

XPath for DL-Lite Ontologies

Egor V. Kostylev¹, Juan L. Reutter², and Domagoj Vrgoč³

¹ University of Oxford egor.kostylev@cs.ox.ac.uk

² PUC Chile jreutter@ing.puc.cl

³ University of Edinburgh domagoj.vrgoc@ed.ac.uk

Abstract. Applications of description logics (DLs) such as OWL 2 and ontology-based data access (OBDA) require understanding of how to pose database queries over DL knowledge bases. While there have been many studies regarding traditional relational query formalisms such as conjunctive queries and their extensions, little attention has been paid to graph database queries, despite the fact that graph databases share the structure of interpretations with DLs; that is they describe essentially the same objects. In particular, not much is known about the interplay between DLs and XPath. The last is a powerful formalism for querying semistructured data: it is in the core of most practical query languages for XML trees, and it is also gaining popularity in theory and practice of graph databases. In this paper we make a step towards coupling knowledge bases and graph databases by studying how to answer powerful XPath-style queries over DL-Lite. We start with adapting the definition of XPath to the DL context, and then proceed to study the complexity of evaluating XPath queries over knowledge bases. Results show that, while query answering is undecidable for the full XPath, by carefully tuning the amount of negation allowed in the queries we can arrive to XPath fragments that have a potential to be used in practical applications.

1 Introduction

Satisfiability and model checking have long been two central problems in the knowledge representation community. But new applications of description logics (DLs for short), such as the Web Ontology Language (OWL) [21] and ontology-based data access (OBDA), are forcing us to develop algorithms solving more complex data manipulation and extraction tasks, and in particular, require understanding how to answer database-style queries over knowledge bases that are specified by DLs [14].

The literature on knowledge bases usually considers relational queries, mostly focusing on conjunctive queries (CQs) and their extensions with union [1, 7], forms of negation [15, 24], or aggregates [10, 18]. It is also natural to look at queries that are designed to operate over graph databases, since these share the same structure with DL knowledge bases; namely they use unary and binary predicates (that is concepts and roles in DL terminology) to represent graph nodes and edges. While the idea of using graph queries in knowledge bases is not new (see e.g., [8]), we are only starting to understand how such queries can be fine tuned for various ontology languages. So far the specific research on this topic has primarily been concerned with the class of *regular path queries* (RPQs; see e.g. [2]), one of the most basic graph query languages. We now know how to evaluate such queries over various description logics [9], how to deal with

some of their extensions [4, 5] and understand how to solve their containment problem in the presence of DL constraints [11].

There are of course many other languages for querying graph and semi-structured data, most notable amongst them being XPath. Originally designed to extract data from XML trees [27], XPath has recently been adapted to work over graph databases [19] and was shown to retain good evaluation properties while at the same time being more powerful than RPQs and many of their extensions. Moreover, XPath also subsumes navigational graph querying features of SPARQL [17], such as property paths, and other commercial graph query languages (see e.g. Neo4j [23]), and thus can also be seen as a querying primitive for more powerful languages. Given that XPath is capable of expressing virtually all relevant querying primitives for graphs, and was tried and tested by XML practitioners, it is natural to see if it will have the same impact as a language for DL knowledge bases.

In this paper we take a decisive step towards the implementation of graph query languages over ontologies and study the problem of evaluating XPath queries over DL. To that extent, we first adapt the XPath query language to work over ontologies, arriving to DLXPath—an expressive language designed specifically for DL knowledge bases. To get a flavour of the interplay between DLXPath and DLs, we then study the problem of evaluating DLXPath queries over $DL-Lite_{\mathcal{R}}$ and $DL-Lite_{core}$ ontologies. The choice for this particular family of DLs is twofold: on the one hand, it is simple enough to start such a research; on the other, this family is of the most importance in practice, since it underlies OWL 2 QL profile [21].

As a simple useful example of an DLXPath query, we can give the following: ‘Can I fly from city A to city B making stops only in cities with UNESCO World Heritage Sites which are endangered?’ Here we assume that the considered ontology has binary predicates (roles) `HasDirectFlight` and `HasUNESCOSite`, as well as unary predicate (concept) `InDanger`. Under this ontology this query cannot be answered by RPQs. But note that it can, however, be expressed as a nested RPQ [4], as this language corresponds to DLXPath without negation. However, if we additionally require the site of the last intermediate stop to be included in the list *not* under cultural criteria (expressed by concept `Cultural`), then such a query cannot be expressed by any known language for knowledge bases, but can be written as DLXPath query.

Regardless of the choice of a particular DL, one cannot expect that the study of DLXPath query answering will come as a straightforward adaptation of known techniques. The language of DLXPath contains expressive primitives such as transitive closure and negation, that is, the operators which are known to bring in conceptual difficulties even when studied in isolation. For example, it has been shown that adding negation to conjunctive queries immediately leads to undecidability of query answering [15], while adding path operators usually implies jumping one or two exponents in the complexity of query evaluation algorithms [9]. Thus, one of the challenges when studying DLXPath is to fine-tune these operators to archive an optimal trade-off between good evaluation properties and expressive power of the query language. To this end, we will distinguish two flavours of the language: the *core* fragment, denoted by $DLXPath_{core}$, and the *regular* fragment, denoted by $DLXPath_{reg}$. The two are designed to match the duality present in the standard XPath query language [27], and while the core fragment

allows transitive closure only over basic role names, the regular fragment lifts this restriction, allowing us to post more general path queries. Regarding negation, we will distinguish full fragments, which allow both unary and binary negation, *path-positive* fragments with only the unary one, and *positive* fragments without any negation.

