

Query Answering with Transitive and Linear-Ordered Data

Antoine Amarilli

LTCI, CNRS, Télécom ParisTech
Université Paris-Saclay

Michael Benedikt

University of Oxford

Pierre Bourhis

CNRS CRISAL, Université Lille 1
INRIA Lille

Michael Vanden Boom

University of Oxford

Abstract

We consider entailment problems involving powerful constraint languages such as *guarded existential rules*, in which additional semantic restrictions are put on a set of distinguished relations. We consider restricting a relation to be transitive, restricting a relation to be the transitive closure of another relation, and restricting a relation to be a linear order. We give some natural generalizations of guardedness that allow inference to be decidable in each case, and isolate the complexity of the corresponding decision problems. Finally we show that slight changes in our conditions lead to undecidability.

1 Introduction

The *query answering problem* (or certain answer problem), abbreviated here as QA, is a fundamental reasoning problem in both knowledge representation and databases. It asks whether a query (e.g. given by an existentially-quantified conjunction of atoms) is entailed by a set of constraints and a set of facts. A common class of constraints used for QA are the *existential rules*, also known as *tuple generating dependencies* (TGDs). Although query answering is known to be undecidable for general TGDs, there are a number of subclasses that admit decidable QA, such as those based on *guardedness*. For instance, *guarded* TGDs require all variables in the body of the dependency to appear in a single body atom (the *guard*). *Frontier-guarded* TGDs (FGTGDs) relax this condition and require only that some guard atom contains the variables that occur in both head and body [Baget *et al.*, 2011]. This includes standard SQL referential constraints as well as important constraint classes (e.g. role inclusions) arising in knowledge representation. Guarded existential rules can be generalized to *guarded logics* that allow disjunction and negation and still enjoy decidable QA, e.g. the guarded fragment of first-order logic (GF) [Andréka *et al.*, 1998] and the Guarded Negation Fragment (GNF) [Bárány *et al.*, 2011].

A key challenge is to extend these results to capture additional semantics of the relations. For example, the property that a binary relation is *transitive* cannot be expressed

directly in guarded logics, and neither can the property that one relation is the *transitive closure* of another. Going beyond transitivity, one cannot express that a binary relation is a *linear order*. Since ordered data is common in applications, this means that a key part of data semantics is being lost.

There has been extensive work on decidability results for guarded logics thus extended with semantic restrictions. We first review known results for the *satisfiability problem*.

Ganzinger *et al.* [1999] showed that satisfiability is not decidable for GF when two relations are restricted to be transitive, even on *arity-two* signatures (i.e. with only unary and binary relations). For linear orders, [Kieronski, 2011] showed GF is undecidable when three relations are restricted to be (non-strict) linear orders, even with only two variables (so on arity-two signatures). Otto [2001] showed satisfiability is decidable for two-variable logic with one relation restricted to be a linear order. For transitive relations, one way to regain decidability for GF satisfiability was shown by Szwast and Tendera [2004]: allow transitive relations *only* in guards.

We now turn to the QA problem. Gottlob *et al.* [2013] showed that query answering for GF with transitive relations only in guards is undecidable, even on arity-two signatures. Baget *et al.* [2015] studied QA with respect to a collection of linear TGDs (those with only a single atom in the body and the head). They showed that the query answering problem is decidable with such TGDs and transitive relations, if the signature is binary or if other additional restrictions are obeyed.

Transitivity has also been studied in description logics, where the signature contains unary relations (concepts) and binary relations (roles). In this arity-two context, QA is decidable for many description logics featuring expressive operators as well as transitivity, such as *ZIQ*, *ZOQ*, *ZOI* [Calvanese *et al.*, 2009], Horn-*SROIQ* [Ortiz *et al.*, 2011], or regular- \mathcal{EL}^{++} [Krötzsch and Rudolph, 2007], but they often restrict the interaction between transitivity and some features such as role inclusions and Boolean role combinations. QA becomes undecidable for more expressive description logics with transitivity such as *ALCOIF** [Ortiz *et al.*, 2010] and *ZOIQ* [Ortiz de la Fuente, 2010], and the problem is open for *SROIQ* and *SHOIQ* [Ortiz and Šimkus, 2012].

The main contribution of this work is to introduce a broad class of constraints over arbitrary-arity vocabularies where

query answering is decidable when additional semantics are imposed on some *distinguished relations*. We show that transitivity restrictions can be handled in guarded and frontier-guarded constraints, as long as these distinguished relations are *not* used as guards — we call this new kind of restriction *base-guardedness*. The base-guarded restriction is orthogonal to prior decidable cases such as transitive guards [Szwast and Tendera, 2004] or linear rules [Baget *et al.*, 2015].

On the one hand, we show that the condition allows us to define very expressive and flexible decidable logics, capable of expressing not only guarded existential rules, but guarded rules with negation and disjunction in the head. They can express not only integrity constraints but also conjunctive queries and their negations. On the other hand, a by-product of our results is new query answering schemes for some previously-studied classes of guarded existential rules with extra semantic restrictions. For example, our base-frontier-guarded constraints encompass all *frontier-one TGDs* (where at most one variable is shared between the body and head) [Baget *et al.*, 2009]. Hence, our results imply that QA is decidable with transitive closure and frontier-one constraints, which answers a question of [Baget *et al.*, 2015]. In fact, beyond transitivity assertions, our results even extend to frontier-one TGDs with distinguished relations that are required to be the transitive closure of other relations.

Our results are shown by arguing that it is enough to consider entailment over “tree-like” sets of facts. By representing the set of witness representations as a tree automaton, we derive upper bounds for the combined complexity of the problem. The sufficiency of tree-like examples also enable a refined analysis of *data complexity* (when the query and constraints are fixed). Further, we use a set of coding techniques to show matching lower bounds within our fragment. We also show that loosening our conditions leads to undecidability.

