

Tractable Queries for Lightweight Description Logics

Meghyn Bienvenu

Laboratoire de Recherche en Informatique
CNRS & Université Paris Sud, France

Magdalena Ortiz

Mantas Šimkus

Guohui Xiao

Institute of Information Systems
Vienna University of Technology, Austria

Abstract

It is a classic result in database theory that conjunctive query (CQ) answering, which is NP-complete in general, is feasible in polynomial time when restricted to acyclic queries. Subsequent results identified more general structural properties of CQs (like bounded treewidth) which ensure tractable query evaluation. In this paper, we lift these tractability results to knowledge bases formulated in the lightweight description logics DL-Lite and $\mathcal{ELH}$. The proof exploits known properties of query matches in these logics and involves a query-dependent modification of the data. To obtain a more practical approach, we propose a concrete polynomial-time algorithm for answering acyclic CQs based on rewriting queries into datalog programs. A preliminary evaluation suggests the interest of our approach for handling large acyclic CQs.

1 Introduction

Conjunctive queries (CQs) form a natural and important class of relational databases queries. In the description logic (DL) research community, there has been increasing interest in the problem of retrieving the answers to a CQ while taking into account the knowledge specified by a DL ontology [Calvanese *et al.*, 2007; Lutz *et al.*, 2009]. The use of an ontology typically leads to an increase in the complexity of CQ answering¹ compared to the relational database setting. Indeed, for the so-called *expressive DLs*, CQ answering is co-NP hard in data complexity (that is, when the ontology and query are considered fixed, and the complexity is measured in the size of the data only), in contrast to the AC_0 upper bound for standard databases. In combined complexity (that is, when the complexity is measured in terms of the combined sizes of the query, ontology, and data), rather than NP-complete, the problem is at least ExpTime-hard, and often requires double-exponential time (cf. the survey [Ortiz and Simkus, 2012]). Since such high complexity is prohibitive for data-rich applications, most work in the area focuses on ontologies formulated in *lightweight DLs* of the DL-Lite [Calvanese *et al.*,

2007] and $\mathcal{EL}$ [Baader *et al.*, 2005] families, which enjoy better computational properties. Indeed, CQ answering for DL-Lite knowledge bases has the same data and combined complexity as for plain databases, whereas for $\mathcal{EL}$, CQ answering is P-complete in data complexity, but remains NP-complete in combined complexity.

A classic result in database theory states that CQ answering becomes feasible in polynomial time when restricted to the class of *acyclic CQs* [Yannakakis, 1981]. Later investigations lead to the identification of more general structural properties, such as bounded treewidth, query width, or hypertree width [Chekuri and Rajaraman, 1997; Gottlob *et al.*, 1999], which guarantee tractable CQ answering. Since the NP-hardness in combined complexity of CQ answering in DLs is a direct consequence of the analogous result for relational databases, it seems natural to ask whether these tractability results also transfer to the DL setting. This would be very desirable since it is likely that most of the queries that will actually occur in applications are acyclic. While there are no collections of real-world CQs that can be used to support this claim in the DL setting, one can find some compelling evidence by looking at the closely related setting of SPARQL queries over RDF data, where it has been reported that acyclic queries (in fact, acyclic conjunctive graph patterns) comprise more than 99% of the queries in a log of around three million queries posed to the DBpedia endpoint [Picalausa and Vansummeren, 2011]. Unfortunately, lifting positive results from databases to the DL setting is often not possible, even when one considers lightweight DLs. For instance, for the logic DL-Lite $\mathcal{R}$, which underlies the QL profile of the OWL 2 standard [OWL Working Group, 2009], CQ answering was recently shown to be NP-hard already for acyclic queries [Kikot *et al.*, 2011].

In this paper, we show that for plain DL-Lite (without role hierarchies) and for $\mathcal{EL}$, the picture is brighter. Specifically, all polynomial-time upper bounds for classes of CQs known from relational databases carry over to DL-Lite. In the case of the $\mathcal{EL}$ family, we get polynomiality even for $\mathcal{ELH}$, thus showing that role hierarchies alone are not the culprit for the loss of tractability. Although this general tractability result relies on known properties of the logics, to our knowledge, it has not been pointed out before. The proof involves a polynomial reduction from the problem of answering a given CQ over a knowledge base $\mathcal{K}$ to answering the same CQ over a database that results from a polynomial expansion of the

¹When talking about the complexity of CQ answering, we mean the complexity of the *query output tuple problem*, that is, to decide whether a given tuple is in the answer of a CQ.

dataset in $\mathcal{K}$. The algorithm arising from this reduction has a disadvantage: it involves a query-dependent expansion of the data, which may be undesirable in many settings. Hence, we also propose an alternative polynomial-time algorithm for acyclic CQs, based on a rewriting into datalog. We have implemented a simple prototype of the approach, and it shows promising results for answering large acyclic CQs.

2 Preliminaries

Description Logics We briefly recall the syntax and semantics of DL-Lite $\mathcal{R}$ [Calvanese *et al.*, 2007], $\mathcal{ELH}$ [Baader *et al.*, 2005], and their sublogics DL-Lite and $\mathcal{EL}$. Let N_C , N_R , and N_I be countably infinite sets of concept names, role names, and individuals, respectively, and let $N_R = N_R \cup \{r^- \mid r \in N_R\}$ be the set of (complex) roles. For $R \in N_R$, R^- denotes r^- if $R = r \in N_R$, and r if $R = r^-$.

An *ABox* is a finite set of *assertions* of the forms $A(b)$ and $r(b, c)$ with $A \in N_C$, $r \in N_R$, and $b, c \in N_I$. A *TBox* is a finite set of *axioms*, whose form depends on the particular DL. In DL-Lite, TBox axioms are *concept inclusions* of the form $C_1 \sqsubseteq C_2$, with $C_1 = B_1$ and $C_2 = (\neg)B_2$ for B_1, B_2 of the form $A \in N_C$ or $\exists R$ with $R \in N_R$. DL-Lite $\mathcal{R}$ TBoxes may also contain *role inclusions* of the form $R_1 \sqsubseteq (\neg)R_2$ with $R_1, R_2 \in N_R$. In $\mathcal{EL}$, TBoxes consist of *concept inclusions* $C_1 \sqsubseteq C_2$, but in this case C_1, C_2 may be *complex concepts* constructed according to the syntax $C := \top \mid A \mid C \sqcap C \mid \exists r.C$. $\mathcal{ELH}$ TBoxes additionally allow role inclusions of the form $r_1 \sqsubseteq r_2$, similarly to DL-Lite $\mathcal{R}$, but with $r_1, r_2 \in N_R$. A DL *knowledge base* (KB) $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ consists of a TBox $\mathcal{T}$ and ABox $\mathcal{A}$.

