), D_{\varphi}, \Sigma) = \emptyset$, and the claim follows. \square

7 Intersection of closed repairs semantics

We now concentrate on $QAns_{ICR}(C_{\perp})$, where C is one of the classes of TGDs in question. The main result of this section follows:

Theorem 7.1. *The t-complexity of $QAns_{ICR}(C_{\perp})$, where $t \in \{c, ba, fp, d\}$ and $C \in \{G, L, A, S\}$, is as shown in Table 4.*

The rest of the section is devoted to establishing the above result. We first show, in Section 7.1, the upper bounds, and then, in Section 7.2, the lower bounds.

7.1 Upper bounds

We can partition the cells of Table 4 into five groups in such a way that the claimed upper bounds can be established in a uniform way:

1. The c-complexity for C_{\perp} , where $C \in \{G, S\}$.
2. The c-complexity for A_{\perp} .
3. The c-complexity for L_{\perp} .
4. The ba-complexity and the d-complexity for C_{\perp} , where $C \in \{G, L, A, S\}$.
5. The fp-complexity for C_{\perp} , where $C \in \{G, L, A, S\}$.

We proceed to give more details for each of the above groups.

Input: database D , set Σ of TGDs and NCs, CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$
Output: accept if $\bar{c} \in \text{cert}_{\text{ICR}}(q, D, \Sigma)$; otherwise, reject

```

 $D^* := \emptyset$ 
foreach  $\alpha \in \text{B}(D, \Sigma)$  do
  | if  $\text{cert}_{\text{AR}}(\alpha, D, \Sigma) \neq \emptyset$  then
  | |  $D^* := D^* \cup \{\alpha\}$ 
if  $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$  then
  | return accept
else
  | return reject

```

Algorithm 5: AlgorithmICR1

7.1.1 The \mathfrak{c} -complexity for \mathcal{C}_\perp , where $\mathcal{C} \in \{\mathbf{G}, \mathbf{S}\}$

The upper bounds are obtained via the procedure AlgorithmICR1, depicted in Algorithm 5, which constructs the intersection of closed repairs D^* , and accepts if the given tuple \bar{c} belongs to $\text{cert}(q, D^*, \Sigma)$; otherwise, it rejects. The intersection of closed repairs D^* is constructed by keeping from $\text{B}(D, \Sigma)$, i.e., the set of all ground atoms that can be formed using constants from $\text{dom}(D)$ and predicates occurring in Σ , only the atoms α that belong to $\text{cl}(D', \tau(\Sigma))$ for each $D' \in \text{reps}(D, \Sigma)$, or, equivalently, for which $\text{cert}_{\text{AR}}(\alpha, D, \Sigma) \neq \emptyset$. The fact that $\text{B}(D, \Sigma)$ consists of exponentially many atoms, allows us to conclude that, for a class \mathcal{C} of TGDs, if $\text{QAns}_{\text{AR}}(\mathcal{C}_\perp)$ is in $\mathcal{C} \in \{\text{EXPTIME}, 2\text{EXPTIME}\}$ in \mathfrak{c} -complexity (and thus, $\text{QAns}(\mathcal{C})$ is in \mathcal{C} in \mathfrak{c} -complexity), and the complexity bound inherited from the algorithm underlying the membership of $\text{QAns}(\mathcal{C})$ in \mathcal{C} in \mathfrak{c} -complexity depends polynomially on the input database, then AlgorithmICR1 shows that $\text{QAns}_{\text{ICR}}(\mathcal{C}_\perp)$ is also in \mathcal{C} in \mathfrak{c} -complexity. By Theorem 5.1, $\text{QAns}_{\text{AR}}(\mathbf{G}_\perp)$ is in 2EXPTIME , and $\text{QAns}_{\text{AR}}(\mathbf{S}_\perp)$ is in EXPTIME in \mathfrak{c} -complexity. Moreover, we know from [10] that the complexity bound inherited from the algorithm underlying the fact that $\text{QAns}(\mathbf{G})$ is in 2EXPTIME in \mathfrak{c} -complexity depends polynomially on the input database. The same holds for the class \mathbf{S} [13], and the desired upper bounds follow.

7.1.2 The \mathfrak{c} -complexity for \mathbf{A}_\perp

For showing that $\text{QAns}_{\text{ICR}}(\mathbf{A}_\perp)$ is in $\text{P}^{\text{NEXPTIME}}$ in \mathfrak{c} -complexity, we rely again on AlgorithmICR1, but we need a more refined complexity analysis than the one given above for the classes \mathbf{G}_\perp and \mathbf{S}_\perp . Since (i) $\text{B}(D, \Sigma)$ consists of exponentially many atoms, (ii) $\text{QAns}_{\text{AR}}(\mathbf{A}_\perp)$ is in $\text{P}^{\text{NEXPTIME}}$ in \mathfrak{c} -complexity by Theorem 5.1, (iii) $\text{QAns}(\mathbf{A})$ is in NEXPTIME in \mathfrak{c} -complexity by Proposition 3.3, and (iv) the complexity bound inherited from the algorithm underlying the fact that $\text{QAns}(\mathbf{A})$ is in NEXPTIME in \mathfrak{c} -complexity depends polynomially on the input database, AlgorithmICR1 allows us to conclude that $\text{QAns}_{\text{ICR}}(\mathbf{A}_\perp)$ is in $\text{NEXPTIME}^{\text{NEXPTIME}}$ in \mathfrak{c} -complexity. We know that $\text{P}^{\text{NEXPTIME}}$ is included in $\text{NEXPTIME}^{\text{NEXPTIME}}$, but we also know from [46] that the two complexity classes coincide if, whenever the NEXPTIME oracle is called, its input is of polynomial size, which gives rise to the complexity class $\text{NEXPTIME}^{\text{NEXPTIME}[\text{poly}]}$ (we borrow the notation from [46]). We now observe that during the execution of AlgorithmICR1, the input to the NEXPTIME oracle, which is responsible for checking whether $\text{cert}_{\text{AR}}(\alpha, D, \Sigma) \neq \emptyset$ for an atom $\alpha \in \text{B}(D, \Sigma)$, is always of polynomial size w.r.t. D and Σ . This implies that $\text{QAns}_{\text{ICR}}(\mathbf{A}_\perp)$ is in $\text{NEXPTIME}^{\text{NEXPTIME}[\text{poly}]}$ (and thus, in $\text{P}^{\text{NEXPTIME}}$) in \mathfrak{c} -complexity, and the claim follows.

7.1.3 The \mathfrak{c} -complexity for \mathbf{L}_\perp

For showing that $\text{QAns}_{\text{ICR}}(\mathbf{L}_\perp)$ is in PSPACE , we need to rely on a refined version of the procedure AlgorithmICR1. Indeed, AlgorithmICR1 only shows that $\text{QAns}_{\text{ICR}}(\mathbf{L}_\perp)$ is in EXPTIME in \mathfrak{c} -complexity, despite the fact that $\text{QAns}_{\text{AR}}(\mathbf{L}_\perp)$ is in PSPACE in \mathfrak{c} -complexity, since $\text{B}(D, \Sigma)$ consists of exponentially many atoms. The key ingredient underlying this refined procedure is the following property of linear TGDs, which is implicit in [11], that essentially states that for computing the certain answers of a CQ, we only need linearly many database atoms.

Lemma 7.1. *Consider a database D , a set $\Sigma \in \mathbf{L}$, a CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$. The following are equivalent:*

Input: database D , set $\Sigma \in \mathbf{L}_\perp$, CQ $q(\bar{x})$, tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$
Output: accept if $\bar{c} \in \text{cert}_{\text{ICR}}(q, D, \Sigma)$; otherwise, reject

```

guess a database  $D^* \subseteq \mathbf{B}(D, \Sigma)$  with  $|D^*| \leq |q|$ 
foreach  $\alpha \in D^*$  do
  | if  $\text{cert}_{\text{AR}}(\alpha, D, \Sigma) = \emptyset$  then
  | | return reject
if  $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$  then
  | return accept
else
  | return reject

```

Algorithm 6: AlgorithmICR2

Input: database D , set Σ of TGDs and NCs, CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$
Output: accept if $\bar{c} \notin \text{cert}_{\text{ICR}}(q, D, \Sigma)$; otherwise, reject

```

guess a database  $D^* \subseteq \text{cl}(D, \tau(\Sigma))$ 
foreach  $\alpha \in \text{cl}(D, \tau(\Sigma)) \setminus D^*$  do
  | guess a database  $D_\alpha \subseteq D$ 
  | if  $\alpha \in \text{cl}(D_\alpha, \tau(\Sigma))$  then
  | | return reject
  | else
  | | foreach  $\beta \in D \setminus D_\alpha$  do
  | | | if there is no  $\sigma \in \nu(\Sigma)$  s.t.  $\text{cert}(q_\sigma, D_\alpha \cup \{\beta\}, \tau(\Sigma)) \neq \emptyset$  then
  | | | | return reject
if  $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$  then
  | return reject
else
  | return accept

```

Algorithm 7: AlgorithmICR3

1. $\bar{c} \in \text{cert}(q, D, \Sigma)$.
2. There exists $D' \subseteq D$ with $|D'| \leq |q|$ such that $\bar{c} \in \text{cert}(q, D', \Sigma)$.