As usual when gauging the usefulness of a new query language for ontologies, we start by considering *data complexity* of the query answering problem, where one assumes that the query and the *terminological knowledge* (TBox) are fixed, while the only input is the *assertional knowledge* (ABox). Although we show that for the most general case the problem is undecidable, by limiting the amount of negation we obtain expressive languages whose queries can be answered in CONP and even NLOGSPACE, both over $DL\text{-}Lite_{\mathcal{R}}$ and $DL\text{-}Lite_{core}$. We then move to studying the *combined complexity*, where both the knowledge base and the query form the input. Here we obtain bounds ranging from NP-complete (thus matching the ones for ordinary CQs), to EXPTIME-complete, and undecidable for the full language. We would like to note that some of the results are obtained using the deep connection between DLs, DLXPath and *propositional dynamic logic*, thus providing some interesting new techniques for answering queries over ontologies.

Proofs Due to the space restrictions, we only give sketches of short proofs and ideas of longer ones.

2 Preliminaries

The language of $DL\text{-}Lite_{\mathcal{R}}$ (and $DL\text{-}Lite_{core}$) [1, 7] contains *individuals* c_1, c_2, \dots , *concept names* A_1, A_2, \dots , and *role names* P_1, P_2, \dots . *Concepts* B and *roles* R are defined by the grammar

$$B ::= A_i \mid \exists R, \quad R ::= P_j \mid P_j^-.$$

A $DL\text{-}Lite_{\mathcal{R}}$ TBox is a finite set of *concept* and *role inclusions* of the form

$$B_1 \sqsubseteq B_2, \quad B_1 \sqcap B_2 \sqsubseteq \perp, \quad R_1 \sqsubseteq R_2, \quad R_1 \sqcap R_2 \sqsubseteq \perp.$$

A $DL\text{-}Lite_{core}$ TBox contains only concept inclusions. An ABox is a finite set of *assertions* of the form $A_i(c_k)$ and $P_j(c_k, c_\ell)$. A *knowledge base* (KB) is a pair $(\mathcal{T}, \mathcal{A})$, where \mathcal{T} is a TBox and \mathcal{A} an ABox.

An *interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ is a nonempty domain $\Delta^{\mathcal{I}}$ of elements with an *interpretation function* $\cdot^{\mathcal{I}}$ that assigns an element $c_k^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ to each individual c_k , a subset $A_i^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$ to each concept name A_i , and a binary relation $P_j^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ to each role name P_j . When dealing with $DL\text{-}Lite$ it is usual to adopt the *unique name assumption* (UNA), and we do so here by requiring that $c_k^{\mathcal{I}} \neq c_\ell^{\mathcal{I}}$, for all individuals $c_k \neq c_\ell$. Our results, however, do not depend on UNA. The interpretation function $\cdot^{\mathcal{I}}$ is extended to roles and concepts in the following standard way:

$$\begin{aligned} (\exists R)^{\mathcal{I}} &= \{d \in \Delta^{\mathcal{I}} \mid \text{there is } d' \in \Delta^{\mathcal{I}} \text{ with } (d, d') \in R^{\mathcal{I}}\}, \\ (P_j^-)^{\mathcal{I}} &= \{(d', d) \in \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \mid (d, d') \in P_j^{\mathcal{I}}\}. \end{aligned}$$

The *satisfaction relation* \models for TBox inclusions and ABox assertions is also standard:

$$\begin{array}{lll}
\mathcal{I} \models B_1 \sqsubseteq B_2 & \text{iff} & B_1^{\mathcal{I}} \subseteq B_2^{\mathcal{I}}, \\
\mathcal{I} \models B_1 \sqcap B_2 \sqsubseteq \perp & \text{iff} & B_1^{\mathcal{I}} \cap B_2^{\mathcal{I}} = \emptyset, \\
\mathcal{I} \models R_1 \sqsubseteq R_2 & \text{iff} & R_1^{\mathcal{I}} \subseteq R_2^{\mathcal{I}}, \\
\mathcal{I} \models R_1 \sqcap R_2 \sqsubseteq \perp & \text{iff} & R_1^{\mathcal{I}} \cap R_2^{\mathcal{I}} = \emptyset, \\
\mathcal{I} \models A_i(c_k) & \text{iff} & c_k^{\mathcal{I}} \in A_i^{\mathcal{I}}, \\
\mathcal{I} \models P_j(c_k, c_\ell) & \text{iff} & (c_k^{\mathcal{I}}, c_\ell^{\mathcal{I}}) \in P_j^{\mathcal{I}}.
\end{array}$$

A KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ is *satisfiable* if there is an interpretation \mathcal{I} satisfying all inclusions of \mathcal{T} and assertions of \mathcal{A} . In this case we write $\mathcal{I} \models \mathcal{K}$ and say that \mathcal{I} is a *model* of \mathcal{K} .