The semantics of DL KBs is defined in terms of (DL) *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where the *domain* $\Delta^{\mathcal{I}}$ is a non-empty set, and the *interpretation function* $\cdot^{\mathcal{I}}$ maps each $a \in N_I$ to an object $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, each $A \in N_C$ to a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, and each $r \in N_R$ to a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The function $\cdot^{\mathcal{I}}$ is extended to complex concepts and roles in the usual way, see [Calvanese *et al.*, 2007; Baader *et al.*, 2005] for details. We say that $\mathcal{I}$ satisfies an axiom $P_1 \sqsubseteq P_2$ if $P_1^{\mathcal{I}} \subseteq P_2^{\mathcal{I}}$, and it satisfies an assertion $A(a)$ (resp. $r(a, b)$) if $a^{\mathcal{I}} \in A^{\mathcal{I}}$ (resp. $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$). Finally, $\mathcal{I}$ is a *model* of a KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ if it satisfies every axiom in $\mathcal{T}$ and assertion in $\mathcal{A}$. $\mathcal{K}$ is *consistent* if it admits some model. We use $\text{Ind}(\mathcal{A})$ for the individuals occurring in $\mathcal{A}$. We let $\mathcal{I}_{\mathcal{A}}$ be the interpretation with $\Delta^{\mathcal{I}} = \text{Ind}(\mathcal{A})$ and such that (i) $c \in A^{\mathcal{I}}$ iff $A(c) \in \mathcal{A}$, and (ii) $(c, d) \in r^{\mathcal{I}}$ iff $r(c, d) \in \mathcal{A}$.

Queries We recall *non-recursive datalog queries* and the more restricted *conjunctive queries*, cf. [Levy and Rousset, 1998]. Let N_V and N_D be countably infinite sets of *variables* and *datalog relations*, respectively. Each $\sigma \in N_D$ has an associated non-negative integer *arity*. *Atoms* are expressions of the form $p(\vec{x})$, where $\vec{x} \in (N_V)^n$, and (i) $p \in N_C$ and $n = 1$, (ii) $p \in N_R$ and $n = 2$, or (iii) $p \in N_D$ and n is the arity of p . Atoms of the form (i) and (ii) are called *DL-atoms*.

A *rule* ρ is an expression of the form $h(\vec{x}) \leftarrow \beta_1, \dots, \beta_m$, where $h(\vec{x}), \beta_1, \dots, \beta_m$ are atoms, h is a datalog relation, and every variable of $\vec{x}$ occurs in $\text{body}(\rho) = \{\beta_1, \dots, \beta_m\}$. Abusing notation, we write $\beta \in \rho$ instead of $\beta \in \text{body}(\rho)$. The variables in $\text{head}(\rho)$ are called the *answer variables* of ρ .

Given a set of rules P , we let $\text{Dep}(P) = (V, E)$ be the directed graph such that: (a) V is the set of all datalog relations occurring in P , and (b) $(p_1, p_2) \in E$ whenever there is a rule $\rho \in P$ such that p_1 is the relation in $\text{head}(\rho)$ and p_2 occurs in $\text{body}(\rho)$. A *non-recursive datalog query* is a pair $Q = (P, q)$, where P is a set of rules such that $\text{Dep}(P)$ has no cycle, and q is a datalog relation. The *arity* of Q is the arity of q .

Given a rule ρ and a DL interpretation $\mathcal{I}$, an *assignment* is a function π that maps every variable of ρ to an object in $\Delta^{\mathcal{I}}$. For a concept atom $A(x) \in \rho$, we write $\mathcal{I} \models_{\pi} A(x)$ if $\pi(x) \in A^{\mathcal{I}}$, and for a role atom $R(x_1, x_2) \in \rho$, we write $\mathcal{I} \models_{\pi} R(x_1, x_2)$ if $(\pi(x_1), \pi(x_2)) \in R^{\mathcal{I}}$. We call π a *match* for ρ in $\mathcal{I}$ if $\mathcal{I} \models_{\pi} \beta$ for all DL-atoms $\beta \in \rho$.

A tuple $\vec{t}$ is an *answer* to a query $Q = (P, h)$ in an interpretation $\mathcal{I}$ if there exists a rule $\rho = h(\vec{x}) \leftarrow \alpha$ in P and a match π for ρ in $\mathcal{I}$ such that (i) $\vec{t} = \pi(\vec{x})$ and (ii) for each non-DL-atom $p(\vec{y}) \in \rho$, $\pi(\vec{y})$ is an answer to (P, p) in $\mathcal{I}$. We use $\text{ans}(Q, \mathcal{I})$ to denote the set of answers to Q in $\mathcal{I}$. We denote by $\text{cert}(Q, \mathcal{K})$ the set of *certain answers* to $Q = (P, q)$ over a KB $\mathcal{K}$, defined as $\text{cert}(Q, \mathcal{K}) = \{\vec{a} \in (N_I)^n \mid (\vec{a})^{\mathcal{I}} \in \text{ans}(Q, \mathcal{I}) \text{ for any model } \mathcal{I} \text{ of } \mathcal{K}\}$, where n is the arity of Q .

A *conjunctive query* (CQ) is a non-recursive datalog query of the form $(\{\rho\}, q)$, such that $\text{body}(\rho)$ contains only DL-atoms. Since the particular relation q is irrelevant, we will use single rules (or rule bodies) to denote CQs.

In this paper, we focus on the decision problem known as the *query output tuple* (QOT) problem, which takes as input a query Q , a KB $(\mathcal{T}, \mathcal{A})$, and a tuple of individuals $\vec{a}$, and consists in deciding whether $\vec{a} \in \text{cert}(Q, (\mathcal{T}, \mathcal{A}))$. Whenever we talk about the *complexity of query answering*, we mean the computational complexity of the QOT problem. We focus on *combined complexity*, which is measured in terms of the size of the whole input $(\vec{a}, Q, \mathcal{T}, \mathcal{A})$.