By Lemma 7.1, we obtain the decision procedure AlgorithmICR2, depicted in Algorithm 6, for $\text{QAns}_{\text{ICR}}(\mathbf{L}_\perp)$ by adapting AlgorithmICR1 as follows: instead of deterministically computing the intersection of closed repairs, we simply guess $|q|$ atoms of $\mathbf{B}(D, \Sigma)$, and then verify that are indeed members of the intersection of closed repairs. Since, by Theorem 5.1, $\text{QAns}_{\text{AR}}(\mathbf{L}_\perp)$ is in PSPACE in c-complexity, AlgorithmICR2 uses polynomial space, and the claim follows.

7.1.4 The ba-complexity and the d-complexity for \mathbf{C}_\perp , where $\mathbf{C} \in \{\mathbf{G}, \mathbf{L}, \mathbf{A}, \mathbf{S}\}$

The upper bounds are obtained via the simple procedure AlgorithmICR3, depicted in Algorithm 6, which is similar in spirit to the procedure AlgorithmIAR1, and checks whether there exists a superset of the intersection of closed repairs that does not entail the given tuple \bar{c} of constants. More precisely, the algorithm guesses a subset D^* of $\text{cl}(D, \tau(\Sigma))$, that is, the set of ground atoms that can be entailed by D and $\tau(\Sigma)$, and then checks that for every atom $\alpha \in \text{cl}(D, \tau(\Sigma)) \setminus D^*$, there exists $D_\alpha \in \text{reps}(D, \Sigma)$ such that $\alpha \notin \text{cl}(D_\alpha, \tau(\Sigma))$, and thus, α is not in the intersection of closed repairs. This implies that D^* is a superset of the intersection of closed repairs. Finally, the algorithm rejects if $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$; otherwise, it accepts. This is correct due to the following lemma that can be shown as Lemma 6.1:

Lemma 7.2. Consider a database D , a set Σ of TGDs and NCs, a CQ $q(\bar{x})$, and a tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$. The following are equivalent:

1. $\bar{c} \notin \text{cert}_{\text{ICR}}(q, D, \Sigma)$.
2. There is $D^* \supseteq \bigcap_{D' \in \text{reps}(D, \Sigma)} \text{cl}(D', \tau(\Sigma))$ such that $\bar{c} \notin \text{cert}(q, D^*, \tau(\Sigma))$.

Input: database D , set Σ of TGDs and NCs, CQ $q(\bar{x})$, tuple $\bar{c} \in \text{dom}(D)^{|\bar{x}|}$
Output: accept if $\bar{c} \in \text{cert}_{\text{ICR}}(q, D, \Sigma)$; otherwise, reject

```

 $D^* := \text{B}(D, \Sigma)$ 
foreach  $\alpha \in \text{B}(D, \Sigma)$  do
  | if there exists  $D_\alpha \in \text{reps}(D, \Sigma)$  such that  $\alpha \notin \text{cl}(D_\alpha, \tau(\Sigma))$  then
  | |  $D^* := D^* \setminus \{\alpha\}$ 
if  $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$  then
  | return accept
else
  | return reject

```

Algorithm 8: AlgorithmICR4

Since we focus on predicates of bounded arity, the non-deterministic procedure AlgorithmICR3 runs in polynomial time, assuming access to an oracle that is powerful enough for solving $\text{QAns}(\mathcal{C})$, where \mathcal{C} is the class from which the input set of TGDs is coming from. Notice that for computing the database $\text{cl}(D, \tau(\Sigma))$ (or $\text{cl}(D_\alpha, \tau(\Sigma))$) we simply need to enumerate the polynomially many ground atoms that can be formed using constants from $\text{dom}(D)$ and predicates occurring in Σ , and for each such atom γ check whether $\text{cert}(\gamma, D, \tau(\Sigma)) \neq \emptyset$. Therefore:

Lemma 7.3. *For a class \mathcal{C} of TGDs, $\text{QAns}_{\text{ICR}}(\mathcal{C}_\perp)$ is in $\text{coNP}^{\mathcal{C}}$ in \mathbf{t} -complexity, where $\mathbf{t} \in \{\text{ba}, \mathbf{d}\}$, assuming that $\text{QAns}(\mathcal{C})$ is in \mathcal{C} in \mathbf{t} -complexity.*

Since, by Lemma 7.2, AlgorithmICR3 is correct, the desired upper bounds for Group 2 are obtained from Propositions 3.1, 3.2, 3.3 and 3.4, Lemma 7.3, and the usual complexity facts that have been discussed in the previous sections.

Remark. Let us observe that we could also employ the procedure AlgorithmICR1 for obtaining the EXP-TIME and $\text{P}^{\text{NEXP-TIME}}$ upper bounds for $\text{QAns}_{\text{ICR}}(\mathcal{G}_\perp)$ and $\text{QAns}_{\text{ICR}}(\mathcal{A}_\perp)$, respectively, since, by Theorem 5.1, $\text{QAns}_{\text{AR}}(\mathcal{G}_\perp)$ is in EXP-TIME , and $\text{QAns}_{\text{AR}}(\mathcal{A}_\perp)$ is in $\text{P}^{\text{NEXP-TIME}}$ in ba -complexity.

7.1.5 The fp-complexity for \mathcal{C}_\perp , where $\mathcal{C} \in \{\mathcal{G}, \mathcal{L}, \mathcal{A}, \mathcal{S}\}$

We finally discuss how the $\Theta_2^P = \text{P}^{\text{NP}[O(\log n)]}$ upper bound for $\text{QAns}_{\text{ICR}}(\mathcal{C}_\perp)$, where $\mathcal{C} \in \{\mathcal{G}, \mathcal{L}, \mathcal{A}, \mathcal{S}\}$, can be established. Actually, this is done by exploiting the procedure AlgorithmICR4, depicted in Algorithm 8, which is an adaptation of AlgorithmIAR2, that constructs the intersection of closed repairs D^* , and accepts if the given tuple \bar{c} belongs to $\text{cert}(q, D^*, \tau(\Sigma))$; otherwise, it rejects. The intersection of closed repairs D^* is constructed by starting from $\text{B}(D, \Sigma)$, i.e., the set of all ground atoms that can be formed using constants from $\text{dom}(D)$ and predicates occurring in Σ , which are polynomially many since the arity is bounded, and removing all the atoms α for which there exists at least one repair $D_\alpha \in \text{reps}(D, \Sigma)$ such that $\alpha \notin \text{cl}(D_\alpha, \tau(\Sigma))$. More precisely, D^* is constructed via polynomially many parallel calls to an NP-oracle. In fact, for each atom $\alpha \in \text{B}(D, \Sigma)$, we call in parallel an NP-oracle that does the following:

1. Guess a database $D_\alpha \subseteq D$.
2. If $\alpha \in \text{cl}(D_\alpha, \tau(\Sigma))$, then reject.
3. For each atom $\beta \in D \setminus D_\alpha$, if there is no $\sigma \in \nu(\Sigma)$ such that $\text{cert}(q_\sigma, D_\alpha \cup \{\beta\}, \tau(\Sigma)) \neq \emptyset$, then return reject; otherwise; return accept.