Finally, we solve the QA problem when the distinguished relations are *linear orders*. We show that it is undecidable even assuming base-guardedness, so we introduce a stronger condition called *base-coveredness*: not only are distinguished relations never used as guards, they are always *covered* by a non-distinguished atom. Our decidability technique works by “compiling away” linear order restrictions, obtaining an entailment problem without any special restrictions. The correctness proof for our reduction to classical QA again relies on the ability to restrict reasoning to sets of facts with tree-like representations. To our knowledge, these are the first decidability results for the QA problem with linear orders. Again, we give tight complexity bounds, and show that weakening the base-coveredness condition leads to undecidability.

Proofs are omitted in this extended abstract.

2 Preliminaries

We work on a *relational signature* σ where each relation $R \in \sigma$ has an associated *arity* (written $\text{arity}(R)$). A *fact* $R(\vec{a})$ (or *R-fact*) consists of a relation $R \in \sigma$ and domain elements \vec{a} , with $|\vec{a}| = \text{arity}(R)$. We denote a (finite or infinite) set of facts over σ by \mathcal{F} . We write $\text{elems}(\mathcal{F})$ for the set of elements mentioned in the facts in \mathcal{F} .

We consider *constraints* and *queries* given in fragments of

first-order logic (FO). For simplicity, we disallow constants in constraints and queries, although our results extend with them. Given a set of facts \mathcal{F} and a sentence φ in FO, we talk of \mathcal{F} *satisfying* φ in the usual way.

The queries that we will use are *conjunctive queries* (CQ), namely, existentially quantified conjunction of atoms, which we restrict for simplicity to be Boolean. We also allow *unions of conjunctive queries* (UCQs), namely, disjunctions of CQs.

Problems considered. Given a *finite* set of facts \mathcal{F}_0 , constraints Σ and query Q (given as FO sentences), we say that \mathcal{F}_0 and Σ *entail* Q if for every (possibly infinite) $\mathcal{F} \supseteq \mathcal{F}_0$ satisfying Σ , \mathcal{F} satisfies Q . This amounts to asking whether $\mathcal{F}_0 \wedge \Sigma \wedge \neg Q$ is satisfiable (by a possibly infinite set of facts). We write $\text{QA}(\mathcal{F}_0, \Sigma, Q)$ for this decision problem, called the *query answering* problem.

In this paper, we study the QA problem when imposing semantic constraints on some *distinguished relations*. We thus split the signature as $\sigma := \sigma_B \sqcup \sigma_D$, where σ_B is the *base signature* (its relations are the *base* relations), and σ_D is the *distinguished* signature. All distinguished relations are required to be binary, and they will be assigned special semantics. We study three kinds of special semantics.

We say \mathcal{F}_0, Σ *entails* Q *over transitive relations*, and write $\text{QATr}(\mathcal{F}_0, \Sigma, Q)$ for the corresponding problem, if $\mathcal{F}_0 \wedge \Sigma \wedge \neg Q$ is satisfied by some set of facts \mathcal{F} where each distinguished relation R_i^+ in σ_D is required to be *transitive*¹ in \mathcal{F} .

We say \mathcal{F}_0, Σ *entails* Q *over transitive closure*, and write $\text{QATc}(\mathcal{F}_0, \Sigma, Q)$ for this problem, if the same holds on some \mathcal{F} where each relation R_i^+ of σ_D is interpreted as the transitive closure of a corresponding binary base relation $R_i \in \sigma_B$.

We say \mathcal{F}_0, Σ *entails* Q *over linear orders*, and write $\text{QAlin}(\mathcal{F}_0, \Sigma, Q)$, if the same holds on some \mathcal{F} where each relation $<_{i \in \sigma_D}$ is required to be a strict linear order on the elements of \mathcal{F} .

We now define the constraint languages (which are all fragments of FO) for which we study these QA problems.

Dependencies. The first constraint languages we study are restricted classes of *tuple-generating dependencies* (TGDs). A TGD is a first-order sentence τ of the form $\forall \vec{x} (\bigwedge_i \gamma_i(\vec{x}) \rightarrow \exists \vec{y} \bigwedge_i \rho_i(\vec{x}, \vec{y}))$ where $\bigwedge_i \gamma_i$ and $\bigwedge_i \rho_i$ are conjunctions of atoms respectively called the *body* and *head* of τ .


We will be interested in TGDs that are *guarded* in various ways. A *guard* for \vec{x} is an atom from σ using every variable in \vec{x} . We will be particularly interested in *base-guards*, which are guards coming from the base relations σ_B .

A *frontier-guarded TGD* (FGTGD) is a TGD whose body contains a guard for the *frontier variables* — variables that occur in both head and body. It is a *base frontier-guarded TGD* (BaseFGTGD) if there is a base-guard for the frontier variables. We allow equality atoms $x = x$ to be guards, so BaseFGTGD subsumes *frontier-one TGDs*, which have one frontier variable. Frontier-guarded and frontier-one TGDs

¹Note that we work with *transitive* relations, which may not be *reflexive*, unlike, e.g., R^* roles in \mathcal{ZOIQ} description logics [Calvanese *et al.*, 2009]. However, our results are easily seen to adapt when replacing transitive predicates (and transitive closure) by transitive and reflexive predicates (and closure).

Fragment	QAttr		QAtc		QAlin	
	data	combin.	data	combin.	data	combin.
BaseGNF	coNP-c	2EXP-c	coNP-c	2EXP-c	undecidable	
BaseCovGNF	coNP-c	2EXP-c	coNP-c	2EXP-c	coNP-c	2EXP-c
BaseFGTGD	in coNP	2EXP-c	coNP-c	2EXP-c	undecidable	
BaseCovFGTGD	P-c	2EXP-c	coNP-c	2EXP-c	coNP-c	2EXP-c

(a) Summary of QA results (for base-covered fragments, queries are also base-covered)



(b) Taxonomy of fragments

have been shown to have an attractive combination of expressivity and computational properties [Baget *et al.*, 2011].