Canonical Models Every consistent DL-Lite $\mathcal{R}$ or $\mathcal{ELH}$ KB $(\mathcal{T}, \mathcal{A})$ possesses a *canonical model* $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. For DL-Lite $\mathcal{R}$, the domain $\Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ consists of all words $aR_1 \dots R_n$ ($n \geq 0$) such that $a \in \text{Ind}(\mathcal{A})$, $R_i \in N_R$, and:

- if $n \geq 1$, then $\mathcal{T}, \mathcal{A} \models \exists R_1(a)$;
- for $1 \leq i < n$, $\mathcal{T} \models \exists R_i^- \sqsubseteq \exists R_{i+1}$ and $R_i^- \neq R_{i+1}$.

We say that $w' \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ is a *child* of $w \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ if $w' = wR$ for some R . The interpretation function is defined as follows:

$$\begin{aligned} a^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} &= a \text{ for all } a \in \text{Ind}(\mathcal{A}) \\ A^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} &= \{a \in \text{Ind}(\mathcal{A}) \mid \mathcal{T}, \mathcal{A} \models A(a)\} \\ &\cup \{aR_1 \dots R_n \mid n \geq 1 \text{ and } \mathcal{T} \models \exists R_n^- \sqsubseteq A\} \\ r^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} &= \{(a, b) \mid r(a, b) \in \mathcal{A}\} \cup \\ &\quad \{(w_1, w_2) \mid w_2 = w_1 S \text{ and } \mathcal{T} \models S \sqsubseteq r\} \cup \\ &\quad \{(w_2, w_1) \mid w_2 = w_1 S \text{ and } \mathcal{T} \models S \sqsubseteq r^-\} \end{aligned}$$

The construction of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ for $\mathcal{ELH}$ KBs is similar, please refer to [Lutz *et al.*, 2009] for details. Note that for both DL-Lite $\mathcal{R}$ and $\mathcal{ELH}$ KBs, $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ is composed of a *core*, which is obtained by restricting $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ to the objects in $\text{Ind}(\mathcal{A})$, and an *anonymous part* consisting of (possibly infinite) trees rooted at objects in the core. It is well-known that $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ can be homomorphically mapped into any model of $\mathcal{T}$ and $\mathcal{A}$, yielding:

Fact 1. Let $\mathcal{K}$ be a consistent DL-Lite or $\mathcal{ELH}$ KB, and let $\mathcal{I}_{\mathcal{K}}$ be its canonical model. Then $\text{cert}(Q, \mathcal{K}) = \text{ans}(Q, \mathcal{I}_{\mathcal{K}})$ for every non-recursive datalog query Q .

3 General Tractability Result

In this section, we observe that for both DL-Lite and $\mathcal{ELH}$, the answers to a CQ ρ over a consistent KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ coincide with the answers to ρ over an interpretation $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ that can be constructed in polynomial time from $\mathcal{K}$ and ρ . Since $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ can be viewed as a relational database, we obtain that any class of CQs that is tractable for plain databases is also tractable for KBs formulated in these DLs.

To establish this result, we rely on the following well-known property of query matches in DL-Lite and $\mathcal{ELH}$ (cf. [Kikot et al., 2012; Lutz et al., 2009]):

Lemma 2. Consider a consistent DL-Lite or $\mathcal{ELH}$ KB $(\mathcal{T}, \mathcal{A})$, a CQ ρ , and an object w in the anonymous part of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. Then all query matches π for ρ that coincide on the set of variables $\{x \mid \pi(x) = w\}$ coincide also on the sets of variables that are matched to the children of w .

It follows from Lemma 2 that if there is a query atom $R(x, y)$ such that $\pi(y)$ is a child of $\pi(x)$ in the anonymous part, then it is uniquely determined which other query variables have to be matched inside the tree rooted at $\pi(x)$, and how these variables are ordered into a tree. We can exploit this property to construct the desired $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ in two steps as follows:

- (1) First, we generate from ρ a polynomial number of tree-shaped queries, which correspond to the different ways that a subquery of ρ can be mapped inside a tree in the anonymous part of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. Each query is generated by selecting an atom $R(x, y) \in \rho$ and considering what happens if x and y were to be mapped respectively to a node w and one of its children in the anonymous part. By repeatedly applying Lemma 2, we can determine which other variables must be matched inside the tree rooted at w , and how the resulting subquery collapses into a tree.
- (2) Only the fragments of the anonymous part of the canonical model into which one of these tree-shaped queries can be homomorphically embedded can participate in query matches. Hence, by appropriately augmenting the core with instantiations of these tree-shaped queries, we obtain an interpretation $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ that is sufficient for retrieving all query answers. Specifically, we attach to each individual a a copy of each tree-shaped query for which there is a match rooted at a . To handle the case of (parts of) queries whose matches may be detached from the core, we also instantiate a disconnected copy of each tree-shaped query which can be mapped inside the anonymous part of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$.

Construction for DL-Lite The following notion of *tree witness* for DL-Lite was defined in [Kontchakov et al., 2010]. Let ρ be a CQ and let $R(x, y)$ be such that either $R(x, y) \in \rho$ or $R^-(y, x) \in \rho$. A *tree witness* for $R(x, y)$ in ρ is a partial map f from the variables in ρ to words over the alphabet $\overline{\mathbb{N}}_{\mathbb{R}}$ such that its domain is minimal (w.r.t. set-theoretic inclusion) and the following conditions hold:

- $f(y) = R$;
- if $f(z) = wS$, $S'(z, z') \in \rho$ or $S'^-(z', z) \in \rho$, and $S' \neq S^-$, then $f(z') = wSS'$; and

- if $f(z) = wS$ and $S(z', z) \in \rho$ or $S^-(z, z') \in \rho$, then $f(z') = w$.

We remark that each tree witness f naturally corresponds to a tree-shaped CQ obtained by restricting the original CQ ρ to the variables in the domain of f and unifying variables z, z' with $f(z) = f(z')$. By definition, there is at most one tree witness for each $R(x, y)$, which we denote by $f_{R(x, y)}$. We say that a tree witness $f_{R(x, y)}$ is *valid w.r.t. $\mathcal{T}$* if for every word $R_1 \dots R_n$ in the range of $f_{R(x, y)}$, and every $1 \leq i < n$, we have $R_{i-1} \neq R_i$ and $\mathcal{T} \models \exists R_{i-1}^- \sqsubseteq \exists R_i$. Existence and validity of tree witnesses can be tested in polynomial time.