3 DLXPath: XPath for Knowledge Bases

As mentioned in the introduction, the family of *DL-Lite* was designed not only to keep satisfiability and model checking problems simple, but mainly to keep the complexity of conjunctive query answering the same as in the case of relational databases, while simultaneously maximising the expressive power of the ontological language. This allows to use *DL-Lite* as a foundation for practical data managing applications, such as OWL 2 QL and OBDA. However, this does not automatically mean that other useful query formalisms with good evaluation properties over databases also have good properties when posed over *DL-Lite* knowledge bases. Hence, each class of queries which can be useful in knowledge base applications requires a separate research on its computational properties.

In this paper we concentrate on an adaptation of XPath query language for XML trees to knowledge bases. Recently it was shown that (a version of) XPath can be successfully used for querying graph databases [19]. Every interpretation of a *DL-Lite* vocabulary can be seen as a graph, and hence every *DL-Lite* KB is an incomplete description of a graph. That is why we expect our adaptation DLXPath for querying knowledge bases to be useful in practical applications.

In what follows, we will consider several fragments of DLXPath. We start with DLXPath_{core}, the fragment which corresponds to core XPath, the theoretical foundation of most practical languages for querying XML trees.⁴

Definition 1. Node formulas φ, ψ and path formulas α, β of DLXPath_{core} are expressions satisfying the grammar

$$\begin{array}{ll}
\varphi, \psi & := A \mid \neg\varphi \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \langle \alpha \rangle, \\
\alpha, \beta & := \varepsilon \mid R \mid [\varphi] \mid \alpha \cup \beta \mid \alpha \cdot \beta \mid \bar{\alpha} \mid R^+,
\end{array} \tag{1}$$

where A ranges over concept names and R ranges over roles (i.e., role names and their inverses).

Semantics $\llbracket \cdot \rrbracket^{\mathcal{I}}$ of DLXPath_{core} for an interpretation \mathcal{I} associates subsets of $\Delta^{\mathcal{I}}$ to node formulas and binary relations on $\Delta^{\mathcal{I}}$ to path formulas as given in Table 1.

⁴ The subscript ‘core’ is used in this paper for two unrelated purposes. This matching is historical and accidental, but we decided to stay with conventional notation despite this undesired collision.

$$\begin{aligned}
\llbracket A \rrbracket^{\mathcal{I}} &= A^{\mathcal{I}} \\
\llbracket \neg \varphi \rrbracket^{\mathcal{I}} &= \Delta^{\mathcal{I}} \setminus \llbracket \varphi \rrbracket^{\mathcal{I}} \\
\llbracket \varphi \wedge \psi \rrbracket^{\mathcal{I}} &= \llbracket \varphi \rrbracket^{\mathcal{I}} \cap \llbracket \psi \rrbracket^{\mathcal{I}} \\
\llbracket \varphi \vee \psi \rrbracket^{\mathcal{I}} &= \llbracket \varphi \rrbracket^{\mathcal{I}} \cup \llbracket \psi \rrbracket^{\mathcal{I}} \\
\llbracket \langle \alpha \rangle \rrbracket^{\mathcal{I}} &= \{d \mid \text{there exists } d' \text{ such that } (d, d') \in \llbracket \alpha \rrbracket^{\mathcal{I}}\} \\
\llbracket \varepsilon \rrbracket^{\mathcal{I}} &= \{(d, d) \mid d \in \Delta^{\mathcal{I}}\} \\
\llbracket R \rrbracket^{\mathcal{I}} &= R^{\mathcal{I}} \\
\llbracket [\varphi] \rrbracket^{\mathcal{I}} &= \{(d, d) \mid d \in \llbracket \varphi \rrbracket^{\mathcal{I}}\} \\
\llbracket \alpha \cup \beta \rrbracket^{\mathcal{I}} &= \llbracket \alpha \rrbracket^{\mathcal{I}} \cup \llbracket \beta \rrbracket^{\mathcal{I}} \\
\llbracket \alpha \cdot \beta \rrbracket^{\mathcal{I}} &= \llbracket \alpha \rrbracket^{\mathcal{I}} \circ \llbracket \beta \rrbracket^{\mathcal{I}} \\
\llbracket \bar{\alpha} \rrbracket^{\mathcal{I}} &= (\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}) \setminus \llbracket \alpha \rrbracket^{\mathcal{I}} \\
\llbracket R^+ \rrbracket^{\mathcal{I}} &\text{ is the transitive closure of } R^{\mathcal{I}}
\end{aligned}$$

Table 1. Semantics of $\text{DLXPath}_{\text{reg}}$. The symbol ‘\’ stands for set-theoretic difference.