The checks $\alpha \in \text{cl}(D_\alpha, \tau(\Sigma))$ and $\text{cert}(q_\sigma, D_\alpha \cup \{\beta\}, \tau(\Sigma)) \neq \emptyset$ are feasible in polynomial time since $\tau(\Sigma)$ and q_σ are fixed, while for each $\mathcal{C} \in \{\mathcal{G}, \mathcal{L}, \mathcal{A}, \mathcal{S}\}$, $\text{QAns}(\mathcal{C})$ is in PTIME in \mathbf{d} -complexity. Therefore, the above oracle is indeed an NP-oracle. It is clear that, for an atom $\alpha \in \text{B}(D, \Sigma)$, if the above oracle returns accept, then α does not belong to the intersection of closed repairs. Consequently, the intersection of closed repairs D^* is constructed by simply removing from D all the atoms α for which the oracle returns accept. Since D^* can be constructed in polynomial time via parallel NP-oracle calls, we can conclude that it can also be constructed in polynomial time via logarithmically many NP-oracle calls; see, e.g., [42]. Once we have D^* in place, we need one more call to an NP-oracle for checking whether $\bar{c} \in \text{cert}(q, D^*, \tau(\Sigma))$; the latter is indeed in NP since, for each $\mathcal{C} \in \{\mathcal{G}, \mathcal{L}, \mathcal{A}, \mathcal{S}\}$, $\text{QAns}(\mathcal{C})$ is in NP in fp -complexity. The claim follows.

7.2 Lower bounds

We now concentrate on the complexity lower bounds claimed in Table 4. The \mathcal{C} -hardness results, where $\mathcal{C} \in \{\text{PSPACE}, \text{EXPTIME}, 2\text{EXPTIME}\}$, are coming for free since $\text{QAns}(\mathcal{C})$ is \mathcal{C} -hard. Therefore, to complete the picture, it suffices to establish the following hardness results:

1. $\text{QAns}_{\text{ICR}}(\mathbf{A}_{\perp})$ is $\text{P}^{\text{NEXP TIME}}$ -hard in **ba**-complexity.
2. $\text{QAns}_{\text{ICR}}(\text{NC})$ is Π_2^P -hard in **ba**-complexity.
3. $\text{QAns}_{\text{ICR}}(\mathbf{C}_{\perp})$, where $\mathbf{C} \in \{\text{L}, \text{A}, \text{S}\}$, is Θ_2^P -hard in **fp**-complexity.
4. $\text{QAns}_{\text{ICR}}(\mathbf{C}_{\perp})$, where $\mathbf{C} \in \{\text{L}, \text{A}, \text{S}\}$, is **coNP**-hard in **d**-complexity.

The rest of the section is devoted to establishing the above lower bounds. But let us first establish an auxiliary lemma, which will be useful for our later analysis. It states that for ground atomic CQs the AR and the ICR semantics coincide:

Lemma 7.4. *Consider a database D , a set Σ of TGDs and NCs, and a ground atom α . Then, $\text{cert}_{\text{AR}}(\alpha, D, \Sigma) \neq \emptyset$ iff $\text{cert}_{\text{ICR}}(\alpha, D, \Sigma) \neq \emptyset$.*

Proof. (\Rightarrow) By hypothesis, for every $D' \in \text{reps}(D, \Sigma)$, $\text{cert}(\alpha, D', \tau(\Sigma)) \neq \emptyset$. This implies that, for every $D' \in \text{reps}(D, \Sigma)$, $\alpha \in \text{cl}(D', \tau(\Sigma))$. Therefore, $\alpha \in \bigcap_{D' \in \text{reps}(D, \Sigma)} \text{cl}(D', \tau(\Sigma))$, which means that $\text{cert}_{\text{ICR}}(\alpha, D, \Sigma) \neq \emptyset$, as needed.

(\Leftarrow) Conversely, assume that $\text{cert}_{\text{AR}}(\alpha, D, \Sigma) = \emptyset$. Hence, there exists $D_{\alpha} \in \text{reps}(D, \Sigma)$ such that $\text{cert}(\alpha, D_{\alpha}, \tau(\Sigma)) = \emptyset$, and thus, $\alpha \notin \text{cl}(D_{\alpha}, \tau(\Sigma))$. Assume now that $\text{cert}(\alpha, D^*, \tau(\Sigma)) \neq \emptyset$ with $D^* = \bigcap_{D' \in \text{reps}(D, \Sigma)} \text{cl}(D', \tau(\Sigma))$. This implies that there exists $D'' \subseteq \text{cl}(D_{\alpha}, \tau(\Sigma))$ such that $\text{cert}(\alpha, D'', \tau(\Sigma)) \neq \emptyset$. But this allows us to conclude that $\alpha \in \text{cl}(D_{\alpha}, \tau(\Sigma))$, which is a contradiction. Therefore, $\text{cert}(\alpha, D^*, \tau(\Sigma)) = \emptyset$, which means that $\text{cert}_{\text{ICR}}(\alpha, D, \Sigma) = \emptyset$. \square

We proceed with the proofs of the claimed lower bounds.

Theorem 7.2. *$\text{QAns}_{\text{ICR}}(\mathbf{A}_{\perp})$ is $\text{P}^{\text{NEXP TIME}}$ -hard in **ba**-complexity.*

Proof. By Lemma 7.4, we can apply the proof for the fact that $\text{QAns}_{\text{AR}}(\mathbf{A}_{\perp})$ is $\text{P}^{\text{NEXP TIME}}$ -hard in **ba**-complexity. Recall that for showing the latter we reduce from the extended exponential tiling problem. In fact, given an extended tiling system \mathcal{E} , we construct a database $D_{\mathcal{E}}$, and a set $\Sigma_{\mathcal{E}} \in \mathbf{A}_{\perp}$ that mentions only predicates of bounded arity, such that \mathcal{E} is valid iff $\text{cert}_{\text{AR}}(\text{Yes}(), D, \Sigma) \neq \emptyset$, where Yes is a 0-ary predicate indicating that \mathcal{E} is indeed valid. By Lemma 7.4, we can conclude that $\text{cert}_{\text{AR}}(\text{Yes}(), D, \Sigma) \neq \emptyset$ iff $\text{cert}_{\text{ICR}}(\text{Yes}(), D, \Sigma) \neq \emptyset$, which shows that $\text{QAns}_{\text{ICR}}(\mathbf{A}_{\perp})$ is $\text{P}^{\text{NEXP TIME}}$ -hard in **ba**-complexity. \square

Theorem 7.3. *$\text{QAns}_{\text{ICR}}(\text{NC})$ is Π_2^P -hard in **ba**-complexity.*

Proof. The proof of Theorem 6.3, showing that $\text{QAns}_{\text{IAR}}(\text{NC})$ is Π_2^P -hard in **ba**-complexity, applies also to the ICR semantics. The reason is that $\text{cl}(D, \Sigma) = D$ for every database D and Σ in **NC**, since Σ does not include any TGD. Hence, the IAR and ICR semantics coincide for the class **NC**. \square

Theorem 7.4. *$\text{QAns}_{\text{ICR}}(\mathbf{C}_{\perp})$, where $\mathbf{C} \in \{\text{L}, \text{A}, \text{S}\}$, is Θ_2^P -hard in **fp**-complexity.*

Proof. We reduce from **Comp3SAT** [32, 33]. Recall that given two sets A and B of 3CNF formulas, this problem asks whether $\#A > \#B$, i.e., whether A contains more satisfiable formulas than B . Recall also that this problem remains Θ_2^P -hard even if $|A| = |B|$, all formulas in A and B are over the same set of variables and have the same number of clauses, and $A = \{\varphi_1, \dots, \varphi_m\}$, $B = \{\psi_1, \dots, \psi_m\}$ are such that φ_{i+1} (resp., ψ_{i+1}) is satisfiable implies φ_i (resp., ψ_i) is satisfiable, for each $i \in \{1, \dots, m-1\}$. Clearly, $\#A > \#B$ iff there exists $i \in \{1, \dots, m\}$ such that φ_i is satisfiable and ψ_i is unsatisfiable.

Given $A = \{\varphi_1, \dots, \varphi_m\}$ and $B = \{\psi_1, \dots, \psi_m\}$, our goal is to construct a database $D_{A,B}$ and a Boolean CQ $q_{A,B}$ such that the following are equivalent:

1. There exists $i \in \{1, \dots, m\}$ such that φ_i is satisfiable and ψ_i is unsatisfiable.
2. $\text{cert}_{\text{ICR}}(q_{A,B}, D_{A,B}, \Sigma) \neq \emptyset$ for some fixed $\Sigma \in \mathbf{C}_{\perp}$, where $\mathbf{C} \in \{\text{L}, \text{A}, \text{S}\}$.