We also introduce the more restricted class of *base-covered frontier-guarded TGDs* (BaseCovFGTGD): they are the BaseFGTGDs where, for every $\sigma_{\mathcal{D}}$ -atom in the body, there is a base guard in the body containing its variables (but each $\sigma_{\mathcal{D}}$ -atom may have a different base guard).

An important special case of frontier-guarded TGDs for applications are *inclusion dependencies* (ID). An ID is a FGTGD where the body and head contain a single atom, and no variable occurs twice in the same atom. A *base inclusion dependency* BaseID is an ID where the body atom is in $\sigma_{\mathcal{B}}$, so it is both base-guarded and base-covered.

Guarded logics. Moving beyond TGDs, the second kind of constraint that we study are *guarded logics*.

The *guarded negation fragment* (GNF) is the fragment of FO which contains all atoms, and is closed under conjunction, disjunction, existential quantification, and the following form of negation: for any GNF formula $\varphi(\vec{x})$ and atom $A(\vec{x}, \vec{y})$ with free variables exactly as indicated, the formula $A(\vec{x}, \vec{y}) \wedge \neg\varphi(\vec{x})$ is in GNF. That is, existential quantification may be unguarded, but the free variables in any negated subformula must be guarded; universal quantification must be expressed with existential quantification and negation. GNF can express all FGTGDs, as well as non-TGD constraints and UCQs. For instance, as it allows disjunction, GNF can express *disjunctive inclusion dependencies*, DIDs, which generalize IDs: their body is a single atom with no repeated variables, and their head is a disjunction of atoms with no repeated variables.

We introduce the *base-guarded negation fragment* BaseGNF over σ : it is defined like GNF, but requires *base guards* instead of guards. The *base-covered guarded negation fragment* BaseCovGNF over σ consists of BaseGNF formulas such that every $\sigma_{\mathcal{D}}$ -atom A that appears negatively (i.e., under the scope of an odd number of negations) appears conjoined with a base guard — i.e., a $\sigma_{\mathcal{B}}$ -atom containing all variables of A . This technical condition is designed to generalize BaseCovFGTGDs. Indeed, a BaseCovFGTGD of the form $\forall\vec{x}(\bigwedge\gamma_i \rightarrow \exists\vec{y}\bigwedge\rho_i)$ can be written in BaseCovGNF as $\neg\exists\vec{x}(\bigwedge\gamma_i \wedge \neg\exists\vec{y}\bigwedge\rho_i)$.

We call a CQ Q *base-covered* if each $\sigma_{\mathcal{D}}$ -atom in Q has a $\sigma_{\mathcal{B}}$ -atom of Q containing its variables. This is the same as saying that $\neg Q$ is in BaseCovGNF. A UCQ is base-covered if each disjunct is.

Examples. We conclude the preliminaries by giving a few examples. Consider a signature with a binary base relation B , a ternary base relation C , and a distinguished relation R^+ .

- $\forall xyz((R^+(x, y) \wedge R^+(y, z)) \rightarrow R^+(x, z))$ is a TGD, but is not a FGTGD since the frontier variables x, z are

not guarded. It cannot even be expressed in GNF.

- $\forall xy(R^+(x, y) \rightarrow B(x, y))$ is an ID, hence a FGTGD. It is not a BaseID or even in BaseGNF, since the frontier variables are not base-guarded.
- $\forall xyz((B(z, x) \wedge R^+(x, y) \wedge R^+(y, z)) \rightarrow R^+(x, z))$ is a BaseFGTGD. It is not a BaseCovFGTGD since there are no base atoms in the body to cover x, y and z .
- $\exists wxyz(R^+(w, x) \wedge R^+(x, y) \wedge R^+(y, z) \wedge R^+(z, w) \wedge C(w, x, y) \wedge C(y, z, w))$ is a base-covered CQ.
- $\exists xy(B(x, y) \wedge \neg(R^+(x, y) \wedge R^+(y, x)) \wedge (R^+(x, y) \vee R^+(y, x)))$ cannot be rewritten as a TGD. But it is in BaseCovGNF.

Our main results are summarized in Table a, and the languages that we study are illustrated in Figure b. In particular, QAttr and QAtc are decidable for BaseGNF. This includes base-frontier-guarded rules, which allow one to use a transitive relation such as “part-of” (or even its transitive closure) whenever only one variable is to be exported to the head. This latter condition holds in the translations of many classical description logics. Our results also imply that QAlin is decidable for BaseCovGNF, which allows constraints that arise from data integration and data exchange over attributes with linear orders — e.g. views defined by selecting rows of a table where some inequality involving the attributes is satisfied.

3 Decidability results for transitivity

We first consider QAtc, where $\sigma_{\mathcal{B}}$ includes binary relations R_1, \dots, R_n , and $\sigma_{\mathcal{D}}$ consists of binary relations R_1^+, \dots, R_n^+ such that R_i^+ is the transitive closure of R_i .

Theorem 1. *We can decide QAtc(\mathcal{F}_0, Σ, Q) in 2EXPTIME, where \mathcal{F}_0 ranges over finite sets of facts, Σ over BaseGNF constraints, and Q over UCQs. In particular, this holds when Σ consists of BaseFGTGDs.*

In order to prove Theorem 1, we give a decision procedure to determine whether $\mathcal{F}_0 \wedge \Sigma \wedge \neg Q$ is satisfiable, when R_i^+ is interpreted as the transitive closure of R_i . When $\Sigma \in$ BaseGNF and Q is a Boolean UCQ, then $\Sigma \wedge \neg Q$ is in BaseGNF. So it suffices to show that BaseGNF satisfiability is decidable, when properly interpreting R_i^+ .