As usual when dealing with ontologies, our interest is not query answering for a particular interpretation, but computing those answers that are true in all possible models of the knowledge base. Formally, let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be a knowledge base and α a DLXPath path formula. The *certain answers* of α over \mathcal{K} , denoted $\text{Certain}(\alpha, \mathcal{K})$, is the set of all pairs (c_1, c_2) of individuals such that $(c_1^{\mathcal{I}}, c_2^{\mathcal{I}}) \in \llbracket \alpha \rrbracket^{\mathcal{I}}$ for all models \mathcal{I} of \mathcal{K} . Similarly, one can define certain answers $\text{Certain}(\varphi, \mathcal{K})$ for a DLXPath node formula φ as the set of all individuals c such that $c^{\mathcal{I}} \in \llbracket \varphi \rrbracket^{\mathcal{I}}$, for all \mathcal{I} models of \mathcal{K} . In the paper all of the results will be stated for path formulas (*queries* from here on), however, they remain unchanged for node formulas.

Example 1. Coming back to the example from the introduction, consider role names HasDirectFlight , which connects cities with a direct flight, and HasUNESCOsite , which connects cities with their UNESCO world heritage sites, as well as concept InDanger denoting that a particular site is endangered. Let KB \mathcal{K} represents the flight destination graph, as well as knowledge about heritage sites, partially explicitly in the ABox, and partially implicitly, by means of TBox inclusions. Then checking whether it is possible to fly from Edinburgh to a city that has an endangered UNESCO world heritage site is equivalent to checking if Edinburgh is in $\text{Certain}(\varphi, \mathcal{K})$, where φ is a node formula

$$\langle \text{HasDirectFlight}^+[\langle \text{HasUNESCOsite}[\text{InDanger}] \rangle] \rangle. \quad \square$$

As already noted, the formalism of core XPath is in the nutshell of the most practical query languages for XML trees. However, in [19] it was shown that the corresponding graph language cannot express certain properties that are deemed essential when querying graphs. Thus, besides $\text{DLXPath}_{\text{core}}$ we also consider its generalisation called *regular DLXPath*, or $\text{DLXPath}_{\text{reg}}$, which extends the core fragment by allowing the use of transitive closure operator $^+$ over arbitrary path formulas. Formally, path formulas of $\text{DLXPath}_{\text{reg}}$ satisfy the grammar

$$\alpha, \beta := \varepsilon \mid R \mid [\varphi] \mid \alpha \cup \beta \mid \alpha \cdot \beta \mid \bar{\alpha} \mid \alpha^+,$$

while node formulas remain the same as for $\text{DLXPath}_{\text{core}}$ in grammar (1). As expected, the semantics $\llbracket \alpha^+ \rrbracket^{\mathcal{I}}$ over an interpretation \mathcal{I} is the transitive closure of $\llbracket \alpha \rrbracket^{\mathcal{I}}$.

Example 2. With full transitive closure we are now able to post more complex queries than in the core fragment. Consider again the knowledge base \mathcal{K} from Example 1. We can now ask if it is possible to fly from Liverpool to Jerusalem, making stopovers only in places that have an endangered UNESCO world heritage site by checking if the pair $(\text{Liverpool}, \text{Jerusalem})$ is in $\text{Certain}(\alpha, \mathcal{K})$, where α is a query

$$(\text{HasDirectFlight}[\langle \text{HasUNESCOsite}[\text{InDanger}] \rangle])^+.$$

Note that here we check if each of the cities along the path has an endangered site. If we want to additionally require that the sites along the route are not included in the UNESCO list under cultural criteria, we need to replace $[\text{InDanger}]$ with $[\text{InDanger} \wedge \neg \text{Cultural}]$ in the query above. \square

Besides these two query languages, we will consider their fragments, which will be introduced as needed.

Before passing on to the complexity of DLXPath query evaluation, we briefly compare our languages with other formalisms. First, DLXPath is clearly incomparable with CQs. However, all tree-shaped CQs and unions of CQs with no more than two free variables can be written as $\text{DLXPath}_{\text{core}}$ queries. On the other hand, all negation-free and $^+$ -free $\text{DLXPath}_{\text{reg}}$ queries can be written as unions of CQs, though with a cost of exponential blow-up. If we allow negation but only over concept and role names, then a query can be written as a union of CQs with safe negation. Second, $\text{DLXPath}_{\text{reg}}$ queries without negation and node tests $[\varphi]$ are *2-way regular path queries (2RPQs)*, the standard formalism for querying graph databases. If, additionally, role inverses are not allowed as path formulas, it becomes plain *RPQs*, that is, essentially, regular expressions. Some fragments of DLXPath can be expressed in other description logic formalisms. For example, it is well-known that a unary tree-shaped CQ with the corresponding tree being rooted and directed, can be written as a concept of \mathcal{EL} , a DL underlying OWL 2 EL profile [21]. Finally, $\text{DLXPath}_{\text{reg}}$ node formulas without path negation are nothing else but *propositional dynamic logic with converse (CPDL)* formulas; we will discuss this connection in more detail later on.

In the following sections we analyse the complexity of evaluating DLXPath queries over DL-Lite knowledge bases.

4 Data Complexity of DLXPath Query Evaluation

As it is widely accepted in theory and proved in practice, the size of the query and TBox is usually much smaller than the size of the ABox (see e.g., [25] for discussion in the relational database context and [7] for DLs). This is why one usually considers *data complexity* of query answering, assuming that the TBox and the query are fixed, and only ABox is part of the input. In this section we study this problem for various fragments of DLXPath . Formally, let \mathcal{T} be a TBox and α a DLXPath query. We are interested in the following family of problems.