If a match π for ρ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ maps x and y to objects w and wR respectively, then there is a tree witness $f_{R(x, y)}$ for $R(x, y)$ in ρ which is valid w.r.t. $\mathcal{T}$ and such that $\pi(z) = wf_{R(x, y)}(z)$ for each variable z in the domain of $f_{R(x, y)}$. Moreover, we may assume that if $\pi(x)$ is minimal (that is, there is no $\pi(x')$ which is a prefix of $\pi(x)$), then $\pi(x)$ is within distance $|\mathcal{T}|$ of the ABox. Hence, the matches for ρ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ coincide with the matches of ρ in the structure $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ obtained by adding to the core:

- (a) the objects aw such that w occurs in the range of a valid tree witness $f_{R(x, y)}$ and $\mathcal{T}, \mathcal{A} \models \exists R(a)$.
- (b) the objects b_iw such that b_i is a fresh object chosen for a valid tree witness $f_{R(x, y)}$, w occurs in the range of $f_{R(x, y)}$, and there is an individual $a \in \text{Ind}(\mathcal{A})$ and chain of roles $R_1, \dots, R_n$ of length at most $|\mathcal{T}|$ such that $\mathcal{T}, \mathcal{A} \models \exists R_1(a), R_n = R$, and for each $1 < i \leq n$, $R_{i-1} \neq R_i$ and $\mathcal{T} \models \exists R_{i-1}^- \sqsubseteq \exists R_i$.

To extend the interpretations of concept and role names to these new objects, we let $wR \in A^{\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}}$ whenever $\mathcal{T} \models \exists R^- \sqsubseteq A$, and $(w, wR) \in R^{\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}}$ for each new wR .

We remark that the existence of a role chain and individual having the properties stated in (b) can be decided in polynomial time by initializing a set *Reach* with all roles S such that $\mathcal{T}, \mathcal{A} \models \exists S(a)$ for some $a \in \text{Ind}(\mathcal{A})$, and then adding U to *Reach* whenever there is $V \in \text{Reach}$ such that $\mathcal{T} \models V^- \sqsubseteq U$. Since there are only polynomially many objects of the forms aw and b_iw as above, and instance checking and TBox reasoning are tractable for DL-Lite KBs [Calvanese et al., 2007], it follows that $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ can be constructed in polynomial time.

Construction for $\mathcal{ELH}$ In what follows, it will prove convenient to use conjunction as a role constructor: if $r_1, r_2 \in \mathbb{N}_{\mathbb{R}}$ are role names, then their conjunction $r_1 \sqcap r_2$ is a role whose interpretation is $(r_1 \sqcap r_2)^{\mathcal{I}} = r_1^{\mathcal{I}} \cap r_2^{\mathcal{I}}$. We denote by $\mathcal{ELH}_{\sqcap}$ the extension of $\mathcal{ELH}$ with role conjunction.

We introduce a notion of tree witness for $\mathcal{ELH}$, inspired by the fork elimination procedure from [Lutz, 2008]. Let ρ be a CQ and $\alpha = r(x, y) \in \rho$. We begin by defining a set D_{α} and equivalence relation $\sim_{\alpha}$ over D_{α} by initializing D_{α} to $\{x, y\}$, and then applying the following rules until convergence:

- if $s(z, z') \in \rho$ and $z \in D_{\alpha}$, then add z' to D_{α}
- if $s(u, u') \in \rho, t(z, z') \in \rho, u, u', z' \in D_{\alpha}$, and $u' \sim_{\alpha} z'$, then add z to D_{α} and put $u \sim_{\alpha} z$

Note that D_{α} and $\sim_{\alpha}$ are uniquely defined and can be computed in linear time in the size of ρ . We let ρ_{α} be the query obtained by restricting ρ to the variables in D_{α} , then replacing each variable z by its equivalence class $[z]$ under $\sim_{\alpha}$. We define a directed graph G_{α} whose nodes are the equivalence

classes in $\sim_\alpha$ and whose edges are the atoms in ρ_α . If G_α contains no (directed) cycles, then we define the *tree witness* for $\alpha = r(x, y)$ in ρ as the map $f : D_\alpha \rightarrow (2^{\mathbb{N}_R} \times 2^{\mathbb{N}_C})^*$ with:

- $f(z) = \epsilon$ if $z \sim_\alpha x$
- $f(z) = M_1 N_1 \dots M_k N_k$ if $[u_0], \dots, [u_k]$ is the unique path in G_α with $u_0 \sim_\alpha x$ and $u_k \sim_\alpha z$, and for every $1 \leq i \leq k$, we have $M_i = \{s \mid s([u_{i-1}], [u_i]) \in G_\alpha\}$ and $N_i = \{A \mid A([u_i]) \in G_\alpha\}$

To every tree witness f , we can naturally associate an $\mathcal{ELH}_\square$ concept $\text{conc}_f(\epsilon)$ as follows: if $f(z) = w$ is a leaf, then $\text{conc}_f(w) = \top$, and if $f(z) = w$ has children $w_1, \dots, w_n$ with $w_i = w M_i N_i$, then $\text{conc}_f(w) = \bigcap_{i=1}^n \exists(\bigcap_{r \in M_i} r) \cdot (\bigcap_{A \in N_i} A \sqcap \text{conc}_f(w_i))$. It follows from results in [Rudolph et al., 2008] that it can be checked in polynomial time whether $\mathcal{T}, \mathcal{A} \models \text{conc}_f(\epsilon)$. Using a reachability construction similar to the one for DL-Lite, we can also test in polynomial time whether $\text{conc}_f(\epsilon)$ is non-empty in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$.

We are now ready to define the interpretation $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$. Its domain is the set of words in $\mathbb{N}_I \cdot (2^{\mathbb{N}_C} \times 2^{\mathbb{N}_R})^*$ that contains each individual occurring in $\mathcal{A}$, each aw such that w is in the domain of a tree witness f for which $(\mathcal{T}, \mathcal{A}) \models \text{conc}_f(\epsilon)(a)$ holds, and each $b_i \cdot w$ such that f_i is a tree witness with $\text{conc}_f(\epsilon)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} \neq \emptyset$ (for b_i a fresh symbol selected for f_i). Concept and role names are interpreted as follows:

$$\begin{aligned} A^{\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}} &= \{a \in \text{Ind}(\mathcal{A}) \mid \mathcal{T}, \mathcal{A} \models A(a)\} \cup \\ &\quad \{c M_1 N_1 \dots M_n N_n \mid \mathcal{T} \models N_n \sqsubseteq A\} \\ r^{\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}} &= \{(a, b) \mid \mathcal{T}, \mathcal{A} \models r(a, b)\} \cup \\ &\quad \{(w_1, w_1 \cdot M C) \mid \mathcal{T} \models s \sqsubseteq r \text{ for some } s \in M\} \end{aligned}$$

and each individual a is interpreted as itself ($a^{\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}} = a$).

