# Regular Path Queries in Lightweight Description Logics: Complexity and Algorithms 

Meghyn Bienvenu<br>MEGHYN@LRI.FR<br>Laboratoire de Recherche en Informatique, CNRS $\xi^{\text {G Uiversité Paris-Sud, France }}$<br>Magdalena Ortiz<br>ORTIZ@KR.TUWIEN.AC.AT<br>Mantas Šimkus<br>SIMKUS@DBAI.TUWIEN.AC.AT<br>Institute of Information Systems,<br>TU Wien, Austria


#### Abstract

Conjunctive regular path queries are an expressive extension of the well-known class of conjunctive queries. Such queries have been extensively studied in the (graph) database community, since they support a controlled form of recursion and enable sophisticated path navigation. Somewhat surprisingly, there has been little work aimed at using such queries in the context of description logic (DL) knowledge bases, particularly for the lightweight DLs that are considered best suited for data-intensive applications. This paper aims to bridge this gap by providing algorithms and tight complexity bounds for answering twoway conjunctive regular path queries over DL knowledge bases formulated in lightweight DLs of the DL-Lite and $\mathcal{E L}$ families. Our results demonstrate that in data complexity, the cost of moving to this richer query language is as low as one could wish for: the problem is NL-complete for DL-Lite and P-complete for $\mathcal{E L}$. The combined complexity of query answering increases from NP- to PSpace-complete, but for two-way regular path queries (without conjunction), we show that query answering is tractable even with respect to combined complexity. Our results reveal two-way conjunctive regular path queries as a promising language for querying data enriched by ontologies formulated in DLs of the DL-Lite and $\mathcal{E L}$ families or the corresponding OWL 2 QL and EL profiles.


## 1. Introduction

Recent years have seen a rapidly growing interest in using description logic (DL) ontologies to query instance data. This setting can be seen as a generalization of the related problem of querying graph databases which, like DL instance data, are sets of ground facts using only unary and binary predicates, i.e., node- and edge-labeled graphs (Consens \& Mendelzon, 1990; Barceló, Libkin, Lin, \& Wood, 2012). The relevance of both problems lies in the fact that in many application areas, data can be naturally represented in such form. This applies, in particular, to XML and RDF data. In the presence of a DL ontology, the domain knowledge expressed in the ontology is exploited when querying the data, which can facilitate query formulation and provide users with more complete answers to their queries. While the DL and database communities share some common research goals, the research agendas they have pursued differ significantly. In the DL research, the focus has been on studying the computational complexity of answering (unions of) plain conjunctive queries (CQs) in the presence ontological constraints expressed in different DLs, and on
the development of efficient algorithms for this setting, see the surveys of Ortiz (2013) and Ortiz and Šimkus (2012). By contrast, the work on graph databases typically does not consider ontological knowledge, but instead aims at supporting expressive navigational query languages.

Regular path queries (RPQs) constitute the basic navigational query language. Formally, an RPQ is given as a regular language (represented as a regular expression or a finite automaton) over the binary predicates in the database facts (arc labels, or roles in DL parlance), and it returns all pairs of objects that are connected by a path whose label is a word belonging to the specified language. A crucial feature of these queries is that they allow for controlled form of recursion that is computationally well behaved yet sufficient for expressing reachability queries and the traversal of paths of unbounded length (Florescu, Levy, \& Suciu, 1998). In two-way RPQs (2RPQs), the regular expressions may use the arc labels in the backwards direction, which allows for more flexible path navigation. Notably, this comes at no computational cost: answering both RPQs and 2RPQs over (plain) graph databases is complete for NL in combined complexity (that is, when the complexity is measured in terms of the whole input, which in this case consists of the query and the data). Conjunctive (two-way) regular path queries (C(2)RPQs), which are one of the most expressive and popular languages for querying graph databases, simultaneously extend plain CQs and (2)RPQs by allowing for conjunctions of atoms that can share variables in arbitrary ways, where the atoms may contain regular expressions that navigate the arcs of the database. If we consider the data complexity measure (in which the complexity is measured only in terms of the size of the data, with all other inputs considered as fixed), then answering C2RPQs is still NL-complete. In combined complexity, the C2RPQ answering problem is NP-complete, which is in fact the lowest complexity that could be expected, given that CQ answering is already NP-hard. We note that the navigational capabilities provided by RPQs and their extensions have long been considered crucial for querying data on the Web. Indeed, navigation along regular paths is at the heart of the XPath language for querying XML data (Berglund, Boag, Chamberlin, Fernández, Kay, Robie, \& Siméon, 2007), and SPARQL 1.1, the language recently recommended by the World Wide Web Consortium (W3C) as the new standard for querying RDF data (Harris \& Seaborne, 2013), adds to the previous standard a feature called property paths, which roughly amounts to extending the 'core' of SPARQL from CQs to C2RPQs.

It comes as a surprise that, despite their advantages and relevance, RPQs and their extensions have received rather little attention in the DL literature. They were first considered in the seminal work of Calvanese, de Giacomo, and Lenzerini (1998) on query answering in the presence of DL ontologies. However, the vast majority of subsequent research has targeted instance queries (IQs), conjunctive queries and unions thereof, and occasionally positive first-order queries. Only a handful of works have considered C2RPQs. Calvanese et al. $(2014,2007,2009)$ showed that C2RPQ answering is 2ExpTimE-complete in combined complexity for the very expressive $\mathrm{DLs}^{1} \mathcal{Z I \mathcal { Q }}, \mathcal{Z I O}$, and $\mathcal{Z O Q}$, which allow for regular expressions as role constructors. The same upper bound was shown for containment of C2RPQs in the presence of rich ontological knowledge (Calvanese et al., 2009). This is noteworthy since it is a well-known fact that most forms of recursion make query answer-

[^0]ing and query containment undecidable in the presence of ontological constraints (Levy \& Rousset, 1996). Moreover, this complexity is not higher than that of other significantly more restricted settings, such as answering plain CQs in the DL $\mathcal{A L C I}$ (Lutz, 2008) or positive first-order queries in $\mathcal{A L C}$ (Ortiz \& Šimkus, 2014). However, hardness for 2ExpTime is nevertheless a prohibitively high complexity for many applications. Even in data complexity, the algorithms underlying the aforementioned results still need exponential time. More recently, algorithms for answering C2RPQs in Horn-SHOIQ and Horn-SROIQ were proposed (Ortiz, Rudolph, \& Šimkus, 2011). These algorithms run in polynomial time in the size of the data, but may still require exponential time in the size of the ontology, which is worst-case optimal for these logics. By contrast, for the lightweight DLs of the DL-Lite (Calvanese, De Giacomo, Lembo, Lenzerini, \& Rosati, 2007) and $\mathcal{E L}$ (Baader, Brandt, \& Lutz, 2005) families, which are the languages of choice for ontology-mediated query answering and notably underlie the OWL 2 QL and EL profiles (Motik, Cuenca Grau, Horrocks, Wu, Fokoue, \& Lutz, 2012), the precise complexity of answering C2RPQs was left open: the ExpTime upper bound in combined complexity for more expressive Horn DLs does not match the NP lower bound stemming of answering plain CQs, and it is not at all apparent how to obtain better upper bounds for C2RPQs by adapting existing techniques. For the DL-Lite family, the data complexity was also left open, with a gap between the NL-hardness of RPQ answering inherited from the graph database setting and the P upper bound that had been established for more expressive Horn DLs.

In this paper, we close these open questions by presenting algorithms and precise complexity bounds for answering (C) 2 RPQs in the $\mathcal{E} \mathcal{L}$ and DL-Lite families of lightweight DLs. Many of our results were first announced in the conference version of this paper (Bienvenu, Ortiz, \& Šimkus, 2013), but here we provide full proofs of these results, additional examples, and a discussion of extensions of our results and their applicability to the OWL 2 profiles. We also strengthen some of the complexity lower bounds by considering more restricted classes of queries (Proposition 4.5) or more succinct representations (Theorem 4.2) and identify interesting restrictions that lead to better combined complexity for C2RPQs (Theorems 6.10 and 6.11). Our main contributions can be summarized as follows:

- We establish a P lower bound in combined complexity for answering 2RPQs in DL-Lite, as well as for RPQs in DL-Lite $\mathcal{R}_{\mathcal{R}}$, which can be contrasted with the NLcompleteness of instance checking in these logics. This result improves upon a similar lower bound from the conference version of this paper by adopting the (less succinct) regular expression representation of queries.
- We present an algorithm for answering 2RPQs over DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$ knowledge bases that runs in polynomial time in combined complexity. This tractability result is extended to all single-atom C2RPQs, including those with existential variables.
- We show that answering CRPQs is PSpace-hard for both DL-Lite and $\mathcal{E L}$. This result had already been shown in the conference version, but here we provide a different proof that holds even for the structurally-restricted class of strongly acyclic CRPQs in which the regular expressions are disjunction-free and of star-height two. This hardness result is interesting when compared with the graph database setting, where

Bienvenu, Ortiz, \& Šimkus

C2RPQ answering is NP-complete (and thus not harder than CQs in the worst case) and becomes feasible in polynomial time when restricted to strongly acyclic C2RPQs.

- We develop a query rewriting procedure for answering C2RPQs in DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E} \mathcal{L H}$, which we use to show that the problem is feasible in PSpace for both logics. The PSpace upper bound for $\mathcal{E L H}$ is especially interesting, since C2RPQs allow for inverse roles and it is well known that adding inverse roles to $\mathcal{E L}$ immediately leads to ExpTime-hardness of query answering, even for instance queries. This result demonstrates that by including inverses in the query language, rather than the ontology language, it is possible to obtain algorithms that use only polynomial space.
- Using the same algorithm, we derive an NL upper bound in data complexity for DLLite $_{\mathcal{R}}$ and a P upper bound for $\mathcal{E} \mathcal{L H}$. In both cases, this is the lowest data complexity that could be expected in light of existing results.
- We also identify cases where C2RPQ answering is feasible in NP, and thus not harder than answering plain CQs. This is the case for queries whose existential variables do not occur in joins (or only occur in joins in a restricted way), or for arbitrary C2RPQs whenever the ontology is guaranteed to have a finite canonical model.

Our new complexity results (and relevant existing results) are summarized in Figure 1.
This paper is organized as follows. We begin in Section 2 by introducing the lightweight description logics $\mathrm{DL}^{-L i t e} \mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$ (and relevant sublogics) and recalling some basic notions related to regular languages and computational complexity. In Section 3, we define the syntax and semantics of the different types of path queries considered in this paper. Section 4 is dedicated to showing lower bounds, first for RPQs and then for CRPQs. In both cases, we start by presenting some easy bounds that follow from known results, before moving on to the main results. Section 5 presents algorithmic techniques and upper bounds for 2RPQs, and then in Section 6, we show how these can be extended to handle C2RPQs. In Section 7, we give a brief overview of related results for similar query languages and other DLs and discuss the applicability of our results to the profiles of the OWL Web Ontology Language. Conclusions and directions for future work are given in Section 8. To improve the readability of the paper, one of the more technical proofs is deferred to the appendix.

## 2. Preliminaries

We briefly recall some basics of description logics, a few computational complexity classes that are relevant in the paper, and some notation for regular languages.

### 2.1 Description Logics

We first recall the syntax and semantics of description logics, focusing on the lightweight families of logics DL-Lite (Calvanese et al., 2007) and $\mathcal{E L}$ (Baader et al., 2005). We also recall the definition of canonical model for these logics.

|  | IQ |  | (2)RPQ |  |
| :---: | :---: | :---: | :---: | :---: |
|  | data | combined | data | combined |
| DL-Lite ${ }_{\text {RDFS }}$ | in $\mathrm{AC}^{0}$ | NL-c | NL-c | NL-c |
|  |  | $\geq$ (A) | $\geq$ (A) | $\leq$ Thm 5.2 |
| DL-Lite $_{(\mathcal{R})}$ | in $\mathrm{AC}^{0}$ | NL-c | NL-c | P-c ${ }^{\dagger}$ |
|  | $\leq$ (G) | $\leq$ (G) | $\leq$ Thm 5.9 | $\geq$ Thm 4.2, $\leq$ Thm 5.9 |
| $\mathcal{E L}(\mathcal{H})$ | P-c | P-c | P-c | P-c |
|  | $\geq$ (E) | $\leq$ (C) | $\leq$ Thm 5.9 | $\leq$ Thm 5.9 |
|  | CQ |  | C(2)RPQ |  |
|  | data | combined | data | combined |
| DL-Lite ${ }_{\text {RDFS }}$ | in $\mathrm{AC}^{0}$ | NP-c | NL-c | NP-c |
|  |  | $\geq$ (B) | $\leq$ Thm 6.8 | $\leq$ Thm 6.8 |
| DL-Lite $_{(\mathcal{R})}$ | in $\mathrm{AC}^{0}$ | NP-c | NL-c | PSpace-c |
|  | $\leq$ (D) | $\leq$ (D) | $\leq$ Thm 6.8 | $\geq \text { Prop } 4.5, \leq \text { Thm } 6.8$ |
| $\mathcal{E L}(\mathcal{H})$ | P-c | NP-c | P-c | PSpace-c |
|  | $\leq$ (F) | $\leq$ (F) | $\leq$ Thm 6.8 | $\geq \text { Prop } 4.5, \leq \text { Thm } 6.8$ |

Figure 1: Complexity of query answering. The 'c' indicates completeness results, and ' $\leq$ ' and ' $\geq$ ' are used for upper and lower bounds respectively. New results are marked in bold. The remaining annotations have the following meanings:
$\dagger$ P-hardness for RPQs applies only to DL-Lite $\mathcal{R}_{\mathcal{R}}$
(A) Easy reduction from the NL-complete directed reachability problem
(B) Follows from NP-hardness of CQ answering over relational databases
(C) Baader et al. (2005)
(D) Calvanese et al. (2007)
(E) Calvanese, De Giacomo, Lembo, Lenzerini, and Rosati (2006)
(F) Rosati (2007), Krisnadhi and Lutz (2007), Krötzsch and Rudolph (2007)
(G) Calvanese et al. (2007), Artale, Calvanese, Kontchakov, and Zakharyaschev (2009)

### 2.1.1 Description Logic Syntax

As usual, we assume countably infinite, mutually disjoint sets $\mathrm{N}_{\mathrm{C}}, \mathrm{N}_{\mathrm{R}}$, and $\mathrm{N}_{\mathrm{I}}$ of concept names, role names, and individuals, respectively. An inverse role takes the form $r^{-}$where $r \in \mathrm{~N}_{\mathrm{R}}$. We will use $\mathrm{N}_{\mathrm{R}}^{ \pm}$to refer to $\mathrm{N}_{\mathrm{R}} \cup\left\{r^{-} \mid r \in \mathrm{~N}_{\mathrm{R}}\right\}$, and if $R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$, we use $R^{-}$to mean $r^{-}$if $R=r$ and $r$ if $R=r^{-}$.

A description logic knowledge base (KB) $\mathcal{K}=(\mathcal{T}, \mathcal{A})$ consists of a TBox $\mathcal{T}$ and an ABox $\mathcal{A}$. The former provides general domain knowledge, while the latter expresses facts

Bienvenu, Ortiz, \& Šimkus
about particular entities. Sometimes we use the generic terms ontology and data(set) in place of TBox and ABox.

Formally, a TBox is a finite set of inclusions, whose form depends on the DL in question. In DL-Lite, TBoxes consist of a set of concept inclusions of the form $B \sqsubseteq C$, with $B$ and $C$ concepts constructed according to the following syntax:

$$
B:=A|\exists R \quad C:=B| \neg B
$$

where $A \in \mathrm{~N}_{\mathrm{C}}$ and $R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$. $\mathrm{DL}^{-L_{i t e}} \mathcal{R}$ additionally allows for role inclusions of the form $R_{1} \sqsubseteq(\neg) R_{2}$, where $R_{1}, R_{2} \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$. The logic DL-Lite ${ }_{\text {RDFS }}$ is obtained from DL-Lite $\mathcal{R}_{\mathcal{R}}$ by disallowing inclusions that contain negation or have existential concepts $(\exists R)$ on the righthand side. $\mathrm{DL}^{-L i t e}{ }_{\mathcal{R}}$ is the basis for the OWL 2 QL profile, and DL-Lite ${ }_{\text {RDFS }}$ corresponds to the fragment of DL-Lite $\mathcal{R}_{\mathcal{R}}$ that is expressible in the RDF Schema ontology language (Brickley \& Guha, 2014).

In $\mathcal{E} \mathcal{L}$, concept inclusions have the form $C_{1} \sqsubseteq C_{2}$ where $C_{1}, C_{2}$ are complex concepts constructed according to the following syntax:

$$
C:=\top|A| C \sqcap C \mid \exists r . C
$$

where $A \in \mathrm{~N}_{\mathrm{C}}$ and $r \in \mathrm{~N}_{\mathrm{R}}$. The $\mathrm{DL} \mathcal{E} \mathcal{L H}$ additionally allows for role inclusions of the form $r_{1} \sqsubseteq r_{2}$, where $r_{1}, r_{2} \in \mathrm{~N}_{\mathrm{R}}$. Note that in $\mathcal{E} \mathcal{L}(\mathcal{H})$ TBoxes, inverse roles are not permitted.

We will use $\operatorname{sig}(\mathcal{T})$ to denote the signature of a $\mathrm{TBox} \mathcal{T}$, that is, the set of all concept and role names appearing in $\mathcal{T}$. It will also prove useful to introduce the set $\mathrm{BC}_{\mathcal{T}}$ of basic concepts for a TBox $\mathcal{T}$, defined as follows: $\mathrm{BC}_{\mathcal{T}}=\left(\mathrm{N}_{\mathrm{C}} \cap \operatorname{sig}(\mathcal{T})\right) \cup\left\{\exists r, \exists r^{-} \mid r \in \mathrm{~N}_{\mathrm{R}} \cap \operatorname{sig}(\mathcal{T})\right\}$ if $\mathcal{T}$ is a $\mathrm{DL}^{-L i t e} \mathcal{R}_{\mathcal{R}}$ TBox, and $\mathrm{BC}_{\mathcal{T}}=\left(\mathrm{N}_{\mathrm{C}} \cap \operatorname{sig}(\mathcal{T})\right) \cup\{\top\}$ if $\mathcal{T}$ is an $\mathcal{E} \mathcal{L} \mathcal{H}$ TBox.

In all of the considered DLs, an ABox is a finite set of concept assertions of the form $A(b)$ and role assertions of the form $r(b, c)$, where $A \in \mathrm{~N}_{\mathrm{C}}, r \in \mathrm{~N}_{\mathrm{R}}$, and $b, c \in \mathrm{~N}_{\mathrm{I}}$. We use $\operatorname{Ind}(\mathcal{A})$ to refer to the set of individuals appearing in the ABox $\mathcal{A}$.

### 2.1.2 Description Logic Semantics

The semantics of DL KBs is defined in terms of interpretations, which take the form $\mathcal{I}=$ $\left(\Delta^{\mathcal{I}},{ }^{\mathcal{I}}\right)$, where $\Delta^{\mathcal{I}}$ is a non-empty set and.$^{\mathcal{I}}$ maps each individual $a \in \mathrm{~N}_{\mathrm{I}}$ to $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, each concept name $A \in \mathrm{~N}_{\mathrm{C}}$ to $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, and each role name $r \in \mathrm{~N}_{\mathrm{R}}$ to $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The function ${ }^{\mathcal{I}}$ is extended to general concepts and roles as follows:

$$
\begin{aligned}
\top^{\mathcal{I}} & =\Delta^{\mathcal{I}} \\
(\neg A)^{\mathcal{I}} & =\Delta^{\mathcal{I}} \backslash A^{\mathcal{I}} \\
(\neg R)^{\mathcal{I}} & =\left(\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}\right) \backslash R^{\mathcal{I}}
\end{aligned}
$$

$$
\left(r^{-}\right)^{\mathcal{I}}=\left\{(c, d) \mid(d, c) \in r^{\mathcal{I}}\right\}
$$

$$
(\exists R)^{\mathcal{I}}=\left\{c \mid \exists d:(c, d) \in R^{\mathcal{I}}\right\}
$$

$$
(\exists r . C)^{\mathcal{I}}=\left\{c \mid \exists d:(c, d) \in r^{\mathcal{I}}, d \in C^{\mathcal{I}}\right\}
$$

An interpretation $\mathcal{I}$ satisfies an inclusion $G \sqsubseteq H$, denoted $\mathcal{I} \models G \sqsubseteq H$, if $G^{\mathcal{I}} \subseteq H^{\mathcal{I}}$. Similarly, $\mathcal{I}$ satisfies an assertion $A(a)$ if $a^{\mathcal{I}} \in A^{\mathcal{I}}$, in symbols $\mathcal{I} \models A(a) ; \mathcal{I}$ satisfies an assertion $r(a, b)$ if $\left(a^{\mathcal{I}}, b^{\mathcal{I}}\right) \in r^{\mathcal{I}}$, in symbols $\mathcal{I} \models r(a, b)$. An interpretation $\mathcal{I}$ is a model of $\mathcal{T}$, if $\mathcal{I}$ satisfies all inclusions in $\mathcal{T}$; it is a model of $\mathcal{A}$ if it satisfies all assertions in $\mathcal{A}$; and it is a model of $(\mathcal{T}, \mathcal{A})$ if it is a model of $\mathcal{T}$ and $\mathcal{A}$. A $\mathrm{KB}(\mathcal{T}, \mathcal{A})$ is satisfiable if it possesses at least one model, else it is unsatisfiable. Note that $\mathcal{E} \mathcal{L} \mathcal{H}$ knowledge bases are


Figure 2: Example ABox $\mathcal{A}_{\text {mob }}$
always satisfiable. If $\mathcal{G}$ is a TBox, ABox, or KB , and $\alpha$ an inclusion or assertion, we say that $\mathcal{G}$ entails $\alpha$, written $\mathcal{G} \models \alpha$, if $\mathcal{I} \models \alpha$ for every model $\mathcal{I}$ of $\mathcal{G}$.

Observe that we do not make the Unique Names Assumption (UNA), as our definition of interpretation allows distinct individuals to be mapped to the same domain element. We remark however that all of the results in this paper hold equally well under the UNA.

Example 2.1. As a motivating example, we consider the domain of public transport and urban mobility. A partial database for this domain is presented in Figure 2. The nodes in the two leftmost columns (shaded in blue) correspond to public transport lines and stations. The arrows between them represent the existing connections, and are labeled according to the type of transport line serving them: sbSub stands for 'served by subway', while sbLFT and sbHFT respectively stand for 'served by low-floor tramway' and 'served by high-floor tramway'. Nodes are also labeled with classes in which they participate. In particular, the labels GLStop and EquipStop indicate two specific kinds of public transport stops: groundlevel stops, and stops that are suitably equipped with ramps and elevators for passengers with restricted mobility. The remaining columns contain places of interest, and the loclnlabeled arrows between them represent the 'located in' relation. Note that this is nothing else but a graphical representation of a DL ABox, which we call $\mathcal{A}_{\text {mob }}$, where each label $A$ on a node $b$ corresponds to the concept assertion $A(b)$, and an arc labeled $r$ from node $b$ to node $c$ corresponds to the role assertion $r(b, c)$.

In our example, we assume that the information in the first two columns is provided and maintained by the local public transport authorities, thus it is complete and well structured. In contrast, the remaining data is crowd-sourced, thus it is likely to be incomplete and may not adhere to a rigid, predefined structure.

In order to better respond to user queries, this data is enriched with domain knowledge expressed by the $\mathcal{E L}$ ontology in Figure 3. The ontology defines subclass relations between concepts like 'stop' and the more specialized 'accessible stop', 'ground-level stop', and it defines new terms not present in the data but which may be useful at query time, such as 'food services'. It also enhances the possibly incomplete data by asserting the existence of other places of interest. For example, a 'family-friendly' location has both dining and playground facilities.
(1) An accessible stop (AccStop) is a public transport stop (Stop). A ground-level stop (GLStop) or a stop that is suitably equipped with ramps and elevators (EquipStop) is an accessible stop.
(2) Restaurants and cafes are food services (FoodServ).
(3) A place that is family friendly (FamFriendly) has some food service and a playground.
(4) A shopping center has a supermarket and a food court.
(5) A food court has some food service.

| AccStop | $\sqsubseteq$ Stop | $(1 \mathrm{a})$ |
| ---: | :--- | ---: |
| GLStop | $\sqsubseteq$ AccStop | $(1 \mathrm{~b})$ |
| EquipStop | $\sqsubseteq$ AccStop | $(1 \mathrm{c})$ |
| Restaurant | $\sqsubseteq$ FoodServ | $(2 \mathrm{a})$ |
| Cafe | $\sqsubseteq$ FoodServ | $(2 \mathrm{~b})$ |
| FamFriendly | $\sqsubseteq \exists$ hasFacility.FoodServ | $(3 \mathrm{a})$ |
| FamFriendly | $\sqsubseteq \exists$ hasFacility.Playground | $(3 \mathrm{~b})$ |
| ShoppingCenter | $\sqsubseteq \exists$ hasFacility.Foodcourt | $(4 \mathrm{a})$ |
| ShoppingCenter | $\sqsubseteq \exists$ hasFacility.Supermarket (4b) |  |
| Foodcourt | $\sqsubseteq \exists$ hasFacility.FoodServ | (5a) |

Figure 3: Example DL-Lite TBox $\mathcal{T}_{\text {mob }}$ expressing domain knowledge

To keep the example compact, we have chosen to write the example ontology in $\mathcal{E} \mathcal{L}$, but the same knowledge can be expressed in $\mathrm{DL}^{-L i t e} \mathcal{R}_{\mathcal{R}}$ by using auxiliary roles to simulate the concept inclusions with qualified existential quantification on the right-hand-side. For instance, the concept inclusion (3a) can be replaced by the following three inclusions in the syntax of DL-Lite:

FamFriendly $\sqsubseteq \exists$ hasFoodServ<br>$\exists$ hasFoodServ ${ }^{-}$「FoodServ<br>hasFoodServ $\sqsubseteq ~ h a s F a c i l i t y ~$

and similarly for $(3 \mathrm{~b}),(4 \mathrm{a}),(4 \mathrm{~b})$, and (5a).