CERTAIN ANSWERS (α, \mathcal{T})

Input: ABox \mathcal{A} and pair (c_1, c_2) of individuals

Question: Is $(c_1, c_2) \in \text{Certain}(\alpha, (\mathcal{A}, \mathcal{T}))$?

As it was previously mentioned, data complexity of CQ query answering over *DL-Lite_R* knowledge bases is the same as over relational databases, that is, in LOGSPACE [7]. At the first glance, one may expect a similar result in the case of DLXPath, where the complexity is NLOGSPACE-complete over graph databases [19]. However, the combination of the open world assumption and allowing negation in queries makes things quite different. In fact, the situation here is more like in the case of CQ with safe negation, where the complexity jump is dramatic: from polynomiality to undecidability [15].

Theorem 1. *There exists a DL-Lite_{core} TBox \mathcal{T} and a DLXPath_{core} query α such that the problem CERTAIN ANSWERS (α, \mathcal{T}) is undecidable.*

The proof uses similar techniques as the proof of Theorem 1 in [15], that shows the undecidability of the problem of computing certain answers for conjunctive queries with safe negation over *DL-Lite_R* knowledge bases.

Having this negative result, a natural direction is to search for DLXPath fragments that have decidable certain answers problem. Based on previous studies for XPath in XML and graph databases (see e.g., [3]), the most problematic primitive seems to be the negation $\bar{\alpha}$ in path formulas. Indeed, we now show that removing this primitive from the syntax leads to decidability of the certain answers problem. We denote DLXPath_{core}^{path-pos} the fragment of DLXPath_{core} that does not allow the negation $\bar{\alpha}$, for α a path formula, and define the class DLXPath_{reg}^{path-pos} accordingly. We then have the following theorem.

Theorem 2. *There exists a DL-Lite_{core} TBox \mathcal{T} and a DLXPath_{core}^{path-pos} query α such that the problem CERTAIN ANSWERS (α, \mathcal{T}) is CONP-hard. The problem is in CONP for any DL-Lite_R TBox \mathcal{T} and DLXPath_{reg}^{path-pos} query α .*

Proof (sketch). We first sketch the CONP-hardness of the problem. Consider a language with role names $P, T_1, T_2, T_3, F_1, F_2$, and F_3 , as well as concept name A . Fix a query

$$\alpha = \left[P \left(\bigvee_{v_1 \neq u_1, v_2 \neq u_2, v_3 \neq u_3} \left(\bigwedge_{i=1,2,3} \langle S_i^{v_i} [A^{u_i}] \rangle \right) \right) \right],$$

where all v_i and u_i range over $\{\top, \perp\}$; $S_i^{v_i}$ is T_i if $v_i = \top$ and F_i if $v_i = \perp$; A^{u_i} is A if $u_i = \top$ and $\neg A$ if $u_i = \perp$.

Consider the complement of the NP-complete problem 3CNF-SAT whose input is a conjunction Φ of clauses of the form $\ell_1 \vee \ell_2 \vee \ell_3$, where the ℓ_i are literals, that is, propositional variables or their negations, and whose answer is ‘yes’ if Φ is not satisfiable and ‘no’ otherwise.

For each variable x in Φ the ABox \mathcal{A} uses an individual c_x . Also, for each clause γ in Φ it uses an individual c_γ . Finally, it uses an individual c .

For each clause γ in Φ the ABox \mathcal{A} contains the assertion $P(c, c_\gamma)$. Also, for each literal ℓ_i in each clause $\gamma = \ell_1 \vee \ell_2 \vee \ell_3$ of Φ the ABox \mathcal{A} contains the assertion $T_i(c_\gamma, c_x)$ if $\ell_i = x$ and the assertion $F_i(c_\gamma, c_x)$ if $\ell_i = \bar{x}$.

It is a technicality to show that (c, c) is a certain answer to α over the KB (\mathcal{A}, \emptyset) if and only if Φ is not satisfiable.

The CONP upper bound can be obtained by a careful analysis of the proof of Theorem 4. \square

In fact, we can see that the reduction for hardness is done for the empty TBox, that is all the complexity is ‘contained’ in the formulation of the fixed query (which does not use the transitive closure operator $^+$).

While this theorem looks positive in light of the general undecidability result, the complexity might still be too high for practical applications, as it leads to algorithms which run in exponential time in the size of the ABox. To lower the complexity and obtain tractable algorithms, one could also consider fragments of DLXPath that do not allow any form of negation, neither in node nor in path formulas. We denote such fragments of $\text{DLXPath}_{\text{core}}$ and $\text{DLXPath}_{\text{reg}}$ by $\text{DLXPath}_{\text{core}}^{\text{pos}}$ and $\text{DLXPath}_{\text{reg}}^{\text{pos}}$, respectively. This last fragment is nothing else but *nested 2RPQs* [22], a language that has already been studied for *DL-Lite* knowledge bases [4], where an NLOGSPACE tight complexity bound was shown for the problem of computing certain answers for both *DL-Lite*_{core} and *DL-Lite* _{\mathcal{R}} . Furthermore, by adapting known results from graph databases, it is not difficult to extend this result to show that the problem is already NLOGSPACE hard even for $\text{DLXPath}_{\text{core}}^{\text{pos}}$ queries (see e.g., [2] for a survey on such techniques).