The following theorem resumes the key properties of the interpretations $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ just described.

Theorem 3. *Let ρ be a CQ, let $(\mathcal{T}, \mathcal{A})$ be a consistent DL-Lite or $\mathcal{ELH}$ knowledge base, and let $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ be the interpretation defined previously. The following statements hold:*

1. $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ can be built in polynomial time in $\rho, \mathcal{T}$ and $\mathcal{A}$.
2. For every tuple $\vec{a}$ of individuals, $\vec{a} \in \text{ans}(\rho, \mathcal{I}_{\mathcal{T}, \mathcal{A}})$ iff $\vec{a} \in \text{ans}(\rho, \mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho})$.

In light of Theorem 3 and Fact 1, to determine whether $\vec{a} \in \text{cert}(\rho, (\mathcal{T}, \mathcal{A}))$, it is sufficient to test the consistency of $(\mathcal{T}, \mathcal{A})$ and then, if $(\mathcal{T}, \mathcal{A})$ is consistent, to decide whether $\vec{a} \in \text{ans}(\rho, \mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho})$. If we view $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ as a plain relational database, the latter check is just a special case of the QOT problem. Hence, we obtain the desired result:

Corollary 4. *Let $\mathcal{Q}$ be a class of CQs for which the query output tuple problem over relational databases is decidable in polynomial time. Then the query output tuple problem for $\mathcal{Q}$ is also tractable for KBs formulated in DL-Lite and $\mathcal{ELH}$.*

The construction of $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ above is easily extended to DL-Lite_R using the notion of tree witnesses proposed in [Kikot et al., 2012]. However, the construction is no longer polynomial since there can be exponentially many tree witnesses for a single atom $R(x, y)$ in ρ [Kikot et al., 2011], and so we do not obtain an analogue of Theorem 3. Indeed, it follows from results by Kikot et al. 2011 that the tuple output problem for tree-shaped CQs is NP-complete for DL-Lite_R KBs.

Therefore, to obtain tractability results to DL-Lite_R, one must impose some syntactic restriction on $\mathcal{T}$ and ρ that ensures a polynomial bound on the number of tree witnesses, e.g. the absence of *twisty roles* proposed in [Kikot et al., 2012].

4 Answering Acyclic Queries

The expansion technique presented in Section 3 allows us to convert any polynomial-time algorithm for evaluating a tractable class of CQs over relational databases into a polynomial-time algorithm for evaluating the same class of queries over DL-Lite and $\mathcal{ELH}$ KBs. However, the resulting algorithm would involve building the structure $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ for each input query ρ , which is clearly undesirable. We now present our main contribution: a polynomial-time procedure for evaluating acyclic CQs which is based upon rewriting CQs into non-recursive datalog programs. We present the approach for DL-Lite KBs, but discuss at the end of the section how the approach can be adapted to DL-Lite_R and $\mathcal{ELH}$.

We begin with some preliminaries. As usual, the query graph $G(\rho)$ of a CQ ρ is defined as the undirected graph whose nodes are the variables of ρ , and that has an edge between x and x' if ρ contains a (body) atom $r(x, x')$ or $r(x', x)$. A CQ ρ is *acyclic* if $G(\rho)$ is acyclic, and *rooted* if every connected component of $G(\rho)$ has at least one answer variable. We consider a slight generalization of acyclicity: we say that a CQ ρ is *acyclic modulo answer variables* (*a-acyclic* for short) if the graph $G^-(\rho)$ obtained by deleting from $G(\rho)$ each edge (x, x') , where x, x' are answer variables is acyclic.

Our rewriting procedure works on queries which are both rooted and a-acyclic (we will discuss later how to handle the non-rooted case). To every rooted a-acyclic CQ ρ , we associate the set of connected components $\{T_1, \dots, T_n\}$ of $G^-(\rho)$. Because ρ is rooted, every T_i is a connected acyclic graph containing at least one vertex which is an answer variable. We select an arbitrary answer variable x_i for each T_i and designate it as the root of T_i , allowing us to view T_i as a tree. Then, given a pair of variables x, y of ρ , we say y is a *child* of x if y is a child of x in the (unique) tree T_i that contains x , and define *descendant* as the transitive closure of child. For a variable x of ρ , we denote by $\vec{x}_\rho$ the tuple $(y_1, \dots, y_k)$, where $y_1, \dots, y_k$ are the answer variables that are descendants of x .

Rewriting procedure for rooted queries Consider a rooted a-acyclic CQ $\rho = q(\vec{x}) \leftarrow \alpha$ and a DL-Lite TBox $\mathcal{T}$, and let X be the set of answer variables which are roots in ρ (see previously). We rewrite the query ρ into the non-recursive datalog program $\text{rew}_{\mathcal{T}}(\rho) = (P, q)$ defined as follows. In addition to q , the program uses the following datalog relations: (i) $(|\vec{x}_\rho| + 1)$ -ary relations q_x, q'_x for every variable x of ρ , and (ii) unary relations q_A and $q_{\exists R}$ for every $A \in \mathbb{N}_C$ and $R \in \overline{\mathbb{N}_R}$ occurring in ρ . We now describe the rules in P . There is a single top-level rule defining q :

$$q(\vec{x}) \leftarrow \bigwedge_{x \in X} q_x(x, \vec{x}_\rho) \wedge \bigwedge_{x_i, x_j \in X \wedge r(x_i, x_j) \in \rho} r(x_i, x_j) \quad (1)$$

For every variable x in ρ , with $Y = \{y_1, \dots, y_n\}$ the set of children of x in ρ , we have the following rule

$$q_x(x, \vec{x}_\rho) \leftarrow \bigwedge_{A(x) \in \rho} q_A(x) \wedge \bigwedge_{y \in Y} q'_y(x, \vec{y}_\rho) \quad (2)$$

and for every $y \in Y$, we also have

$$q'_y(x, \vec{y}_\rho) \leftarrow \bigwedge_{r(x,y) \in \rho} r(x, y) \wedge q_y(y, \vec{y}_\rho) \quad (3)$$

and for all $y \in Y$ satisfying the following conditions:

- (i) there is an atom $R(x, y) \in \rho$ or $R^-(y, x) \in \rho$ and the tree witness $f_{R(x,y)}$ exists and is valid
 - (ii) for every u in domain of $f_{R(x,y)}$ with $f_{R(x,y)}(u) = wS$ and $A(u) \in \rho$, we have $\mathcal{T} \models \exists S^- \sqsubseteq A$
 - (iii) the set $Z = \{z \mid f_{R(x,y)}(z) = \varepsilon \wedge z \neq x\}$ contains all answer variables in the domain of $f_{R(x,y)}$
- we additionally have the rule

$$q'_y(x, \vec{y}) \leftarrow q_{\exists R}(x) \wedge \bigwedge_{z \in Z} q_z(x, \vec{z}_\rho) \quad (4)$$

Finally, for every $B \in \mathbf{N}_C \cup \{\exists R \mid R \in \overline{\mathbf{N}_R}\}$ with q_B a datalog relation in P , we have the rules

$$\begin{aligned} q_B(x) &\leftarrow A(x) && \text{for all } A \in \mathbf{N}_C \text{ such that } \mathcal{T} \models A \sqsubseteq B \\ q_B(x) &\leftarrow s(x, y) && \text{for all } s \in \mathbf{N}_R \text{ such that } \mathcal{T} \models \exists s \sqsubseteq B \quad (5) \\ q_B(x) &\leftarrow s(y, x) && \text{for all } s \in \mathbf{N}_R \text{ such that } \mathcal{T} \models \exists s^- \sqsubseteq B \end{aligned}$$

Intuitively, the relation q_x corresponds to the query $\rho|x$ whose answer variables are $\{x\} \cup \vec{x}_\rho$ and whose body is obtained by restricting the body of ρ to the atoms whose arguments among x and its descendants; whereas the relation q'_y corresponds to the query $\rho|xy$ (with y a child of x) obtained by adding to $\rho|x$ the role atoms linking x and y . Rule (1) stipulates that a tuple is in the answer to ρ if it makes true all of the role atoms linking two root variables and each of the queries $\rho|x$ associated with a root variable x of ρ . Then rule (2) states that to make $\rho|x$ hold at an individual, we must satisfy the concept atoms for x and the query $\rho|xy$ for each child y of x . Rules (3) and (4) provide two ways of satisfying $\rho|xy$. The first way, captured by rule (3), is to map y to an ABox individual, in which case the role atoms between x and y must occur in the ABox, and the query $\rho|y$ must hold at this individual. The second possibility, treated by rule (4), is that y is mapped to an element of the anonymous part of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ which is a child of x . For this to occur, several conditions must be verified. First, if y is an R -successor of x , then the tree witness $f_{R(x,y)}$ must exist and be valid w.r.t. $\mathcal{T}$. Second, we must ensure that all concept atoms concerning variables that are mapped inside the anonymous part by $f_{R(x,y)}$ are satisfied (this is checked in item (ii)). Finally, since answer variables cannot be mapped inside the anonymous part, we need condition (iii), which checks that every answer variable z in the domain of $f_{R(x,y)}$ is such that $f_{R(x,y)}(z) = \varepsilon$. If all of these conditions are met, then rule (4) states that the query $\rho|xy$ can be satisfied by making $\exists R$ hold at x (thereby guaranteeing the existence of the required paths in the anonymous part of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$) and by satisfying the remainder of the query $\rho|xy$ (i.e. the query obtained by removing the atoms mapped inside the anonymous part). The latter corresponds precisely to the union of the queries $\rho|z$ where z is a descendant of x with $f_{R(x,y)}(z) = \varepsilon$. Finally, the rules in (5) provide the standard rewriting of concepts A and $\exists R$ w.r.t. $\mathcal{T}$.

We establish the correctness of our rewriting procedure.

Theorem 5. *Let ρ be a rooted a-acyclic CQ and $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ a DL-Lite KB. Then $\text{cert}(\rho, \mathcal{K}) = \text{ans}(\text{rew}_{\mathcal{T}}(\rho), \mathcal{I}_{\mathcal{A}})$.*

Proof idea. The key step in the proof is showing that for every variable x in ρ , $\text{cert}(\rho|x, \mathcal{K}) = \text{ans}((P, q_x), \mathcal{I}_{\mathcal{A}})$. This can be proven by induction on the number of variables in $\rho|x$, utilizing Fact 1 and properties of tree witnesses. $\square$

The next theorem shows that our rewriting procedure provides a polynomial-time algorithm for evaluating rooted a-acyclic conjunctive queries.

Theorem 6. *Given a rooted a-acyclic CQ ρ and a DL-Lite KB $(\mathcal{T}, \mathcal{A})$, the program $\text{rew}_{\mathcal{T}}(\rho)$ can be computed in polynomial time in the size of ρ and $\mathcal{T}$, and $\vec{a} \in \text{ans}(\text{rew}_{\mathcal{T}}(\rho), \mathcal{I}_{\mathcal{A}})$ can be tested in polynomial time in the size of Q and $\mathcal{A}$.*

Proof. For the first point, we observe that the number of relations in $\text{rew}_{\mathcal{T}}(\rho)$ is linear in the number of atoms in ρ and that each relation is defined using linearly many rules in the size of ρ and $\mathcal{T}$. We also note that testing the conditions for rules of type (4) can be done in polynomial time in the size of ρ and $\mathcal{T}$ (cf. Section 3). The second statement is true because once the answer variables in $\text{rew}_{\mathcal{T}}(\rho)$ have been instantiated with the individuals in $\vec{a}$, we have a non-recursive datalog program (with constants) where every rule contains at most two variables. It is known that datalog programs of this form can be evaluated in polynomial time (cf. [Dantsin *et al.*, 2001]). $\square$

Handling non-rooted queries. We now return to the case of non-rooted a-acyclic CQs, and show that such queries can be answered via a reduction to the rooted query case. Given an a-acyclic CQ ρ over a KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, the set $\text{cert}(\rho, \mathcal{K})$ can be computed using the following steps:

1. If ρ is rooted, then return $\text{cert}(\rho, \mathcal{K})$.
2. Choose a maximally connected component β of ρ , such that β contains no answer variable.
3. If β has a match fully in the anonymous part of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, then drop β from ρ and go to step 1.
4. If there is a variable x in β such that the rooted query $g(x) \leftarrow \beta$ has a non-empty answer over $\mathcal{K}$, then drop β from ρ and go to step 1. Otherwise, return $\emptyset$.