### 2.1.3 Normal Form for $\mathcal{E L H}$ TBoxes

To simplify the presentation, we will assume throughout the paper that $\mathcal{E} \mathcal{L H}$ TBoxes are in normal form, meaning that all concept inclusions are of one of the following forms:

$$
A \sqsubseteq B \quad A \sqsubseteq \exists r . B \quad A_{1} \sqcap A_{2} \sqsubseteq B \quad \exists r . B \sqsubseteq A
$$

with $A, A_{1}, A_{2}, B \in \mathrm{~N}_{\mathrm{C}} \cup\{\top\}$. The following well-known property (see (Baader et al., 2005)) shows that this assumption is without loss of generality.

Proposition 2.2. For every $\mathcal{E} \mathcal{L H}$ TBox $\mathcal{T}$, one can construct in polynomial time an $\mathcal{E} \mathcal{L H}$ TBox $\mathcal{T}^{\prime}$ in normal form (possibly using new concept names) such that $\mathcal{T}^{\prime}$ is a model conservative extension of $\mathcal{T}$, that is, every model of $\mathcal{T}^{\prime}$ is a model of $\mathcal{T}$, and for every model $\mathcal{I}$ of $\mathcal{T}$, there is a model $\mathcal{I}^{\prime}$ of $\mathcal{T}^{\prime}$ such that $\mathcal{I}$ and $\mathcal{I}^{\prime}$ have the same domain and coincide on the interpretation of all concept and role names except those in $\operatorname{sig}\left(\mathcal{T}^{\prime}\right) \backslash \operatorname{sig}(\mathcal{T})$.

### 2.1.4 Canonical Models

Canonical models are a key technical tool used to study lightweight description logics, and we will use them in the proofs of many of our results.

We recall the definition of the canonical model $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ (alternatively denoted $\mathcal{I}_{\mathcal{K}}$ ) of a satisfiable DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E L H} \operatorname{KB} \mathcal{K}=(\mathcal{T}, \mathcal{A})$. The domain $\Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ consists of sequences of the form $a R_{1} C_{1} \ldots R_{n} C_{n}(n \geq 0)$, where $a \in \operatorname{Ind}(\mathcal{A})$, each $C_{i}$ is a concept, and each $R_{i}$ is a (possibly inverse) role. The exact definition depends on which logic we consider:
(A) When $\mathcal{T}$ is a DL-Lite $\mathcal{R}_{\mathcal{R}}$ TBox, the domain $\Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ contains exactly those sequences $a R_{1} \exists R_{1}^{-} \ldots R_{n} \exists R_{n}^{-}$which satisfy:

- if $n \geq 1$, then $\mathcal{T}, \mathcal{A} \models \exists R_{1}(a)$
- for $1 \leq i<n, \mathcal{T} \models \exists R_{i}^{-} \sqsubseteq \exists R_{i+1}$.
(B) When $\mathcal{T}$ is an $\mathcal{E L H}$ TBox, ${ }^{2}$ the domain $\Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ contains exactly those sequences $a r_{1} A_{1} \ldots r_{n} A_{n}$ for which each $r_{i} \in \mathrm{~N}_{\mathrm{R}}$, and:
- if $n \geq 1$, then $\mathcal{T}, \mathcal{A} \models \exists r_{1} . A_{1}(a)$;
- for $1 \leq i<n, \mathcal{T} \models A_{i} \sqsubseteq \exists r_{i+1} . A_{i+1}$.

For elements $e \in \Delta^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}} \backslash \operatorname{Ind}(\mathcal{A})$, we will use the notation $\operatorname{Tail}(e)$ to denote the final concept in $e$. The set $\mathrm{TC}_{\mathcal{T}}$ of tail concepts for a TBox $\mathcal{T}$ is defined as follows: $\mathrm{TC}_{\mathcal{T}}=\left\{\exists r, \exists r^{-} \mid r \in\right.$ $\left.\mathrm{N}_{\mathrm{R}} \cap \operatorname{sig}(\mathcal{T})\right\}$ if $\mathcal{T}$ is a DL-Lite $\mathcal{R}_{\mathcal{R}}$ TBox, and $\mathrm{TC}_{\mathcal{T}}=\mathrm{BC}_{\mathcal{T}}=\left(\mathrm{N}_{\mathrm{C}} \cap \operatorname{sig}(\mathcal{T})\right) \cup\{\mathrm{T}\}$ if $\mathcal{T}$ is an $\mathcal{E} \mathcal{L H}$ TBox. Clearly, if $e \in \Delta^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}} \backslash \operatorname{Ind}(\mathcal{A})$, then $\operatorname{Tail}(e) \in \mathrm{TC}_{\mathcal{T}}$.

To complete the definition of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, we must fix the interpretation of the individual names, concept names, and role names. This is done as follows:

$$
\begin{aligned}
A^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}= & \{a \in \operatorname{Ind}(\mathcal{A}) \mid \mathcal{T}, \mathcal{A} \models A(a)\} \cup\left\{e \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} \backslash \operatorname{Ind}(\mathcal{A}) \mid \mathcal{T} \models \operatorname{Tail}(e) \sqsubseteq A\right\} \\
r^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}= & \{(a, b) \mid \mathcal{T}, \mathcal{A} \models r(a, b)\} \cup\left\{\left(e_{1}, e_{2}\right) \mid e_{2}=e_{1} S C \text { and } \mathcal{T} \models S \sqsubseteq r\right\} \cup \\
& \left\{\left(e_{2}, e_{1}\right) \mid e_{2}=e_{1} S C \text { and } \mathcal{T} \models S \sqsubseteq r^{-}\right\} \\
a^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}= & a \quad \text { for all } a \in \operatorname{Ind}(\mathcal{A})
\end{aligned}
$$

Note that $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ is composed of a core consisting of the ABox individuals and an anonymous part consisting of (possibly infinite) trees rooted at the ABox individuals. We will use $\left.\mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{e}$ to denote the submodel of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ obtained by restricting the domain to those elements containing $e$ as a prefix. Observe that, by construction, $\left.\mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{e}$ and $\left.\mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{e^{\prime}}$ are isomorphic whenever $\operatorname{Tail}(e)=\operatorname{Tail}\left(e^{\prime}\right)$.

It can be verified that $\mathcal{I}_{\mathcal{K}} \models \mathcal{K}$ for every satisfiable $\mathrm{KB} \mathcal{K}$. Moreover, it is well known that the canonical model $\mathcal{I}_{\mathcal{K}}$ can be homomorphically embedded into any model of $\mathcal{K}$.

Example 2.3. The canonical model of $\left(\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}\right)$ from Example 2.1 is depicted in Figure 4. To satisfy the existential restrictions in the inclusions (3a) - (5a) of $\mathcal{T}_{\text {mob }}$, the

[^1]

Figure 4: Canonical model $\mathcal{I}_{\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}}$ of the $\operatorname{KB}\left(\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}\right)$
canonical model contains the following five anonymous elements $e_{1}, \ldots e_{5}$, which form two tree-shaped structures rooted at the nodes park:huberPark and shopping:cityMall:

```
e}\mp@subsup{e}{1}{}=\mathrm{ park:huberPark hasFacility FoodServ
e}\mp@subsup{e}{2}{}=\mathrm{ park:huberPark hasFacility Playground
e}=\mathrm{ shopping:cityMall hasFacility Foodcourt
e}\mp@subsup{4}{4}{}=\mathrm{ shopping:cityMall hasFacility Supermarket
e
```

Example 2.4. To illustrate why canonical models can be infinite, we present in Figure 5 a simple DL-Lite $\mathcal{R}_{\mathcal{R}}$ knowledge base $(\mathcal{T}, \mathcal{A})$ and a depiction of its canonical model $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. As in the preceding example, we use names $e_{i}$ as abbreviations for the anonymous objects in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. Note that the tree rooted at $b_{2}$ is infinite, since every object that belongs to the concept $B$ has an $r$-child that also belongs to $B$.

### 2.2 Regular Languages

We assume the reader is familiar with regular languages, represented either by regular expressions or nondeterministic finite state automata (NFAs). Regular expressions $\mathcal{E}$ over alphabet $\Sigma$ are defined by the grammar

$$
\mathcal{E} \rightarrow \varepsilon|\sigma| \mathcal{E} \cdot \mathcal{E}|\mathcal{E} \cup \mathcal{E}| \mathcal{E}^{*}
$$

where $\sigma \in \Sigma$ and $\varepsilon$ denotes the empty word (i.e., the sequence of length 0 ). An NFA over an alphabet $\Sigma$ is a tuple $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$, where $S$ is a finite set of states, $\delta \subseteq S \times \Sigma \times S$ the

$$
\begin{aligned}
& \mathcal{T}=\{ B \sqsubseteq \exists r, \quad \exists r^{-} \sqsubseteq B, \\
&\left.B \sqsubseteq \exists r_{1}, \quad r_{1} \sqsubseteq r_{2}^{-} \quad\right\} \\
& \mathcal{A}=\left\{r(a, b), r_{2}(b, c), D(b)\right\}
\end{aligned}
$$



Figure 5: Example DL-Lite $\mathcal{R}_{\mathcal{R}}$ knowledge base $(\mathcal{T}, \mathcal{A})$ and its canonical model $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$
transition relation, $s_{0} \in S$ the initial state, and $F \subseteq S$ the set of final states. We use $L(\mathcal{E})$ (resp. $L(\alpha)$ ) to denote the language defined by the regular expression $\mathcal{E}$ (resp. the NFA $\alpha$ ).

We recall that NFAs are exponentially more succinct than regular expressions, in that there exists a polynomial translation of regular expressions into equivalent NFAs, while the translation from NFAs to regular expressions may incur an exponential blowup (Ehrenfeucht \& Zeiger, 1974). Thus, to ensure that our complexity results hold regardless of the chosen representation, we will prove our complexity lower bounds using the regular expression representation, and for our upper bounds, we will adopt the NFA representation.

### 2.3 Computational Complexity

We assume familiarity with standard complexity classes, such as NL (problems solvable in non-deterministic logarithmic space), P (problems solvable in polynomial time), NP (problems solvable in non-deterministic polynomial time), PSpace (problems solvable in polynomial space), and NPSPACE (problems solvable in non-deterministic polynomial space). We recall that by Savitch's theorem, we have NPspace $=$ PSPace. We shall also consider the oracle classes $\mathrm{NL}^{\mathrm{NL}}$ and $\mathrm{NL}^{\mathrm{P}}$ consisting of problems solvable in non-deterministic logarithmic space when given access to an NL (respectively, P ) oracle. It is well-known that $\mathrm{NL}^{\mathrm{NL}}=\mathrm{NL}$ and $\mathrm{NL}^{\mathrm{P}}=\mathrm{P}$. The circuit complexity class $\mathrm{AC}^{0}$ that was mentioned in Figure 1 comprises problems that can be computed by a family of unbounded fan-in circuits of constant depth and polynomial size. The preceding classes are ordered as follows:

$$
\mathrm{AC}^{0} \subsetneq \mathrm{NL} \subseteq \mathrm{P} \subseteq \mathrm{NP} \subseteq \mathrm{PSPACE}
$$

For precise definitions of these complexity classes, and other standard notions of computational complexity, we refer the reader to the recent textbook of Arora and Barak (2009) and references therein.

## 3. Path Queries

In this section, we introduce the different query languages considered in this paper and define the relevant computational problems.

$$
\begin{aligned}
& q_{1}(x, y)=\text { AccStop? } \cdot\left(\left(\operatorname{sbSub} \cdot \mathrm{sbSub}^{-}\right) \cup\left(\mathrm{sbLFT} \cdot \mathrm{sbLFT}^{-}\right)\right)^{*} \cdot \operatorname{AccStop} ?(x, y) \\
& q_{2}(x, y)=\exists z_{1}, z_{2} \cdot \operatorname{AccStop} ? \cdot\left(\left(\text { sbSub } \cdot \text { sbSub }^{-}\right) \cup\left(\text { sbLFT } \cdot \operatorname{sbLFT}^{-}\right)\right)^{*} \cdot \operatorname{AccStop} ?(x, y) \\
& \wedge \text { locln } \cdot\left(\text { locln }^{-}\right)^{*} \cdot \text { hasFacility }{ }^{*} \cdot \text { FoodServ? }\left(x, z_{1}\right) \\
& \wedge \text { locln } \cdot\left(\text { locln }{ }^{-}\right)^{*} \cdot \text { hasFacility }{ }^{*} \cdot \text { Playground } ?\left(y, z_{2}\right)
\end{aligned}
$$

Figure 6: Example queries

### 3.1 Syntax of Path Queries

A conjunctive (two-way) regular path query (C2RPQ) has the form $q(\vec{x})=\exists \vec{y} \cdot \varphi$ where $\vec{x}$ and $\vec{y}$ are disjoint tuples of variables, and $\varphi$ is a conjunction of atoms of the forms:
(i) $A(t)$, where $A \in \mathrm{~N}_{\mathrm{C}}$ and $t \in \mathrm{~N}_{\mathrm{I}} \cup \vec{x} \cup \vec{y}$
(ii) $\Lambda\left(t, t^{\prime}\right)$, where $\Lambda$ is an NFA or regular expression defining a regular language over $\mathrm{N}_{\mathrm{R}}^{ \pm} \cup\left\{A ? \mid A \in \mathrm{~N}_{\mathrm{C}}\right\}$, and $t, t^{\prime} \in \mathrm{N}_{\mathrm{I}} \cup \vec{x} \cup \vec{y}$

As usual, variables and individuals are called terms, the variables in $\vec{x}$ are called answer variables, and the variables in $\vec{y}$ are called quantified variables. We use terms $(q)$, vars $(q)$, $\operatorname{avars}(q)$, and quars $(q)$ to refer respectively to the sets of terms, variables, answer variables, and quantified variables appearing in query $q$. A query with no answer variables is called a Boolean query. Note that where convenient we will treat a query as its set of atoms.

Conjunctive (one-way) regular path queries (CRPQs) are obtained by disallowing symbols from $\mathrm{N}_{\mathrm{R}}^{ \pm} \backslash \mathrm{N}_{\mathrm{R}}$ in atoms of type (ii), and conjunctive queries (CQs) result from only allowing type (ii) atoms of the form $r\left(t, t^{\prime}\right)$ with $r \in \mathbf{N}_{\mathrm{R}}$. Two-way regular path queries (2RPQs) consist of a single atom of type (ii) such that $t$ and $t^{\prime}$ are both answer variables. Regular path queries ( $R P Q s$ ) are 2 RPQs that do not use any symbols from $N_{R}^{ \pm} \backslash N_{R}$. Finally, instance queries (IQs) take the form $A(x)$ with $A \in \mathbf{N}_{\mathrm{C}}$, or $r(x, y)$ with $r \in \mathrm{~N}_{\mathrm{R}}$.

Note that it will sometimes prove convenient to treat queries as sets of atoms, using the notation $\alpha \in q$ to indicate that $\alpha$ is an atom of $q$.

Example 3.1. Figure 6 shows two example queries. The $2 \mathrm{RPQ} q_{1}$ retrieves pairs $x, y$ of public transport stops that have an accessible connection, that is, both $x$ and $y$ are accessible stops, and there is a public transport route between them that uses only subway and low-floor tramway. The query $q_{2}$ retrieves pairs $x, y$ of public transport stops that have an accessible connection (as in $q_{1}$ ) and such that there is a place to eat at the location of $x$ and a playground at the location of $y$.

Note that by using the Kleene star $\left(^{*}\right)$, we can query for services such as restaurants and playgrounds available at some location without having to take care of the different ways in which places can be related. For example, a restaurant could be at the same location as a stop, or it could be that it is in a food court inside some shopping center that is itself at the same location as a stop. In both cases, $q_{2}$ correctly identifies it as a food service at that location. This is a very useful feature of RPQs and their extensions, particularly in cases where the data does not comply to a rigid schema.

### 3.2 Semantics of Path Queries

We now proceed to define the semantics of C2RPQs. Given an interpretation $\mathcal{I}$, a path from $e_{0}$ to $e_{n}$ in $\mathcal{I}$ is a sequence $e_{0} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ with $n \geq 0$ such that every $e_{i}$ is an element from $\Delta^{\mathcal{I}}$, every $u_{i}$ is a symbol from $\mathrm{N}_{\mathrm{R}}^{ \pm} \cup\left\{A\right.$ ? $\left.\mid A \in \mathrm{~N}_{\mathrm{C}}\right\}$, and for every $1 \leq i \leq n$ :

- If $u_{i}=A$ ?, then $e_{i-1}=e_{i} \in A^{\mathcal{I}}$;
- If $u_{i}=R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$, then $\left(e_{i-1}, e_{i}\right) \in R^{\mathcal{I}}$.

The label $\lambda(p)$ of path $p=e_{0} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ is the word $u_{1} u_{2} \ldots u_{n}$. Note that if $p=e_{0}$, then we define $\lambda(p)$ to be $\varepsilon$.

Then for every language $L$ over $\mathrm{N}_{\mathrm{R}}^{ \pm} \cup\left\{A\right.$ ? $\left.\mid A \in \mathrm{~N}_{\mathrm{C}}\right\}$, the semantics of $L$ w.r.t. interpretation $\mathcal{I}$ is defined as follows:

$$
L^{\mathcal{I}}=\left\{\left(e_{0}, e_{n}\right) \mid \text { there is some path } p \text { from } e_{0} \text { to } e_{n} \text { with } \lambda(p) \in L\right\}
$$

A match for a C2RPQ $q$ in an interpretation $\mathcal{I}$ is a mapping $\pi$ from the terms in $q$ to elements in $\Delta^{\mathcal{I}}$ such that

- $\pi(c)=c^{\mathcal{I}}$ for each $c \in \mathbf{N}_{\mathbf{I}}$,
- $\pi(t) \in A^{\mathcal{I}}$ for each atom $A(t)$ in $q$, and
- $\left(\pi(t), \pi\left(t^{\prime}\right)\right) \in L(\Lambda)^{\mathcal{I}}$ for each $\Lambda\left(t, t^{\prime}\right)$ in $q$.

Given a C2RPQ $q$ with answer variables $x_{1}, \ldots, x_{k}$, we say that a tuple of individuals $\left(a_{1}, \ldots, a_{k}\right)$ from $\operatorname{Ind}(\mathcal{A})$ is a certain answer to $q$ w.r.t. the $\operatorname{KB} \mathcal{K}=(\mathcal{T}, \mathcal{A})$ just in the case that in every model $\mathcal{I}$ of $\mathcal{K}$ there is a match $\pi$ for $q$ such that $\pi\left(v_{i}\right)=a_{i}^{\mathcal{I}}$ for every $1 \leq i \leq k$. We use $\operatorname{cert}(q, \mathcal{K})$ to denote the set of certain answers to $q$ w.r.t. the KB $\mathcal{K}$. Note that if $q$ is a Boolean query, then either $\operatorname{cert}(q, \mathcal{K})=\{()\}$ (where () denotes the empty tuple) or $\operatorname{cert}(q, \mathcal{K})=\emptyset$. In the former case, we say that $q$ is entailed from $\mathcal{K}$, and we write $\mathcal{K} \models q$.

We remark that the normal form for $\mathcal{E L H}$ TBoxes can also be assumed without loss of generality for query answering. Indeed, we can always assume that the fresh symbols in the TBox $\mathcal{T}^{\prime}$ in normal form do not occur in $q$, and it follows from Proposition 2.2 and the definition of certain answers that $\operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))=\operatorname{cert}\left(q,\left(\mathcal{T}^{\prime}, \mathcal{A}\right)\right)$ for every $\mathrm{C} 2 \mathrm{RPQ} q$ that does not use any symbols in $\operatorname{sig}\left(\mathcal{T}^{\prime}\right) \backslash \operatorname{sig}(\mathcal{T})$.

While the definition of certain answers involves all models of a KB, for the DLs considered in this paper, it is in fact sufficient to look for matches in the canonical model.
Lemma 3.2. For every $D L$-Lite $_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H} K B \mathcal{K}=(\mathcal{T}, \mathcal{A}), C 2 R P Q q(\vec{x})$ of arity $k$, and $k$-tuple $\vec{a}$ of individuals from $\mathcal{A}: \vec{a} \in \operatorname{cert}(q, \mathcal{K})$ if and only if there is a match $\pi$ for $q$ in $\mathcal{I}_{\mathcal{K}}$ such that $\pi(\vec{x})=\vec{a}$.

Proof sketch. It is well known that the canonical model $\mathcal{I}_{\mathcal{K}}$ can be homomorphically embedded into any model of $\mathcal{K}$ (Calvanese et al., 2007; Rosati, 2007; Krisnadhi \& Lutz, 2007; Krötzsch \& Rudolph, 2007). It follows that whenever query matches are preserved under homomorphisms, the existence of a match in $\mathcal{I}_{\mathcal{K}}$ implies the existence of a match in every model. This has been often observed for CQs, and it applies equally well to C2RPQs (Calvanese et al., 2014; Ortiz et al., 2011). Since the converse is trivially true, certain answers coincide with the answers over the canonical model.


Figure 7: A match witnessing $a \in \operatorname{cert}(q, \mathcal{K})$ for $q$ and $\mathcal{K}$ from Example 3.4

Example 3.3. In our urban mobility example, the stop at the cathedral square is not known to be accessible (i.e., AccStop(stop:cathSq) is not entailed by ( $\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}$ )), hence stop:cathSq cannot participate in any match for $q_{1}$. The stop at the theater is accessible, but it is only connected to other stops via high-floor tramway. Thus stop:Volkstheater only participates in one mapping for $q_{1}$ in $\mathcal{I}_{\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}}$, namely $\pi(x)=\pi(y)=$ stop:Volkstheater. Indeed, the path

## stop:Volkstheater AccStop? stop:Volkstheater AccStop? stop:Volkstheater

witnesses that (stop:Volkstheater, stop:Volkstheater) $\in L_{1}^{\mathcal{T}}$ for the language $L_{1}$ specified in $q_{1}$, and there is no longer path starting or ending at stop:Volkstheater whose label belongs to $L_{1}$. The stops stop:opera, stop:cityPark, and stop:trainStation are all accessible and mutually connected via accessible public transport (i.e., subway and low-floor tramway lines). Hence, we can find a path between any pair of them whose label is in $L_{1}$, and all such pairs are certain answers to $q_{1}$. Thus $\operatorname{cert}\left(q_{1},\left(\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}\right)\right)$ contains (stop:Volkstheater, stop:Volkstheater), and all pairs of stops involving stop:opera, stop:cityPark, and stop:trainStation.

The certain answers to $q_{2}$ are precisely those pairs $\left(s_{1}, s_{2}\right)$ of stops that are an answer to $q_{1}$ such that there is some food service at the location of $s_{1}$ and a playground at the location of $s_{2}$. Since in $\mathcal{I}_{\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}}$, we find both some food service and a playground at the location of stop:Volkstheater, we have (stop:Volkstheater, stop:Volkstheater) $\in \operatorname{cert}\left(q_{2},\left(\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}\right)\right)$. We also find food services at the locations of stop:opera and stop:trainStation, and a playground at the location of stop:cityPark, hence $\operatorname{cert}\left(q_{2},\left(\mathcal{T}_{\text {mob }}, \mathcal{A}_{\text {mob }}\right)\right)$ also contains the pairs (stop:opera, stop:cityPark) and (stop:trainStation, stop:cityPark).