Note that from this result we conclude that nesting and inverse in regular expressions, as well as fixed *DL-Lite* _{\mathcal{R}} TBoxes do not increase the tractable data complexity of query evaluation.

5 Combined complexity of DLXPath Query Evaluation

Even if data complexity is the most important measure in practice, *combined complexity*, that is complexity under the assumption that both TBox, ABox and the query are given as input, allows us to get a better understanding of the query answering problem, and often provides a blueprint of how to solve the problem in practice. That is why we continue our study in this direction. Formally, we consider the following family of problems, where X ranges over $\{\text{core}, \mathcal{R}\}$, y over $\{\text{core}, \text{reg}\}$, and z is either nothing, or ‘path-pos’, or ‘pos’.

DL-Lite _{X} CERTAIN ANSWERS OVER DLXPath_y^z
Input: *DL-Lite* _{X} KB \mathcal{K} , DLXPath_y^z query α ,
and pair (c_1, c_2) of individuals
Question: Is $(c_1, c_2) \in \mathcal{C}ertain(\alpha, \mathcal{K})$?

As it immediately follows from Theorem 1, the problem remains undecidable for full $\text{DLXPath}_{\text{core}}$ and, hence, for full $\text{DLXPath}_{\text{reg}}$. However, the last result also follows from the connection of DLXPath with propositional dynamic logic (PDL) and some well known properties of PDL. Since this connection will be heavily used in the remainder of the paper we now discuss it in more detail.

First of all, we note that in PDL community a different terminology is used: for example, interpretations are called *Kripke structures*, concept names are called *propositional letters* or *variables*, role names are *atomic programs*, and inverse operator is *converse*. Though, to be consistent with the rest of the paper we stay with the DL terminology.

As already mentioned, node formulas of $\text{DLXPath}_{\text{reg}}^{\text{path-pos}}$ are, essentially, PDL with converse (CPDL) formulas (see [16] for a good introduction on the topic). Formally, the syntax of plain *propositional dynamic logic* (PDL) is the same as $\text{DLXPath}_{\text{reg}}^{\text{path-pos}}$ (i.e., CPDL), except that it does not allow the inverse P_j^- of role names as path expressions. The standard problem in the PDL community is the *satisfiability* of a node formula φ ; that is, checking whether there exists an interpretation \mathcal{I} and an element d in its domain such that φ holds in d (written, $\mathcal{I}, d \models \varphi$).

Lutz et al. showed that the satisfiability of PDL formulas extended with arbitrary path negation (such a logic is denoted PDL^\neg) is undecidable [20], which already implies the undecidability of the certain answers problem $\text{DLXPath}_{\text{reg}}$: indeed, a PDL^\neg formula φ is satisfiable if and only if the $\text{DLXPath}_{\text{reg}}$ query $[\neg\varphi]$ has the certain answer (c, c) over the empty KB with individual c in the language. Note, however, that it does not immediately imply undecidability in data complexity shown in Theorem 1.

Turning our attention to the certain answers problem for $\text{DLXPath}^{\text{path-pos}}$ queries, we again use the connection with the theory of PDL. It is well-known that the satisfiability problem for PDL is EXPTIME-complete [16]. It remains in EXPTIME even if these formulas are allowed to use the transitive closure operator $^+$ only over role names. The same holds for CPDL [16] and $\text{PDL}^{(\neg)}$, that is, the extension of PDL which allows path negation in a limited form—only on role names [20]. Similarly to the previous undecidability result, the EXPTIME lower bounds for these classes of PDL formulas already imply EXPTIME-hardness of the certain answers problem for $\text{DLXPath}_{\text{core}}^{\text{path-pos}}$, even for empty KBs. Also, in [4] hardness was established even for $\text{DLXPath}_{\text{reg}}^{\text{pos}}$.

The upper bound is, however, much more challenging. To deal with *DL-Lite \mathcal{R}* knowledge bases, in particular with role inclusions, we join the aforementioned extensions of PDL with inverse and negation on role names, and consider the language $\text{CPDL}^{(\neg)}$ whose node formulas obey the same grammar (1) as PDL (and DLXPath), and path formulas are defined as follows:

$$\alpha, \beta := \varepsilon \mid R \mid [\varphi] \mid \alpha \cup \beta \mid \alpha \cdot \beta \mid \overline{R} \mid \alpha^+.$$

We need the following result to establish complexity of the certain answers problem.

Theorem 3. *Checking satisfiability of a $\text{CPDL}^{(\neg)}$ node formula can be done in EXPTIME.*

The proof makes use of ideas from [20] and [26]. Although it is not strictly related to description logics and knowledge representation, we state the result explicitly as we believe it might be of interest to the PDL community.

Of course, satisfiability results do not transfer directly to query answering over KBs. However, widening and recasting the ideas from [12] and [13], we obtain the desired upper bound.

Lemma 1. *The problem DL-Lite_R CERTAIN ANSWERS OVER DLXPath_{reg}^{path-pos} is in EXPTIME.*

Summing up, we obtain the following theorem (the hardness results follows from [16] and [4] and are included just for completeness).