As mentioned in the introduction, our proofs rely heavily on the fact that in query answering problems for these constraint languages, one can restrict to sets of facts that have a “tree-like” structure. We now make this notion precise. A *tree decomposition* of \mathcal{F} consists of a tree (T, Child) and a labelling function λ associating each node of the tree T to a set of facts of \mathcal{F} , called the *bag* of that node, that satisfies the following conditions: (i) each fact of \mathcal{F} must be in the image

of λ ; (ii) for each element $e \in \text{elems}(\mathcal{F})$, the set of nodes whose bag uses e is a connected subset of T . It is \mathcal{F}_0 -rooted if the root node is associated with \mathcal{F}_0 . It has *width* $k - 1$ if each bag other than the root mentions at most k elements.

For a number k , a σ sentence φ is said to have *transitive-closure friendly k -tree-like witnesses* if: for every finite set of facts \mathcal{F}_0 , if there is an $\mathcal{F} \supseteq \mathcal{F}_0$ satisfying φ in which each R^+ is interpreted as the transitive closure of R , then there is such an \mathcal{F} that has an \mathcal{F}_0 -rooted $(k - 1)$ -width tree decomposition. We can show that BaseGNF sentences have this kind of k -tree-like witness for an easily computable k . The proof uses a standard technique, involving an unravelling based on “guarded negation bisimulation” [Bárány *et al.*, 2011]:

Proposition 1. *Every sentence φ in BaseGNF has transitive-closure friendly k -tree-like witnesses, where $k \leq |\varphi|$.*

Here k can be taken to be the “width” of φ [Bárány *et al.*, 2011], which is roughly the maximum number of free variables in any subformula. Hence, it suffices to test satisfiability for BaseGNF restricted to sets of facts with tree decompositions of width $|\varphi| - 1$. It is well known that sets of facts of bounded tree-width can be encoded as trees over a finite alphabet that depends only on the signature and the tree-width. This makes the problem amenable to tree automata techniques, since we can design a tree automaton that runs on representations of these tree decompositions and checks whether some sentence holds in the corresponding set of facts.

Theorem 2. *Let φ be a sentence in BaseGNF, and let \mathcal{F}_0 be a finite set of facts. We can construct in 2EXPTIME a 2-way alternating parity tree automaton $\mathcal{A}_{\varphi, \mathcal{F}_0}$ such that*

$$\mathcal{F}_0 \wedge \varphi \text{ is satisfiable} \quad \text{iff} \quad L(\mathcal{A}_{\varphi, \mathcal{F}_0}) \neq \emptyset$$

when each $R_i^+ \in \sigma_{\mathcal{D}}$ is interpreted as the transitive closure of $R_i \in \sigma_{\mathcal{B}}$. The number of states of $\mathcal{A}_{\varphi, \mathcal{F}_0}$ is exponential in $|\varphi| \cdot |\mathcal{F}_0|$ and the number of priorities is linear in $|\varphi|$.

The construction can be viewed as an extension of [Calvanese *et al.*, 2005], and incorporates ideas from automata for guarded logics (see, e.g., [Grädel and Walukiewicz, 1999]).

Because 2-way tree automata emptiness is decidable in time exponential in the number of states and priorities [Vardi, 1998], this yields the 2EXPTIME bound for Theorem 1.

Consequences for QAtr. We can derive results for QAtr by observing that the QAtr problem subsumes it: to enforce that $R^+ \in \sigma_{\mathcal{D}}$ is transitive, simply interpret it as the transitive closure of a relation R that is never otherwise used. Hence:

Corollary 1. *We can decide QAtr(\mathcal{F}_0, Σ, Q) in 2EXPTIME, where \mathcal{F}_0 ranges over finite sets of facts, Σ over BaseGNF constraints (in particular, BaseFGTGD), and Q over UCQs.*

In particular, this result holds for *frontier-one TGDs* (those with a single frontier variable), as a single variable is always base-guarded. This answers a question of [Baget *et al.*, 2015].

Data complexity. Our results in Theorem 1 and Corollary 1 show upper bounds on the *combined complexity* of the QAtr and QAtrc problems. We now turn to the complexity when the query and constraints are fixed but the initial set of facts varies — the *data complexity*.

We first show a CoNP data complexity upper bound for QAtrc for BaseGNF constraints. The algorithm uses the fact

that a counterexample to QAtrc can be taken to have a \mathcal{F}' -rooted tree decomposition, for some \mathcal{F}' that does not add new elements to \mathcal{F}_0 , only new facts. While such a decomposition could be large, it suffices to guess \mathcal{F}' and annotations describing, for each $|\varphi|$ -tuple \vec{c} in \mathcal{F}' , sufficiently many formulas holding in the subtree that interfaces with \vec{c} . The technique generalizes an analogous result in [Bárány *et al.*, 2012].

Theorem 3. *For any fixed BaseGNF constraints Σ and UCQ Q , given a finite set of facts \mathcal{F}_0 , we can decide QAtrc(\mathcal{F}_0, Σ, Q) in CoNP data complexity.*

For FGTGDs, the data complexity of QA is in PTIME [Baget *et al.*, 2011]. We can show that the same holds, but only for BaseCovFGTGDs, and for QAtr rather than QAtrc:

Theorem 4. *For any fixed BaseCovFGTGD constraints Σ and base-covered UCQ Q , given a finite set of facts \mathcal{F}_0 , we can decide QAtr(\mathcal{F}_0, Σ, Q) in PTIME data complexity.*

The proof uses a reduction to the standard QA problem for FGTGDs, and then applies the PTIME result of [Baget *et al.*, 2011]. The reduction again makes use of tree-likeness to show that we can replace the requirement that the R_i^+ are transitive by the weaker requirement of transitivity within small sets (intuitively, within bags of a decomposition). We will also use this idea for linear orders (see Proposition 3).