Example 3.4. We also give an example of a query over the $\operatorname{KB}(\mathcal{T}, \mathcal{A})$ in Figure 5:

$$
q(x)=\exists y, z \cdot r^{*} \cdot r_{1} \cdot r_{2} \cdot r^{-}(x, y) \wedge r^{-} \cdot r^{-}(y, z) \wedge D(z)
$$

We have $\operatorname{cert}(q, \mathcal{K})=\{a, b\}$. To see why $a$ is a certain answer, consider the mapping $\pi(x)=a, \pi(y)=e_{11}$, and $\pi(z)=b$. The path arbre ${ }_{1} r_{11} r e_{111} r_{1} e_{1112} r_{2} e_{111} r^{-} e_{11}$ witnesses that $\left(a, e_{11}\right) \in L\left(r^{*} \cdot r_{1} \cdot r_{2} \cdot r^{-}\right)^{\mathcal{I} \mathcal{T}, \mathcal{A}}$, while the path $e_{11} r^{-} e_{1} r^{-} b$ witnesses that $\left(e_{11}, b\right) \in$
$L\left(r^{-} \cdot r^{-}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. Since $b \in D^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$, the mapping $\pi$ is a match for $q$. This match is depicted in Figure 7. To see that $b$ is also a certain answer, consider $\pi^{\prime}(x)=\pi^{\prime}(z)=b$ and $\pi^{\prime}(y)=$ $e_{11}$, and observe that $\pi^{\prime}$ is also a match because bre $e_{1} r e_{11} r e_{111} r_{1} e_{1112} r_{2} e_{111} r^{-} e_{11}$ witnesses $\left(b, e_{11}\right) \in L\left(r^{*} \cdot r_{1} \cdot r_{2} \cdot r^{-}\right)^{\mathcal{I}_{\tau, \mathcal{A}}}$.

Note that in both matches, $y$ is mapped to an element in the anonymous part, and there is no match mapping $y$ to an individual. This illustrates that anonymous elements may play a decisive role in query answering, and every complete query answering algorithm must consider possible matches into the possibly infinite anonymous part of canonical models.

### 3.3 Computational Problems

In this paper, we will be interested in the problem of computing the certain answers to C2RPQs, and more precisely, the associated decision problem of determining whether a given tuple is a certain answer to a query. In what follows, for a query language $Q \in$ $\{\mathrm{IQ}, \mathrm{CQ}, \mathrm{RPQ}, \mathrm{CRPQ}, 2 \mathrm{RPQ}, \mathrm{C} 2 \mathrm{RPQ}\}$, we will use the term $Q$ answering to refer to the problem of deciding given a KB $\mathcal{K}$, tuple $\vec{a}$, and query $q$ from $Q$, whether $\vec{a} \in \operatorname{cert}(q, \mathcal{K})$.

There are different ways of measuring the complexity of query answering, depending on which of the three parameters of the problem $(\mathcal{T}, \mathcal{A}$, and $q)$ are considered as inputs and which are considered fixed. In this work, we consider the two most commonly used measures: combined complexity and data complexity. Combined complexity treats all three parameters as inputs, so the complexity is measured with respect to the total size $|\mathcal{T}|+|\mathcal{A}|+|q|$ (we use $|\cdot|$ to denote the size of an object, e.g. the length of its string representation according to some suitable encoding). Data complexity takes $\mathcal{A}$ as input and assumes $\mathcal{T}$ and $q$ to be fixed, so the complexity is measured only with respect to $|\mathcal{A}|$, with $|\mathcal{T}|$ and $|q|$ treated as constants.

## 4. Lower Bounds

In this section, we establish the required complexity lower bounds. We begin with some lower bounds for RPQs that can be straightforwardly obtained from existing results.

Proposition 4.1. $R P Q$ answering is

## 1. NL-hard in data complexity for DL-Lite ${ }_{\text {RDFS }}$;

2. P-hard in data complexity for $\mathcal{E L}$;

Proof. Statement (1) follows from the analogous result for graph databases (Consens \& Mendelzon, 1990). It can be shown by a simple reduction from the NL-complete directed reachability problem: vertex $b$ is reachable from vertex $a$ in a directed graph $G$ if and only if $(a, b)$ is a certain answer to the RPQ $r^{*}(x, y)$ w.r.t. the $\mathrm{KB}\left(\emptyset, \mathcal{A}_{G}\right)$, where

$$
\mathcal{A}_{G}=\left\{r\left(v_{1}, v_{2}\right) \mid \text { there is a directed edge from } v_{1} \text { to } v_{2} \text { in } G\right\} .
$$

Statement (2) is a direct consequence of the P-hardness in data complexity of instance checking in $\mathcal{E L}$ (Calvanese et al., 2006), since the instance query $A(x)$ can be computed using the RPQ $A ?(x, y)$.

In the case of DL-Lite, we establish a P lower bound for 2RPQs, which contrasts with the NL-completeness of instance checking. We remark that a similar result was given in

Bienvenu, Ortiz, \& Šimkus
the conference version of this paper (Bienvenu et al., 2013), but the reduction required an NFA representation of the regular language in the query. The complexity of 2RPQs using the (less succinct) regular expression representation was left open and is resolved by the following theorem.

Theorem 4.2. 2RPQ answering is P-hard in combined complexity for DL-Lite.
Proof. We give a reduction from the P-complete entailment problem for propositional definite Horn theories. Without loss of generality, we suppose that we are given a propositional Horn theory $\Psi$ over variables $v_{1}, \ldots, v_{n}$ that consists of

- a set of rules $\rho_{i}=v_{i_{1}} \wedge v_{i_{2}} \rightarrow v_{i_{3}}(1 \leq i \leq m)$
- a single "initialization" fact: $v_{1}$, with $v_{1} \neq v_{n}$

Indeed, any arbitrary propositional definite Horn theory $\Psi^{\prime}$ can be transformed into a theory of the preceding form as follows: take a fresh variable $v_{1}$ not appearing in $\Psi^{\prime}$, add $v_{1}$ to the body of every rule in $\Psi^{\prime}$, add the fact $v_{1}$, and finally perform standard syntactic manipulations (possibly introducing additional fresh variables) to ensure that all rules other than the initialization fact $v_{1}$ contain exactly two body variables.

In what follows, we show how to construct, given a propositional Horn theory $\Psi$ of the form above, a DL-Lite $\mathrm{KB} \mathcal{K}=(\mathcal{T}, \mathcal{A})$ and Boolean 2RPQ $q$ such that $\mathcal{K} \models q$ if and only if $\Psi \models v_{n}$. We first provide an informal description of the reduction, and then present it formally. As is well known, $\Psi \models v_{n}$ just in the case that there exists a proof tree for $v_{n}$ from $\Psi$, which can be defined as a binary tree $T$ where each node is labeled with a variable from $v_{1}, \ldots, v_{n}$ such that the following conditions are satisfied: (i) the root is labeled with $v_{n}$, (ii) the leaves are labeled with $v_{1}$, and (iii) for any inner node $d$, if $d$ is labeled with $v_{k}$, then there is a rule $\rho_{i} \in \Psi$ such that $v_{k}=v_{i_{3}}$ and the two children of $d$ are labeled with $v_{i_{1}}$ and $v_{i_{2}}$, respectively. The existence of a node-labeled proof tree $T$ as just described is equivalent to the existence of an edge-labeled proof tree $T^{\prime}$, defined as a sibling-ordered binary tree whose edges are labeled with rules from $\Psi$ as follows. First, the two edges outgoing from the root are labeled with a rule $\rho_{i}$ with $v_{i_{3}}=v_{n}$. For a non-root node $d_{\ell}$ which is the $\ell$-th child of its parent $d$, where $\ell \in\{1,2\}$, if the edge $\left(d, d_{\ell}\right)$ is labeled with $\rho_{i}$, then either $v_{i \ell}=v_{1}$, or $d_{\ell}$ has two outgoing edges $E_{1}$ and $E_{2}$ that are both labeled with a rule $\rho_{j}$ such that $v_{j_{3}}=v_{i_{\ell}}$.

We will next show how to construct $\mathcal{K}$ and $q$ such that $\mathcal{K} \models q$ if and only if there exists a edge-labeled proof tree $T^{\prime}$ for $v_{n}$ from $\Psi$. Roughly speaking, we use $\mathcal{K}$ to generate in the anonymous part a tree that contains all possible edge-labeled proof trees. Since such proof trees are based upon sibling-ordered trees, we need to distinguish between the first and second children of a node, and so we use two roles $r_{i, 1}$ and $r_{i, 2}$ for each rule $\rho_{i}$. An edge-labeled proof tree $T^{\prime}$ will thus map into a subtree of $\mathcal{I}_{\mathcal{K}}$ of the same structure, but with label $\rho_{i}$ replaced by either $r_{i, 1}$ or $r_{i, 2}$, depending on whether the edge leads to the first or second child of the parent node. We then use a $2 \mathrm{RPQ} q$ to determine whether $\mathcal{I}_{\mathcal{K}}$ actually contains such a subtree. The intuition is that every path that witnesses the satisfaction of $q$ corresponds to the complete depth-first traversal of (the representation of) a valid edgelabeled proof tree that starts and ends at the root, and in which the left subtree of a node is always visited before the right one.

$$
\mathcal{E}=\left(\bigcup_{1 \leq i \leq m} r_{i, 1}\right) \cdot\left(\bigcup_{1 \leq i \leq m} r_{i, 1} \cup \bigcup_{k \in F_{1}}\left(r_{k, 1}^{-} \cdot r_{k, 2}\right) \cup \bigcup_{k \in F_{2}}\left(r_{k, 2}^{-} \cdot\left(\bigcup_{1 \leq i \leq m} r_{i, 2}^{-}\right)^{*} \cdot\left(\bigcup_{1 \leq i \leq m}\left(r_{i, 1}^{-} \cdot r_{i, 2}\right) \cup \varepsilon\right)\right)\right)^{*}
$$

Figure 8: Regular expression used in the proof of Theorem 4.2.

The ABox $\mathcal{A}$ consists of a single assertion $A(a)$, and the TBox $\mathcal{T}$ contains the following concept inclusions:

- $A \sqsubseteq \exists r_{i, \ell}$, where $\ell \in\{1,2\}$ and $1 \leq i \leq m$ with $v_{i_{3}}=v_{n}$
- $\exists r_{i, \ell}^{-} \sqsubseteq \exists r_{k, j}$, where $\ell, j \in\{1,2\}$ and $1 \leq i, k \leq m$ such that $v_{i_{\ell}}=v_{k_{3}}$

For $\ell \in\{1,2\}$, we define the set

$$
F_{\ell}=\left\{k \mid 1 \leq k \leq m, v_{k_{\ell}}=v_{1}\right\} .
$$

Intuitively, $F_{\ell}$ contains the index of each rule where we can 'turn back' since its $\ell$-th variable is the initial variable $v_{1}$, (and so the corresponding child node in the proof tree would be a leaf). We use $F_{1}$ and $F_{2}$ to define the regular expression $\mathcal{E}$ in Figure 8, which we then use to define the following 2 RPQ :

$$
q=\mathcal{E}(a, a)
$$

We now prove the correctness of the reduction.
$(\Rightarrow)$ Suppose that $\mathcal{K} \models q$. Then by Lemma 3.2, there is a match for $q$ in the canonical model $\mathcal{I}_{\mathcal{K}}$ of $\mathcal{K}$. This means that $(a, a)=\left(a^{\mathcal{I}_{\mathcal{K}}}, a^{\mathcal{I}_{\mathcal{K}}}\right) \in L(\mathcal{E})^{\mathcal{I}_{\mathcal{K}}}$, and so there exists a path $e_{0} \sigma_{1} \ldots \sigma_{p} e_{p}$ in $\mathcal{I}_{\mathcal{K}}$ whose label is in $L(\mathcal{E})$ such that $e_{0}=e_{p}=a$.
Claim 1. For every $2 \leq j \leq p$ :

1. If $\sigma_{j}=r_{i, 1}$, then there exists $j^{\prime}>j$ with $\sigma_{j^{\prime}}=r_{i, 2}^{-}$.
2. If $\sigma_{j}=r_{i, 2}^{-}$, then there exists $j^{\prime}<j$ such that $\sigma_{j^{\prime}}=r_{i, 1}^{-}$.
3. If $\sigma_{j}=r_{i, 1}^{-}$and $i \notin F_{1}$, then $\sigma_{j-1}=r_{i^{\prime}, 2}^{-}$where $v_{i_{3}^{\prime}}=v_{i_{1}}$.
4. If $\sigma_{j}=r_{i, 2}^{-}$and $i \notin F_{2}$, then $\sigma_{j-1}=r_{i^{\prime}, 2}^{-}$where $v_{i_{3}^{\prime}}=v_{i_{2}}$.

Proof of claim. For Point 1, we remark that in $\mathcal{I}_{\mathcal{K}}$ all roles are directed away from the ABox; formally, for every role name $s$, if $\left(g, g^{\prime}\right) \in s^{\mathcal{I}_{\mathcal{K}}}$, then $g^{\prime}=g s \exists s^{-}$. It follows that if $\sigma_{j}=r_{i, 1}$, then $e_{j}=e_{j-1} r_{i, 1} \exists r_{i, 1}^{-}$. Since the sequence of elements $e_{j}, \ldots, e_{p}$ defines a path in $\mathcal{I}_{\mathcal{K}}$ from $e_{j}$ to $e_{p}=a$, by continuity, there must be some $j_{0}>j$ such that $e_{j_{0}-1}=e_{j}$ and $e_{j_{0}}=e_{j-1}$, in which case we must have $\sigma_{j_{0}}=r_{i, 1}^{-}$. Next note that the structure of $\mathcal{E}$ ensures that every occurrence of $r_{i, 1}^{-}$is immediately followed by $r_{i, 2}$. Then repeating the same argument, substituting $r_{i, 2}$ for $r_{i, 1}$, we can find some $j^{\prime}>j_{0}$ such that $\sigma_{j^{\prime}}=r_{i, 2}^{-}$.

For Point 2, we again use the fact that roles in $\mathcal{I}_{\mathcal{K}}$ are directed away from the ABox. Thus, if $\sigma_{j}=r_{i, 2}^{-}$, then $e_{j-1}=e_{j} r_{i, 2} \exists r_{i, 2}^{-}$. The sequence $e_{0}, \ldots, e_{j-1}$ of elements forms a
path in $\mathcal{I}_{\mathcal{K}}$ from $e_{0}=a$ to $e_{j} r_{i, 2} \exists r_{i, 2}^{-}$. Thus, by continuity, there must exist some $j^{\prime}<j$ such that $e_{j^{\prime}}=e_{j}, e_{j^{\prime}+1}=e_{j} r_{i, 2} \exists r_{i, 2}^{-}$, and $\sigma_{j^{\prime}+1}=r_{i, 2}$. By examining the structure of $\mathcal{E}$, we can see that any occurrence of $r_{i, 2}$ must be immediately preceded by $r_{i, 1}^{-}$, and so $\sigma_{j^{\prime}}=r_{i, 1}^{-}$.

To show Point 3, suppose that $\sigma_{j}=r_{i, 1}^{-}$and $i \notin F_{1}$. The structure of $\mathcal{E}$ ensures that the preceding symbol $\sigma_{j-1}$ takes the form $r_{i^{\prime}, 2}^{-}$. We thus have $\left(e_{j}, e_{j-1}\right) \in r_{i, 1}^{\mathcal{I}_{\mathcal{K}}}$ and $\left(e_{j-1}, e_{j-2}\right) \in r_{i^{\prime}, 2}^{\mathcal{I}_{\mathcal{K}}}$. Once again using the fact that elements in $\mathcal{I}_{\mathcal{K}}$ do not contain inverse role names, we obtain $e_{j-2}=e_{j} r_{i, 1} \exists r_{i, 1}^{-} r_{i^{\prime}, 2} \exists r_{i^{\prime}, 2}^{-}$. It follows that $\mathcal{T} \models \exists r_{i, 1}^{-} \sqsubseteq \exists r_{i^{\prime}, 2}$, which can only be the case if $v_{i_{3}^{\prime}}=v_{i_{1}}$.

Finally, for Point 4, suppose that $\sigma_{j}=r_{i, 2}^{-}$and $i \notin F_{2}$. Examining the structure of $\mathcal{E}$, it is clear that the preceding symbol $\sigma_{j-1}$ must be of the form $r_{i^{\prime}, 2}^{-}$, and so we have $\left(e_{j}, e_{j-1}\right) \in r_{i, 2}^{\mathcal{I}_{\mathcal{K}}}$ and $\left(e_{j-1}, e_{j-2}\right) \in r_{i^{\prime}, 2}^{\mathcal{I}_{\mathcal{K}}}$. This means that $\mathcal{T} \models \exists r_{i, 2}^{-} \sqsubseteq \exists r_{i^{\prime}, 2}$, and hence $v_{i_{3}^{\prime}}=v_{i_{2}}$. (end proof of claim)

It is easy to see that the first symbol $\sigma_{1}$ must have the form $r_{i, 1}$, and so we have $\left(e_{0}, e_{1}\right) \in r_{i, 1}^{\mathcal{I}_{\mathcal{K}}}$. Since $e_{0}=a$, we have $\left(a, a r_{i, 1} \exists r_{i, 1}^{-}\right) \in r_{i, 1}^{\mathcal{I}_{\mathcal{K}}}$. This implies that $\mathcal{T} \models A \sqsubseteq \exists r_{i, 1}$, and hence that $v_{i_{3}}=v_{n}$. Applying Points 1 and 2 of the preceding claim, we can find $j, k$ such that $\sigma_{j}=r_{i, 1}^{-}$and $\sigma_{k}=r_{i, 2}^{-}$. To complete the proof of this direction, we establish the following claim.

Claim 2. For every $1 \leq j \leq p$ and $\ell \in\{1,2\}$, if $\sigma_{j}=r_{i, \ell}^{-}$, then $\Psi \models v_{i_{\ell}}$.
Proof of claim. We proceed by induction on $j$. For the base case, suppose that $\sigma_{j}=r_{i, \ell}^{-}$, and there is no $j^{\prime}<j$ with $\sigma_{j^{\prime}}=r \quad r_{i^{\prime}, \ell^{\prime}}^{-}$. It follows from Claim 1 that $i \in F_{1}$ and $\ell=1$. We thus have $v_{i_{\ell}}=v_{1}$, so $\Psi \models v_{i_{\ell}}$ trivially holds.

For the induction step, suppose that claim holds for all $j<k$, and consider $\sigma_{k}=r_{i, \ell}^{-}$. We consider three cases:

- Case 1: $\sigma_{k}=r_{i, \ell}^{-}$and $i \in F_{\ell}$

Since $i \in F_{\ell}$, we have $v_{i_{\ell}}=v_{1}$, so $\Psi \models v_{i_{\ell}}$ follows immediately.

- Case 2: $\sigma_{k}=r_{i, 1}^{-}$and $i \notin F_{1}$

By Point 3 of Claim 1, $\sigma_{k-1}=r_{i^{\prime}, 2}^{-}$where $v_{i_{3}^{\prime}}=v_{i_{1}}$. Applying the induction hypothesis to $\sigma_{k-1}$, we obtain $\Psi \models v_{i_{2}^{\prime}}$. From Point 2 of Claim 2, there exists some $j<k$ such that $\sigma_{j}=r_{i^{\prime}, 1}^{-}$. A second application of the induction hypothesis yields $\Psi \models v_{i_{1}^{\prime}}$. The rule $v_{i_{1}^{\prime}} \wedge v_{i_{2}^{\prime}} \rightarrow v_{i_{3}^{\prime}}$ belongs to $\Psi$, so we must also have $\Psi \models v_{i_{3}^{\prime}}$. Then since $v_{i_{3}^{\prime}}=v_{i_{1}}$, we obtain $\Psi \models v_{i_{1}}$.

- Case 3: $\sigma_{k}=r_{i, 2}^{-}$and $i \notin F_{2}$

We can use almost the same argument as for Case 2, except that we must use Point 4 of Claim 1, rather than Point 3. (end proof of claim)

As $\Psi$ contains the rule $v_{i_{1}} \wedge v_{i_{2}} \rightarrow v_{i_{3}}$, it follows from Claim 2 that $\Psi \models v_{n}$.
$(\Leftarrow)$ If $\Psi \models v_{n}$, then there must be exist an edge-labeled proof tree $T^{\prime}$ for $v_{n}$ from $\Psi$ as described at the beginning of the proof. We define a mapping $f$ from the nodes of $T^{\prime}$ to domain elements in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ as follows:

- $f(d)=a$ for $d$ the root of $T^{\prime}$;
- for every non-root node $d$, if $d$ is the first (resp., second) child of its parent $d_{p}$ and $\left(d_{p}, d\right)$ is labeled $\rho_{i}$, then $f(d)=f\left(d_{p}\right) r_{i, 1} \exists r_{i, 1}^{-}$(resp. $\left.f(d)=f\left(d_{p}\right) r_{i, 2} \exists r_{i, 2}^{-}\right)$.

We show the following claim:
Claim 3: For every non-leaf node $d$ in $T^{\prime},(f(d), f(d)) \in L(\mathcal{E})^{\mathcal{L}_{\mathcal{T}}, \mathcal{A}}$.
Proof of claim. For every non-leaf node $d$, we consider a depth-first traversal of the subtree of $T^{\prime}$ rooted at $d$ that always visits the left subtree of a node before visiting the right one, and then returns to the root. Let $d_{1}, d_{2}, \ldots, d_{n}$ be the sequence of nodes visited in this traversal, with $d_{1}=d_{n}=d$, and for every $2 \leq i \leq n$, let $\sigma_{i-1, i}$ be as follows:

- $\sigma_{i-1, i}=r_{j, 1}$ if $d_{i}$ is the first child of $d_{i-1}$ and $\left(d_{i-1}, d_{i}\right)$ is labeled with $\rho_{j}$;
- $\sigma_{i-1, i}=r_{j, 2}$ if $d_{i}$ is the second child of $d_{i-1}$ and $\left(d_{i-1}, d_{i}\right)$ is labeled with $\rho_{j}$;
- $\sigma_{i-1, i}=r_{j, 1}^{-}$if $d_{i-1}$ is the first child of $d_{i}$ and $\left(d_{i}, d_{i-1}\right)$ is labeled with $\rho_{j}$;
- $\sigma_{i-1, i}=r_{j, 2}^{-}$if $d_{i-1}$ is the second child of $d_{i}$ and $\left(d_{i}, d_{i-1}\right)$ is labeled with $\rho_{j}$.

We now define a path $p_{d}$ as follows:

$$
p_{d}=f\left(d_{1}\right) \sigma_{1,2} f\left(d_{2}\right) \sigma_{2,3} \ldots \sigma_{n-1, n} f\left(d_{n}\right)
$$

To show the claim, it suffices to show the following for every non-leaf node $d$ :
$(*) \quad p_{d}$ is a path in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\lambda\left(p_{d}\right) \in L(\mathcal{E})$.
This can be shown by induction on the minimal distance of $d$ to a leaf in $T^{\prime}$. An important observation is that, for $\ell \in\{1,2\}$, if the $\ell$ th child $d^{\ell}$ of a node $d$ of $T^{\prime}$ is a leaf, and the edge from $d$ to $d^{\ell}$ is labeled $\rho_{i}$, then by definition $i \in F_{\ell}$.