Theorem 4. *The problem DL-Lite_R CERTAIN ANSWERS OVER DLXPath_{reg}^{path-pos} is EXPTIME-complete. It remains EXPTIME-hard for DLXPath_{reg}^{pos} queries with DL-Lite_{core} KBs, and for DLXPath_{core}^{path-pos} even with empty KBs.*

The only remaining fragment is that of DLXPath_{core}^{pos} queries, that is, the restriction of the DLXPath_{core} language that use neither binary nor unary negation. In the previous section we saw that data complexity of answering DLXPath^{pos} queries is the same as answering RPQs and 2RPQs, regardless of whether we made use of the regular or the core fragment. For combined complexity, the case is now different. We have already seen that query answering remains EXPTIME-hard for DLXPath^{pos}. We now show that the restriction to the core fragment limits the complexity by almost one exponential.

Theorem 5. *The problem DL-Lite_R CERTAIN ANSWERS OVER DLXPath_{core}^{pos} is NP-complete. It remains NP-hard for DL-Lite_{core}.*

Proof (sketch). For designing an NP algorithm we rethink the ideas from the proof of Theorem 4.4 in [3], where a similar problem was solved in the context of XML. First, we show that a pair (c_1, c_2) is a certain answer to a query α over a satisfiable KB \mathcal{K} if and only if $(c_1, c_2) \in \llbracket \alpha \rrbracket^{Can(\mathcal{K})}$, where $Can(\mathcal{K})$ is the canonical model of \mathcal{K} (see e.g., [7] for definition). Second, if (c_1, c_2) is an answer, then this is witnessed by a sub-interpretation of $Can(\mathcal{K})$ of polynomial size. Hence, the algorithm can just guess this sub-interpretation.

To show hardness we give a reduction from the NP-complete problem 3CNF-SAT. Suppose we are given a conjunction Φ of clauses of the form $\ell_1 \vee \ell_2 \vee \ell_3$, where ℓ_k are literals, that is propositional variables or their negations (we can assume that all literals in each clause are distinct). Let x_1, \dots, x_n be the variables in Φ .

Consider the KB \mathcal{K} with the ABox $\mathcal{A} = \{A(c)\}$ and the TBox \mathcal{T} consisting of the inclusions

$$\begin{aligned} A &\sqsubseteq \exists(X_1^v)^-, \text{ for } v \in \{\top, \perp\}, \\ \exists X_{i-1}^v &\sqsubseteq \exists(X_i^u)^-, \text{ for } v, u \in \{\top, \perp\}, 1 < i \leq n. \end{aligned}$$

For every variable x_i and truth value $v \in \{\top, \perp\}$ let $\psi_{x_i}^v$ be the node formula

$$\langle (X_n^\top \cup X_n^\perp) \cdots (X_{i+1}^\top \cup X_{i+1}^\perp) \cdot X_i^v \cdot (X_{i-1}^\top \cup X_{i-1}^\perp) \cdots (X_1^\top \cup X_1^\perp) \rangle.$$

Let φ be the node formula $\langle \alpha[\psi] \rangle$, where

$$\alpha = ((X_1^\top)^- \cup (X_1^\perp)^-) \cdot \dots \cdot ((X_n^\top)^- \cup (X_n^\perp)^-),$$

and ψ is a conjunction of node formulas of the following form, for each clause γ in Φ , using variables x, y, z :

$$\bigvee_{\substack{(v_x, v_y, v_z) \text{ assignment of } (x, y, z) \\ \text{satisfying } \gamma}} \psi_x^{v_x} \wedge \psi_y^{v_y} \wedge \psi_z^{v_z}.$$

Complexity	DLXPath ^{pos}	DLXPath ^{path-pos}	DLXPath
Data	NLOGSPACE-c	CONP-c	undec.
Combined	NP-c* / EXPTIME-c [†]	EXPTIME-c	undec.

Table 2. A summary of the complexity results. Here “-c” stands for “-complete” and “undec.” for “undecidable”. All the results hold for both $DL-Lite_{core}$ and $DL-Lite_{\mathcal{R}}$, as well as for both $DLXPath_{core}$ and $DLXPath_{reg}$; the only exception is that * holds just for $DLXPath_{core}$ and [†] just for $DLXPath_{reg}$. The new results of this paper are set off in bold.

One can show that (c, c) is a certain answer to $[\varphi]$ over \mathcal{K} iff Φ is satisfiable. \square

It is worth to mention, that the result holds even if the queries are not allowed to use the transitive closure operator $^+$ at all, being, essentially, very restricted form of unions of CQs but with the same complexity of query answering. Hence, the border between tractable and intractable combined complexity of the problem lies somewhere very close to here, leaving 2RPQs on one side and $DLXPath_{core}^{pos}$ with CQs on the other.

6 Conclusions and Future Work

In this paper we conducted a detailed study about using the XPath language to query ontologies of the DL-Lite family. The results, summarised in Table 2, show that although the problem is generally undecidable, by limiting the amount of negation we can get decidable and even tractable fragments. The deep connection between XPath, DL and PDL allowed us to use ideas developed in other areas and gave insight on how query evaluation can be affected by certain aspects of the language. Although this connection was explored previously (see e.g., discussion in Chapters 13 and 14 in [6]), we believe that the time is ripe for a comprehensive survey describing how techniques from one area can be transferred to another and hope that the results of this paper can motivate such a survey.