Restricting to QAtr is in fact essential to make data complexity tractable, as hardness holds otherwise.

Hardness. We now show complexity lower bounds. We already know that all our variants of QA are 2EXPTIME-hard in combined complexity, and CoNP-hard in data complexity, when GNF constraints are allowed: this follows from existing bounds on GNF reasoning [Bárány *et al.*, 2012] even without distinguished predicates. However, in the case of the QAtrc problem, we can show the same result for the much weaker language of BaseIDs.

We do this via a reduction from QA with *disjunctive inclusion dependencies*, which is known to be 2EXPTIME-hard in combined complexity [Bourhis *et al.*, 2013, Theorem 2] and CoNP-hard in data complexity [Calvanese *et al.*, 2006; Bourhis *et al.*, 2013], even without distinguished relations. We use the transitive closure to emulate disjunction (as was already suggested in the description logic context [Horrocks and Sattler, 1999]) by creating an R_i^+ -fact and limiting the length of a witness R_i -path (this limit is imposed by Q'). The choice of the length of the witness path among two possibilities is used to mimic the disjunction. We thus show:

Theorem 5. *For any finite set of facts \mathcal{F}_0 , DIDs Σ , and UCQ Q on a signature σ , we can compute in PTIME a set of facts \mathcal{F}'_0 , BaseIDs Σ' , and a base-covered CQ Q' on a signature σ' (with a single distinguished relation), such that QA(\mathcal{F}_0, Σ, Q) iff QAtrc($\mathcal{F}'_0, \Sigma', Q'$).*

This implies the following, contrasting with Theorem 4:

Corollary 2. *The QAtrc problem with BaseIDs and base-covered CQs is CoNP-hard in data complexity and 2EXPTIME-hard in combined complexity.*

In fact, the data complexity lower bound for QAtrc even holds in the absence of constraints:

Proposition 2. *There is a base-covered CQ Q such that the data complexity of QAtrc($\mathcal{F}_0, \emptyset, Q$) is CoNP-hard.*

We prove this by reducing the problem of 3-coloring a directed graph, known to be NP-hard, to the complement of QAtc. It is well-known how to do this using TGDs that have disjunction in the head. As in the proof of Theorem 5, we simulate this disjunction by using a choice of the length of paths that realize transitive closure facts asserted in \mathcal{F}_0 .

In all of these hardness results, we first prove them for UCQs, and then show how the use of disjunction can be eliminated, using a prior “trick” (see, e.g., [Gottlob and Papadimitriou, 2003]) to code the intermediate truth values of disjunctions within a CQ.

4 Decidability results for linear orders

We now move to QAlin, the setting where the distinguished relations $<_i$ of $\sigma_{\mathcal{D}}$ are *linear* (total) strict orders, i.e., they are transitive, irreflexive, and total. We consider constraints and queries that are base-covered. We prove the following result.

Theorem 6. *We can decide $\text{QAlin}(\mathcal{F}_0, \Sigma, Q)$ in 2EXPTIME, where \mathcal{F}_0 ranges over finite sets of facts, Σ over BaseCovGNF, and Q over base-covered UCQs. In particular, this holds when Σ consists of BaseCovFGTGDs.*

Our technique here is to reduce this to traditional QA where no additional restrictions (like being transitive or a linear order) are imposed. Starting with BaseCovGNF constraints, we reduce to a traditional QA problem with GNF constraints, and hence prove decidability in 2EXPTIME using [Bárány *et al.*, 2012]. However, the reduction is quite simple, and hence could be applicable to other constraint classes.

The idea behind the reduction is to include additional constraints that enforce the linear order conditions. However, we cannot express transitivity or totality in GNF. Hence, we will only add a weakening of these properties that is expressible in GNF, and then argue that this is sufficient for our purposes.

The reduction is described in the following proposition.

Proposition 3. *For any finite set of facts \mathcal{F}_0 , constraints $\Sigma \in \text{BaseCovGNF}$, and base-covered UCQ Q , we can compute \mathcal{F}'_0 and $\Sigma' \in \text{BaseGNF}$ in PTIME such that $\text{QAlin}(\mathcal{F}_0, \Sigma, Q)$ iff $\text{QA}(\mathcal{F}'_0, \Sigma', Q)$.*

In particular, \mathcal{F}'_0 is \mathcal{F}_0 together with facts $G(a, b)$ for every pair $a, b \in \text{elems}(\mathcal{F}_0)$, where G is some fresh binary base relation. Σ' is Σ together with the k -guardedly-linear axioms for each distinguished relation $<$, where k is $|\Sigma \wedge \neg Q|$; namely: (i) guardedly total: $\forall xy((\text{guarded}_{\sigma_{\mathcal{B}} \cup \{G\}}(x, y) \wedge x \neq y) \rightarrow x < y \vee y < x)$ (ii) k -guardedly transitive: $\forall xy((\text{guarded}_{\sigma_{\mathcal{B}} \cup \{G\}}(x, y) \wedge x <^k y) \rightarrow x < y)$ (iii) k -cycle free: $\neg \exists x(x <^k x)$, where $\text{guarded}_{\sigma_{\mathcal{B}} \cup \{G\}}(x, y)$ is the formula expressing that x, y is base-guarded (an existentially-quantified disjunction over all possible base-guards containing x and y) and $x <^k y$ expresses that there is some $<$ -path of at most k elements that starts at x and ends at y . Note that $x <^k y$ can be written in BaseGNF so, unlike the property of being a linear order, the k -guardedly-linear axioms can be expressed in BaseGNF.