In fact, the construction is along the lines of the one given in Section 6 for showing that $\text{QAns}_{\text{IAR}}(\mathbb{G}_\perp)$ is Θ_2^P -hard in fp-complexity. We first construct a database D_A and a CQ $q_A(x)$ such that, for each $i \in \{1, \dots, m\}$, φ_i is satisfiable iff $\langle f_i \rangle \in q_A(D_A)$; the constant f_i should be understood as the identifier for the formula φ_i . Moreover, we construct a database D_B and a CQ $q_B(x)$ such that, for $i \in \{1, \dots, m\}$, ψ_i is unsatisfiable iff $\langle f_i \rangle \in \text{cert}_{\text{ICR}}(q_B, D_B, \Sigma) \neq \emptyset$ for some fixed set $\Sigma \in \mathbb{C}_\perp$, for $\mathbb{C} \in \{\text{L}, \text{A}, \text{S}\}$; here, f_i is the identifier of ψ_i . Once we have the above in place, we can easily construct $D_{A,B}$ and $q_{A,B}$. We assume that all the formulas in A and B are over the variables x_1, \dots, x_n and have k clauses.

The database D_A and the CQ q_A . Actually, for D_A and q_A we can use exactly the same construction as in the proof of the fact that $\text{QAns}_{\text{IAR}}(\mathbb{G}_\perp)$ is Θ_2^P -hard in fp-complexity given in Section 6.

The database D_B and the CQ q_B . Let us now explain the construction of D_B and q_B , which is significantly different (at least the construction of D_B) than the one given in the previous section. Given a 3CNF formula $\psi_i = C_{i,1} \wedge \dots \wedge C_{i,k}$ from B with $C_{i,j} = (\ell_{i,j}^1 \vee \ell_{i,j}^2 \vee \ell_{i,j}^3)$, the database D_B essentially stores all the possible truth assignments for each clause of ψ_i . To this end, we use a 5-ary predicate $\text{BClause}_{b_1 b_2 b_3}^{s_1 s_2 s_3}$, where $s_1, s_2, s_3 \in \{\text{p}, \text{n}\}$ and $b_1, b_2, b_3 \in \{0, 1\}$. For example, given the clause $C_{i,j} = x_{i_1} \vee x_{i_2} \vee \neg x_{i_3}$, the atom

$$\text{BClause}_{101}^{\text{ppn}}(f_i, c_j, x_{i_1}, x_{i_2}, x_{i_3})$$

encodes the clause itself, with the superscript **ppn** indicating that the variable of the first (resp., second, third) literal appears positively (resp., positively, negatively) in $C_{i,j}$, and at the same time encodes the truth assignment that sets x_{i_1} to true, x_{i_2} to false, and x_{i_3} to true. We proceed to formally define D_B .

Recall that for a literal ℓ , $\text{sign}(\ell) = \text{p}$ (resp., $\text{sign}(\ell) = \text{n}$) if $\ell = x$ (resp., $\ell = \neg x$). For a clause $C_{i,j}$ of φ , we write $s_{i,j}^r$ for $\text{sign}(\ell_{i,j}^r)$, where $r \in \{1, 2, 3\}$. Recall that $\text{var}(\ell)$ is the variable of the literal ℓ . The database D_B is defined as

$$\bigcup_{1 \leq i \leq m} \bigcup_{1 \leq j \leq k} \bigcup_{b_1, b_2, b_3 \in \{0,1\}} \left\{ \text{BClause}_{b_1 b_2 b_3}^{s_{i,j}^1 s_{i,j}^2 s_{i,j}^3} (f_i, c_j, \text{var}(\ell_{i,j}^1), \text{var}(\ell_{i,j}^2), \text{var}(\ell_{i,j}^3)) \right\}.$$

Regarding the CQ $q_B(x)$, is defined as the atomic query $\text{Unsat}(x)$, which simply asks whether there exists a formula in B that is unsatisfiable.

We claim that there exists a set $\Sigma \in \mathbb{C}_\perp$, for $\mathbb{C} \in \{\text{L}, \text{A}, \text{S}\}$, such that, for each $i \in \{1, \dots, m\}$, ψ_i is unsatisfiable iff $\langle f_i \rangle \in \text{cert}_{\text{ICR}}(q_B, D_B, \Sigma)$. In particular, $\Sigma = \Sigma_{\text{cons}} \cup \Sigma_{\text{unsat}}$ with Σ_{cons} being a set of NCs that performs a consistency check on the truth assignment for ψ_i , i.e., each variable of ψ_i is assigned exactly one value, and Σ_{unsat} being a set of TGDs that entails the ground atom $\text{Unsat}(f_i)$ whenever a clause of ψ_i evaluates to false. More precisely, Σ_{cons} consists of the following NCs; a \star symbol in the superscript is a placeholder for **p** or **n**, a \star in the subscript is a placeholder for 0 or 1, and, as usual, $_$ is a “don’t care” variable:

$$\begin{aligned} \text{BClause}_{1\star\star}^{\star\star\star}(x, _, y, _, _) , \text{BClause}_{0\star\star}^{\star\star\star}(x, _, y, _, _) &\rightarrow \perp \\ \text{BClause}_{1\star\star}^{\star\star\star}(x, _, y, _, _) , \text{BClause}_{\star 0\star}^{\star\star\star}(x, _, _, y, _) &\rightarrow \perp \\ \text{BClause}_{1\star\star}^{\star\star\star}(x, _, y, _, _) , \text{BClause}_{\star\star 0}^{\star\star\star}(x, _, _, _, y) &\rightarrow \perp \\ \text{BClause}_{\star 1\star}^{\star\star\star}(x, _, _, y, _) , \text{BClause}_{0\star\star}^{\star\star\star}(x, _, y, _, _) &\rightarrow \perp \\ \text{BClause}_{\star 1\star}^{\star\star\star}(x, _, _, y, _) , \text{BClause}_{\star 0\star}^{\star\star\star}(x, _, _, y, _) &\rightarrow \perp \\ \text{BClause}_{\star 1\star}^{\star\star\star}(x, _, _, y, _) , \text{BClause}_{\star\star 0}^{\star\star\star}(x, _, _, _, y) &\rightarrow \perp \\ \text{BClause}_{\star\star 1}^{\star\star\star}(x, _, _, _, y) , \text{BClause}_{0\star\star}^{\star\star\star}(x, _, y, _, _) &\rightarrow \perp \\ \text{BClause}_{\star\star 1}^{\star\star\star}(x, _, _, _, y) , \text{BClause}_{\star 0\star}^{\star\star\star}(x, _, _, y, _) &\rightarrow \perp \\ \text{BClause}_{\star\star 1}^{\star\star\star}(x, _, _, _, y) , \text{BClause}_{\star\star 0}^{\star\star\star}(x, _, _, _, y) &\rightarrow \perp. \end{aligned}$$

Moreover, the set of TGDs Σ_{unsat} consists of:

$$\begin{aligned}
\text{BClause}_{111}^{\text{nnn}}(x, _, _, _, _) &\rightarrow \text{Unsat}(x) \\
\text{BClause}_{110}^{\text{nnp}}(x, _, _, _, _) &\rightarrow \text{Unsat}(x) \\
\text{BClause}_{101}^{\text{npn}}(x, _, _, _, _) &\rightarrow \text{Unsat}(x) \\
\text{BClause}_{100}^{\text{ppp}}(x, _, _, _, _) &\rightarrow \text{Unsat}(x) \\
\text{BClause}_{011}^{\text{pnn}}(x, _, _, _, _) &\rightarrow \text{Unsat}(x) \\
\text{BClause}_{010}^{\text{ppp}}(x, _, _, _, _) &\rightarrow \text{Unsat}(x) \\
\text{BClause}_{001}^{\text{ppn}}(x, _, _, _, _) &\rightarrow \text{Unsat}(x) \\
\text{BClause}_{000}^{\text{ppp}}(x, _, _, _, _) &\rightarrow \text{Unsat}(x).
\end{aligned}$$

Observe that Σ_{unsat} falls in \mathbf{C} , for each $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$. We proceed to show that:

Lemma 7.5. *For $i \in \{1, \dots, m\}$, ψ_i is unsatisfiable iff $\langle f_i \rangle \in \text{cert}_{\text{ICR}}(q_B, D_B, \Sigma)$.*

Proof. (\Rightarrow) Consider an arbitrary repair $D \in \text{reps}(D_B, \Sigma)$. It is easy to verify that, due to Σ_{cons} , D encodes a consistent assignment μ_D of truth values to the variables of ψ_i , i.e., for each clause $C_{i,j}$ of ψ_i , D contains exactly one atom of the form

$$\text{BClause}_{b_1 b_2 b_3}^{s_i^1 s_i^2 s_i^3}(f_i, c_j, \text{var}(\ell_{i,j}^1), \text{var}(\ell_{i,j}^2), \text{var}(\ell_{i,j}^3)).$$

Since, by hypothesis, ψ_i is unsatisfiable, there exists a clause $C_{i,j}$ of ψ_i such that, according to μ_D , evaluates to false. This implies that a TGD of Σ_{unsat} will be triggered, and thus, $\text{Unsat}(f_i) \in \text{cl}(D, \tau(\Sigma))$. Hence,

$$\text{Unsat}(f_i) \in \bigcap_{D' \in \text{reps}(D_B, \Sigma)} \text{cl}(D', \tau(\Sigma)),$$

which in turn implies that $\langle f_i \rangle \in \text{cert}_{\text{ICR}}(q_B, D_B, \Sigma)$.