In order to be able to more easily argue that some words belong to $L(\mathcal{E})$, we give names to the relevant subexpressions of $\mathcal{E}$ :

$$
\begin{aligned}
\mathcal{E}_{1} & =\bigcup_{1 \leq i \leq m} r_{i, 1} \\
\mathcal{E}_{2} & =\bigcup_{k \in F_{1}}\left(r_{k, 1}^{-} \cdot r_{k, 2}\right) \\
\mathcal{E}_{3} & =\left(\bigcup_{1 \leq i \leq m} r_{i, 2}^{-}\right)^{*} \\
\mathcal{E}_{4} & =\left(\bigcup_{1 \leq i \leq m}\left(r_{i, 1}^{-} \cdot r_{i, 2}\right) \cup \varepsilon\right) \\
\mathcal{E}_{5} & =\bigcup_{k \in F_{2}}\left(r_{k, 2}^{-} \cdot\left(\bigcup_{1 \leq i \leq m} r_{i, 2}^{-}\right)^{*} \cdot\left(\bigcup_{1 \leq i \leq m}\left(r_{i, 1}^{-} \cdot r_{i, 2}\right) \cup \varepsilon\right)\right. \\
\mathcal{E}_{6} & =\left(\bigcup_{1 \leq i \leq m} r_{i, 1} \cup \bigcup_{k \in F_{1}}\left(r_{k, 1}^{-} \cdot r_{k, 2}\right) \cup \bigcup_{k \in F_{2}}\left(r_{k, 2}^{-} \cdot\left(\bigcup_{1 \leq i \leq m} r_{i, 2}^{-}\right)^{*} \cdot\left(\bigcup_{1 \leq i \leq m}\left(r_{i, 1}^{-} \cdot r_{i, 2}\right) \cup \varepsilon\right)\right)\right)^{*}
\end{aligned}
$$

Note that we have $\mathcal{E}_{5}=\bigcup_{k \in F_{2}}\left(r_{k, 2}^{-} \cdot \mathcal{E}_{3} \cdot \mathcal{E}_{4}\right), \mathcal{E}_{6}=\left(\mathcal{E}_{1} \cup \mathcal{E}_{2} \cup \mathcal{E}_{5}\right)^{*}$, and $\mathcal{E}=\mathcal{E}_{1} \cdot \mathcal{E}_{6}$.
Now we are ready to prove (*). For the base case, when both children of $d$ are leaves, let $\rho_{i}$ be the label of the edges from $d$ to its children. By construction of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}, f(d)$ has an $r_{i, 1}$-child and an $r_{i, 2}$-child, so $p_{d}$ is a path in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\lambda\left(p_{d}\right)=r_{i, 1} \cdot r_{i, 1}^{-} \cdot r_{i, 2} \cdot r_{i, 2}^{-}$. Both children of $d$ are leaves, so $i \in F_{1} \cap F_{2}$. We thus have $r_{i, 1} \in L\left(\mathcal{E}_{1}\right), r_{i, 1}^{-} \cdot r_{i, 2} \in L\left(\mathcal{E}_{2}\right)$, and $r_{i, 2}^{-} \in L\left(\mathcal{E}_{5}\right)=L\left(\bigcup_{k \in F_{2}}\left(r_{k, 2}^{-} \cdot \mathcal{E}_{3} \cdot \mathcal{E}_{4}\right)\right)$ (for the latter, observe that $\varepsilon \in L\left(\mathcal{E}_{3}\right)$ and $\left.\varepsilon \in L\left(\mathcal{E}_{4}\right)\right)$. Using $\mathcal{E}_{6}=\left(\mathcal{E}_{1} \cup \mathcal{E}_{2} \cup \mathcal{E}_{5}\right)^{*}$, we get $r_{i, 1}^{-} \cdot r_{i, 2} \cdot r_{i, 2}^{-} \in L\left(\mathcal{E}_{6}\right)$, and using $\mathcal{E}=\mathcal{E}_{1} \cdot \mathcal{E}_{6}$, we obtain $\lambda\left(p_{d}\right)=r_{i, 1} \cdot r_{i, 1}^{-} \cdot r_{i, 2} \cdot r_{i, 2}^{-} \in L(\mathcal{E})$.

For the induction step, let $d^{L}$ and $d^{R}$ be the left and right children of $d$ respectively, and let $\rho_{i}$ be the rule that is used to label the edges from $d$ to $d^{L}$ and from $d$ to $d^{R}$. By construction of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ and $f$, we know that $f\left(d^{L}\right)$ is an $r_{i, 1}$-child of $f(d)$ and that $f\left(d^{R}\right)$ is an $r_{i, 2}$-child of $f(d)$. We distinguish three cases:

- If neither of the children $d^{L}$ and $d^{R}$ of $d$ is a leaf, then we know from the induction hypothesis that $p_{d^{L}}$ and $p_{d^{R}}$ are paths in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ which start and end at $f\left(d^{L}\right)$ and $f\left(d^{R}\right)$ respectively and are such that $\left\{\lambda\left(p_{d^{L}}\right), \lambda\left(p_{d^{R}}\right)\right\} \subseteq L(\mathcal{E})$. It follows that $p_{d}=$ $f(d) r_{i, 1} p_{d^{L}} r_{i, 1}^{-} f(d) r_{i, 2} p_{d^{R}} r_{i, 2}^{-} f(d)$ is a path in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. We let $\rho_{k^{L}}$ (resp. $\rho_{k^{R}}$ ) be the label linking $d^{L}$ (resp. $d^{R}$ ) to its two children. Note that by construction, $\lambda\left(p_{d^{L}}\right)$ (resp. $\left.\lambda\left(p_{d^{L}}\right)\right)$ ends with $r_{k^{L}, 1}^{-}\left(\right.$resp. $\left.r_{k^{L}, 2}^{-}\right)$. It follows that $\lambda\left(p_{d^{L}}\right) \in L\left(\mathcal{E}_{1} \cdot \mathcal{E}_{6} \cdot \mathcal{E}_{5}\right)$ and that for the subexpression $\mathcal{E}_{4}=\left(\bigcup_{1 \leq i \leq m}\left(r_{i, 1}^{-} \cdot r_{i, 2}\right) \cup \varepsilon\right)$ of the final $\mathcal{E}_{5}$, we must select $\varepsilon$. By choosing to instantiate $\mathcal{E}_{4}$ with $r_{i, 1}^{-} r_{i, 2}$ instead, we can show that $\lambda\left(p_{d^{L}}\right) \cdot r_{i, 1}^{-} \cdot r_{i, 2} \in$ $L\left(\mathcal{E}_{6}\right)$. A similar argument for $\lambda\left(p_{d^{R}}\right)$ can be used to show that $\lambda\left(p_{d^{R}}\right) \cdot r_{i, 2}^{-} \in L\left(\mathcal{E}_{6}\right)$. We therefore obtain $\lambda\left(p_{d}\right)=r_{i, 1} \cdot \lambda\left(p_{d_{1}}\right) \cdot r_{i, 1}^{-} \cdot r_{i, 2} \cdot \lambda\left(p_{d_{1}}\right) \cdot r_{i, 2}^{-} \in L\left(\mathcal{E}_{1} \cdot \mathcal{E}_{6} \cdot \mathcal{E}_{6}\right) \subseteq L(\mathcal{E})$.
- If $d^{L}$ is not a leaf but $d^{R}$ is, then we can apply the induction hypothesis to infer that $p_{d^{L}}$ is a path in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ that starts and ends at $f\left(d^{L}\right)$ and is such that $\lambda\left(p_{d^{L}}\right) \in L(\mathcal{E})$. By using the same reasoning as in the previous case, we obtain that $\lambda\left(p_{d^{L}}\right) \cdot r_{i, 1}^{-} \cdot r_{i, 2} \in$ $L\left(\mathcal{E}_{6}\right)$. As $d^{R}$ is a leaf, it follows that $i \in F_{2}$, and hence $r_{i, 2}^{-} \in L\left(\mathcal{E}_{5}\right)$ (here again we choose $\varepsilon$ to satisfy the subexpressions $\mathcal{E}_{3}$ and $\mathcal{E}_{4}$ of $\mathcal{E}_{5}$ ). Putting this together, we find that $\lambda\left(p_{d}\right)=r_{i, 1} \lambda\left(p_{d_{1}}\right) r_{i, 1}^{-} r_{i, 2} r_{i, 2}^{-} \in L\left(\mathcal{E}_{1} \cdot \mathcal{E}_{6} \cdot \mathcal{E}_{5}\right) \subseteq L(\mathcal{E})$.
- If $d^{L}$ is a leaf but $d^{R}$ is not, the argument is analogous to the previous case.
(end proof of claim)
Since $f(d)=a$ for the root $d$ of $T^{\prime}$, it follows from the preceding claim that $(a, a) \in$ $L(\mathcal{E})^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, hence $\mathcal{K} \models q$.

In DL-Lite ${ }_{\mathcal{R}}$, we can strengthen Theorem 4.2 by using role inclusions to eliminate inverse roles in the query.

Corollary 4.3. $R P Q$ answering is P -hard in combined complexity in $D L-$ Lite $_{\mathcal{R}}$.
Proof. Let $q$ be a 2 RPQ and $(\mathcal{T}, \mathcal{A})$ be a DL-Lite $\mathcal{R}_{\mathcal{R}}$ KB. For each inverse role $r^{-}$appearing in $q$, we introduce a new role name $r_{\text {inv }}$. We then let $q^{\prime}$ be the RPQ obtained by replacing every occurrence of $r^{-}$in the query by $r_{\text {inv }}$, and let $\mathcal{T}^{\prime}$ be the extension of $\mathcal{T}$ with the role inclusions $r^{-} \sqsubseteq r_{\text {inv }}$ and $r_{\text {inv }} \sqsubseteq r^{-}$for each new role name $r_{\text {inv }}$. It is easy to see that $\operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))=\operatorname{cert}\left(q^{\prime},\left(\mathcal{T}^{\prime}, \mathcal{A}\right)\right)$.

We leave open whether RPQ answering in DL-Lite is P-hard in combined complexity.
We next provide combined complexity lower bounds for CRPQs. As CRPQs generalize CQs, we inherit an NP lower bound from the well-known NP-hardness in combined complexity of CQ answering in relational databases (see (Abiteboul, Hull, \& Vianu, 1995)):

Proposition 4.4. CRPQ answering is NP-hard in combined complexity for DL-Lite $e_{\text {RDFs }}$.
For DL-Lite and $\mathcal{E L}$, we show that CRPQ answering is PSpace-hard for combined complexity, in contrast to CQ answering which is NP-complete. Interestingly, PSpacehardness holds even under strong restrictions. In particular, we consider the following restrictions on the shape of the query and the form of the regular languages in query atoms:

- (Strong) acyclicity: a C2RPQ $q$ is acyclic if its associated undirected graph $G_{q}=$ $\left\{\left\{t, t^{\prime}\right\} \mid L\left(t, t^{\prime}\right) \in q\right\}$ is acyclic. It is strongly acyclic if additionally (i) it does not contain any atoms of the form $L(t, t)$ and (ii) for every pair of distinct atoms $L_{1}\left(t_{1}, t_{1}^{\prime}\right), L_{2}\left(t_{2}, t_{2}^{\prime}\right)$, we have $\left\{t_{1}, t_{1}^{\prime}\right\} \neq\left\{t_{2}, t_{2}^{\prime}\right\}$.
- Disjunction-freeness and star-height of regular expressions: a regular expression is disjunction-free if it does not contain $\cup$. The star-height of a regular expression is defined as the maximum nesting depth of stars appearing in the regular expression.

We point out that in the graph database setting, (strong) acyclicity leads to tractability: acyclic C2RPQs can be evaluated in polynomial time in combined complexity (Barceló et al., 2012), and for strongly C2RPQs, evaluation can even be done in linear time in the size of the database (Barceló, 2013). By contrast, the following result shows that strong acyclicity has no impact on the worst-case complexity of CRPQ answering in our setting:

Theorem 4.5. CRPQ answering is PSpace-hard in combined complexity for DL-Lite and $\mathcal{E} \mathcal{L}$. This result applies even under the restriction to strongly acyclic CRPQs whose regular languages are given by disjunction-free regular expressions of star-height two.

Proof. We give a reduction from the problem of emptiness of the intersection of an arbitrary number of regular languages, which is known to be PSPACE-hard. This result was first shown by Kozen (1977) for regular languages given as deterministic NFAs. More recently, Bala (2002) proved that this problem is also PSpace-hard when the regular languages are given as disjunction-free regular expressions of star-height two. So, let $\mathcal{E}_{1}, \ldots, \mathcal{E}_{n}$ be disjunctionfree regular expressions of star-height two over the alphabet $\Sigma=\left\{\sigma_{1}, \ldots, \sigma_{m}\right\}$. We will use the symbols in $\Sigma$ as role names, and we will also use concept names $A$ and $B$. For our reduction, we consider the following Boolean CRPQ:

$$
q=\exists x_{1}, \ldots, x_{n}, y \cdot A\left(x_{1}\right) \wedge \ldots A\left(x_{n}\right) \wedge \mathcal{E}_{1}\left(x_{1}, y\right) \wedge \ldots \wedge \mathcal{E}_{n}\left(x_{n}, y\right)
$$

Observe that $q$ satisfies the restrictions in the proposition statement. For the KB, we use the ABox $\mathcal{A}=\{A(a)\}$ and a TBox $\mathcal{T}$ whose form depends on the logic in question. For DL-Lite, we use

$$
\mathcal{T}=\left\{A \sqsubseteq \exists \sigma_{i} \mid \sigma_{i} \in \Sigma\right\} \cup\left\{\exists \sigma_{i}^{-} \sqsubseteq \exists \sigma_{j} \mid \sigma_{i}, \sigma_{j} \in \Sigma\right\}
$$

and for $\mathcal{E L}$, we use instead:

$$
\mathcal{T}=\{A \sqsubseteq B\} \cup\left\{B \sqsubseteq \exists \sigma_{i} . B \mid \sigma_{i} \in \Sigma\right\} .
$$

Notice that in both cases the canonical model $\mathcal{I}_{\mathcal{K}}$ of $\mathcal{K}=(\mathcal{T}, \mathcal{A})$ consists of an infinite tree rooted at $a$ such that every element in the interpretation has a unique $\sigma_{i}$-child for each $\sigma_{i} \in \Sigma$ (and no other children). Thus, we can associate to every domain element in $\mathcal{I}_{\mathcal{K}}$ the word over $\Sigma$ given by the sequence of role names encountered along the unique path from $a$, and moreover, for every word $w \in \Sigma^{*}$ we can find an element $e_{w}$ such that the sequence of role names on the path from $a$ to $e_{w}$ is exactly $w$.

We claim that $L\left(\mathcal{E}_{1}\right) \cap \ldots \cap L\left(\mathcal{E}_{n}\right)$ is non-empty if and only if $\mathcal{K} \models q$. To see why, first note that if $w \in L\left(\mathcal{E}_{1}\right) \cap \ldots \cap L\left(\mathcal{E}_{n}\right)$, then we can define a match for $q$ in the canonical model by mapping the variables $x_{1}, \ldots, x_{n}$ to $a$ and $y$ to $e_{w}$. Conversely, if $q$ is entailed from $\mathcal{K}$, then there is a match $\pi$ of $q$ in $\mathcal{I}_{\mathcal{K}}$. Since $A^{\mathcal{I}_{\mathcal{K}}}=\{a\}$, we must have $\pi\left(x_{i}\right)=a$ for every $1 \leq i \leq n$. It follows that the unique path from $a$ to $\pi(y)$ in $\mathcal{I}_{\mathcal{K}}$ is a word that belongs to every $L\left(\mathcal{E}_{i}\right)$, which means that $L\left(\mathcal{E}_{1}\right) \cap \ldots \cap L\left(\mathcal{E}_{n}\right)$ is non-empty.

We note that the proof of the preceding theorem from the conference version (Bienvenu et al., 2013) uses a simpler query in which the variables $x_{i}$ are replaced by $a$ and the atoms $A\left(x_{i}\right)$ are dropped. However, that query is not strongly acyclic. We also remark that a similar proof had already been used to establish PSpace hardness of CQs in an extension of $\mathcal{E} \mathcal{L H}$ that allows for regular role hierarchies (Krötzsch \& Rudolph, 2007).

## 5. Upper Bounds for 2RPQs

In the next two sections, we provide concrete query answering algorithms for the considered classes of path queries and DLs, which we leverage to derive matching upper bounds to the complexity lower bounds from Section 4. The technical developments are presented in stages. We begin this section by giving a simple algorithm for answering 2RPQs in DLLiterergs. The remainder of the section is then devoted to showing how this algorithm can $^{\text {a }}$ be extended to handle DL-Lite $\mathcal{R}^{2}$ and $\mathcal{E} \mathcal{L H}$. Afterwards, in Section 6, we introduce a query rewriting procedure that, when combined with the algorithms from the present section, yields a method for answering C2RPQs.

In this and the following section, we assume that all binary atoms take the form $\alpha\left(t, t^{\prime}\right)$, where $\alpha$ is an NFA over $\mathrm{N}_{\mathrm{R}}^{ \pm} \cup\left\{A\right.$ ? $\left.\mid A \in \mathrm{~N}_{\mathrm{C}}\right\}$. This is without loss of generality, since every regular expression can transformed into an equivalent NFA; an examination of the standard technique for constructing NFAs from regular expressions (Thompson, 1968) reveals that the transformation can in fact be performed by a deterministic logspace transducer.

It will also be useful to introduce notation for NFAs that result from changing the initial and final states of some other NFA. In what follows, given an NFA $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$, we will use $\alpha_{s, G}$ to denote the NFA $(S, \Sigma, \delta, s, G)$, i.e., the NFA with the same states and transitions as $\alpha$ but with initial state $s$ and final states $G$. When there is a single final state $s^{\prime}$, we write $\alpha_{s, s^{\prime}}$ in place of $\alpha_{s,\left\{s^{\prime}\right\}}$.

## Algorithm BasicEval

InPut: NFA $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$ with $\Sigma \subseteq \mathrm{N}_{\mathrm{R}}^{ \pm} \cup\left\{A ? \mid A \in \mathrm{~N}_{\mathrm{C}}\right\}$, DL-Lite DDFS $\operatorname{KB}(\mathcal{T}, \mathcal{A})$, $(a, b) \in \operatorname{Ind}(\mathcal{A}) \times \operatorname{Ind}(\mathcal{A})$

1. Initialize current $=\left(a, s_{0}\right)$ and count $=0$. Set $\max =|\mathcal{A}| \cdot|S|$.
2. While count $<\max$ and current $\notin\left\{\left(b, s_{f}\right) \mid s_{f} \in F\right\}$
(a) Let current $=(c, s)$.
(b) Guess a pair $\left(d, s^{\prime}\right) \in \operatorname{Ind}(\mathcal{A}) \times S$ and a transition $\left(s, \sigma, s^{\prime}\right) \in \delta$

- If $\sigma \in \mathrm{N}_{\mathrm{R}}^{ \pm}$, then verify that $\mathcal{T}, \mathcal{A} \models \sigma(c, d)$, and return no if not.
- If $\sigma=A$ ?, then verify that $c=d$ and $\mathcal{T}, \mathcal{A} \models A(c)$, and return no if not.
(c) Set current $=\left(d, s^{\prime}\right)$ and increment count.

3. If current $=\left(b, s_{f}\right)$ for some $s_{f} \in F$, return yes. Else return no.

Figure 9: Non-deterministic algorithm for 2 RPQ answering in DL-LiterdFs.

### 5.1 Warm-up: 2RPQ Answering in DL-Lite RDFS

A standard technique for answering 2RPQs in the absence of an ontology is to nondeterministically guess a path between a pair of individuals that is labeled by a word from the specified regular language. Such a procedure can be made to run in logarithmic space by only keeping a small portion of the path in memory at any time.

In Figure 9, we present a simple non-deterministic algorithm BasicEval for answering 2 RPQs over DL-Lite ${ }_{\text {RDFS }}$ knowledge bases that implements this idea. The algorithm takes as input an NFA $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$, a DL-Lite ${ }_{\text {RDFS }} \mathrm{KB} \mathcal{K}=(\mathcal{T}, \mathcal{A})$, and a pair of individuals $(a, b)$ from $\mathcal{A}$, and it decides whether $(a, b) \in \operatorname{cert}(\alpha(x, y), \mathcal{K})$. In Step 1, we initialize current with the pair $\left(a, s_{0}\right)$ and the counter count to 0 . We also compute the maximum value max of the counter, which corresponds to the largest length of path that needs to be considered. At every iteration of the while loop (Step 2), we start with a single pair $(c, s)$ stored in current and then proceed to guess a new pair $\left(d, s^{\prime}\right)$ together with a transition of the form $\left(s, \sigma, s^{\prime}\right)$. The idea is that we would like to append $\sigma d$ to the path guessed so far, but to do so, we must ensure that the conditions of paths are satisfied. This is the purpose of the entailment checks in Step 2(b). If the applicable check succeeds, then we place $\left(d, s^{\prime}\right)$ in current and increment count. We exit the while loop once we have reached the maximum counter value or the pair in count takes the form $\left(b, s_{f}\right)$ with $s_{f}$ a final state. In the latter case, we have managed to guess a path with the required properties, and so the algorithm returns yes.

The use of the counter ensures that the algorithm terminates, and the following proposition proves that it always outputs the correct result.

Proposition 5.1. For every $2 R P Q q=\alpha(x, y)$, $D L$-Lite RDFS $^{\text {R }} K B \mathcal{K}=(\mathcal{T}, \mathcal{A})$, and pair of individuals $(a, b)$ from $\operatorname{Ind}(\mathcal{A}):(a, b) \in \operatorname{cert}(q, \mathcal{K})$ if and only if there is some execution of BasicEval $(\alpha, \mathcal{K},(a, b))$ that returns yes.

Proof. Consider a $2 \mathrm{RPQ} q=\alpha(x, y)$ with $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$, a DL-Lite ${ }_{\text {RDFs }}$ KB $\mathcal{K}=$ $(\mathcal{T}, \mathcal{A})$, and a pair of individuals $(a, b)$ from $\operatorname{Ind}(\mathcal{A})$.

First suppose that $(a, b) \in \operatorname{cert}(q, \mathcal{K})$. Then there is a path $p=e_{0} u_{1} \ldots u_{n} e_{n}$ in $\mathcal{I}_{\mathcal{K}}$ such that $e_{0}=a, e_{n}=b$, and $\lambda(p) \in L(\alpha)$. We may assume without loss of generality that there is no path of shorter length than $p$ that satisfies these conditions. As $\mathcal{T}$ is a DL-Lite ${ }_{\text {RDFS }}$ TBox, we know that $\Delta^{\mathcal{I}_{\mathcal{K}}}=\operatorname{Ind}(\mathcal{A})$, so every $e_{i} \in \operatorname{Ind}(\mathcal{A})$. Since $\lambda(p) \in L(\alpha)$, we can find a sequence of states $s_{0} s_{1} \ldots s_{n}$ from $S$ such that $s_{n} \in F$ and for every $1 \leq i \leq n$, $\left(s_{i-1}, u_{i}, s_{i}\right) \in \delta$. Because of our minimality assumption, we know that $\left(e_{i}, s_{i}\right) \neq\left(e_{j}, s_{j}\right)$ for $i \neq j$, so the sequence of pairs $\left(e_{0}, s_{0}\right)\left(e_{1}, s_{1}\right) \ldots\left(e_{n}, s_{n}\right)$ has length at most $|\mathcal{A}| \cdot|S|$. It is easily verified that by guessing this sequence of pairs, together with the corresponding transitions $\left(s_{i-1}, u_{i}, s_{i}\right)$, we obtain an execution of BasicEval that returns yes.

For the other direction, suppose that there is some execution of $\operatorname{BasicEval}(\alpha, \mathcal{K},(a, b))$ that returns yes, and let $\left(c_{0}, s_{0}\right)\left(a_{1}, s_{1}\right) \ldots\left(c_{n}, s_{n}\right)$ be the sequence of pairs that were guessed by this execution. Then we must have $c_{0}=a, c_{n}=b$, and $s_{n} \in F$. Moreover, for every $1 \leq i \leq n$, there must exist a transition $\left(s_{i-1}, u_{i}, s_{i}\right) \in \delta$ such that $\mathcal{T}, \mathcal{A} \models u_{i}\left(c_{i-1}, c_{i}\right)$ if $u_{i} \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$, or $c_{i-1}=c_{i}$ and $\mathcal{T}, \mathcal{A} \models A\left(c_{i}\right)$ if $u_{i}=A$ ?. It follows that the sequence $p=c_{0} u_{1} c_{1} \ldots u_{n} c_{n}$ is a path in $\mathcal{I}_{\mathcal{K}}$ with $\lambda(p) \in L(\alpha)$, and so $\left(c_{0}, c_{n}\right)=(a, b) \in \operatorname{cert}(q, \mathcal{K})$.

By analyzing the complexity of the procedure BasicEval, we obtain an NL upper bound, which matches the NL lower bound from Proposition 4.1.

Theorem 5.2. $2 R P Q$ answering is in NL in combined complexity for $D L$-LiterdFs.
Proof. By Proposition 5.1, we know that BasicEval is a decision procedure for 2 RPQ answering in DL-LiterdFs. To see why BasicEval runs in non-deterministic logarithmic space, we note that (i) by utilizing a binary encoding, only logarithmic space is needed to store the value of count, the old and new values of current, and the guessed transition from $\delta$, and (ii) the entailment checks in 2 (b) can be performed in non-deterministic logarithmic space. The latter follows from the fact that instance checking is in NL in combined complexity in the superlogic DL-Lite $\mathcal{R}_{\mathcal{R}}$ (Calvanese et al., 2007).

### 5.2 2RPQ Answering in DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$

We now turn to the problem of answering 2RPQs over DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$ knowledge bases. The following example shows that the basic evaluation algorithm we used for DL-Lite RDFS is incomplete for these logics. Intuitively, the problem lies in the fact that the algorithm only considers paths along ABox individuals, whereas to satisfy the query it may be necessary to consider paths that pass through the anonymous part.