We now sketch the argument for the correctness of the reduction. The easy direction is where we assume $\text{QA}(\mathcal{F}'_0, \Sigma', Q)$ holds, so any $\mathcal{F}' \supseteq \mathcal{F}'_0$ satisfying Σ' must satisfy Q . Now consider $\mathcal{F} \supseteq \mathcal{F}_0$ that satisfies Σ and where

all $<$ in $\sigma_{\mathcal{D}}$ are strict linear orders. We must show that \mathcal{F} satisfies Q . First, observe that \mathcal{F} satisfies Σ' since the k -guardedly-linear axioms for $<$ are clearly satisfied for all k when $<$ is a strict linear order. Now consider the extension of \mathcal{F} to \mathcal{F}' with facts $G(a, b)$ for all $a, b \in \text{elems}(\mathcal{F}_0)$. This must still satisfy Σ' : adding these facts means there are additional k -guardedly-linear requirements on the elements from \mathcal{F}_0 , but these requirements already hold since $<$ is a strict linear order. Hence, by our initial assumption, \mathcal{F}' must satisfy Q . Since Q does not mention G , the restriction of \mathcal{F}' back to \mathcal{F} still satisfies Q as well. Therefore, $\text{QAlin}(\mathcal{F}_0, \Sigma, Q)$ holds.

For the harder direction, suppose for the sake of contradiction that $\text{QA}(\mathcal{F}'_0, \Sigma', Q)$ does not hold, but $\text{QAlin}(\mathcal{F}_0, \Sigma, Q)$ does. Then there is some $\mathcal{F}' \supseteq \mathcal{F}'_0$ such that \mathcal{F}' satisfies $\Sigma' \wedge \neg Q$. We will again rely on the ability to restrict to tree-like \mathcal{F}' , but with a slightly different notion of tree-likeness.

We say a set E of elements from $\text{elems}(\mathcal{F})$ are *base-guarded* in \mathcal{F} if there is some $\sigma_{\mathcal{B}}$ -fact in \mathcal{F} that mentions all of the elements in E . A *base-guarded-interface tree decomposition* $(T, \text{Child}, \lambda)$ for \mathcal{F} is a tree decomposition satisfying the following additional property: for all nodes n_1 that are not the root of T , if n_2 is a child of n_1 and E is the set of elements mentioned in both n_1 and n_2 , then E is base-guarded in \mathcal{F} . A sentence φ has *base-guarded-interface k -tree-like witnesses* if for any finite set of facts \mathcal{F}_0 , if there is some $\mathcal{F} \supseteq \mathcal{F}_0$ satisfying φ then there is such an \mathcal{F} with a \mathcal{F}_0 -rooted $(k-1)$ -width base-guarded-interface tree decomposition.

Although the transformation from Σ to Σ' makes the formula larger, it does not increase the “width” that controls the bag size of tree-like witnesses. Hence, we can show:

Lemma 1. *The sentence $\Sigma' \wedge \neg Q$ has base-guarded-interface k -tree-like witnesses for some $k \leq |\Sigma \wedge \neg Q|$.*

Using this lemma, we can assume that we have some $\mathcal{F}' \supseteq \mathcal{F}'_0$ which has a $(k-1)$ -width base-guarded-interface tree decomposition and witnesses $\Sigma' \wedge \neg Q$. If every $<$ in $\sigma_{\mathcal{D}}$ is a strict linear order in \mathcal{F}' , then restricting \mathcal{F}' to the set of σ -facts yields some \mathcal{F} that would satisfy $\Sigma \wedge \neg Q$, a contradiction. Hence, there are some distinguished relations $<$ that are not strict linear orders in \mathcal{F}' . We can show that such an \mathcal{F}' can actually be extended to some \mathcal{F}'' that still satisfies $\Sigma' \wedge \neg Q$ but where all $<$ in $\sigma_{\mathcal{D}}$ are strict linear orders, which we already argued is impossible.

The crucial part of the argument is thus about extending k -guardedly-linear counterexamples to genuine linear orders:

Lemma 2. *If there is $\mathcal{F}' \supseteq \mathcal{F}'_0$ that satisfies $\Sigma' \wedge \neg Q$ and has a \mathcal{F}'_0 -rooted base-guarded-interface $(k-1)$ -width tree decomposition, then there is $\mathcal{F}'' \supseteq \mathcal{F}'$ that satisfies $\Sigma' \wedge \neg Q$ where each distinguished relation is a strict linear order.*

The proof of Lemma 2 proceeds by showing that sets of facts that have $(k-1)$ -width base-guarded-interface tree decompositions and satisfy k -guardedly-linear axioms must already be cycle-free with respect to $<$. Hence, by taking the transitive closure of $<$ in \mathcal{F} , we get a new set of facts where every $<$ is a strict *partial* order. Any strict partial order can be further extended to a strict linear order using known techniques, so we can obtain $\mathcal{F}'' \supseteq \mathcal{F}'$ where $<$ is a strict partial order. This \mathcal{F}'' may have more $<$ facts than \mathcal{F}' , but the

k -guardedly linear axioms ensure that these new $<$ -facts are only about pairs of elements that are not base-guarded.

It remains to show that \mathcal{F}'' satisfies $\Sigma' \wedge \neg Q$. It is clear that \mathcal{F}'' still satisfies the k -guardedly-linear axioms, but it could no longer satisfy $\Sigma \wedge \neg Q$. However, this is where the base-covered assumption is used: it can be shown that satisfiability of $\Sigma \wedge \neg Q$ in BaseCovGNF is not affected by adding new $<$ facts about pairs of elements that are not base-guarded.

Data complexity. Again, the result of Theorem 6 is a combined complexity upper bound. However, as it works by reducing to traditional QA in PTIME, data complexity upper bounds follow from [Bárány *et al.*, 2012].