(\Leftarrow) Conversely, assume that ψ_i is satisfiable, and let μ be a satisfying assignment that witnesses this fact. It is easy to verify that there is a repair $D_\mu \in \text{reps}(D_B, \Sigma)$ that encodes μ . Since μ is a satisfying assignment, all the clauses of ψ_i evaluate to true. This means that none of the TGDs of Σ_{unsat} will be triggered, and thus, $\langle f_i \rangle \notin \text{cert}_{\text{AR}}(q_B, D_B, \Sigma)$. By Lemma 7.4, we can conclude that $\langle f_i \rangle \notin \text{cert}_{\text{ICR}}(q_B, D_B, \Sigma) = \emptyset$, and the claim follows. \square

The database $D_{A,B}$ and the CQ $q_{A,B}$. We can now easily construct the database $D_{A,B}$ and the Boolean CQ $q_{A,B}$ with the desired property:

$$D_{A,B} = D_A \cup D_B \quad \text{and} \quad q_{A,B} = \exists x(q_A(x) \wedge q_B(x)).$$

With Σ being the set devised above, it is not difficult to verify that

$$\bigcap_{D \in \text{reps}(D_{A,B}, \Sigma)} \text{cl}(D, \Sigma) = D_A \cup \left(\bigcap_{D \in \text{reps}(D_B, \Sigma)} \text{cl}(D, \Sigma) \right)$$

since D_A is consistent with Σ . By exploiting this observation, and Lemmas 6.7 and 7.5, we can provide a proof that mimics the one of Lemma 6.9, and show that:

Lemma 7.6. *The following are equivalent:*

1. *There exists $i \in \{1, \dots, m\}$ such that φ_i is satisfiable and ψ_i is unsatisfiable.*
2. $\text{cert}_{\text{ICR}}(q_{A,B}, D_{A,B}, \Sigma) \neq \emptyset$.

With Lemma 7.6 in place, we can conclude that

$$\#A > \#B \iff \text{cert}_{\text{ICR}}(q_{A,B}, D_{A,B}, \Sigma) \neq \emptyset$$

for a fixed $\Sigma \in \mathbf{C}_\perp$, for $\mathbf{C} \in \{\mathbf{L}, \mathbf{A}, \mathbf{S}\}$, which implies that $\text{QAns}_{\text{ICR}}(\mathbf{C}_\perp)$ is Θ_2^P -hard in fp-complexity , and the claim follows. \square

Theorem 7.5. $\text{QAns}_{\text{ICR}}(\mathcal{C}_\perp)$, where $\mathcal{C} \in \{\text{L}, \text{A}, \text{S}\}$, is coNP -hard in d -complexity.

The proof of the above exploits the coNP -hard problem of deciding whether a 3CNF formula is unsatisfiable:

PROBLEM : 3UNSAT
INPUT : A 3CNF Boolean formula φ .
QUESTION : Is φ unsatisfiable?

We now proceed with the proof of Theorem 7.5.

Proof of Theorem 7.5. The construction of the database is essentially the same as the one given in the previous proof for the database D_B , with the difference that we have to deal only with one formula and not a set of formulas. Although it is easy to modify the construction given above, we give it here for the sake of completeness.

For a 3CNF formula $\varphi = C_1 \wedge \dots \wedge C_k$ over the variables x_1, \dots, x_n with $C_i = (\ell_i^1 \vee \ell_i^2 \vee \ell_i^3)$, the database D_φ stores all the possible truth assignments for a clause of φ . To this end, we use a 4-ary predicate $\text{Clause}_{b_1 b_2 b_3}^{s_1 s_2 s_3}$, where $s_1, s_2, s_3 \in \{\text{p}, \text{n}\}$ and $b_1, b_2, b_3 \in \{0, 1\}$. For example, for $C_i = x_{i_1} \vee x_{i_2} \vee \neg x_{i_3}$, the atom

$$\text{Clause}_{101}^{\text{ppn}}(c_i, x_{i_1}, x_{i_2}, x_{i_3})$$

encodes the clause itself, with the superscript ppn indicating that the variable of the first (resp., second, third) literal appears positively (resp., positively, negatively) in C_i , and at the same time encodes the truth assignment that sets x_{i_1} to true, x_{i_2} to false, and x_{i_3} to true. We proceed to formally define D_φ .

As usual, for a literal ℓ , $\text{sign}(\ell) = \text{p}$ (resp., $\text{sign}(\ell) = \text{n}$) if $\ell = x$ (resp., $\ell = \neg x$) for some variable x . For a clause C_i of φ , we write s_i^j for $\text{sign}(\ell_i^j)$, where $j \in \{1, 2, 3\}$. Recall that $\text{var}(\ell)$ is the variable of the literal ℓ . The database D_φ is

$$\bigcup_{1 \leq i \leq k} \bigcup_{b_1, b_2, b_3 \in \{0, 1\}} \left\{ \text{Clause}_{b_1 b_2 b_3}^{s_i^1 s_i^2 s_i^3}(c_i, \text{var}(\ell_i^1), \text{var}(\ell_i^2), \text{var}(\ell_i^3)) \right\}.$$

This completes the definition of D_φ .

We claim that there exists a set Σ of TGDs and NCs such that φ is unsatisfiable iff $\text{cert}_{\text{ICR}}(\text{Unsat}(), D_\varphi, \Sigma) \neq \emptyset$. In particular, $\Sigma = \Sigma_{\text{cons}} \cup \Sigma_{\text{unsat}}$ with Σ_{cons} being a set of NCs that performs a consistency check on the truth assignment for φ , i.e., each variable of φ is assigned exactly one value, and Σ_{unsat} being a set of TGDs that entails the ground atom $\text{Unsat}()$ whenever a clause of φ evaluates to false. More precisely, Σ_{cons} consists of the following NCs; a \star symbol in the superscript is a placeholder for p or n , a \star in the subscript is a placeholder for 0 or 1, and, as usual, $_$ is a “don’t care” variable that occurs only once:

$$\begin{aligned} \text{Clause}_{1\star\star}^{\star\star\star}(_, x, _, _) , \text{Clause}_{0\star\star}^{\star\star\star}(_, x, _, _) &\rightarrow \perp \\ \text{Clause}_{1\star\star}^{\star\star\star}(_, x, _, _) , \text{Clause}_{\star 0\star}^{\star\star\star}(_, _, x, _) &\rightarrow \perp \\ \text{Clause}_{1\star\star}^{\star\star\star}(_, x, _, _) , \text{Clause}_{\star\star 0}^{\star\star\star}(_, _, _, x) &\rightarrow \perp \\ \text{Clause}_{\star 1\star}^{\star\star\star}(_, _, x, _) , \text{Clause}_{0\star\star}^{\star\star\star}(_, x, _, _) &\rightarrow \perp \\ \text{Clause}_{\star 1\star}^{\star\star\star}(_, _, x, _) , \text{Clause}_{\star 0\star}^{\star\star\star}(_, _, x, _) &\rightarrow \perp \\ \text{Clause}_{\star 1\star}^{\star\star\star}(_, _, x, _) , \text{Clause}_{\star\star 0}^{\star\star\star}(_, _, _, x) &\rightarrow \perp \\ \text{Clause}_{\star\star 1}^{\star\star\star}(_, _, _, x) , \text{Clause}_{0\star\star}^{\star\star\star}(_, x, _, _) &\rightarrow \perp \\ \text{Clause}_{\star\star 1}^{\star\star\star}(_, _, _, x) , \text{Clause}_{\star 0\star}^{\star\star\star}(_, _, x, _) &\rightarrow \perp \\ \text{Clause}_{\star\star 1}^{\star\star\star}(_, _, _, x) , \text{Clause}_{\star\star 0}^{\star\star\star}(_, _, _, x) &\rightarrow \perp. \end{aligned}$$