Example 5.3. In the remaining of the paper, we will consider the $\operatorname{KB}(\mathcal{T}, \mathcal{A})$ with $\mathcal{T}=$ $\left\{B \sqsubseteq \exists r, \exists r^{-} \sqsubseteq B, B \sqsubseteq \exists r_{1}, r_{1} \sqsubseteq r_{2}^{-}\right\}$and $\mathcal{A}=\left\{r(a, b), r_{2}(b, c), D(b)\right\}$ from Example 2.4. As an example 2RPQ, we consider $q^{\prime}(x, y)=\alpha(x, y)$, where $\alpha=\left\langle S, \Sigma, \delta, s_{0},\left\{s_{f}\right\}\right\rangle$ with

$$
\begin{aligned}
S & =\left\{s_{0}, s_{1}, s_{2}, s_{f}\right\} \\
\Sigma & =\left\{r, r_{1}, r_{2}, r^{-}\right\} \\
\delta & =\left\{\left(s_{0}, r, s_{0}\right),\left(s_{0}, r_{1}, s_{1}\right),\left(s_{1}, r_{2}, s_{2}\right),\left(s_{2}, r^{-}, s_{f}\right)\right\}
\end{aligned}
$$



Figure 10: A match witnessing $(a, a) \in \operatorname{cert}(q, \mathcal{K})$ for $q(x, y)=\alpha(x, y)$

The NFA $\alpha$ is depicted in the upper right-hand-side of Figure 10. Observe that $L(\alpha)=$ $L\left(r^{*} \cdot r_{1} \cdot r_{2} \cdot r^{-}\right)$(i.e., the same language as in the first atom of the C2RPQ $q$ of Example 3.4). Note that $(a, a) \in \operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))$, but this is only witnessed by the path $\operatorname{arbr}_{1} e_{2} r_{2} b r^{-} a$, which passes by the element $e_{2}$ in the anonymous part of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. On this input, the algorithm BasicEval would start from $\left(a, s_{0}\right)$. In the first iteration of the while loop, it could guess the pair $\left(b, s_{0}\right)$ and the transition $\left(s_{0}, r, s_{0}\right)$, and the entailment check in the first item of Step 2(b) would succeed since $\mathcal{T}, \mathcal{A} \models r(a, b)$. However, in the next iteration the only transitions from $s_{0}$ are $\left(s_{0}, r, s_{0}\right),\left(s_{0}, r_{1}, s_{1}\right)$, but there is no $d \in \operatorname{Ind}(\mathcal{A})$ such that $\mathcal{T}, \mathcal{A} \models r(b, d)$ or $\mathcal{T}, \mathcal{A} \models r_{1}(b, d)$. Hence there is no good pair to guess in Step 2(b) and the algorithm would fail. Since there are no other possible guesses in the first iteration that satisfy the entailment check, the algorithm incorrectly returns no.

Our aim is to modify the evaluation algorithm to take into account detours through the anonymous part. We observe that any path between two ABox individuals in $\mathcal{I}_{\mathcal{K}}$ can be decomposed into a sequence of paths of two types:

- paths whose elements all belong to the ABox
- paths that begin and end with the same ABox individual and whose intermediate points all belong to the anonymous part

Paths of the first type are already handled by the evaluation algorithm. To handle paths of the second type, we will show how to check whether there is a path of this form that starts and ends at a given individual $a$ and takes the query automaton from state $s$ to state $s^{\prime}$. If such a "loop" exists at the individual $a$, then during query evaluation, we will be allowed to jump directly from $(a, s)$ to $\left(a, s^{\prime}\right)$. By modifying the evaluation algorithm to allow such shortcuts in addition to normal transitions, we can ensure that all possible paths in the canonical model are taken into account.

A key observation is that to decide whether a loop is available at a given ABox individual $a$ it is sufficient to consider the basic concepts that hold at $a$. This leads us to define a table $\mathrm{ALoop}_{\alpha}$ in which the entry $\mathrm{ALoop}_{\alpha}\left[s, s^{\prime}\right]$ contains the set of all basic concepts $C$ that force the existence of a path of the second type whose label takes the query automaton from state $s$ to state $s^{\prime}$. In order to define $\mathrm{ALoop}_{\alpha}$, we will require a second table Loop ${ }_{\alpha}$ that will contain for each pair of states $\left(s, s^{\prime}\right)$, a set of tail concepts that guarantee the existence of
a path from an anonymous element $e$ to itself that takes the query automaton from $s$ to $s^{\prime}$ while never leaving the subtree of $\mathcal{I}_{\mathcal{K}}$ rooted at $e$ (note that we do allow $e$ to occur multiple times along the path).

Let us now proceed to the definition of the tables ALoop $_{\alpha}$ and Loop ${ }_{\alpha}$. If $\mathcal{T}$ is a DL-Lite $\mathcal{R}_{\mathcal{R}}$ TBox, then Loop $_{\alpha}$ is defined inductively using the following rules:
(L1) For every $s \in S: \operatorname{Loop}_{\alpha}[s, s]=\mathrm{TC}_{\mathcal{T}}$.
(L2) If $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$ and $C \in \operatorname{Loop}_{\alpha}\left[s_{2}, s_{3}\right]$, then $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{3}\right]$.
(L3) If $C \in \mathrm{TC}_{\mathcal{T}}, \mathcal{T} \models C \sqsubseteq A$, and $\left(s_{1}, A\right.$ ?, $\left.s_{2}\right) \in \delta$, then $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$.
(L4) If $C \in \mathrm{TC}_{\mathcal{T}}, \mathcal{T} \models C \sqsubseteq \exists R, \mathcal{T} \models R \sqsubseteq R^{\prime}, \mathcal{T} \models R^{-} \sqsubseteq R^{\prime \prime},\left(s_{1}, R^{\prime}, s_{2}\right) \in \delta, \exists R^{-} \in$ $\operatorname{Loop}_{\alpha}\left[s_{2}, s_{3}\right]$, and $\left(s_{3}, R^{\prime \prime}, s_{4}\right) \in \delta$, then $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{4}\right]$.
and the table $\mathrm{ALoop}_{\alpha}$ is constructed from Loop $_{\alpha}$ using the rule:
(L5) If $C \in \mathrm{BC}_{\mathcal{T}}, \mathcal{T} \models C \sqsubseteq \exists R, \mathcal{T} \models R \sqsubseteq R^{\prime}, \mathcal{T} \models R^{-} \sqsubseteq R^{\prime \prime},\left(s_{1}, R^{\prime}, s_{2}\right) \in \delta, \exists R^{-} \in$ $\operatorname{Loop}_{\alpha}\left[s_{2}, s_{3}\right]$, and $\left(s_{3}, R^{\prime \prime}, s_{4}\right) \in \delta$, then $C \in \operatorname{ALoop}_{\alpha}\left[s_{1}, s_{4}\right]$.

For $\mathcal{E} \mathcal{L H}$, we use the same definitions, except that the rules 4 and 5 are replaced by:
(L4') If $C \in \mathrm{TC}_{\mathcal{T}}, \mathcal{T} \models C \sqsubseteq \exists r . D, \mathcal{T} \models r \sqsubseteq r^{\prime}, \mathcal{T} \models r \sqsubseteq r^{\prime \prime},\left(s_{1}, r^{\prime}, s_{2}\right) \in \delta, D \in$ $\operatorname{Loop}_{\alpha}\left[s_{2}, s_{3}\right]$, and $\left(s_{3}, r^{\prime \prime-}, s_{4}\right) \in \delta$, then $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{4}\right]$.
(L5') If $C \in \mathrm{BC}_{\mathcal{T}}, \mathcal{T} \models C \sqsubseteq \exists r . D, \mathcal{T} \models r \sqsubseteq r^{\prime}, \mathcal{T} \models r \sqsubseteq r^{\prime \prime},\left(s_{1}, r^{\prime}, s_{2}\right) \in \delta, D \in$ $\operatorname{Loop}_{\alpha}\left[s_{2}, s_{3}\right]$, and $\left(s_{3}, r^{\prime \prime-}, s_{4}\right) \in \delta$, then $C \in \operatorname{ALoop}_{\alpha}\left[s_{1}, s_{4}\right]$.

Note that since $\mathrm{TC}_{\mathcal{T}}=\mathrm{BC}_{\mathcal{T}}$ in $\mathcal{E L H}$, the only difference between ( $\mathrm{L} 4^{\prime}$ ) and ( $\mathrm{L} 5^{\prime}$ ) is that the former adds concepts to the table $\mathrm{Loop}_{\alpha}$, while the latter adds concepts to ALoop ${ }_{\alpha}$.

The following example illustrates the construction of the tables Loop ${ }_{\alpha}$ and ALoop ${ }_{\alpha}$.
Example 5.4. Observe that in our running example $\mathrm{TC}_{\mathcal{T}}=\left\{\exists r, \exists r^{-}, \exists r_{1}, \exists r_{1}^{-}, \exists r_{2}, \exists r_{2}^{-}\right\}$ and $\mathrm{BC}_{\mathcal{T}}=\{B\} \cup \mathrm{TC}_{\mathcal{T}}$. In the first step of the loop computation for the $2 \mathrm{RPQ} q^{\prime}$ from Example 5.3, we get $\operatorname{Loop}_{\alpha}[s, s]=\mathrm{TC}_{\mathcal{T}}$ for every $s \in\left\{s_{0}, s_{1}, s_{2}, s_{f}\right\}$. We can infer that $\exists r^{-} \in \operatorname{Loop}_{\alpha}\left[s_{0}, s_{2}\right]$ using rule L4 and the following facts:

$$
\begin{array}{lll}
\mathcal{T} \models \exists r^{-} \sqsubseteq \exists r_{1} & \mathcal{T} \models r_{1} \sqsubseteq r_{1} & \mathcal{T} \models r_{1}^{-} \sqsubseteq r_{2} \\
\left(s_{0}, r_{1}, s_{1}\right) \in \delta & \exists r_{1}^{-} \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{1}\right] & \left(s_{1}, r_{2}, s_{2}\right) \in \delta
\end{array}
$$

Intuitively, every element $e$ that satisfies $\exists r^{-}$has an $r_{1}$-child $e^{\prime}$ (by $\mathcal{T} \models \exists r^{-} \sqsubseteq \exists r_{1}$ ), and $e^{\prime}$ is in turn an $r_{2}$-child of $e\left(\right.$ by $\left.r_{1}^{-} \sqsubseteq r_{2}\right)$. Hence, starting from such an element $e$, we can always use the transition $\left(s_{0}, r_{1}, s_{1}\right)$ to go to the child $e^{\prime}$ and then return to $e$ using the transition $\left(s_{1}, r_{2}, s_{2}\right)$, i.e., there is a loop from $s_{0}$ to $s_{2}$ at $e$.

In a further step, we can infer $\exists r^{-} \in \operatorname{Loop}_{\alpha}\left[s_{0}, s_{3}\right]$ by using rule L4 together with:

$$
\begin{array}{llr}
\mathcal{T} \models \exists r^{-} \sqsubseteq \exists r & \mathcal{T} \models r \sqsubseteq r & \mathcal{T} \models r^{-} \sqsubseteq r^{-} \\
\left(s_{0}, r, s_{0}\right) \in \delta & \exists r^{-} \in \operatorname{Loop}_{\alpha}\left[s_{0}, s_{2}\right] & \left(s_{2}, r^{-}, s_{f}\right) \in \delta
\end{array}
$$

This reflects that whenever an element satisfies $\exists r^{-}$, there is a loop from $s_{0}$ to $s_{3}$ as follows: we can move to some $r$-child (which exists by $\mathcal{T} \models \exists r^{-} \sqsubseteq \exists r$ ) staying in state $s_{0}$ with $\left(s_{0}, r, s_{0}\right)$, then use the previously computed loop to jump to $s_{2}$ at the same element, and then go back up to $e$ with the transition $\left(s_{2}, r^{-}, s_{f}\right)$.

One can verify that no further loops can be inferred with the rules, so we obtain:

$$
\operatorname{Loop}_{\alpha}\left[s_{i}, s_{i}\right]=\mathrm{TC}_{\mathcal{T}} \text { for } 0 \leq i \leq 3 \quad \operatorname{Loop}_{\alpha}\left[s_{0}, s_{2}\right]=\left\{\exists r^{-}\right\} \quad \operatorname{Loop}_{\alpha}\left[s_{0}, s_{3}\right]=\left\{\exists r^{-}\right\}
$$

Now we compute ALoop $_{\alpha}$. First we note that there are applications of L5 analogous to the two described applications of L 4 , which respectively result in $\exists r^{-} \in \operatorname{ALoop}_{\alpha}\left[s_{0}, s_{2}\right]$ and $\exists r^{-} \in$ ALoop $_{\alpha}\left[s_{0}, s_{3}\right]$. Moreover, similar applications but using $\mathcal{T} \models B \sqsubseteq \exists r_{1}$ instead of $\mathcal{T} \models \exists r^{-} \sqsubseteq \exists r_{1}$ and $\mathcal{T} \models B \sqsubseteq \exists r$ instead of $\mathcal{T} \models \exists r^{-} \sqsubseteq \exists r$ yield $B \in$ ALoop $_{\alpha}\left[s_{0}, s_{2}\right]$ and $B \in$ ALoop $_{\alpha}\left[s_{0}, s_{3}\right]$. Further applications of L5 yield no new loops, hence we obtain:

$$
\operatorname{ALoop}_{\alpha}\left[s_{0}, s_{2}\right]=\left\{\exists r^{-}, B\right\} \quad \operatorname{ALoop}_{\alpha}\left[s_{0}, s_{3}\right]=\left\{\exists r^{-}, B\right\}
$$

We observe that the tables Loop $_{\alpha}$ and ALoop $_{\alpha}$ can be constructed in polynomial time in $|\mathcal{T}|$ and $|\alpha|$ since entailment of inclusions is in P for both DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E} \mathcal{L H}$ (Calvanese et al., 2007; Baader et al., 2005). The following propositions show that Loop ${ }_{\alpha}$ and ALoop ${ }_{\alpha}$ have the desired meaning:

Proposition 5.5. For every $D L$-Lite $_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H} K B \mathcal{K}=(\mathcal{T}, \mathcal{A})$ and $d \in \Delta^{\mathcal{I}_{\mathcal{K}}} \backslash \operatorname{lnd}(\mathcal{A})$, the following are equivalent:

1. Tail $(d) \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$;
2. There is a path $p=e_{0} u_{1} e_{1} \ldots u_{n} e_{n}$ in $\mathcal{I}_{\mathcal{K}}$ such that $\lambda(p) \in L\left(\alpha_{s, s^{\prime}}\right), e_{0}=e_{n}=d$, and $\left.e_{i} \in \Delta^{\mathcal{I}_{\mathcal{K}}}\right|_{d}$ for every $0 \leq i \leq n$.

Proof. Consider a $\mathrm{KB} \mathcal{K}=(\mathcal{T}, \mathcal{A})$ and an automaton $\alpha=(S, \Sigma, \delta, s, F)$. We begin by proving that the first statement implies the second. Fix a sequence of applications of the rules $\mathrm{L} 1, \mathrm{~L} 2, \mathrm{~L} 3$, and L 4 ( or L4') which generates the full table Loop ${ }_{\alpha}$, and let $k$ be the length of this sequence. It then suffices to show the following claim for all $1 \leq i \leq k$ :

Claim: If $C$ is inserted into $\operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$ on the $i$-th rule application and $d \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} \backslash \operatorname{lnd}(\mathcal{A})$ is such that $\operatorname{Tail}(d)=C$, then there is a path $p=e_{0} u_{1} e_{1} \ldots u_{n} e_{n}$ in $\mathcal{I}_{\mathcal{K}}$ such that $\lambda(p) \in$ $L\left(\alpha_{s, s^{\prime}}\right), e_{0}=e_{n}=d$, and $e_{i} \in \Delta^{\left.\mathcal{I}_{\mathcal{K}}\right|_{d}}$ for every $0 \leq i \leq n$.
Proof of claim. The proof is by induction on $i$. First suppose that $C$ is inserted into $\operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$ with the first rule application and $d \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} \backslash \operatorname{Ind}(\mathcal{A})$ is such that $\operatorname{Tail}(d)=C$. Then either rule L1 or rule L3 must have been applied. In the former case, we have $s=s^{\prime}$, so the path $p=d$ satisfies the required conditions (recall that in this case $\lambda(p) \in \varepsilon$ ). If instead it was rule L3 that was applied, then we must have $\mathcal{T} \models C \sqsubseteq A$ and $\left(s, A ?, s^{\prime}\right) \in \delta$ for some concept name $A$. Since $\operatorname{Tail}(d)=C$, we must have $d \in A^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. It follows that $p=d A ? d$ is a path satisfying the required conditions.

For the induction step, suppose that the statement holds for all $1 \leq i<k$, and let $d \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} \backslash \operatorname{Ind}(\mathcal{A})$ be such that $C=\operatorname{Tail}(d)$ is inserted into $\operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$ on the $k$-th rule application. The first possibility is that the $k$-th rule application involves rules L1 or L3, in which case we proceed as in the base case. The next possibility is that rule L2 was
applied. Then there must exist $s^{\prime \prime}$ such that after the first $k-1$ rule applications, we have $C \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime \prime}\right]$ and $C \in \operatorname{Loop}_{\alpha}\left[s^{\prime \prime}, s^{\prime}\right]$. Applying the induction hypothesis, we find paths $p_{1}$ and $p_{2}$ that both begin and end with $d$, contain only elements from $\Delta^{\left.\mathcal{I}_{\mathcal{K}}\right|_{d}}$, and are such that $\lambda\left(p_{1}\right) \in L\left(\alpha_{s, s^{\prime \prime}}\right)$ and $\lambda\left(p_{2}\right) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime}}\right)$. Let $p_{3}$ be the path obtained by taking $p_{1}$ then adding $p_{2}$ with its first occurrence of $d$ removed. Then $p_{3}$ begins and ends at $d$, contains only elements from $\Delta^{\left.\mathcal{I}_{\mathcal{K}}\right|_{d}}$, and is such that $\lambda\left(p_{3}\right) \in L\left(\alpha_{s, s^{\prime}}\right)$.

The final possibility is that the $k$-th rule application involves rule L4. Here the proof differs depending on whether $\mathcal{T}$ is formulated in DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E L H}$. We give the proof only for DL-Lite $\mathcal{R}_{\mathcal{R}}$; the proof for $\mathcal{E L H}$ proceeds analogously. We first note that since an application of rule L4 leads to the insertion of $C$ into $\operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$ at stage $k$, it must be the case that we can find $R, R^{\prime}, R^{\prime \prime} \in \mathrm{N}_{\mathrm{R}}^{ \pm}$and $s^{\prime \prime}, s^{\prime \prime \prime} \in S$ such that

- $\mathcal{T} \models C \sqsubseteq \exists R, \mathcal{T} \models R \sqsubseteq R^{\prime}, \mathcal{T} \models R^{-} \sqsubseteq R^{\prime \prime}$,
- $\left(s, R^{\prime}, s^{\prime \prime}\right) \in \delta$ and $\left(s^{\prime \prime \prime}, R^{\prime \prime}, s^{\prime}\right) \in \delta$, and
- $\exists R^{-} \in \operatorname{Loop}_{\alpha}\left[s^{\prime \prime}, s^{\prime \prime \prime}\right]$ (after $k-1$ rule applications).

As $\operatorname{Tail}(d)=C$ and $\mathcal{T} \models C \sqsubseteq \exists R$, the element $d^{\prime}=d R \exists R^{-}$must belong to $\Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. Then by applying the induction hypothesis, we can infer that there is a path $p^{\prime}$ that begins and ends at $d^{\prime}$, contains only elements from $\Delta^{\left.\mathcal{I}_{\mathcal{K}}\right|_{d^{\prime}}}$, and is such that $\lambda\left(p^{\prime}\right) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime \prime \prime}}\right)$. It follows that the path $p=d R^{\prime} p^{\prime} R^{\prime \prime} d$ satisfies all requirements, and in particular, is such that $\lambda(p) \in L\left(\alpha_{s, s^{\prime}}\right)$. (end proof of claim)

To show the other direction, we proceed by induction on the length of the path $p=$ $e_{0} u_{1} e_{1} \ldots u_{n} e_{n}$. The first base case is when $n=0$, i.e., when $\lambda(p)=\varepsilon$. Then $\varepsilon \in L\left(\alpha_{s, s^{\prime}}\right)$, which implies that $s=s^{\prime}$. By rule L1 of the definition of Loop ${ }_{\alpha}$, we must have Tail $(d) \in$ $\operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$. The second base case is when $n=1$, i.e., $p=d A$ ? $d$. Since $p$ is a path, we must have $d \in A^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$, which means that $\mathcal{T} \models \operatorname{Tail}(d) \sqsubseteq A$. We also know that $\lambda(p)=A$ ? $\in$ $L\left(\alpha_{s, s^{\prime}}\right)$, which implies that $\left(s, A ?, s^{\prime}\right) \in \delta$. We have thus shown that the conditions of rule L 3 are satisfied, and so $\operatorname{Tail}(d) \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$.

For the induction step, suppose that the second direction holds for all $0 \leq \ell<k$, and suppose that there is a path $p=e_{0} u_{1} e_{1} \ldots u_{k} e_{k}$ in $\mathcal{I}_{\mathcal{K}}$ with $k>2$ such that $\lambda(p) \in L\left(\alpha_{s, s^{\prime}}\right)$, $e_{0}=e_{k}=d$, and $e_{i} \in \Delta^{\left.\mathcal{I}_{\mathcal{K}}\right|_{d}}$ for every $0 \leq i \leq k$. First suppose that there exists some $e_{j}$ with $0<j<k$ such that $e_{j}=d$. Let $p_{1}=e_{0} u_{1} e_{1} \ldots e_{j}$ and $p_{2}=e_{j} u_{j} \ldots e_{k}$. We know that $\lambda(p)=\lambda\left(p_{1}\right) \lambda\left(p_{2}\right) \in L\left(\alpha_{s, s^{\prime}}\right)$, so there must exist some state $s^{\prime \prime}$ such that $\lambda\left(p_{1}\right) \in L\left(\alpha_{s, s^{\prime \prime}}\right)$ and $\lambda\left(p_{2}\right) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime}}\right)$. Applying the induction hypothesis to $p_{1}$ and $p_{2}$, we obtain $\operatorname{Tail}(d) \in \operatorname{Loop}_{\alpha}\left[s^{\prime \prime}, s^{\prime}\right]$ and Tail $(d) \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime \prime}\right]$. Hence, by rule L2 of the construction of $\operatorname{Loop}_{\alpha}$, we must have $\operatorname{Tail}(d) \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$.

Now let us consider the second possibility, which is that $e_{j} \neq d$ for all $0<j<k$. Since $\left.\mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{d}$ is tree-shaped, and $p$ is a path, it must be the case that $e_{1}=e_{k-1}=d R C \in \Delta^{\left.\mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{d}}$. At this point, the proof slightly differs depending on whether we are in DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E L H}$. We present the proof for the case of $\mathcal{E L} \mathcal{H}$, in which case we have $R=r \in \mathrm{~N}_{\mathrm{R}}$ and $\mathcal{T} \models$ $\operatorname{Tail}(d) \sqsubseteq \exists r . C$. As $\lambda(p)=u_{1} \ldots u_{k} \in L\left(\alpha_{s, s^{\prime}}\right), e_{0}=e_{k}=d$, and $e_{1}=e_{k-1}=d R C$, it must be the case that

- $u_{1} \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$and $\mathcal{T} \models r \sqsubseteq u_{1}$, and
- $u_{k} \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$and $\mathcal{T} \models r \sqsubseteq t$ where $u_{k}=t^{-}$.

We also know that there must exist states $s^{\prime \prime}, s^{\prime \prime \prime} \in S$ such that

- $\left(s, u_{1}, s^{\prime \prime}\right) \in \delta$ and $\left(s^{\prime \prime \prime}, u_{k}, s^{\prime}\right) \in \delta$, and
- $u_{2} \ldots u_{k-1} \in L\left(\alpha_{s^{\prime \prime}, s^{\prime \prime \prime}}\right)$.

We can apply the induction hypothesis to $p^{\prime}=e_{1} u_{1} \ldots u_{k-1} e_{k-1}$ to infer that Tail $\left(e_{1}\right)=$ $C \in \operatorname{Loop}_{\alpha}\left[s^{\prime \prime}, s^{\prime \prime \prime}\right]$. We thus satisfy all of the required conditions for applying rule L4' to obtain $\operatorname{Tail}(d) \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$. The proof for DL-Lite $\mathcal{R}_{\mathcal{R}}$ is analogous.

Proposition 5.6. For every $D L-$ Lite $_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H} K B \mathcal{K}=(\mathcal{T}, \mathcal{A})$, NFA $\alpha$ containing states $s, s^{\prime}$, and $a \in \operatorname{Ind}(\mathcal{A})$, the following statements are equivalent:

1. There is a concept $C \in \operatorname{ALoop}_{\alpha}\left[s, s^{\prime}\right]$ such that $\mathcal{T}, \mathcal{A} \models C(a)$.
2. There is a path $p=e_{0} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ in $\mathcal{I}_{\mathcal{K}}$ such that $\lambda(p) \in L\left(\alpha_{s, s^{\prime}}\right), e_{0}=e_{n}=a$, $n>1$, and $e_{i} \notin \operatorname{Ind}(\mathcal{A})$ for $0<i<n$.