For testing whether β satisfies the condition in step 3, it suffices to consider the tree witnesses f of the subquery β whose domain are all the variables in β . For each such witness f , we need to test: (i) validity of f , (ii) whether $\mathcal{T} \models \exists R^- \sqsubseteq A$ for all x with $f(x)$ of the form wR and all atoms $A(x)$ in β , and (iii) whether there exists a chain of roles $R_1, \dots, R_n$ that is implied by the knowledge base and that reaches the first role in f , as in item (b) of the construction of $\mathcal{I}_{\mathcal{T}, \mathcal{A}, \rho}$ for DL-Lite. Since all these steps are feasible in polynomial time, we obtain a polynomial-time procedure for solving the QOT problem for arbitrary a-acyclic CQs over DL-Lite KBs.

Adapting the rewriting for DL-Lite_R and $\mathcal{ELH}$ We now discuss how to modify the rewriting to handle DL-Lite_R and $\mathcal{ELH}$ KBs. First, since both DLs support role inclusions, we must replace atoms $r(x, y)$ by atoms $q_r(x, y)$, and add the corresponding rules $q_r(x, y) \leftarrow S(x, y)$ with $\mathcal{T} \models S \sqsubseteq r$.

Then the algorithm can be directly employed for DL-Lite_R, but using the notion of tree witness in [Kikot *et al.*, 2012]. Similarly as in Section 3, the algorithm may not be polynomial since there can be exponentially many tree witnesses.

For $\mathcal{ELH}$, apart from handling role inclusions as above, rule (4) must be modified to use the $\mathcal{ELH}$ version of tree witnesses. In particular, instead of an atom $q_{\exists R}(x)$, we use $q_{C_f}(x)$, where C_f is the concept induced by the tree witness f for $r(x, y)$. Note that there are only linearly many tree witnesses, hence only linearly many new relations q_{C_f} . Assuming that the $\mathcal{ELH}$ TBox is in normal form (cf. [Baader *et al.*, 2005]), we only need to consider atomic concepts A which entail C , yielding a linear number of rules of the form $q_{C_f}(x) \leftarrow q_A(x)$. Finally, we must modify the rules in (5) defining the relations q_A to capture entailment in $\mathcal{ELH}$, which may require the use of recursive datalog rules (see e.g. [Eiter *et al.*, 2012]). Importantly, these rules have at most two variables each, and so they do not impact the polynomial upper bound argument. Thus, by using this modified rewriting, and handling non-rooted CQs in a similar way to DL-Lite, we obtain a polynomial-time algorithm for deciding the QOT problem for a-acyclic CQs over $\mathcal{ELH}$ KBs.

5 Preliminary Evaluation

We developed a prototype rewriting system that takes as input a rooted acyclic CQ ρ and a DL-Lite_R TBox $\mathcal{T}$, and outputs an SQL statement expressing the resulting non-recursive datalog program $\text{rew}_{\mathcal{T}}(\rho)$ (using common table expressions). We evaluated the result over ABoxes stored in a relational database, using the PostgreSQL database system, and Owlgres [Stocker and Smith, 2008] for loading the data.

To test our prototype, we used the university ontology LUBM₂₀ described in [Lutz *et al.*, 2012], which adds concept inclusions with additional concept names, and with existential concepts on the right-hand side, to the original LUBM ontology [Guo *et al.*, 2005]. We considered three acyclic queries from the benchmark in [Lutz *et al.*, 2012] (q_2 , q_4 , and q_5), which are rather small (at most 4 atoms), and created four additional large acyclic queries, with 13 to 34 atoms, and 7 to 17 variables. The new queries are also significantly larger than the ones in the REQUIEM test suite (≤ 7 atoms). The importance of handling such larger queries in practice has been previously argued in [Rosati and Almatelli, 2010].

We compared our rewriting procedure with REQUIEM [Pérez-Urbina *et al.*, 2009] and IQAROS [Venetis *et al.*, 2012] which, like most of the existing query rewriting systems for the DL-Lite family, generate unions of CQs rather than non-recursive datalog programs. For the four large queries, both REQUIEM and IQAROS did not terminate (within ten minutes). Even for the small q_4 and q_5 , the generated rewritings were too large to be posed directly to an off-the-shelf RDBMS: REQUIEM generated tens of thousands of queries for both, and IQAROS almost 15 thousand for q_4 , and almost one thousand for q_5 . By contrast, for our approach, the rewriting times were negligible for all queries (under half a second) and the produced rewritings have less than 30 rules, disregarding the rules (5) of the algorithm (since the latter are independent of the query, we computed them separately and

#Uni	q_2	q_4	q_5	q_7	q_8	q_9	q_{10}
20	2.5	3.0	4.2	1.5	1.0	0.0	0.0
50	9.0	7.3	4.6	2.0	3.3	0.0	0.0
100	20.5	15.0	9.4	4.2	7.2	0.0	0.0
150	25.6	21.8	14.1	6.6	11.6	15.2	14.9
200	33.5	>600	27.0	15.2	26.9	31.2	30.3

Table 1: Scalability of our system (runtime in seconds)

stored them using a database view per concept/role name).

We also tested the feasibility of evaluating our rewritings over large ABoxes. For this, we used the modified LUBM data generator [Lutz *et al.*, 2012] (with 5% incompleteness). Each university has approximately 17k individuals, 28K concept assertions, and 47K role assertions. We carried out our experiments on ABoxes with 20 – 200 universities, resulting in very large ABoxes (up to ca. 1.5 GB on disk). The results reported in Table 1 show that the algorithm scales well.

Finally, we note that a performance comparison with PRESTO [Rosati and Almatelli, 2010], which also outputs non-recursive datalog programs, was not possible because this system is not publicly available. However, we can observe that its underlying algorithm may produce exponential-size rewritings for acyclic CQs, as witnessed by the family of queries $q(x) \leftarrow r(x, y) \wedge \bigwedge_{0 \leq i \leq n} p(y, z_i) \wedge p(u_i, z_i)$ coupled with the empty TBox. Intuitively, the PRESTO algorithm generates a separate rule for every possible way of selecting a collection of variables from $\{u_1, \dots, u_n\}$ and identifying them with y . This exponential blow-up suggests that our positive results are not merely an artifact of the datalog representation, but derive also from our use of acyclicity.