Moreover, the set of TGDs Σ_{unsat} consists of:

$$\begin{aligned} \text{Clause}_{111}^{\text{nnn}}(_, _, _, _) &\rightarrow \text{Unsat}() \\ \text{Clause}_{110}^{\text{nnp}}(_, _, _, _) &\rightarrow \text{Unsat}() \\ \text{Clause}_{101}^{\text{npn}}(_, _, _, _) &\rightarrow \text{Unsat}() \\ \text{Clause}_{100}^{\text{npp}}(_, _, _, _) &\rightarrow \text{Unsat}() \\ \text{Clause}_{011}^{\text{pnn}}(_, _, _, _) &\rightarrow \text{Unsat}() \end{aligned}$$

	c-complexity	ba-complexity	fp-complexity	d-complexity
F	EXPTIME	NP	NP	P TIME
WG	2EXPTIME	EXPTIME	EXPTIME	EXPTIME
WA	2EXPTIME	2EXPTIME	NP	P TIME
WS	2EXPTIME	2EXPTIME	NP	P TIME

Table 5: Complexity of $\text{QAns}(\mathcal{C})$, where $\mathcal{C} \in \{F, WG, WA, WS\}$; these are completeness results.

$$\begin{aligned} \text{Clause}_{010}^{\text{pnp}}(_, _, _, _) &\rightarrow \text{Unsat}() \\ \text{Clause}_{001}^{\text{ppn}}(_, _, _, _) &\rightarrow \text{Unsat}() \\ \text{Clause}_{000}^{\text{ppp}}(_, _, _, _) &\rightarrow \text{Unsat}(). \end{aligned}$$

Observe that Σ_{unsat} falls in \mathcal{C} , for each $\mathcal{C} \in \{L, A, S\}$. This completes the definition of Σ . By providing a proof that mimics the one for Lemma 7.5, we can show that φ is unsatisfiable iff $\text{cert}_{\text{ICR}}(\text{Unsat}(), D_\varphi, \Sigma) \neq \emptyset$, and the claim follows. \square

Remark. Observe that in the above reduction the query is a ground atomic CQ, that is, $\text{Unsat}()$. This fact, together with Lemma 7.4, implies that φ is unsatisfiable iff $\text{cert}_{\text{AR}}(\text{Unsat}(), D_\varphi, \Sigma) \neq \emptyset$. Thus, the above reduction provides an alternative proof for the fact that $\text{QAns}_{\text{AR}}(\mathcal{C}_\perp)$, where $\mathcal{C} \in \{L, A, S\}$, is coNP-hard in d-complexity. However, the proof given in Section 5 shows that the coNP-hardness holds even without TGDs, but it exploits a more complex CQ.

8 Full dependencies and beyond

A central class of TGDs, which is incomparable (at the syntax level) to all the classes that we have seen so far, is the class of *full TGDs*, i.e., TGDs without existentially quantified variables, which we denote by F . Indeed, this class forms a powerful language for modeling ontologies that has been used in several different scenarios. For example, it is known that the logical core of the RL profile of OWL 2, which is aimed at applications that require efficient reasoning without sacrificing too much expressive power, corresponds to full TGDs.⁷ Interestingly, the main classes of TGDs that we have seen in the previous sections based on the notions of guardedness, acyclicity and stickiness, come with their “weakly” version that incorporates full TGDs: *weakly-guarded* (WG) [10], *weakly-acyclic* (WA) [21], and *weakly-sticky* (WS), respectively. The definition of all these “weakly” versions follows the same principle: the underlying syntactic condition is relaxed in such a way that only certain “harmful” variables are taken into account; for details we refer the reader to the references given above.

The complexity of $\text{QAns}(\mathcal{C})$, where $\mathcal{C} \in \{F, WG, WA, WS\}$, is by now well-understood and is summarized in Table 5. The results for $\text{QAns}(F)$ are coming from the Datalog literature since a set of full TGDs is essentially a Datalog program [17]. For all the other classes, we refer the reader to the references mentioned above. But what about the complexity of consistent query answering under the semantics that we have seen so far, when the above classes of TGDs are combined with NCs? It turned out that the analysis performed in the previous sections for the less expressive classes of TGDs allows us to easily complete the picture.

Theorem 8.1. *The t-complexity of $\text{QAns}_s(\mathcal{C}_\perp)$, where $t \in \{c, ba, fp, d\}$, $s \in \{\text{AR}, \text{IAR}, \text{ICR}\}$, and $\mathcal{C} \in \{F, WG, WA, WS\}$, is as shown in Table 6.*

Let us briefly summarize how the above complexity results are obtained by exploiting the algorithms and the reductions devised in the previous sections:

AR semantics. The upper bounds are obtained via the procedure `AlgorithmAR`. The 2EXPTIME and EXPTIME lower bounds are inherited from $\text{QAns}(\mathcal{C})$, while the Π_2^P and coNP lower bounds are inherited from $\text{QAns}_{\text{AR}}(\text{NC})$.

IAR semantics. All the upper bounds, apart from the Θ_2^P ones, are obtained via `AlgorithmIAR1`, while the Θ_2^P upper bounds via `AlgorithmIAR2`. As above, the 2EXPTIME and EXPTIME lower bounds are inherited from $\text{QAns}(\mathcal{C})$, while the Π_2^P one from $\text{QAns}_{\text{IAR}}(\text{NC})$. For the Θ_2^P and coNP lower bounds,

⁷https://www.w3.org/TR/owl2-profiles/#OWL_2_RL

		c-complexity	ba-complexity	fp-complexity	d-complexity
F_{\perp}	AR	EXPTIME	Π_2^P	Π_2^P	coNP
	IAR	EXPTIME	Π_2^P	Θ_2^P	coNP
	ICR	EXPTIME	Π_2^P	Θ_2^P	coNP
WG_{\perp}	AR	2EXPTIME	EXPTIME	EXPTIME	EXPTIME
	IAR	2EXPTIME	EXPTIME	EXPTIME	EXPTIME
	ICR	2EXPTIME	EXPTIME	EXPTIME	EXPTIME
WA_{\perp}	AR	2EXPTIME	2EXPTIME	Π_2^P	coNP
	IAR	2EXPTIME	2EXPTIME	Θ_2^P	coNP
	ICR	2EXPTIME	2EXPTIME	Θ_2^P	coNP
WS_{\perp}	AR	2EXPTIME	2EXPTIME	Π_2^P	coNP
	IAR	2EXPTIME	2EXPTIME	Θ_2^P	coNP
	ICR	2EXPTIME	2EXPTIME	Θ_2^P	coNP

Table 6: Complexity of $QAns_s(C_{\perp})$, where $C \in \{F, WG, WA, WS\}$. For each class, the first (resp., second, third) row corresponds to AR (resp., IAR, ICR); these are completeness results.




Figure 1: Complete picture of the relationships among inconsistent-tolerant semantics [4]. The semantics to which the arrow points is a complete approximation of the semantics from where the arrow starts; e.g., both AR and ICAR entail all the ICR answers, and ICR all the IAR answers.

it suffices to observe that the proof for the fact that $QAns_{IAR}(G_{\perp})$ is Θ_2^P -hard in fp-complexity and coNP-hard in d-complexity exploits only full TGDs.

ICR semantics. The c-complexity upper bounds are obtained via AlgorithmICR1. The ba- and d-complexity upper bounds, as well as the fp-complexity upper bound in the case of WG_{\perp} , are obtained via AlgorithmICR3. The Θ_2^P upper bounds are established by using AlgorithmICR4. The 2EXPTIME and EXPTIME lower bounds are inherited from $QAns(C)$, while the Π_2^P one from $QAns_{ICR}(NC)$. Finally, for the Θ_2^P and coNP lower bounds, observe that the proofs for showing that $QAns_{ICR}(C'_{\perp})$ is Θ_2^P -hard in fp-complexity and coNP-hard in d-complexity, for $C' \in \{L, A, S\}$, use only full TGDs.

9 Related work

There has been an extensive body of work on querying inconsistent knowledge bases in the context of DL and existential rule languages. Arguably, as discussed in the introduction, the AR, IAR, and ICR semantics have been the most prominent inconsistency-tolerant semantics. The AR semantics (known in the database literature as consistent query answering) was first developed for relational databases

in [1], and then applied to several DLs in [28, 31]. Intuitively, the AR semantics entails the set of answers that are classically entailed in every possible repair. The intractability of the AR semantics was first established in [31], which showed that ontological UCQ answering is coNP-complete in data complexity. This result was then strengthened in [28], which showed that the coNP-hardness holds even for ground atomic queries and when the knowledge base is expressed in $DL-Lite_{core}$ (the least expressive logic in the $DL-Lite$ family). The work of [44] studied both the data and the combined complexity for a wide spectrum of DLs, while [3] identified cases for simple ontologies (within the $DL-Lite$ family) for which tractable data complexity results can be obtained. In [37, 38, 40], the data and different types of combined complexity of the AR semantics have been studied for ontologies modeled via existential rules and negative constraints.