Proof. Fix a DL-Lite $\mathcal{R}$ TBox $\mathcal{T}$ and an NFA $\alpha=(S, \Sigma, \delta, s, F)$ (the proof for $\mathcal{E} \mathcal{L H}$ is similar and left to reader). For the first direction, suppose that there is a concept $C \in \operatorname{ALoop}_{\alpha}\left[s, s^{\prime}\right]$ such that $\mathcal{T}, \mathcal{A} \models C(a)$. It follows from the definition of ALoop $_{\alpha}$ that there exists roles $R, R^{\prime}, R^{\prime \prime} \in \mathrm{N}_{\mathrm{R}}^{ \pm}$and states $s^{\prime \prime}, s^{\prime \prime \prime} \in S$ such that $\mathcal{T} \models C \sqsubseteq \exists R, \mathcal{T} \models R \sqsubseteq R^{\prime}, \mathcal{T} \models R^{-} \sqsubseteq$ $R^{\prime \prime},\left(s, R^{\prime}, s^{\prime \prime}\right) \in \delta, \exists R^{-} \in \operatorname{Loop}_{\alpha}\left[s^{\prime \prime}, s^{\prime \prime \prime}\right]$, and $\left(s^{\prime \prime \prime}, R^{\prime \prime}, s^{\prime}\right) \in \delta$. Since $\mathcal{T}, \mathcal{A} \models C(a)$ and $\mathcal{T} \models C \sqsubseteq \exists R$, it follows from the definition of $\mathcal{I}_{\mathcal{K}}$ that $a R \exists R^{-} \in \Delta^{\mathcal{I}_{\mathcal{K}}}$. By Proposition 5.5, $\exists R^{-} \in \operatorname{Loop}_{\alpha}\left[s^{\prime \prime}, s^{\prime \prime \prime}\right]$ implies that there is a path $p=e_{0} u_{0} \ldots u_{n} e_{n}$ in $\left.\mathcal{I}_{\mathcal{K}}\right|_{a R \exists R^{-}}$with $e_{0}=$ $e_{n}=a R \exists R^{-}$and $\lambda(p) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime \prime \prime}}\right)$. Then it can be verified that by taking $p^{\prime}=a R^{\prime} p R^{\prime \prime} a$, we obtain a path in $\mathcal{I}_{\mathcal{K}}$ that satisfies the conditions of the second statement. In particular, since $\left(s, R^{\prime}, s^{\prime \prime}\right) \in \delta, \lambda(p) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime \prime \prime}}\right)$, and $\left(s^{\prime \prime \prime}, R^{\prime \prime}, s^{\prime}\right) \in \delta$, we have $\lambda\left(p^{\prime}\right) \in L\left(\alpha_{s, s^{\prime}}\right)$.

For the second direction, suppose that there is a path $p=e_{0} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ in $\mathcal{I}_{\mathcal{K}}$ such that $\lambda(p) \in L\left(\alpha_{s, s^{\prime}}\right), e_{0}=e_{n}=a, n>1$, and $e_{i} \notin \operatorname{Ind}(\mathcal{A})$ for $0<i<n$. Since $n>1$, $e_{1} \notin \operatorname{Ind}(\mathcal{A})$, and $p$ is a path, it must be the case that $e_{1}=a R \exists R^{-}$for some $R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$. Moreover, since $\left.\mathcal{I}_{\mathcal{K}}\right|_{a}$ is a tree, we must have $e_{n-1}=e_{1}$. Then by the definition of paths, it must be the case that

- $u_{1} \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$and $\mathcal{T} \models R \sqsubseteq u_{1}$, and
- $u_{n} \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$and $\mathcal{T} \models R^{-} \sqsubseteq u_{n}$.

We also know that there must exist states $s^{\prime \prime}, s^{\prime \prime \prime} \in S$ such that:

- $\left(s, u_{1}, s^{\prime \prime}\right) \in \delta$ and $\left(s^{\prime \prime \prime}, u_{n}, s^{\prime}\right) \in \delta$
- $u_{2} \ldots u_{n-1} \in L\left(\alpha_{s^{\prime \prime}, s^{\prime \prime \prime}}\right)$

By applying Proposition 5.5 to $p^{\prime}=e_{1} u_{1} \ldots u_{n-1} e_{n-1}$, we can infer that $\exists R^{-} \in \operatorname{Loop}_{\alpha}\left[s^{\prime \prime}, s^{\prime \prime \prime}\right]$. Since $a R \exists R^{-} \in \Delta^{\mathcal{I}_{\mathcal{K}}}$, the definition of $\mathcal{I}_{\mathcal{K}}$ ensures that there is some $C \in \mathrm{BC}_{\mathcal{T}}$ such that $a \in$ $C^{\mathcal{I}_{\mathcal{K}}}$ and $\mathcal{T} \models C \sqsubseteq \exists R$. The conditions of Rule 5 are thus satisfied, so $C \in \operatorname{ALoop}_{\alpha}\left[s, s^{\prime}\right]$.

Algorithm EvalAtom
InPut: NFA $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$ with $\Sigma \subseteq \mathrm{N}_{\mathrm{R}}^{ \pm} \cup\left\{A ? \mid A \in \mathrm{~N}_{\mathrm{C}}\right\}$, DL-Lite D $_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L} \mathcal{H} \mathrm{KB}$ $(\mathcal{T}, \mathcal{A}),(a, b) \in \operatorname{Ind}(\mathcal{A}) \times \operatorname{Ind}(\mathcal{A})$

1. Test whether $(\mathcal{T}, \mathcal{A})$ is satisfiable, output yes if not.
2. Initialize current $=\left(a, s_{0}\right)$ and count $=0$. Set $\max =|\mathcal{A}| \cdot|S|+1$.
3. While count $<\max$ and current $\notin\left\{\left(b, s_{f}\right) \mid s_{f} \in F\right\}$
(a) Let current $=(c, s)$.
(b) Guess a pair $\left(d, s^{\prime}\right) \in \operatorname{Ind}(\mathcal{A}) \times S$ together with either $\left(s, \sigma, s^{\prime}\right) \in \delta$ or $B \in$ $\mathrm{ALoop}_{\alpha}\left[s, s^{\prime}\right]$.
(c) If $\left(s, \sigma, s^{\prime}\right)$ was guessed

- If $\sigma \in \mathrm{N}_{\mathrm{R}}^{ \pm}$, then verify that $\mathcal{T}, \mathcal{A} \models \sigma(c, d)$, and return no if not.
- If $\sigma=A$ ?, then verify that $c=d$ and $\mathcal{T}, \mathcal{A} \models A(c)$, and return no if not.
(d) If $B$ was guessed, then verify that $c=d$ and $\mathcal{T}, \mathcal{A} \models B(c)$, and return no if not.
(e) Set current $=\left(d, s^{\prime}\right)$ and increment count.

4. If current $=\left(b, s_{f}\right)$ for some $s_{f} \in F$, return yes. Else return no.

Figure 11: Non-deterministic algorithm for 2 RPQ answering in DL-Lite $_{\mathcal{R}}$ and $\mathcal{E} \mathcal{L H}$.

Now that we have a means to determine which loops through the anonymous part are available from a given ABox individual, we are ready to present the extended evaluation algorithm EvalAtom in Figure 11 that handles DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$ KBs. The algorithm EvalAtom differs from BasicEval in two respects. First, because DL-Lite $\mathcal{R}_{\mathcal{R}}$ KBs may contain contradictions, there is an initial consistency check in Step 1 to determine whether the input KB is satisfiable (this step can be skipped for $\mathcal{E L H} \mathrm{KBs}$, which are always satisfiable). If the KB is shown to be unsatisfiable, then the query trivially holds, so the algorithm outputs yes. The second difference occurs in Step 3(b) within the while loop, where we now have the choice between guessing a pair $\left(d, s^{\prime}\right) \in \operatorname{Ind}(\mathcal{A}) \times S$ (as before) and guessing a concept $B \in \operatorname{ALoop}_{\alpha}\left[s, s^{\prime}\right]$. The first option corresponds to taking a step in the ABox, whereas the second corresponds to a shortcut through the anonymous part. If we choose the second option, then we must check that the selected concept is entailed at the current individual. The exit conditions for the while loop and the criterion for outputting yes in Step 4 remain unchanged.

Example 5.7. Algorithm EvalAtom correctly returns yes on the input ( $a, a$ ) together with our example query and KB, in contrast to BasicEval (as was shown in Example 5.3). Indeed, a successful execution (illustrated pictorially in Figure 12) starts from ( $a, s_{0}$ ) and guesses the pair $\left(b, s_{0}\right)$ and the transition $\left(s_{0}, r, s_{0}\right)$ in the first iteration of the while loop. The checks in Step 3(c) succeed as $\mathcal{T}, \mathcal{A} \models r(a, b)$. In the next iteration, it can guess ( $b, s_{2}$ ) and, as $\exists r^{-} \in \operatorname{ALoop}_{\alpha}\left[s_{0}, s_{2}\right]$, it can also guess the concept $\exists r^{-}$. The checks in Step 3(d) succeed because $\mathcal{T}, \mathcal{A} \models \exists r^{-}(b)$. In the third iteration, it guesses the pair ( $a, s_{f}$ ) and the transition


Figure 12: Establishing $(a, a) \in \operatorname{cert}(q, \mathcal{K})$ using EvalAtom for $q(x, y)=\alpha(x, y)$
$\left(s_{2}, r^{-}, s_{f}\right)$. As $\mathcal{T}, \mathcal{A} \models r^{-}(b, a)$, the checks in Step 3(c) succeed again. Since $s_{f} \in F$, this is the last iteration, and in Step 4 the algorithm returns yes.

Proposition 5.8. For every $2 R P Q q=\alpha(x, y)$, DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E L H} K B \mathcal{K}=(\mathcal{T}, \mathcal{A})$, and pair of individuals $(a, b)$ from $\operatorname{Ind}(\mathcal{A}):(a, b) \in \operatorname{cert}(q, \mathcal{K})$ if and only if there is some execution of EvalAtom $(\alpha, \mathcal{K},(a, b))$ that returns yes.

Proof. Consider a $2 \mathrm{RPQ} q=\alpha(x, y)$ with $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$, a DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E L H}$ KB $\mathcal{K}=(\mathcal{T}, \mathcal{A})$, and a pair of individuals $(a, b)$ from $\operatorname{Ind}(\mathcal{A})$.

For the first direction, suppose that $(a, b) \in \operatorname{cert}(q, \mathcal{K})$. Then there is a path $p=$ $e_{0} u_{1} \ldots u_{n} e_{n}$ in $\mathcal{I}_{\mathcal{K}}$ such that $e_{0}=a, e_{n}=b$, and $\lambda(p) \in L(\alpha)$. We may assume without loss of generality that there is no shorter path with these properties. Since $\lambda(p) \in L(\alpha)$, we can find a sequence of states $s_{0} s_{1} \ldots s_{n}$ from $S$ such that $s_{n} \in F$ and for every $1 \leq i \leq n$, $\left(s_{i-1}, u_{i}, s_{i}\right) \in \delta$. Let $j_{1}<\ldots<j_{m}$ be all of the indices $i$ such that $e_{i} \in \operatorname{Ind}(\mathcal{A})$, and consider the sequence of pairs $\Gamma=\left(e_{j_{1}}, s_{j_{1}}\right)\left(e_{j_{2}}, s_{j_{2}}\right) \ldots\left(e_{j_{m}}, s_{j_{m}}\right)$. Observe that $j_{1}=0$ and $j_{m}=n$, so we have $e_{j_{1}}=a, e_{j_{m}}=b$, and $s_{j_{m}} \in F$. Also observe that we must have $\left(e_{j_{\ell}}, s_{j_{\ell}}\right) \neq\left(e_{j_{k}}, s_{j_{k}}\right)$ whenever $\ell \neq k$, since otherwise, we could construct a shorter path with the same properties as $p$, contradicting our minimality assumption. It follows that the sequence $\Gamma$ contains at most $|\mathcal{A}| \cdot|S|$ pairs. Thus, to prove that the sequence $\Gamma$ leads to an execution of EvalAtom that returns yes, it only remains to show that it is always possible to guess a transition or concept in 3 (b) such that the checks in 3 (c) and $3(\mathrm{~d})$ succeed. Thus, let us suppose that current $=\left(e_{j_{\ell}}, s_{j_{\ell}}\right)$ and we guess the pair $\left(e_{j_{\ell+1}}, s_{j_{\ell+1}}\right)$ in $3(\mathrm{~b})$. There are three cases to consider:

- Case 1: $j_{\ell+1}=j_{\ell}+1$ and $u_{j_{\ell+1}}=R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$. From earlier, we have that $\left(s_{j_{\ell}}, R, s_{j_{\ell+1}}\right) \in \delta$, and since $p$ is a path, we know that $\left(e_{j_{\ell}}, e_{j_{\ell+1}}\right) \in R^{\mathcal{I}_{\mathcal{K}}}$. The latter implies that $\mathcal{T}, \mathcal{A} \models R\left(e_{j_{\ell}}, e_{j_{\ell+1}}\right)$, so by choosing the transition $\left(s_{j_{\ell}}, R, s_{j_{\ell+1}}\right)$ in $3(\mathrm{~b})$, we can ensure that the entailment check will succeed in 3(c).
- Case 2: $j_{\ell+1}=j_{\ell}+1$ and $u_{j_{\ell+1}}=A$ ?. From earlier, we have that $\left(s_{j_{\ell}}, A\right.$ ?, $\left.s_{j_{\ell+1}}\right) \in \delta$, and since $p$ is a path, we know that $e_{j_{\ell}}=e_{j_{\ell+1}} \in A^{\mathcal{I}_{\mathcal{K}}}$. By choosing the transition ( $s_{j_{\ell}}, A$ ?, $s_{j_{\ell+1}}$ ), we can ensure that the conditions of 3(c) are satisfied.
- Case 3: $j_{\ell+1}>j_{\ell}+1$. In this case, the path $p^{\prime}=e_{j_{\ell}} u_{j_{\ell}+1} \ldots u_{j_{\ell+1}} e_{j_{\ell+1}}$ is such that $\lambda\left(p^{\prime}\right) \in L\left(\alpha_{s_{j_{\ell}}, s_{j_{\ell+1}}}\right), e_{j_{\ell}}=e_{j_{\ell+1}} \in \operatorname{Ind}(\mathcal{A})$, and for every $j_{\ell}<i<j_{\ell+1}$, we have
$e_{i} \notin \operatorname{Ind}(\mathcal{A})$. We can thus apply Lemma 5.6 to find a concept $C \in \operatorname{ALoop}_{\alpha}\left[s_{j_{\ell}}, s_{j_{\ell+1}}\right]$ such that $\mathcal{T}, \mathcal{A} \models C\left(e_{j_{\ell}}\right)$. By choosing $C$ in $3(\mathrm{~b})$, we can be sure that the entailment check in $3(\mathrm{~d})$ will succeed.

For the other direction, consider some execution of $\operatorname{Eval} \operatorname{Atom}(\alpha, \mathcal{K},(a, b))$ that returns yes, and let $\left(c_{0}, s_{0}\right)\left(a_{1}, s_{1}\right) \ldots\left(c_{n}, s_{n}\right)$ be the sequence of pairs that were guessed by this execution. Then we must have $c_{0}=a, c_{n}=b$, and $s_{n} \in F$. To complete the proof, it suffices to establish the following claim:

Claim: For every $0 \leq i \leq n$, there is a path $p_{i}$ from $c_{0}$ to $c_{i}$ in $\mathcal{I}_{\mathcal{K}}$ such that $\lambda\left(p_{i}\right) \in L\left(\alpha_{s_{0}, s_{i}}\right)$. Proof of claim. The proof is by induction on $i$. For the base case $(i=0)$, we can simply take the path $p_{0}=c_{0}$ since $\lambda\left(p_{0}\right)=\varepsilon$ and $\varepsilon \in L\left(\alpha_{s_{0}, s_{0}}\right)$. For the induction step, we suppose that $p_{k-1}$ is a path from $c_{0}$ to $c_{k-1}$ such that $\lambda\left(p_{k-1}\right) \in L\left(\alpha_{s_{0}, s_{k-1}}\right)$, and we show how to construct a path $p_{k}$ with the required properties. There are three cases to consider depending on which transition or concept was guessed together with ( $c_{k}, s_{k}$ ) in Step 3(b).

- Case 1: the transition $\left(s_{k-1}, R, s_{k}\right) \in \delta$ was guessed. Since the check in 3(c) succeeded, we have $\mathcal{T}, \mathcal{A} \models R\left(c_{k_{1}}, c_{k}\right)$. Then $p_{k}=p_{k-1} R c_{k}$ is a path in $\mathcal{I}_{\mathcal{K}}$ from $c_{0}$ to $c_{k}$ such that $\lambda\left(p_{k}\right) \in L\left(\alpha_{s_{0}, s_{k}}\right)$.
- Case 2: the transition $\left(s_{k-1}, A\right.$ ?, $\left.s_{k}\right) \in \delta$ was guessed. Then since the check in 3(c) succeeded, we have $c_{k}=c_{k-1}$ and $\mathcal{T}, \mathcal{A} \models A\left(c_{k_{1}}\right)$. Then $p_{k}=p_{k-1} A ? c_{k}$ is a path in $\mathcal{I}_{\mathcal{K}}$ from $c_{0}$ to $c_{k}$ such that $\lambda\left(p_{k}\right) \in L\left(\alpha_{s_{0}, s_{k}}\right)$.
- Case 3: the concept $B \in \operatorname{ALoop}_{\alpha}\left[s_{k-1}, s_{k}\right]$ was guessed. Since the check in 3(d) succeeded, we have $c_{k}=c_{k-1}$ and $\mathcal{T}, \mathcal{A} \models B\left(c_{k-1}\right)$. By applying Lemma 5.6, we can find a path $p^{\prime}=e_{0} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ in $\mathcal{I}_{\mathcal{K}}$ such that $\lambda\left(p^{\prime}\right) \in L\left(\alpha_{s_{k-1}, s_{k}}\right)$ and $e_{0}=$ $e_{n}=c_{k-1}$. Then $p_{k}=p_{k-1} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ is a path from $c_{0}$ to $c_{k}$ with $\lambda\left(p_{k}\right)=$ $\lambda\left(p_{k-1}\right) \lambda\left(p^{\prime}\right) \in L\left(\alpha_{s_{0}, s_{k}}\right)$.

By analyzing the complexity of our modified evaluation procedure, we obtain upper bounds that match the lower bounds from Section 4.

Theorem 5.9. $2 R P Q$ answering is in

1. NL in data complexity for DL-Lite $\mathcal{R}_{\mathcal{R}}$;
2. P in combined complexity for $D L$-Lite $\mathcal{R}_{\mathcal{R}}$;
3. P in combined and data complexity for $\mathcal{E} \mathcal{L H}$.

Proof. The procedure EvalAtom involves three different types of checks: the consistency check performed in Step 1 and the entailment and loop checks that take place in Step 3. The cost of these checks depends on the choice of the DL and the complexity measure. Aside from these checks, the base procedure runs in non-deterministic logarithmic space in combined complexity, as only logarithmic space is needed to keep track of the value of count, the old and new values of current, and the guesses of transitions and concepts.

For DL-Lite $\mathcal{R}^{\text {, }}$, we know from existing results (Calvanese et al., 2007) that the consistency and entailment checks can be performed in non-deterministic logarithmic space in combined
complexity (and hence also for data complexity). We have also seen that the table ALoop ${ }_{\alpha}$ can be constructed in polynomial time in $|\mathcal{T}|$ and $\alpha$, so the loop checks can be performed in constant time in $|\mathcal{A}|$ and in polynomial time w.r.t. the whole input. It follows that EvalAtom runs in non-deterministic logarithmic space w.r.t. $|\mathcal{A}|$, yielding Statement (1). Regarding Statement (2), we note that EvalAtom can be viewed as an NL procedure that uses a $P$ oracle to handle the loop checks. Since $N L^{P}=P$, this yields a $P$ upper bound on the combined complexity of 2 RPQ answering in DL-Lite $\mathcal{R}_{\mathcal{R}}$. The proof of Statement (3) is similar. We simply note that when the input TBox is formulated in $\mathcal{E L H}$, all three types of checks can be performed in polynomial time w.r.t. the whole input (Baader et al., 2005). Using $\mathrm{NL}^{\mathrm{P}}=\mathrm{P}$, we may conclude that EvalAtom provides a polynomial-time procedure for 2 RPQ answering in $\mathcal{E} \mathcal{L H}$.

## 6. Upper Bounds for C2RPQs

The main objective of this section is to define a procedure for deciding, given a $\mathrm{KB} \mathcal{K}=$ $(\mathcal{T}, \mathcal{A})$, C2RPQ $q(\vec{x})$ of arity $k$, and $k$-tuple $\vec{a}$ of individuals from $\mathcal{A}$, whether there is a match $\pi$ for $q$ in $\mathcal{I}_{\mathcal{K}}$ such that $\pi(\vec{x})=\vec{a}$. A naïve approach might consist in guessing a mapping $\pi$ from the query variables to the individuals in the core of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ and then checking that $\pi$ is a match by running the EvalAtom algorithm on $\alpha\left(\pi(t), \pi\left(t^{\prime}\right)\right)$ for every query atom $\alpha\left(t, t^{\prime}\right)$. Such an algorithm would properly take into account paths between individuals that pass by the anonymous part, but because it does not consider matches that send variables to anonymous objects, it would still be incomplete. In particular, such a procedure would not return $a$ or $b$ as answers in Example 3.4. To regain completeness, one could instead guess matches into the entire domain of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, but this would not yield a decision procedure since the latter may be infinite. Moreover, since matches of C2RPQs can involve domain elements that are arbitrarily far apart, it is not apparent how to identify a suitable finite subset of the domain that is guaranteed to contain a match for the query if one exists. To address these challenges, the procedure we propose in this section comprises two main steps. As a first step, we rewrite the input query $q$ into a set $Q$ of C2RPQs such that there is a match $\pi$ for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\pi(\vec{x})=\vec{a}$ if and only if there is a match $\pi^{\prime}$ for some $q^{\prime} \in Q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\pi^{\prime}(\vec{x})=\vec{a}$. The advantage of the rewritten queries is that we will only need to consider matches $\pi^{\prime}$ which map query variables to $\operatorname{Ind}(\mathcal{A})$. The second step decides the existence of such restricted matches for the rewritten queries using the EvalAtom procedure defined in Section 5.

For the purposes of this section, it will prove convenient to work with DL-Lite $\mathcal{R}_{\mathcal{R}}$ TBoxes that satisfy the following condition: for every role name $r \in \operatorname{sig}(\mathcal{T})$, there exists concept names $A_{r}, A_{r}^{-}$such that $\mathcal{T}$ contains the inclusions $A_{r} \sqsubseteq \exists r, \exists r \sqsubseteq A_{r}, A_{r^{-}} \sqsubseteq \exists r^{-}, \exists r^{-} \sqsubseteq$ $A_{r^{-}}$, and these are the only inclusions in $\mathcal{T}$ involving the concept names $A_{r}$ and $A_{r^{-}}$. Note that if $\mathcal{T}$ does not satisfy this condition, then we can simply choose fresh concepts $A_{r}, A_{r^{-}}$ for each role name $r \in \operatorname{sig}(\mathcal{T})$ and add the corresponding inclusions to $\mathcal{T}$. The resulting TBox $\mathcal{T}^{\prime}$ is a model conservative of extension of $\mathcal{T}$, hence for every ABox $\mathcal{A}$ and every $\mathrm{C} 2 \operatorname{RPQ} q$ with $\operatorname{sig}(q) \cap\left(\operatorname{sig}\left(\mathcal{T}^{\prime}\right) \backslash \operatorname{sig}(\mathcal{T})\right)=\emptyset$, we have $\operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))=\operatorname{cert}\left(q,\left(\mathcal{T}^{\prime}, \mathcal{A}\right)\right)$. We may therefore assume without any loss of generality that all DL-Lite $\mathcal{R}_{\mathcal{R}}$ TBoxes considered in this section satisfy this syntactic condition. In what follows, when we use a concept name $A_{R}$, with $R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$, we will assume the TBox contains the inclusions $A_{R} \sqsubseteq \exists R$ and $\exists R \sqsubseteq A_{R}$.

Bienvenu, Ortiz, \& Šimkus

### 6.1 Query Rewriting

Our aim is to rewrite our query in such a way that we do not need to map any variables to the anonymous part of the model. We draw our inspiration from a query rewriting procedure for Horn-SHIQ introduced by Eiter, Ortiz, Šimkus, Tran, and Xiao (2012). The main intuition is as follows. Suppose we have a match $\pi$ for $q$ which maps some variable $y$ to the anonymous part, and no other variable is mapped below $\pi(y)$. Then we modify $q$ in such a way that the resulting query $q^{\prime}$ has a match $\pi^{\prime}$ that is the same as $\pi$ except that those variables mapped to $\pi(y)$ by $\pi$ are now mapped by $\pi^{\prime}$ to the (unique) parent of $\pi(y)$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. The delicate point is that we must "split" atoms of the form $\alpha\left(t, t^{\prime}\right)$ with $y \in\left\{t, t^{\prime}\right\}$ into the parts which are satisfied in the subtree $\left.\mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{\pi(y)}$, and those which occur above $\pi(y)$, whose satisfaction still needs to be determined and thus must be incorporated into the new query. With each iteration of the rewriting procedure, we obtain a query which has a match which maps variables "closer" to the core of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, until eventually we obtain a query that has a match which maps all terms to $\operatorname{Ind}(\mathcal{A})$.