6 Future Work

We plan to generalize our rewriting technique to larger tractable classes of CQs, like bounded treewidth CQs. Another direction is to identify suitable restrictions for more expressive DLs that allow for tractable answering of acyclic queries. Our proof-of-concept implementation raises hopes that efficient evaluation of large acyclic queries is feasible, but many challenges must still be addressed. For example, we observe that breaking down the queries into small rules as is done by our rewriting may lead to a loss of structure that could be used by database management systems for optimized evaluation. There are many other aspects, not directly related to the rewriting technique, that must also be taken into account to achieve practicable query answering, such as exploring more efficient forms of representing data in ABoxes, using different kinds of indexes, and considering different translations of our programs into SQL. Using semantic indexes [Rodriguez-Muro and Calvanese, 2012] for handling the rules of type (5) appears particularly promising.

Acknowledgements The authors were supported by a Université Paris-Sud Attractivité starting grant and ANR project PAGODA ANR-12-JS02-007-01 (Bienvenu), FWF project T515-N23 (Ortiz), FWF project P25518-N23 and WWTF project ICT12-015 (Šimkus), Vienna PhD School of informatics and EU project Optique FP7-318338 (Xiao).

References

- [Baader *et al.*, 2005] Franz Baader, Sebastian Brandt, and Carsten Lutz. Pushing the EL envelope. In *Proc. of IJCAI*, pages 364–369, 2005.
- [Calvanese *et al.*, 2007] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *J. Automated Reasoning*, 39(3):385–429, 2007.
- [Chekuri and Rajaraman, 1997] Chandra Chekuri and Anand Rajaraman. Conjunctive query containment revisited. In *Proc. of ICDT*, pages 56–70. Springer, 1997.
- [Dantsin *et al.*, 2001] Evgeny Dantsin, Thomas Eiter, Georg Gottlob, and Andrei Voronkov. Complexity and expressive power of logic programming. *ACM Computing Survey*, 33(3):374–425, September 2001.
- [Eiter *et al.*, 2012] Thomas Eiter, Magdalena Ortiz, Mantas Simkus, Trung-Kien Tran, and Guohui Xiao. Query rewriting for Horn-SHIQ plus rules. In *Proc. of AAAI*. AAAI Press, 2012.
- [Gottlob *et al.*, 1999] Georg Gottlob, Nicola Leone, and Francesco Scarcello. Hypertree decompositions and tractable queries. In *Proc. of PODS*, pages 21–32. ACM Press, 1999.
- [Guo *et al.*, 2005] Yuanbo Guo, Zhengxiang Pan, and Jeff Heflin. LUBM: A benchmark for OWL knowledge base systems. *J. Web Semantics*, 3(2-3):158–182, 2005.
- [Kikot *et al.*, 2011] Stanislav Kikot, Roman Kontchakov, and Michael Zakharyashev. On (in)tractability of OBDA with OWL 2 QL. In *Proc. of DL*. CEUR-WS.org, 2011.
- [Kikot *et al.*, 2012] Stanislav Kikot, Roman Kontchakov, and Michael Zakharyashev. Conjunctive query answering with OWL 2 QL. In *Proc. of KR*. AAAI Press, 2012.
- [Kontchakov *et al.*, 2010] Roman Kontchakov, Carsten Lutz, David Toman, Frank Wolter, and Michael Zakharyashev. The combined approach to query answering in DL-Lite. In *Proc. of KR*. AAAI Press, 2010.
- [Levy and Rousset, 1998] Alon Y. Levy and Marie-Christine Rousset. Combining Horn rules and description logics in CARIN. *Artificial Intelligence*, 104(1-2):165–209, 1998.
- [Lutz *et al.*, 2009] Carsten Lutz, David Toman, and Frank Wolter. Conjunctive query answering in the description logic EL using a relational database system. In *Proc. of IJCAI*, pages 2070–2075, 2009.
- [Lutz *et al.*, 2012] Carsten Lutz, Inanc Seylan, David Toman, and Frank Wolter. The combined approach to OBDA: Taming role hierarchies using filters (with appendix). In *Proc. of SSWS+HPCSW*, 2012.
- [Lutz, 2008] Carsten Lutz. The complexity of conjunctive query answering in expressive description logics. In *Proc. of IJCAR*, pages 179–193. Springer, 2008.
- [Ortiz and Simkus, 2012] Magdalena Ortiz and Mantas Simkus. Reasoning and query answering in description logics. In *Reasoning Web*, pages 1–53. Springer, 2012.
- [OWL Working Group, 2009] W3C OWL Working Group. *OWL 2 Web Ontology Language: Document Overview*. W3C Recommendation, 2009. Available at <http://www.w3.org/TR/owl2-overview/>.
- [Pérez-Urbina *et al.*, 2009] Héctor Pérez-Urbina, Boris Motik, and Ian Horrocks. A comparison of query rewriting techniques for DL-Lite. In *Proc. of DL*. CEUR-WS.org, 2009.
- [Picalausa and Vansummeren, 2011] François Picalausa and Stijn Vansummeren. What are real SPARQL queries like? In *Proc. of SWIM*. ACM, 2011.
- [Rodriguez-Muro and Calvanese, 2012] Mariano Rodriguez-Muro and Diego Calvanese. High performance query answering over DL-Lite ontologies. In *Proc. of KR*, pages 308–318. AAAI Press, 2012.
- [Rosati and Almatelli, 2010] Riccardo Rosati and Alessandro Almatelli. Improving query answering over DL-Lite ontologies. In *Proc. of KR*. AAAI Press, 2010.
- [Rudolph *et al.*, 2008] Sebastian Rudolph, Markus Krötzsch, and Pascal Hitzler. Cheap boolean role constructors for description logics. In *Proc. of JELIA*, pages 362–374, 2008.
- [Stocker and Smith, 2008] Markus Stocker and Michael Smith. Owlgres: A scalable OWL reasoner. In *Proc. of OWLED*. CEUR-WS.org, 2008.
- [Venetis *et al.*, 2012] Tassos Venetis, Giorgos Stoilos, and Giorgos B. Stamou. Incremental query rewriting for OWL 2 QL. In *Proc. of DL*. CEUR-WS.org, 2012.
- [Yannakakis, 1981] Mihalis Yannakakis. Algorithms for acyclic database schemes. In *Proc. of VLDB*, pages 82–94. IEEE Computer Society, 1981.