From a practical point of view, we would like to find and test the classes of queries that are of particular practical interest and have either tractable general algorithms or reliable heuristics. Here we primarily want to tackle positive and path-positive fragments of XPath, as these queries were implemented and tested by the XML community, and many good heuristics have been developed over the years.

Acknowledgements. We thank Diego Figueira for his precious comments. Reutter is funded by FONDECYT grant 11130648 and the Millennium Nucleus Center for Semantic Web Research under Grant NC120004.

References

1. A. Artale, D. Calvanese, R. Kontchakov, and M. Zakharyashev. The DL-Lite family and relations. *J. Artif. Intell. Res. (JAIR)*, 36:1–69, 2009.
2. P. Barceló. Querying graph databases. In *PODS*, pages 175–188, 2013.

3. M. Benedikt, W. Fan, and F. Geerts. XPath satisfiability in the presence of DTDs. *J. ACM*, 55(2), 2008.
4. M. Bienvenu, D. Calvanese, M. Ortiz, and M. Šimkus. Nested regular path queries in description logics. *In KR*, 2014.
5. M. Bienvenu, M. Ortiz, and M. Šimkus. Conjunctive regular path queries in lightweight description logics. *In IJCAI*, 2013.
6. P. Blackburn, J. F. A. K. v. Benthem, and F. Wolter. *Handbook of Modal Logic, Volume 3 (Studies in Logic and Practical Reasoning)*. Elsevier Science Inc., New York, NY, USA, 2006.
7. D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *J. of Automated Reasoning*, 39(3):385–429, 2007.
8. D. Calvanese, G. De Giacomo, M. Lenzerini, and M. Vardi. Containment of conjunctive regular path queries with inverse. *In 7th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 176–185, 2000.
9. D. Calvanese, T. Eiter, and M. Ortiz. Answering regular path queries in expressive description logics: An automata-theoretic approach. *In AAAI*, pages 391–396, 2007.
10. D. Calvanese, E. Kharlamov, W. Nutt, and C. Thorne. Aggregate queries over ontologies. In R. Elmasri, M. Doerr, M. Brochhausen, and H. Han, editors, *ONISW*, pages 97–104. ACM, 2008.
11. D. Calvanese, M. Ortiz, and M. Šimkus. Containment of regular path queries under description logic constraints. *In IJCAI*, pages 805–812, 2011.
12. G. De Giacomo and M. Lenzerini. Boosting the Correspondence between Description Logics and Propositional Dynamic Logics. *In AAAI*, pages 205–212, 1994.
13. G. De Giacomo and M. Lenzerini. TBox and ABox Reasoning in Expressive Description Logics. *In KR*, pages 316–327, 1996.
14. B. Glimm, C. Ogbuji, S. Hawke, I. Herman, B. Parsia, A. Polleres, and A. Seaborne. SPARQL 1.1 entailment regimes, 2013. W3C Recommendation 21 March 2013, <http://www.w3.org/TR/2013/REC-sparql11-entailment-20130321/>.
15. V. Gutiérrez-Basulto, Y. A. Ibáñez-García, R. Kontchakov, and E. V. Kostylev. Conjunctive queries with negation over dl-lite: A closer look. *In RR*, pages 109–122, 2013.
16. D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. MIT Press, 2000.
17. S. Harris and A. Seaborne. Sparql 1.1 query language. *W3C working draft*, 14, 2010.
18. E. V. Kostylev and J. L. Reutter. Answering counting aggregate queries over ontologies of the DL-Lite family. *In AAAI*, 2013.
19. L. Libkin, W. Martens, and D. Vrgoč. Querying graph databases with XPath. *In ICDT*, pages 129–140, 2013.
20. C. Lutz and D. Walther. PDL with Negation of Atomic Programs. *Journal of Applied Non-Classical Logics*, 15(2):189–213, 2005.
21. B. Motik, B. Cuenca Grau, I. Horrocks, Z. Wu, A. Fokoue, and C. Lutz. *OWL 2 Web Ontology Language Profiles (2nd Edition)*, 2012. W3C Recommendation.
22. J. Pérez, M. Arenas, and C. Gutierrez. nSPARQL: A navigational language for RDF. *Web Semantics: Science, Services and Agents on the World Wide Web*, 8(4):255 – 270, 2010.
23. I. Robinson, J. Webber, and E. Eifrem. *Graph databases*. ” O’Reilly Media, Inc.”, 2013.
24. R. Rosati. The limits of querying ontologies. *In Proc. of the 11th Int. Conf. on Database Theory (ICDT)*, volume 4353 of LNCS, pages 164–178. Springer, 2007.
25. M. Y. Vardi. The complexity of relational query languages (extended abstract). *In STOC*, pages 137–146, 1982.
26. M. Y. Vardi and P. Wolper. Automata-Theoretic Techniques for Modal Logics of Programs. *J. Comput. Syst. Sci.*, 32(2):183–221, 1986.
27. XML Path Language (XPath) 2.0 (Second Edition). www.w3.org/TR/xpath20, 2010.