In Figure 13, we implement this intuition by defining an algorithm OneStep that performs a single (non-deterministic) rewriting step. We illustrate the functioning of OneStep in the following examples.

Example 6.1. Recall the query $q(x)=\exists y, z \cdot r^{*} \cdot r_{1} \cdot r_{2} \cdot r^{-}(x, y) \wedge r^{-} \cdot r^{-}(y, z) \wedge D(z)$ from Example 3.4, and the $\operatorname{KB}(\mathcal{T}, \mathcal{A})$ from Example 2.4.

To illustrate the rewriting algorithm, we first disregard the first atom and consider the simpler Boolean query $q_{1}=\exists y, z \cdot \beta(y, z) \wedge D(z)$, where $\beta$ is the NFA with language $r^{-} \cdot r^{-}$ depicted in Figure 14. The figure shows a match $\pi$ for $q_{1}$ with $\pi(y)=e_{11}$ and $\pi(z)=b$. Since $y$ is a leaf of the image of $q_{1}$ under $\pi$, we want to modify $q_{1}$ into a query $q_{1}^{\prime}$ that has a match $\pi^{\prime}$ which only differs from $\pi$ in having $\pi^{\prime}(y)=e_{1}$, where $e_{1}$ is the parent of $\pi(y)=e_{11}$. Intuitively, this is done using OneStep by choosing Leaf $=\{y\}$ as the only variable to be 'moved up'. Then we choose the concept $B$ because if it holds at $e_{1}$, then this enforces the existence of an $r^{-}$role from a child of $e_{1}$ (namely, $e_{11}$ ) to $e_{1}$. Hence, by checking that $B$ holds at $e_{1}$, we implicitly check that the first $r^{-}$needed to satisfy $\beta$ holds below $e_{1}$. This rewriting step is illustrated in the upper part of Figure 14. Formally, in Step 1 we choose Leaf $=\{y\}$. Then in Step 2 we choose $\exists r^{-}$(since, intuitively, this is the tail concept that 'causes' everything that holds at $e_{11}$ ). In Step 3, we simply take the final state $s_{f}^{\prime}$, so the atom $\beta(z, y)$ remains the same, and in Step 4 we do nothing. These two steps are so simple in this example because there are no complex paths in our query match that need to be deeper than $e_{11}$. Next, in Step 5 , we choose $B$, since $\mathcal{T} \models \exists r$, that is, $B$ enforces the existence of a node that satisfies the tail concept $\exists r^{-}$that we guessed above. In Step 5 b we only need to take care of the atom $\beta(y, z)$. We choose $s_{1}^{\prime}$ and in Step 6, we replace $\beta$ by $\beta_{s_{1}^{\prime}, s_{f}^{\prime}}$. In Step 7, we add the atom $B(y)$, and we output the query $B(y) \wedge \beta_{s_{1}^{\prime}, s_{f}^{\prime}}(y, z) \wedge D(z)$.

The lower part of the figure illustrates a successive application of OneStep in which we proceed analogously, dropping the second $r^{-}$of $\beta$, and adding again an atom $B(y)$. This results in the query $q_{1}^{\prime \prime}=\exists y, z \cdot B(y) \beta_{s_{f}^{\prime}, s_{f}^{\prime}}(y, z) \wedge D(z)$, which has a match $\pi^{\prime \prime}(z)=\pi^{\prime \prime}(y)=b$ ranging over individuals only. Note that $L\left(\beta_{s_{f}^{\prime}, s_{f}^{\prime}}\right)=\{\varepsilon\}$ so this query is equivalent to $\exists y . B(y) \wedge D(y)$.

## Algorithm OneStep

Input: C2RPQ $q$ with all binary atoms specified by NFAs, DL-Lite $\mathcal{R}^{\text {or }} \mathcal{E} \mathcal{L H}$ TBox $\mathcal{T}$

1. Guess a non-empty set Leaf $\subseteq \operatorname{qvars}(q)$ and $y \in$ Leaf.

Rename all variables in Leaf to $y$.
2. Guess $C \in \mathrm{TC}_{\mathcal{T}}$ such that $\mathcal{T} \models C \sqsubseteq B$ for every atom $B(y) \in q$. Drop all such atoms from $q$.
3. For each atom $\alpha\left(t, t^{\prime}\right)$ where $\alpha=(S, \Sigma, \delta, s, F)$ is an NFA and $y \in\left\{t, t^{\prime}\right\}$,

- Guess a sequence $s_{1}, \ldots, s_{n-1}$ of distinct states from $S$ and a state $s_{n} \in F$.
- Replace $\alpha\left(t, t^{\prime}\right)$ by the atoms $\alpha_{s, s_{1}}(t, y), \alpha_{s_{1}, s_{2}}(y, y), \ldots, \alpha_{s_{n-2}, s_{n-1}}(y, y)$, $\alpha_{s_{n-1}, s_{n}}\left(y, t^{\prime}\right)$.

4. Drop all atoms $\alpha_{s, s^{\prime}}(y, y)$ such that $C \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$.
5. Guess some $D \in \mathrm{BC}_{\mathcal{T}}$ and $R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$such that:
(a) $C=\exists R^{-}$and $\mathcal{T} \models D \sqsubseteq \exists R$ (if $\mathcal{T}$ is a DL-Lite $\mathcal{R}_{\mathcal{R}}$ TBox) $R \in \mathrm{~N}_{\mathrm{R}}$ and $\mathcal{T} \models D \sqsubseteq \exists R . C \quad$ (if $\mathcal{T}$ is an $\mathcal{E} \mathcal{L H}$ TBox)
(b) For each atom $\alpha(y, t)$ with $\alpha=(S, \Sigma, \delta, s, F)$, there exists $s^{\prime} \in S$ and $U \in \mathbf{N}_{\mathrm{R}}^{ \pm}$ with $\left(s, U^{-}, s^{\prime}\right) \in \delta$ and $\mathcal{T} \models R \sqsubseteq U$.
(c) For each atom $\alpha(t, y)$ with $\alpha=(S, \Sigma, \delta, s, F)$, there exists $s^{\prime \prime} \in S, s_{f} \in F$, and $U \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$with $\left(s^{\prime \prime}, U, s_{f}\right) \in \delta$ and $\mathcal{T} \models R \sqsubseteq U$.

For atoms of the form $\alpha(y, y)$, conditions (b) and (c) must both be satisfied.
6. Replace

- each atom $\alpha(y, t)$ with $t \neq y$ by $\alpha_{s^{\prime}, F}(y, t)$
- each atom $\alpha(t, y)$ with $t \neq y$ by $\alpha_{s, s^{\prime \prime}}(t, y)$
- each atom $\alpha(y, y)$ by $\alpha_{s^{\prime}, s^{\prime \prime}}(y, y)$
with $s, s^{\prime}, s^{\prime \prime}, F$ as in Step 5.

7. If $D \in \mathrm{~N}_{\mathrm{C}}$ is the concept chosen in Step 5 , add $D(y)$ to $q$. If $D=\exists P^{-}$was the chosen concept, add $A_{P^{-}}(y)$ to $q$. Output $q$.

Figure 13: Non-deterministic query rewriting algorithm OneStep.

Example 6.2. Now we illustrate two rewriting steps for $q(x)=\exists y, z \cdot r^{*} \cdot r_{1} \cdot r_{2} \cdot r^{-}(x, y) \wedge$ $r^{-} \cdot r^{-}(y, z) \wedge D(z)$ (see Figure 15). First, recall the match $\pi$ from Figure 7, which is also reproduced in the top part of Figure 15. Observe that $\pi(y)=e_{11}$ is a leaf in the image of $\pi$. To move the match up one step, we make the same initial choices as in Example 6.1, setting Leaf $=\{y\}$ in Step 1 and selecting the tail concept $\exists r^{-}$that characterises $e_{11}$ in


Figure 14: Two successive applications of OneStep to $q_{1}=\exists y, z \cdot r^{-} \cdot r^{-}(y, z) \wedge D(z)$

Step 2. However, because the path that satisfies $\alpha$ goes deeper than $\pi(y)=e_{11}$, in Step 3, we will need to separate the path into the parts that occur below $e_{11}$ and those that occur above it. We choose the sequence of states $s_{0}, s_{f}$ as these correspond to the states of the automata when it visits $e_{11}$, and we replace $\alpha(x, y)$ by $\alpha_{s_{0}, s_{0}}(x, y)$ and $\alpha_{s_{0}, s_{f}}(y, y)$. Note that this intermediate query is illustrated at the top of Figure 15. Since $\exists r^{-} \in \operatorname{Loop}_{\alpha}\left[s_{0}, s_{f}\right]$ (see Example 5.4), we drop the second atom in Step 4. Now we can 'move up' $y$ similarly as above, but considering simultaneously atoms $\alpha_{s_{0}, s_{0}}(x, y)$ and $\beta(y, z)$. In Step 5, we again choose the concept $B$, since it satisfies $\mathcal{T} \models B \sqsubseteq \exists r$. For the atom $\beta(y, z)$ we choose $s_{1}^{\prime}$ and $U=r($ Step $5(\mathrm{~b}))$ and for $\alpha_{s_{0}, s_{0}}(x, y)$, we choose $s_{0}$ and $U=r$ (Step 5(c)). In Step 6, we replace the two query atoms by $\alpha_{s_{0}, s_{0}}(x, y)$ and $\beta_{s_{1}^{\prime}, s_{f}^{\prime}}(y, z)$, and in Step 7, we add the atom $B(y)$ and output the resulting query $\alpha_{s_{0}, s_{0}}(x, y) \wedge B(y) \wedge \beta_{s_{1}^{\prime}, s_{f}^{\prime}}(y, z)$, which is displayed in the middle of Figure 15. The lower part of this figure depicts a second, similar application of OneStep that outputs a query that has a match ranging over the individuals only.

Slightly abusing notation, we will use $\operatorname{OneStep}(q, \mathcal{T})$ to denote the set of queries that are output by some execution of OneStep on input $(q, \mathcal{T})$. We then consider the set $\operatorname{Rewrite}(q, \mathcal{T})$ consisting of all queries that can be obtained from $(q, \mathcal{T})$ by zero or more applications of OneStep. Formally, we define $\operatorname{Rewrite}(q, \mathcal{T})$ as the smallest set that contains the initial query $q$ and is closed under applications of OneStep, i.e., if $q^{\prime} \in \operatorname{Rewrite}(q, \mathcal{T})$ and $q^{\prime \prime} \in$ OneStep $\left(q^{\prime}, \mathcal{T}\right)$, then $q^{\prime \prime} \in \operatorname{Rewrite}(q, \mathcal{T})$.

The next proposition shows that using $\operatorname{Rewrite}(q, \mathcal{T})$, we can reduce the problem of finding an arbitrary query match to finding a match involving only ABox individuals.


Figure 15: Two rewriting steps for $q(x)=\exists y, z \cdot r^{*} \cdot r_{1} \cdot r_{2} \cdot r^{-}(x, y) \wedge r^{-} \cdot r^{-}(y, z) \wedge D(z)$

Proposition 6.3. For every satisfiable $D L-$ Lite $_{\mathcal{R}}$ or $\mathcal{E L H} K B(\mathcal{T}, \mathcal{A})$ and $C 2 R P Q$ : $\vec{a} \in \operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))$ if and only if there exists a query $q^{\prime}(\vec{x}) \in \operatorname{Rewrite}(q, \mathcal{T})$ and a match $\pi$ for $q^{\prime}$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ such that $\pi(\vec{x})=\vec{a}$ and $\pi(z) \in \operatorname{Ind}(\mathcal{A})$ for every variable $z$ in $q^{\prime}$.

We split Proposition 6.3 into two lemmas, the first showing completeness of Rewrite, and the second showing its correctness.

Lemma 6.4. If $\vec{a} \in \operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))$, then there exists a query $q^{\prime}(\vec{x}) \in \operatorname{Rewrite}(q, \mathcal{T})$ and $a$ match $\pi$ for $q^{\prime}$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ such that $\pi(\vec{x})=\vec{a}$ and $\pi(z) \in \operatorname{Ind}(\mathcal{A})$ for every variable $z$ in $q^{\prime}$.

Proof. Consider a knowledge base $(\mathcal{T}, \mathcal{A})$ and its canonical model $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. For every element $e \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, we define the distance $\operatorname{dist}(e)$ of $e$ from the core of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ as follows: $\operatorname{dist}\left(a R_{1} C_{1} \ldots R_{n} C_{n}\right)=n$. Observe that $\operatorname{dist}(e)=0$ implies $e \in \operatorname{Ind}(\mathcal{A})$. Using this notion of distance, we can define the cost of a match $\mu$ of some query in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ as $\Sigma_{v \in \operatorname{dom}(\mu)} \operatorname{dist}^{\mathcal{T}, \mathcal{A}}(\mu(v))$, where $\operatorname{dom}(\mu)$ the domain of $\mu$. We remark that a match has cost equal to zero just in the case that it maps all query variables to ABox individuals.

Suppose that $\vec{a} \in \operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))$, and $\pi$ is a match for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ such that $\pi(\vec{x})=\vec{a}$. If $\pi$ maps all variables in $q$ to $\operatorname{Ind}(\mathcal{A})$, then we are done. Otherwise, there must exist some variable $y$ such that $\pi(y)=e S C$ and there is no $z \in \operatorname{vars}(q)$ such that $\pi(y)$ is a proper prefix of $\pi(z)$. We aim to construct a match $\pi^{\prime}$ for some query $q^{\prime} \in \operatorname{OneStep}(q, \mathcal{T})$ such that
(c1) $\pi^{\prime}(t)=\pi(t)$ for every $t \in \operatorname{terms}(q)$ such that $\pi(t) \neq \pi(y)$; and
(c2) $\pi^{\prime}(t)=e$ for every $t \in \operatorname{terms}(q)$ with $\pi(t)=\pi(y)$.
In other words, $\pi^{\prime}$ is essentially the same as $\pi$ except that it maps each $t$ with $\pi(t)=\pi(y)$ one step closer to the ABox. Observe that (c1) and (c2) together ensure that the cost of $\pi^{\prime}$ is strictly inferior to the cost of $\pi$. Thus, by repeatedly applying this operation, we will eventually obtain a query $q^{*} \in \operatorname{Rewrite}(q, \mathcal{T})$ and a match $\pi^{*}$ for $q^{*}$ that has cost zero, i.e., such that $\pi^{*}(z) \in \operatorname{Ind}(\mathcal{A})$ for every variable $z$ in $q^{*}$.

We now show how to obtain a query $q^{\prime} \in \operatorname{OneStep}(q, \mathcal{T})$ and match $\pi^{\prime}$ with properties (c1) and (c2). In Step 1, we set Leaf $=\{t \in \operatorname{terms}(q) \mid \pi(t)=\pi(y)=e S C\}$. Note that Leaf $\subseteq q \operatorname{vars}(q)$ since $e S C \notin \operatorname{Ind}(\mathcal{A})$. We define a function $\sigma$ as follows: $\sigma(t)=t$ if $\pi(t) \neq \pi(y)$, else $\sigma(t)=y$. At the end of Step 1, we have the query:

$$
\left\{B(\sigma(t)) \mid B(t) \in q_{0}\right\} \cup\left\{\alpha\left(\sigma(t), \sigma\left(t^{\prime}\right)\right) \mid \alpha\left(t, t^{\prime}\right) \in q_{0}\right\}
$$

In Step 2, we choose the concept $C$ (note that $C \in \mathrm{TC}_{\mathcal{T}}$ because $\pi(y)=e S C$ ). Consider some atom $B(y)$ that is present at the end of Step 1 . The existence of atom $B(y)$ at the end of this step means that there must have existed an atom $B(v) \in q$ with $v \in$ Leaf. Since $\pi$ is a match for $q$, we know that $\pi(v)=\pi(y) \in B^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. As $\pi(y)=e S C$, it follows from the definition of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ that $\mathcal{T} \models C \sqsubseteq B$, as required by Step 2 .

Next we show how to select states $s_{1}, \ldots, s_{n}$ in Step 3. Consider an atom $\alpha\left(t_{1}, t_{2}\right)$ which is present in the query at the start of Step 3 , such that $y \in\left\{t_{1}, t_{2}\right\}$ and $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$. Then we know that there is an atom $\alpha\left(t_{1}^{\prime}, t_{2}^{\prime}\right)$ in the input query $q$ such that $t_{1}=\sigma\left(t_{1}^{\prime}\right)$ and $t_{2}=\sigma\left(t_{2}^{\prime}\right)$. Since $\pi$ is a match for $q$ with $\pi\left(t_{1}\right)=\pi\left(t_{1}^{\prime}\right)$ and $\pi\left(t_{2}\right)=\pi\left(t_{2}^{\prime}\right)$, we know that $\left(\pi\left(t_{1}\right), \pi\left(t_{2}\right)\right) \in L(\alpha)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$. It follows that we can find a path $p=e_{0} u_{1} e_{1} \ldots u_{m} e_{m}$ such that $e_{0}=\pi\left(t_{1}\right), e_{m}=\pi\left(t_{2}\right)$, and $\lambda(p) \in L(\alpha)$. We assume without loss of generality that $m$ is minimal, i.e., we cannot find a path satisfying the same properties but of shorter length. Now let $j_{1}<\ldots<j_{n-1}$ be all of the indices $0<\ell<m$ such that $e_{\ell}=\pi(y)$, and consider the following paths:

- $p_{1}=e_{0} u_{1} \ldots u_{j_{1}} e_{j_{1}}$
- $p_{i}=e_{j_{i-1}} u_{j_{i-1}+1} \ldots u_{j_{i}} e_{j_{i}}$ for $1<i<n$
- $p_{n}=e_{j_{n-1}} u_{j_{n-1}+1} \ldots u_{m} e_{m}$

We also define a sequence of states $s_{1}, \ldots, s_{n}$ such that

- $\lambda\left(p_{1}\right) \in L\left(\alpha_{s_{0}, s_{1}}\right)$,
- $\lambda\left(p_{i}\right) \in L\left(\alpha_{s_{i-1}, s_{i}}\right)$ for $1<i \leq n$, and
- $s_{n} \in F$.

Note that such a sequence of states must exist since $\lambda(p)=\lambda\left(p_{1}\right) \lambda\left(p_{2}\right) \ldots \lambda\left(p_{n}\right) \in L(\alpha)$ and $\alpha$ has start state $s_{0}$ and final states $F$. Using the fact that $e_{j_{i}}=\pi(y)$ for $1 \leq i<n$, we have the following:

- $\left(\pi\left(t_{1}\right), \pi(y)\right)=\left(\pi\left(t_{1}^{\prime}\right), \pi(y)\right) \in L\left(\alpha_{s_{0}, s_{1}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$.
- $(\pi(y), \pi(y)) \in L\left(\alpha_{s_{i-1}, s_{i}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, for $2 \leq i<n$.
- $\left(\pi(y), \pi\left(t_{2}\right)\right)=\left(\pi(y), \pi\left(t_{2}^{\prime}\right)\right) \in L\left(\alpha_{s_{n-1}, s_{n}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$.

We aim to show that $s_{i} \neq s_{j}$ for every $1 \leq k<\ell<n$. Suppose for a contradiction that $s_{k}=s_{\ell}$ for some $1 \leq k<\ell<n$, and consider the path $p^{\prime}=e_{0} u_{1} \ldots u_{j_{k}} e_{j_{k}} u_{j_{\ell}+1} e_{j_{\ell}+1} \ldots u_{m} e_{m}$. As $\lambda\left(p^{\prime}\right)=u_{1} \ldots u_{j_{k}} u_{j_{\ell}+1} \ldots u_{m}, u_{1} \ldots u_{j_{k}} \in L\left(\alpha_{s_{0}, s_{k}}\right), u_{j_{\ell}+1} \ldots u_{m} \in L\left(\alpha_{s_{\ell}, s_{n}}\right)=L\left(\alpha_{s_{k}, s_{n}}\right)$, and $s_{n} \in F$, it follows that $\lambda\left(p^{\prime}\right) \in L(\alpha)$. However, this means that $p^{\prime}$ satisfies the same conditions as $p$ but has strictly shorter length, contradicting our minimality assumption. Hence all states in the sequence $s_{1}, \ldots, s_{n-1}$ must be distinct. We can thus choose this sequence of states in Step 3, and replace the atom $\alpha\left(t_{1}, t_{2}\right)=\alpha\left(\sigma\left(t_{1}^{\prime}\right), \sigma\left(t_{2}^{\prime}\right)\right)$ with the atoms:

$$
\alpha_{s_{0}, s_{1}}\left(\sigma\left(t_{1}^{\prime}\right), y\right), \alpha_{s_{1}, s_{2}}(y, y), \alpha_{s_{n-2}, s_{n-1}}(y, y), \alpha_{s_{n-1}, s_{n}}\left(y, \sigma\left(t_{2}^{\prime}\right)\right) .
$$

The final choices to be made occur in Step 5, where we must choose a concept $D \in B C_{\mathcal{T}}$ and role $R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$such that conditions (a), (b), and (c) are satisfied. We set $R=S$ (recall that $\pi(y)=e S C)$. If $e \notin \operatorname{Ind}(\mathcal{A})$, then we let $D$ be the unique concept such that $e=e^{\prime} P D$. Note that if we are in DL-Lite $\mathcal{R}_{\mathcal{R}}$, then $D=\exists P^{-}$. It follows from the definition of canonical models that $\mathcal{T} \models D \sqsubseteq \exists S$ (if we are in DL-Lite $\mathcal{R}_{\mathcal{R}}$ ) or $\mathcal{T} \models D \sqsubseteq \exists S . C$ (for $\mathcal{E} \mathcal{L H}$ ), so condition (a) holds. If instead we have $e \in \operatorname{Ind}$, then the definition of canonical models, together with our normal form for $\mathcal{E L H}$ TBoxes, guarantees that there is some $D \in \mathrm{BC}_{\mathcal{T}}$ such that $e \in D^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ and $\mathcal{T} \models D \sqsubseteq \exists S$ (if we are in DL-Lite $\mathcal{R}_{\mathcal{R}}$ ) or $\mathcal{T} \models D \sqsubseteq \exists S . C$ (for $\mathcal{E} \mathcal{L H})$. Note that in the case of DL-Lite $\mathcal{R}_{\mathcal{R}}, \mathcal{T}, \mathcal{A} \models \exists S(e)$ implies that one of the following holds: (i) there is a concept assertion $A(e) \in \mathcal{A}$ such that $\mathcal{T} \models A \sqsubseteq \exists S$, (ii) there is a role assertion $S^{\prime}\left(e, e^{\prime}\right) \in \mathcal{A}$ such that $\mathcal{T} \models \exists S^{\prime} \sqsubseteq \exists S$, or (iii) there is a role assertion $S^{\prime}\left(e^{\prime}, e\right) \in \mathcal{A}$ such that $\mathcal{T} \models \exists\left(S^{\prime}\right)^{-} \sqsubseteq \exists S$. Thus, it is always possible to choose $D$ such that $\mathcal{A} \models D(e)$, and we will assume in what follows that $D$ has this property.