Corollary 3. *For any BaseCovGNF constraints Σ and base-covered UCQ Q , given a finite set of facts \mathcal{F}_0 , we can decide $\text{QAlin}(\mathcal{F}_0, \Sigma, Q)$ in CoNP data complexity.*

This is similar to the way data complexity bounds were shown for QAttr (in Theorem 4). However, unlike for the QAttr problem, the constraint rewriting in this section introduces disjunction, so rewriting a QAlin problem for BaseCovFGTGDs does not produce a classical query answering problem for FGTGDs. Thus the rewriting does not imply a PTIME data complexity upper bound for BaseCovFGTGD. Indeed, we will see in Proposition 4 that this is CoNP-hard.

Hardness. QAlin for BaseCovGNF constraints is again immediately CoNP-hard in data complexity, and 2EXPTIME-hard in combined complexity, from the corresponding bounds on GNF [Bárány *et al.*, 2012]. However, we can show that hardness holds for the much weaker constraint language BaseID, by a reduction from DID reasoning, as in Section 3.

Theorem 7. *For any finite set of facts \mathcal{F}_0 , DIDs Σ , and UCQ Q on a signature σ , we can compute in PTIME a set of facts \mathcal{F}'_0 , BaseIDs Σ' , and CQ Q' on a signature σ' (with a single distinguished relation), such that $\text{QA}(\mathcal{F}_0, \Sigma, Q)$ iff $\text{QAlin}(\mathcal{F}'_0, \Sigma', Q')$.*

The reduction allows us to transfer hardness results for DID from [Calvanese *et al.*, 2006; Bourhis *et al.*, 2013], exactly as was done in Theorem 5, to conclude:

Corollary 4. *The QAlin problem with BaseID and base-covered CQs is CoNP-hard in data complexity and 2EXPTIME-hard in combined complexity.*

Again, as in the previous section, the data complexity lower bound even holds in the absence of constraints:

Proposition 4. *There is a base-covered CQ Q such that the data complexity of $\text{QAlin}(\mathcal{F}, \emptyset, Q)$ is CoNP-hard.*

5 Undecidability results for transitivity

We have shown in Section 3 that query answering is decidable with transitive relations (even with transitive closure), BaseFGTGDs, and UCQs (Theorem 1). Removing our base-guarded condition leads to undecidability of QAtc, even when constraints are inclusion dependencies:

Theorem 8. *There is a signature $\sigma = \sigma_B \sqcup \sigma_D$ with a single distinguished predicate S^+ in σ_D , a set Σ of IDs on σ , and a CQ Q on σ_B , such that the following problem is undecidable: given a finite set of facts \mathcal{F}_0 , decide $\text{QAtc}(\mathcal{F}_0, \Sigma, Q)$.*

The proof is by reduction from a tiling problem. The constraints use a transitive successor relation to define a grid of integer pairs. It then uses transitive closure to emulate disjunction, as in Theorem 5, and relies on Q to test for forbidden adjacent tile patterns.

An analogous result can be shown for QAttr, using (non-base-guarded) disjunctive inclusion dependencies:

Theorem 9. *There is an arity-two signature $\sigma = \sigma_B \sqcup \sigma_D$ with a single distinguished predicate S^+ in σ_D , a set Σ of DIDs on σ , a CQ Q on σ_B , such that the following problem is undecidable: given a finite set of facts \mathcal{F}_0 , decide $\text{QAttr}(\mathcal{F}_0, \Sigma, Q)$.*

These results complement the undecidability results of [Gottlob *et al.*, 2013, Theorem 2], which showed that, on arity-two signatures, QAttr is undecidable with guarded TGDs and atomic CQs, even when transitive relations occur only in guards. Our results also contrast with the decidability results of [Baget *et al.*, 2015] which apply to QAttr: our Theorem 8 shows that their results cannot extend to QAtc.

6 Undecidability results for linear orders

Section 4 has shown that QAlin is decidable for base-covered CQs and BaseCovGNF constraints. Dropping the base-covered requirement on the query leads to undecidability:

Theorem 10. *There is a signature $\sigma = \sigma_B \sqcup \sigma_D$ where σ_D is a single strict linear order relation, a CQ Q on σ , and a set Σ of inclusion dependencies on σ_B (i.e., not mentioning the linear order, so in particular base-covered), such that the following problem is undecidable: given a finite set of facts \mathcal{F}_0 , decide $\text{QAlin}(\mathcal{F}_0, \Sigma, Q)$.*

This result is close to [Gutiérrez-Basulto *et al.*, 2013, Theorem 3], which deals not with a linear order, but inequalities in queries, which we can express with a linear order. However, this requires a UCQ. As in our prior hardness and undecidability results, we can adapt the technique to use a CQ.

By letting $\Sigma' := \Sigma \wedge \neg Q$ where Σ and Q are as in the previous theorem, we obtain base-guarded constraints for which QAlin is undecidable. In fact, Σ' can be expressed as a set of BaseFGTGDs. This implies that the base-covered requirement is necessary for the constraint language:

Corollary 5. *There is a signature $\sigma = \sigma_B \sqcup \sigma_D$ where σ_D is a single strict linear order relation, and a set Σ' of BaseFGTGD constraints, such that, letting \top be the tautological query, the following problem is undecidable: given a finite set of facts \mathcal{F}_0 , decide $\text{QAlin}(\mathcal{F}_0, \Sigma', \top)$.*

7 Conclusion

We have given a detailed picture of the impact of transitivity, transitive closure, and linear order restrictions on query answering problems for a broad class of guarded constraints. We have shown that transitive relations and transitive closure restrictions can be handled in guarded constraints as long as they are not used in guards. For linear orders, the same is true if order atoms are covered by base atoms. This implies the analogous results for frontier-guarded TGDs, in particular frontier-one. But in the linear order case we show that PTIME data complexity cannot always be preserved.

We leave open the question of entailment over *finite* sets of facts. There are few techniques for deciding entailment over finite sets of facts for logics where it does not coincide with general entailment (and for the constraints considered here it does not coincide). An exception is [Bárány and Bojańczyk, 2012], but it is not clear if the techniques there can be extended to our constraint languages.