The IAR semantics was introduced in [28] as a sound (under-)approximation of AR, as it entails the set of answers that are classically entailed from the intersection of all repairs. The work of [28] showed that ontological UCQ answering is in PTIME in data complexity for $DL-Lite_A$. On the other hand, [29] showed that ontological CQ answering under the IAR semantics is first-order rewritable for DLs of the $DL-Lite$ family. The combined complexity of the IAR semantics for ontology languages of the $DL-Lite$ family was investigated in [9]. The work of [6] analyzed the data and combined complexity of ontological query answering under the AR and IAR semantics for different notions of maximal repairs focusing on the lightweight logic $DL-Lite_{\mathcal{R}}$. Practical implementations of the AR and IAR semantics have been developed in [7, 30].

The ICR semantics was introduced in [3], where it was also shown that ontological CQ answering is in PTIME in data complexity for simple DL ontologies. The ICR semantics entails the set of classical answers obtained from the intersection of the logical closure of all possible repairs, and it is an over-approximation of IAR (i.e., IAR answers are ICR answers, but the reverse does not hold) and an under-approximation of AR. The complexity of ontological query answering under the IAR and ICR semantics for a wide range of existential rule languages and for different complexity measures has been investigated in [34]. The work of [36] investigated the complexity of ontological query answering under the AR, IAR, and ICR semantics for several existential rule languages and complexity measures when repairs are cardinality-maximal.

The work of [28] also introduces other semantics that under and over approximate AR, namely CAR and ICAR, which stand for *Closed ABox Repairs* and *Intersection of Closed ABox Repairs*, respectively. The rationale for the CAR semantics is that AR is dependant of the syntax of the database, which means that logically equivalent knowledge bases may yield different answers under the AR semantics. CAR is an over-approximation of AR based on repairs that are computed from the consistent closure of the database with respect to the (DL-based or rule-based) ontology Σ ; intuitively, the consistent closure is the set of all atoms that can be consistently derived from the database and Σ , i.e., such that no negative axiom is violated in the derivation. A repair is now any subset of the consistent closure that “maximally preserves” the content of the original database. The CAR semantics corresponds then to the set of answers that are classically entailed from every closed repair, and contains all AR answers. Therefore, it is a complete approximation, but there are answers that are true under the CAR semantics that are not true under the AR semantics. In [3, 44] it is shown that ontological query answering for $DL-Lite_{core}$ under CAR is in PTIME in data complexity when we focus on atomic queries, and coNP-complete for UCQs. It has been also shown that for the DL EL , ontological UCQ answering under CAR is DP-complete in data complexity. Analogously to IAR, [28] defines the ICAR semantics, a sound approximation to CAR that computes the answers from the intersection of all closed repairs.

Following these families of under and over approximations to the AR semantics, further semantics were developed trying to formalize more granular conflict resolution techniques. The notion of k -lazy consistent answers, proposed in [38], provides an alternative semantics that offers a compromise between quality of answers and computation time. Lazy answers are based on a “budget” (the parameter k) that restricts the size of removals that need to be made in an inconsistent set of facts in order to make it consistent; if the budget is large enough, then all possible ways of resolving the conflicts within the budget are considered, but if it is not enough then the whole inconsistent set is removed. If we think of the problem of querying inconsistent KBs as a reasoning task for an intelligent agent, then the value of the budget would be a bound on its reasoning capabilities (more complex reasoning can thus be afforded with higher budgets). The k -lazy semantics is non-monotonic with respect to k in the sense that the answers obtained under the semantics with parameter k may not be a subset of those obtained with $k + 1$; nevertheless, the union- k -lazy extension proposed in [41] allows to monotonically expand the set of consistent answers. Although the k -lazy approach is not strictly based on the same notion of repair, it was shown that there always exists a value k for which k -lazy and AR coincide.

The k -support semantics [9] increasingly produces more fine-grained under (sound) approximations of the AR semantics. On the other hand, [9] also proposed the k -defeater semantics, which provides

increasingly tighter upper (complete) approximations of the AR semantics. The k -support semantics restricts the number of distinct supports (i.e., consistent derivations for an answer) that can cover all the repairs; with $k = 1$, the same support must be present in every repair, so it coincides with IAR – increasing parameter k yields larger sets of answers until AR is reached. On the other hand, in the k -defeater semantics, only sets of size k that create a contradiction w.r.t. every minimal support for an answer are considered; clearly, when $k = 0$ the semantics coincides with the brave semantics, i.e., the set of answers that can be obtained from some repair. As k increases, larger defeater sets are considered and the set of answers are incrementally reduced until the set of AR answers is reached. Both semantics enjoy desirable computational properties; however, note that k -defeater may entail answers that are conflicting among each other (even for the same value of k). Figure 1, reproduced from [4], summarizes the relationship among the main inconsistency-tolerant semantics defined so far in the literature. We refer the reader to [4] for more details on the semantics and complexity results for several families of DLs.

The AR semantics was extended to the *generalized repair (GR)* semantics, and its computational complexity was analyzed in [20]. In the GR semantics, not only atoms from the database, but also ontological axioms may be removed and considered as part of the repairs; notice, however, that some database atoms and axioms may be specified to be non-removable. The generalized repair semantics was applied to the IAR and ICR semantics in [34], where its complexity was analyzed for different existential rule languages and complexity measures.

Recently, inconsistency-tolerant semantics for ontological query answering has also been considered from an explanation perspective. A (minimal) explanation for an ontological query can be defined in different ways. However, the literature has lately focused on a quite natural definition in which a (minimal) explanation for an ontological query is a (minimal) set of facts that, together with the ontology, entail the query (see [7, 15] for the DL setting, and [14] for the existential rule setting, and the references therein). This concept of explanation has been extended to the inconsistency-tolerant semantics, in which, intuitively, an explanation in terms of (sets of) facts is provided to justify why a query is entailed under the AR, IAR, ICR, semantics (see [7] for the DL setting, and [35] for the existential rule setting).

10 Conclusion

We performed a thorough complexity analysis of consistent query answering under the main classes of TGDs based on the notions of guardedness, linearity, acyclicity, and stickiness, and extensions thereof. In our analysis, we focused on the standard inconsistency-tolerant semantics (AR semantics), as well as the main sound approximations of it (IAR and ICR semantics), and we considered different complexity measures with the aim of understanding how the complexity is affected when key parameters of the input are considered to be fixed.

Another goal of our analysis, apart from clarifying the complexity landscape, was to understand whether the IAR and ICR semantics have the desired effect on the data complexity of our problem, i.e., whether they lead to tractability. It turned out that this is not the case, as the problem remains coNP-hard in most of the cases. The only exceptions are the classes based on linearity, acyclicity, and stickiness when we focus on the IAR semantics. In these cases, we show that the problem is in AC_0 in data complexity. This is established via FO-rewritability, which in turn relies on the fact that the classes in question are UCQ-rewritable.

As for future work, apart from performing a complexity analysis with other semantics, we believe that it is important to empirically evaluate the performance of the various inconsistency-tolerant semantics with respect to their expressive power. Many of the semantics proposed in the literature have been designed to trade off expressive power for computational tractability. Many also maintain soundness with respect to the AR semantics as a kind of quality guarantee that is difficult to evaluate in practice. One way to compare the performance of the various alternative semantics is therefore to design quality metrics that yield objective comparisons of expressive power in practice. This would also shed light on the performance of semantics that go beyond the classical concept of data repair, which are typically not designed to be sound with respect to the AR semantics. The main challenges in this line of work involve selecting adequate real-world datasets, as well as designing well-founded methods to synthetically generate datasets.

Acknowledgments

We would like to thank Marco Calautti for bringing to our attention Corollary 4 from [46], which was crucial for showing that $\text{QAns}_{\text{ICR}}(\mathbf{A}_{\perp})$ is in $\text{P}^{\text{NEXPTIME}}$ in \mathbf{c} -complexity. This work was supported by the UK EPSRC grants EP/J008346/1, EP/R013667/1, EP/L012138/1, and EP/M025268/1, by the AXA Research Fund, and by the Alan Turing Institute under the EPSRC grant EP/N510129/1. M.V. Martinez and G.I. Simari were partially supported by Proyecto PIP-CONICET 112-201101-01000. A. Pieris was also supported by the EPSRC grant EP/S003800/1.