Acknowledgments

Amarilli was partly funded by the Télécom ParisTech Research Chair on Big Data and Market Insights. Bourhis was supported by CPER Nord-Pas de Calais/FEDER DATA Advanced Data Science and Technologies 2015-2020 and the ANR Aggreg Project ANR-14-CE25-0017, INRIA Northern European Associate Team Integrated Linked Data. Benedikt’s work was sponsored by the Engineering and Physical Sciences Research Council of the United Kingdom (EPSRC), grants EP/M005852/1 and EP/L012138/1. Vanden Boom was partially supported by EPSRC grant EP/L012138/1.

References

- [Andréka *et al.*, 1998] Hajnal Andréka, István Németi, and Johan van Benthem. Modal languages and bounded fragments of predicate logic. *J. Philosophical Logic*, 27(3), 1998.
- [Baget *et al.*, 2009] Jean-François Baget, Michel Leclère, Marie-Laure Mugnier, and Eric Salvat. Extending decidable cases for rules with existential variables. In *IJCAI*, 2009.
- [Baget *et al.*, 2011] Jean-François Baget, Marie-Laure Mugnier, Sebastian Rudolph, and Michaël Thomazo. Walking the complexity lines for generalized guarded existential rules. In *IJCAI*, 2011.
- [Baget *et al.*, 2015] Jean-François Baget, Meghyn Bienvenu, Marie-Laure Mugnier, and Swan Rocher. Combining existential rules and transitivity: Next steps. In *IJCAI*, 2015.
- [Bárány and Bojańczyk, 2012] Vince Bárány and Mikołaj Bojańczyk. Finite satisfiability for guarded fixpoint logic. *IPL*, 112(10):371–375, 2012.
- [Bárány *et al.*, 2011] Vince Bárány, Balder ten Cate, and Luc Segoufin. Guarded negation. In *ICALP*, 2011.
- [Bárány *et al.*, 2012] Vince Bárány, Balder ten Cate, and Martin Otto. Queries with guarded negation. *PVLDB*, 5(11):1328–1339, 2012.
- [Bourhis *et al.*, 2013] Pierre Bourhis, Michael Morak, and Andreas Pieris. The impact of disjunction on query answering under guarded-based existential rules. In *IJCAI*, 2013.
- [Calvanese *et al.*, 2005] Diego Calvanese, Giuseppe De Giacomo, and Moshe Y. Vardi. Decidable containment of recursive queries. *Theor. Comput. Sci.*, 336(1):33–56, 2005.
- [Calvanese *et al.*, 2006] Diego Calvanese, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Data complexity of query answering in description logics. In *KR*, 2006.
- [Calvanese *et al.*, 2009] Diego Calvanese, Thomas Eiter, and Magdalena Ortiz. Regular path queries in expressive description logics with nominals. In *IJCAI*, 2009.
- [Ganzinger *et al.*, 1999] Harald Ganzinger, Christoph Meyer, and Margus Veanes. The two-variable guarded fragment with transitive relations. In *LICS*, 1999.
- [Gottlob and Papadimitriou, 2003] Georg Gottlob and Christos Papadimitriou. On the complexity of single-rule datalog queries. *Inf. Comp.*, 183, 2003.
- [Gottlob *et al.*, 2013] Georg Gottlob, Andreas Pieris, and Lidia Tendera. Querying the guarded fragment with transitivity. In *ICALP*, 2013.
- [Grädel and Walukiewicz, 1999] Erich Grädel and Igor Walukiewicz. Guarded fixed point logic. In *LICS*, 1999.
- [Gutiérrez-Basulto *et al.*, 2013] Victor Gutiérrez-Basulto, Yazmin Ibañez Garcia, Roman Kontchakov, and Egor V. Kostylev. Conjunctive queries with negation over DL-Lite: A closer look. In *Web Reasoning and Rule Systems*, 2013.
- [Horrocks and Sattler, 1999] Ian Horrocks and Ulrike Sattler. A description logic with transitive and inverse roles and role hierarchies. *Journal of Logic and Computation*, 9(3):385–410, 1999.
- [Kieronski, 2011] Emanuel Kieronski. Decidability issues for two-variable logics with several linear orders. In *CSL*, 2011.
- [Krötzsch and Rudolph, 2007] Markus Krötzsch and Sebastian Rudolph. Conjunctive queries for \mathcal{EL} with role composition. In *DL*, 2007.
- [Ortiz and Šimkus, 2012] Magdalena Ortiz and Mantas Šimkus. Reasoning and query answering in description logics. *Reasoning Web. Semantic Technologies for Advanced Query Answering*, pages 1–53, 2012.
- [Ortiz *et al.*, 2010] Magdalena Ortiz, Sebastian Rudolph, and Mantas Šimkus. Query answering is undecidable in DLs with regular expressions, inverses, nominals, and counting. Technical report, Technische Universität Wien, 2010.
- [Ortiz *et al.*, 2011] Magdalena Ortiz, Sebastian Rudolph, and Mantas Šimkus. Query answering in the Horn fragments of the description logics SHOIQ and SROIQ. In *IJCAI*, 2011.
- [Ortiz de la Fuente, 2010] Maria Magdalena Ortiz de la Fuente. *Query answering in expressive description logics*. PhD thesis, Technischen Universität Wien, 2010.
- [Otto, 2001] Martin Otto. Two variable first-order logic over ordered domains. *J. Symb. Log.*, 66(2):685–702, 2001.
- [Szwast and Tendera, 2004] Wiesław Szwast and Lidia Tendera. The guarded fragment with transitive guards. *Annals of Pure and Applied Logic*, 128(13):227 – 276, 2004.
- [Vardi, 1998] Moshe Y. Vardi. Reasoning about the past with two-way automata. In *ICALP*, 1998.