References

- [1] M. Arenas, L. E. Bertossi, J. Chomicki, Consistent query answers in inconsistent databases, in: PODS, 1999, pp. 68–79.
- [2] C. Beeri, M. Y. Vardi, The implication problem for data dependencies, in: ICALP, 1981, pp. 73–85.
- [3] M. Bienvenu, On the complexity of consistent query answering in the presence of simple ontologies, in: AAI, 2012, pp. 705–711.
- [4] M. Bienvenu, Inconsistency handling in ontology-mediated query answering: A progress report, in: DL, 2019.
- [5] M. Bienvenu, C. Bourgaux, Inconsistency-tolerant querying of description logic knowledge bases, in: Reasoning Web, 2016, pp. 156–202.
- [6] M. Bienvenu, C. Bourgaux, F. Goasdoué, Querying inconsistent description logic knowledge bases under preferred repair semantics, in: AAI, 2014, pp. 996–1002.
- [7] M. Bienvenu, C. Bourgaux, F. Goasdoué, Computing and explaining query answers over inconsistent dl-lite knowledge bases, *J. Artif. Intell. Res.* 64 (2019) 563–644.
- [8] M. Bienvenu, R. Rosati, New inconsistency-tolerant semantics for robust ontology-based data access, in: DL, 2013, pp. 53–64.
- [9] M. Bienvenu, R. Rosati, Tractable approximations of consistent query answering for robust ontology-based data access, in: IJCAI, 2013, pp. 775–781.
- [10] A. Cali, G. Gottlob, M. Kifer, Taming the infinite chase: Query answering under expressive relational constraints, *J. Artif. Intell. Res.* 48 (2013) 115–174.
- [11] A. Cali, G. Gottlob, T. Lukasiewicz, A general Datalog-based framework for tractable query answering over ontologies, *J. Web Sem.* 14 (2012) 57–83.
- [12] A. Cali, G. Gottlob, T. Lukasiewicz, B. Marnette, A. Pieris, Datalog+/-: A family of logical knowledge representation and query languages for new applications, in: LICS, 2010, pp. 228–242.
- [13] A. Cali, G. Gottlob, A. Pieris, Towards more expressive ontology languages: The query answering problem, *Artif. Intell.* 193 (2012) 87–128.
- [14] Í. Í. Ceylan, T. Lukasiewicz, E. Malizia, A. Vaiceničius, Explanations for query answers under existential rules, in: IJCAI, 2019, pp. 1639–1646.
- [15] Í. Í. Ceylan, T. Lukasiewicz, E. Malizia, A. Vaiceničius, Explanations for ontology-mediated query answering in description logics, in: ECAI, 2020, pp. 672–679.
- [16] J. Chomicki, J. Marcinkowski, Minimal-change integrity maintenance using tuple deletions, *Inf. Comput.* 197 (2005) 90–121.
- [17] E. Dantsin, T. Eiter, G. Gottlob, A. Voronkov, Complexity and expressive power of logic programming, *ACM Comput. Surv.* 33 (2001) 374–425.
- [18] E. Dantsin, A. Voronkov, Complexity of query answering in logic databases with complex values, in: LFCS, 1997, pp. 56–66.

- [19] A. Deutsch, A. Nash, J. B. Remmel, The chase revisited, in: PODS, 2008, pp. 149–158.
- [20] T. Eiter, T. Lukasiewicz, L. Predoiu, Generalized consistent query answering under existential rules, in: KR, 2016, pp. 359–368.
- [21] R. Fagin, P. G. Kolaitis, R. J. Miller, L. Popa, Data exchange: Semantics and query answering, *Theor. Comput. Sci.* 336 (2005) 89–124.
- [22] G. Gottlob, E. Malizia, Achieving new upper bounds for the hypergraph duality problem through logic, in: CSL-LICS, 2014, pp. 43:1–43:10.
- [23] G. Gottlob, M. Manna, A. Pieris, Polynomial combined rewritings for existential rules, in: KR, 2014, pp. 268–277.
- [24] G. Gottlob, G. Orsi, A. Pieris, Query rewriting and optimization for ontological databases, *ACM Trans. Database Syst.* 39 (2014) 25:1–25:46.
- [25] G. Greco, E. Malizia, L. Palopoli, F. Scarcello, On the complexity of core, kernel, and bargaining set, *Artif. Intell.* 175 (2011) 1877–1910.
- [26] L. A. Hemachandra, The strong exponential hierarchy collapses, *J. Comput. Syst. Sci.* 39 (1989) 299–322.
- [27] D. S. Johnson, A catalog of complexity classes, in: J. van Leeuwen (Ed.), *Handbook of Theoretical Computer Science (Vol. A)*, Elsevier, 1990, pp. 67–161.
- [28] D. Lembo, M. Lenzerini, R. Rosati, M. Ruzzi, D. F. Savo, Inconsistency-tolerant semantics for description logics, in: RR, 2010, pp. 103–117.
- [29] D. Lembo, M. Lenzerini, R. Rosati, M. Ruzzi, D. F. Savo, Query rewriting for inconsistent dl-lite ontologies, in: RR, 2011, pp. 155–169.
- [30] D. Lembo, M. Lenzerini, R. Rosati, M. Ruzzi, D. F. Savo, Inconsistency-tolerant query answering in ontology-based data access, *J. Web Semant.* 33 (2015) 3–29.
- [31] D. Lembo, M. Ruzzi, Consistent query answering over description logic ontologies, in: RR, 2007, pp. 194–208.
- [32] T. Lukasiewicz, E. Malizia, On the complexity of mCP -nets, in: AAI, 2016, pp. 558–564.
- [33] T. Lukasiewicz, E. Malizia, A novel characterization of the complexity class θ_k^P based on counting and comparison, *Theor. Comput. Sci.* 694 (2017) 21–33.
- [34] T. Lukasiewicz, E. Malizia, C. Molinaro, Complexity of approximate query answering under inconsistency in Datalog+/-, in: IJCAI, 2018, pp. 1921–1927.
- [35] T. Lukasiewicz, E. Malizia, C. Molinaro, Explanations for inconsistency-tolerant query answering under existential rules, in: AAI, 2020, pp. 2909–2916.
- [36] T. Lukasiewicz, E. Malizia, A. Vaicenavicius, Complexity of inconsistency-tolerant query answering in datalog+/- under cardinality-based repairs, in: AAI, 2019, pp. 2962–2969.
- [37] T. Lukasiewicz, M. V. Martinez, A. Pieris, G. I. Simari, From classical to consistent query answering under existential rules, in: AAI, 2015, pp. 1546–1552.
- [38] T. Lukasiewicz, M. V. Martinez, G. I. Simari, Inconsistency handling in Datalog+/- ontologies, in: ECAI, 2012, pp. 558–563.
- [39] T. Lukasiewicz, M. V. Martinez, G. I. Simari, Inconsistency-tolerant query rewriting for linear Datalog+/-, in: Datalog 2.0, 2012, pp. 123–134.
- [40] T. Lukasiewicz, M. V. Martinez, G. I. Simari, Complexity of inconsistency-tolerant query answering in Datalog+/-, in: ODBASE, 2013, pp. 488–500.
- [41] M. V. Martinez, G. I. Simari, Explanation-friendly query answering under uncertainty, in: Reasoning Web, 2019, pp. 65–103.

- [42] C. H. Papadimitriou, *Computational Complexity*, Addison-Wesley, 1994.
- [43] A. Poggi, D. Lembo, D. Calvanese, G. De Giacomo, M. Lenzerini, R. Rosati, Linking data to ontologies, *J. Data Semantics* 10 (2008) 133–173.
- [44] R. Rosati, On the complexity of dealing with inconsistency in description logic ontologies, in: *IJCAI*, 2011, pp. 1057–1062.
- [45] M. Schaefer, Graph ramsey theory and the polynomial hierarchy, *J. Comput. Syst. Sci.* 62 (2001) 290–322.
- [46] U. Schöning, K. W. Wagner, Collapsing oracle hierarchies, census functions and logarithmically many queries, in: *STACS*, 1988, pp. 91–97.
- [47] M. Y. Vardi, On the complexity of bounded-variable queries, in: *PODS*, 1995, pp. 266–276.