It remains to show that conditions (b) and (c) are verified when $R=S$. For (b), consider some binary atom $\beta(y, t)$ which belongs to the query at the start of Step 5 . Then we know that there must exist an atom $\alpha\left(t_{1}^{\prime}, t_{2}^{\prime}\right) \in q$ with $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$ such that $\beta(y, t)$ is equal to one of the following atoms which replaced $\alpha\left(\sigma\left(t_{1}^{\prime}\right), \sigma\left(t_{2}^{\prime}\right)\right)$ during Step 3: $\alpha_{s_{0}, s_{1}}\left(\sigma\left(t_{1}^{\prime}\right), y\right), \alpha_{s_{1}, s_{2}}(y, y), \ldots, \alpha_{s_{n-2}, s_{n-1}}(y, y), \alpha_{s_{n-1}, s_{n}}\left(y, \sigma\left(t_{2}^{\prime}\right)\right)$. Thus, we have an atom of the form $\alpha_{s_{i-1}, s_{i}}(y, t)$. Using property ( $\star$ ) and considering the different possible values for $t$, we can infer that $(\pi(y), \pi(t)) \in L\left(\alpha_{s_{i-1}, s_{i}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, as witnessed by the path $p_{i}=$
$e_{j_{i-1}} u_{j_{i-1}+1} \ldots u_{j_{i}} e_{j_{i}}$. We also know that $e_{\ell} \neq \pi(y)$ for every $j_{i-1}<\ell<j_{i}$. In particular, this means that either the path $p_{i}$ is entirely contained in the subtree rooted at $\pi(y)$ or it never visits any element below $\pi(y)$. The former option cannot hold, since it would imply that $t=y$ and that $C \in \operatorname{Loop}_{\alpha}\left[s_{i}, s_{i+1}\right]$, so the atom would have been removed in Step 4. Thus, it must be the case that the first "step" in the path $p_{i}$ goes from $\pi(y)$ to its parent $e$. It follows that $u_{j_{i-1}+1}=U^{-}$for some $U \in \mathbf{N}_{\mathrm{R}}^{ \pm}$such that $(e, \pi(y)) \in U^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$. Since $\pi(y)=e S C$, we must have $\mathcal{T} \models S \sqsubseteq U$. Since $\lambda\left(p_{i}\right) \in L\left(\alpha_{s_{i-1}, s_{i}}\right)$, there must exist a state $s^{\prime} \in S$ such that $\left(s_{i-1}, U^{-}, s^{\prime}\right) \in \delta$ and $u_{j_{i-1}+2} \ldots u_{j_{i}} \in L\left(\alpha_{s^{\prime}, s_{i}}\right)$. This shows that condition (b) is satisfied, and also that $(e, \pi(t)) \in L\left(\alpha_{s^{\prime}, s}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$

We now consider condition (c). Take some atom of the form $\beta(t, y)$ which appears in the query at the start of Step 5. Then we know from earlier that we can find some atom $\alpha\left(t_{1}^{\prime}, t_{2}^{\prime}\right) \in q\left(\right.$ where $\left.\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)\right)$ such that $\beta(t, y)$ is equal to one of the following atoms that replaced $\alpha\left(\sigma\left(t_{1}^{\prime}\right), \sigma\left(t_{2}^{\prime}\right)\right)$ during Step 3: $\alpha_{s_{0}, s_{1}}\left(\sigma\left(t_{1}^{\prime}\right), y\right), \alpha_{s_{1}, s_{2}}(y, y)$, $\ldots, \alpha_{s_{n-2}, s_{n-1}}(y, y), \alpha_{s_{n-1}, s_{n}}\left(y, \sigma\left(t_{2}^{\prime}\right)\right)$. It follows that $\beta(t, y)$ has the form $\alpha_{s_{i}, s_{i+1}}(t, y)$. Using property $(\star)$, and considering the two possible values for $t$, we can deduce that $(\pi(t), \pi(y)) \in L\left(\alpha_{s_{i}, s_{i+1}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$, as witnessed by the path $p_{i}=e_{j_{i-1}} u_{j_{i-1}+1} \ldots u_{j_{i}} e_{j_{i}}$. Arguing as we did for (b), we can show that the path $p_{i}$ is entirely contained in the subtree rooted at $\pi(y)$ or the path never visits any element below $\pi(y)$. The former option would imply that $t=y$ and that $C \in \operatorname{Loop}_{\alpha}\left[s_{i}, s_{i+1}\right]$, so the atom would have been removed in Step 5. Thus, it must be the case that the last "step" in the path is from $e$ to $\pi(y)$, so $u_{j_{i}}=U$ for some $U \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$such that $(e, \pi(y)) \in U^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. Since $\pi(y)=e S C$, we must have $\mathcal{T} \models S \sqsubseteq U$. Since $\lambda\left(p_{i}\right) \in L\left(\alpha_{s_{i}, s_{i+1}}\right)$, we also know that there must exist a state $s^{\prime \prime}$ such that $\left(s^{\prime \prime}, U, s_{i+1}\right) \in \delta$ and $u_{j_{i-1}+1} \ldots u_{j_{i}-1} \in L\left(\alpha_{s_{i}, s^{\prime \prime}}\right)$. This shows that condition (c) is satisfied, and also that $(\pi(t), e) \in L\left(\alpha_{s_{i}, s^{\prime \prime}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. It is also important to note that if $t=y$, then we can apply the arguments for conditions (b) and (c) together to show that $(e, e) \in L\left(\alpha_{s^{\prime}, s^{\prime \prime}} \mathcal{I}^{\mathcal{I}} \mathcal{T}, \mathcal{A}\right.$ (with state $s^{\prime}$ as in (b), and $s^{\prime \prime}$ as required for (c)).

Now let $q^{\prime}$ be the query we obtain at the end of Step 7 when all non-deterministic choices are made in the manner that we have described. Consider the mapping $\pi^{\prime}$ defined as follows:

- $\pi^{\prime}(t)=\pi(t)$ for every $t \in \operatorname{terms}(q)$ with $\pi(t) \neq \pi(y)$.
- $\pi^{\prime}(t)=e$ for every $t \in \operatorname{terms}(q)$ with $\pi(t)=\pi(y)$.

Note that the mapping $\pi^{\prime}$ satisfies the properties (c1) and (c2). It remains to show that $\pi^{\prime}$ is a match for $q^{\prime}$.

Consider first a concept atom $B(t) \in q^{\prime}$. There are two possibilities. Either $B(t)$ appears in $q$ and $t \notin$ Leaf, or $B(t)$ was introduced in Step 7. In the former case, we know that $\pi$ satisfies $B(t)$, and since $\pi^{\prime}(t)=\pi(t)$ (since $t \notin$ Leaf), the same is true of $\pi^{\prime}$. In the latter case, we must have $t=y$ and either $B=D$ if $D \in \mathrm{~N}_{\mathrm{C}}$ or $B=A_{P^{-}}$if $D=\exists P^{-}$. If $B=D$, then we can use the fact that $\pi^{\prime}(t)=e$ and $D$ was chosen so that $e \in D^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$ to infer that $\pi^{\prime}$ satisfies $B(t)$. If $B=A_{P^{-}}$, then we have that $e \in\left(\exists P^{-}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$. Since $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ is a model of $\mathcal{T}$ and $\exists P^{-} \sqsubseteq A_{P^{-}} \in \mathcal{T}$, we also have $e \in A_{P_{-}}^{\mathcal{I}_{\mathcal{A}}, \mathcal{A}}$, which means that $\pi^{\prime}$ satisfies $B(t)$.

Now consider an atom $\gamma\left(t^{\prime}, t^{\prime \prime}\right) \in q^{\prime}$. If $y \notin\left\{t^{\prime}, t^{\prime \prime}\right\}$, then $\gamma\left(t^{\prime}, t^{\prime \prime}\right) \in q$. As $\pi$ is a match for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, it must be the case that $\left(\pi\left(t^{\prime}\right), \pi\left(t^{\prime \prime}\right)\right) \in L(\gamma)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. Since $\pi^{\prime}\left(t^{\prime}\right)=\pi\left(t^{\prime}\right)$ and $\pi^{\prime}\left(t^{\prime \prime}\right)=\pi\left(t^{\prime \prime}\right)$, the same holds for $\pi^{\prime}$, and so the atom $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ is satisfied by $\pi^{\prime}$. Next
suppose that $y \in\left\{t^{\prime}, t^{\prime \prime}\right\}$. An examination of Rewrite shows that $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ must have replaced some atom during Step 6. We distinguish three cases:

- Case 1: $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ replaces $\alpha_{s_{i}, s_{i+1}}(y, t)$ with $t \neq y$. Then $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ must have the form $\alpha_{s^{\prime}, s_{i+1}}(y, t)$, where $s^{\prime}$ is the state that was chosen to ensure condition (b) in Step 5. We recall that $s^{\prime}$ is such that $(e, \pi(t)) \in L\left(\alpha_{s^{\prime}, s_{i+1}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$. Since $t \neq y$, we know that $t \notin$ Leaf, and so $\pi(t)=\pi^{\prime}(t)$. It follows that $\left(\pi^{\prime}(y), \pi^{\prime}(t)\right) \in L\left(\alpha_{s^{\prime}, s_{i+1}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$, so the atom $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ is satisfied by $\pi^{\prime}$.
- Case 2: $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ replaces $\alpha_{s_{i}, s_{i+1}}(t, y)$ with $t \neq y$. Then $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ must have the form $\alpha_{s_{i}, s^{\prime \prime}}(t, y)$, where $s^{\prime \prime}$ is the state that was used in condition $5(\mathrm{c})$. We showed earlier when examining condition (c) that $(\pi(t), e) \in L\left(\alpha_{s_{i}, s^{\prime \prime}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$. Using the fact that $\pi^{\prime}(t)=\pi(t)$ and $\pi^{\prime}(y)=e$, we can infer that $\left(\pi^{\prime}(t), \pi^{\prime}(y)\right) \in L\left(\alpha_{s_{i}, s^{\prime \prime}} \mathcal{I}_{\mathcal{T}, \mathcal{A}}\right.$, hence $\pi^{\prime}$ satisfies the atom $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$.
- Case 3: $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ replaces $\alpha_{s_{i}, s_{i+1}}(y, y)$. Then $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$ must have the form $\alpha_{s^{\prime}, s^{\prime \prime}}(y, y)$, where $s^{\prime}$ is the state from $5(\mathrm{~b})$ and $s^{\prime \prime}$ is the state from $5(\mathrm{c})$. We have $\left(\pi^{\prime}(y), \pi^{\prime}(y)\right)=$ $(e, e) \in L\left(\alpha_{s^{\prime}, s^{\prime \prime}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, which means that $\pi^{\prime}$ satisfies the atom $\gamma\left(t^{\prime}, t^{\prime \prime}\right)$.

As we have shown that every atom in $q^{\prime}$ is satisfied by the mapping $\pi^{\prime}$, it follows that $\pi^{\prime}$ is a match for $q^{\prime}$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, which completes the proof.

Lemma 6.5. If $\vec{a} \in \operatorname{cert}\left(q^{\prime},(\mathcal{T}, \mathcal{A})\right)$ for some $q^{\prime} \in \operatorname{Rewrite}(q, \mathcal{T})$, then $\vec{a} \in \operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))$.
Proof. It is sufficient to show that if $q^{\prime} \in \operatorname{OneStep}(q, \mathcal{T})$ and $\vec{a} \in \operatorname{cert}\left(q^{\prime},(\mathcal{T}, \mathcal{A})\right)$, then $\vec{a} \in$ $\operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))$. Fix a C2RPQ $q$ and a DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H} \operatorname{TBox} \mathcal{T}$, and let $q^{\prime} \in \operatorname{OneStep}(q, \mathcal{T})$ be such that $\vec{a} \in \operatorname{cert}\left(q^{\prime},(\mathcal{T}, \mathcal{A})\right)$. By Lemma 3.2, there exists a match $\pi^{\prime}$ for $q^{\prime}$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ such that $\pi^{\prime}(\vec{x})=\vec{a}$, with $\vec{x}$ the answer variables of $q^{\prime}$.

Consider the execution of $\operatorname{OneStep}(q, \mathcal{T})$ that leads to the query $q^{\prime}$ being output. Let Leaf be the non-empty subset of $\operatorname{qvars}(q)$ that was selected in Step 1, let $y$ be the variable from Leaf chosen in Step 1, let $C \in \mathrm{TC}_{\mathcal{T}}$ be the concept selected in Step 2, and let $D \in \mathrm{BC}_{\mathcal{T}}$ and $R$ be the concept and role selected in Step 5 . In Step 7 , if $D \in \mathrm{~N}_{\mathrm{C}}$, then $D(y)$ was added, and if $D=\exists P^{-}$, then $A_{P^{-}}(y)$ was added. In the former case, we know that $\pi^{\prime}(y) \in D^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ and that either (i) $\mathcal{T} \models D \sqsubseteq \exists R$ and $C=\exists R^{-}$, or (ii) $\mathcal{T} \models D \sqsubseteq \exists R . C$, hence there must exist an $R$-child $e$ of $\pi^{\prime}(y)$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\operatorname{Tail}(e)=C$. In the latter case, we have $\pi^{\prime}(y) \in A_{P-}^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. By our assumption on DL-Lite $\mathcal{R}_{\mathcal{R}}$ TBoxes, $\mathcal{T}$ must contain the inclusion $A_{P-} \sqsubseteq \exists P^{-}$. Since $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ is a model of $\mathcal{T}$, this yields $\pi^{\prime}(y) \in\left(\exists P^{-}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, and hence $\pi^{\prime}(y) \in D^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. We then use the fact that $\mathcal{T} \models D \sqsubseteq \exists R$ where $C=\exists R^{-}$to find an $R$-child $e$ of $\pi^{\prime}(y)$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\operatorname{Tail}(e)=C$.

We define a mapping $\pi$ : terms $(q) \rightarrow \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ by setting $\pi(t)=e$ for every $t \in$ Leaf and setting $\pi(t)=\pi^{\prime}(t)$ for every $t \in \operatorname{terms}\left(q^{\prime}\right) \backslash\{y\}$. This mapping is well-defined since every term in $q$ either belongs to Leaf or appears in $q^{\prime}$. Observe that $\pi(\vec{x})=\vec{a}$ since $\pi^{\prime}(\vec{x})=\vec{a}$ and Leaf does not contain answer variables. We aim to show that $\pi$ is a match for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. To this end, consider some concept atom $B(t) \in q$. First suppose that $t \in$ Leaf. Then we know that the concept $C$ selected in Step 2 is such that $\mathcal{T} \models C \sqsubseteq B$. We then use the fact that since $t \in$ Leaf, we have $\pi(t)=e \in C^{\mathcal{L}_{\mathcal{T}, \mathcal{A}}}$. If $t \notin$ Leaf, then $B(t) \in q^{\prime}$. As $\pi^{\prime}$ is a match for $q^{\prime}$, we have $\pi^{\prime}(t) \in B^{\mathcal{I}_{\mathcal{I}}, \mathcal{A}}$. Since $\pi^{\prime}(t)=\pi(t)$, we get $\pi(t) \in B^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$.

Now consider some atom of the form $\alpha\left(t, t^{\prime}\right) \in q$, where $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$. If both $t \notin$ Leaf and $t^{\prime} \notin$ Leaf, then it can be verified that $\alpha\left(t, t^{\prime}\right) \in q^{\prime}$. As $\pi^{\prime}$ is a match for $q^{\prime}$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, it must be the case that $\left(\pi^{\prime}(t), \pi^{\prime}\left(t^{\prime}\right)\right) \in L(\alpha)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. Since $\pi^{\prime}(t)=\pi(t)$ and $\pi^{\prime}(t)=\pi(t)$, the same holds for $\pi$. Let us next consider the more interesting case in which $\left\{t, t^{\prime}\right\} \cap$ Leaf $\neq \emptyset$. In Step 3, we have a query containing $\alpha\left(\sigma(t), \sigma\left(t^{\prime}\right)\right)$, where the mapping $\sigma$ is defined as follows: $\sigma\left(t^{\prime \prime}\right)=t^{\prime \prime}$ for $t^{\prime \prime} \notin$ Leaf and $\sigma\left(t^{\prime \prime}\right)=y$ for $t^{\prime \prime} \in$ Leaf. Note that since $\left\{t, t^{\prime}\right\} \cap$ Leaf $\neq \emptyset$, at least one of $\sigma(t)$ and $\sigma\left(t^{\prime}\right)$ must be $y$. It follows that in Step 3, we will guess a sequence $s_{1}, \ldots, s_{n-1}$ of distinct states from $S$ and a state $s_{n} \in F$, and we will replace $\alpha\left(\sigma(t), \sigma\left(t^{\prime}\right)\right)$ by the atoms: $\alpha_{s_{0}, s_{1}}(\sigma(t), y), \alpha_{s_{1}, s_{2}}(y, y), \ldots, \alpha_{s_{n-2}, s_{n-1}}(y, y)$, $\alpha_{s_{n-1}, s_{n}}\left(y, \sigma\left(t^{\prime}\right)\right)$. Let us denote this set of atoms by $Q_{\alpha}$. Slightly abusing terminology, we use the phrase match for $Q_{\alpha}$ to refer to a match for the Boolean query given by the conjunctions of all atoms in $Q_{\alpha}$. We now establish the following claim:
Claim 1. If $\pi$ is a match for $Q_{\alpha}$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, then $\pi$ is a match for $\alpha\left(t, t^{\prime}\right)$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$.
Proof of claim. Suppose that $\pi$ is a match for the atoms in $Q_{\alpha}$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. This means that

- $(\pi(\sigma(t)), \pi(y)) \in L\left(\alpha_{s_{0}, s_{1}}\right) \mathcal{I}_{\mathcal{T}, \mathcal{A}}$,
- $(\pi(y), \pi(y)) \in L\left(\alpha_{s_{i}, s_{i+1}}\right)^{\mathcal{L}_{\mathcal{T}, \mathcal{A}}}$ for every $1 \leq i<n-1$, and
- $\left(\pi(y), \pi\left(\sigma\left(t^{\prime}\right)\right)\right) \in L\left(\alpha_{s_{n-1}, s_{n}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$.

We then remark that the language consisting of all words $w_{1} \ldots w_{n}$ such that $w_{1} \in L\left(\alpha_{s_{0}, s_{1}}\right)$, $w_{i} \in L\left(\alpha_{s_{i}, s_{i+1}}\right)$ for every $1 \leq i<n-1$, and $w_{n} \in L\left(\alpha_{s_{n-1}, s_{n}}\right)$ is a subset of the language $L\left(\alpha_{s_{0}, s_{n}}\right)$, and hence of $L(\alpha)$. Thus, by composing the paths witnessing each of the statements in the preceding list, we can show that $\left(\pi\left(\sigma(t), \pi\left(\sigma\left(t^{\prime}\right)\right) \in L(\alpha)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}\right.\right.$. Then to complete the proof, we simply note that $\pi(\sigma(t))=\pi(t)$ and $\pi\left(\sigma\left(t^{\prime}\right)\right)=\pi(t)$, because of the way we defined $\pi$ and $\sigma$. (end proof of claim)

Because of Claim 1, to complete the proof that $\pi$ is a match for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, it is sufficient to show that $\pi$ is a match for $Q_{\alpha}$, which is established by the following claim:

Claim 2. For every $\alpha_{s, s^{\prime}}\left(u, u^{\prime}\right) \in Q_{\alpha}:\left(\pi(u), \pi\left(u^{\prime}\right)\right) \in L\left(\alpha_{s, s^{\prime}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$.
Proof of claim. Take some $\alpha_{s, s^{\prime}}\left(u, u^{\prime}\right) \in Q_{\alpha}$. We start with the case where $u=u^{\prime}=y$ and $C \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$. As $\pi(y)=e$ and Tail $(e)=C$, we have Tail $(\pi(y))=C$. We can thus apply Proposition 5.5 to infer that $(\pi(y), \pi(y)) \in L\left(\alpha_{s, s^{\prime}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, which yields the desired result given that $u=u^{\prime}=y$. Next suppose that either $u \neq y, u^{\prime} \neq y$, or $C \notin \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$. Then we will not remove $\alpha_{s, s^{\prime}}\left(u, u^{\prime}\right)$ in Step 4, so it will still be present in Step 5. There are three cases depending on which of $u$ and $u^{\prime}$ equals $y$. We treat each case separately:

- Case 1: $u=y$ and $u^{\prime} \neq y$. It follows that $\pi(u)=e$ and $\pi\left(u^{\prime}\right)=\pi^{\prime}\left(u^{\prime}\right)$. In Step 6, we will replace $\alpha_{s, s^{\prime}}\left(u, u^{\prime}\right)$ with $\alpha_{s^{\prime \prime}, s^{\prime}}\left(u, u^{\prime}\right)$ where $s^{\prime \prime} \in S$ is such that $\left(s, U^{-}, s^{\prime \prime}\right) \in \delta$ for some $U \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$with $\mathcal{T} \models R \sqsubseteq U$. Using the facts that $\left(\pi^{\prime}(y), e\right) \in R^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}, \pi(u)=e, \mathcal{T} \models$ $R \sqsubseteq U$, and $\left(s, U^{-}, s^{\prime \prime}\right) \in \delta$, we can infer that $\left(\pi(u), \pi^{\prime}(y)\right) \in L\left(\alpha_{s, s^{\prime \prime}}\right)^{\mathcal{I}}, \mathcal{A}$. We also know that the atom $\alpha_{s^{\prime \prime}, s^{\prime}}\left(u, u^{\prime}\right)$ belongs to $q^{\prime}$, and so it must be satisfied by $\pi^{\prime}$, which yields $\left(\pi^{\prime}(u), \pi^{\prime}\left(u^{\prime}\right)\right) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime}}\right)^{\mathcal{T}}, \mathcal{A}$. By combining $\left(\pi(u), \pi^{\prime}(y)\right) \in L\left(\alpha_{s, s^{\prime \prime}} \mathcal{I}^{\mathcal{T}}, \mathcal{A}\right.$ and $\left(\pi^{\prime}(u), \pi^{\prime}\left(u^{\prime}\right)\right) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$, and using the fact that $\left(\pi^{\prime}(u), \pi^{\prime}\left(u^{\prime}\right)\right)=\left(\pi^{\prime}(y), \pi\left(u^{\prime}\right)\right)$, we can infer that $\left(\pi(u), \pi\left(u^{\prime}\right)\right) \in L\left(\alpha_{s, s^{\prime}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$.
- Case 2: $u \neq y$ and $u^{\prime}=y$. It follows that $\pi(u)=\pi^{\prime}(u)$ and $\pi\left(u^{\prime}\right)=e$. In Step 6 , we will replace $\alpha_{s, s^{\prime}}\left(u, u^{\prime}\right)$ with an atom $\alpha_{s, s^{\prime \prime}}\left(u, u^{\prime}\right)$ where $s^{\prime \prime} \in S$ is such that $\left(s^{\prime \prime}, U, s^{\prime}\right) \in \delta$ for some $U \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$with $\mathcal{T} \models R \sqsubseteq U$. Using similar arguments to Case 1, we can show that $\left(\pi(u), \pi^{\prime}(y)\right) \in L\left(\alpha_{s, s^{\prime \prime}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$ and $\left(\pi^{\prime}(y), \pi\left(u^{\prime}\right)\right) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$, from which we can deduce that $\left(\pi(u), \pi\left(u^{\prime}\right)\right) \in L\left(\alpha_{s, s^{\prime}}\right)^{\mathcal{I}}, \mathcal{A}$.
- Case 3: $u=u^{\prime}=y$. It follows that $\pi(u)=\pi\left(u^{\prime}\right)=e$. In Step 6, we will replace $\alpha_{s, s^{\prime}}\left(u, u^{\prime}\right)$ with an atom $\alpha_{s^{\prime \prime}, s^{\prime \prime \prime}}\left(u, u^{\prime}\right)$ where $s^{\prime \prime}, s^{\prime \prime \prime}$ are such that $\left(s, U^{-}, s^{\prime \prime}\right) \in \delta$ and $\left(s^{\prime \prime}, U^{\prime}, s^{\prime}\right) \in \delta$ for some roles $U, U^{\prime} \in \mathrm{N}_{\mathrm{R}}^{ \pm}$with $\mathcal{T} \models R \sqsubseteq U$ and $\mathcal{T} \models R \sqsubseteq U^{\prime}$. By applying similar reasoning to that used in Cases 1 and 2, we can show that $\left(\pi(u), \pi^{\prime}(y)\right) \in$ $L\left(\alpha_{s, s^{\prime \prime}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}},\left(\pi^{\prime}(y), \pi^{\prime}(y)\right) \in L\left(\alpha_{s^{\prime \prime}, s^{\prime \prime \prime}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, and $\left(\pi^{\prime}(y), \pi\left(u^{\prime}\right)\right) \in L\left(\alpha_{s^{\prime \prime \prime}, s^{\prime}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$. From this, we can infer that $\left(\pi(u), \pi\left(u^{\prime}\right)\right) \in L\left(\alpha_{s, s^{\prime}}\right)^{\mathcal{I} \mathcal{T}, \mathcal{A}}$.

Together, Lemmas 6.4 and 6.5 establish Proposition 6.3. We remark that the number of possible atoms appearing in the queries in $\operatorname{Rewrite}(q, \mathcal{T})$ is polynomially bounded by $|q|+|\mathcal{T}|$. This is the key property used to show the following:

Proposition 6.6. There are only exponentially many queries in $\operatorname{Rewrite}(q, \mathcal{T})$, each having size polynomial in $|q|+|\mathcal{T}|$.

Proof. Consider a DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E L H} \operatorname{TBox} \mathcal{T}$, a $\operatorname{C2RPQ} q$, and $q^{\prime} \in \operatorname{Rewrite}(q, \mathcal{T})$. We first note that OneStep never introduces any fresh variables, so vars $\left(q^{\prime}\right) \subseteq \operatorname{vars}(q)$. We next note that OneStep does not introduce any fresh concept names, and it only introduces binary atoms whose NFAs are obtained from one of the original NFAs by changing the initial and finite states. Thus, every atom in $q^{\prime}$ takes one of the following forms:

- $A(v)$, where $A \in \mathrm{~N}_{\mathrm{C}} \cap \operatorname{sig}(\mathcal{T})$ and $v \in \operatorname{vars}(q)$
- $\alpha_{s, s^{\prime}}\left(v, v^{\prime}\right)$, where $\alpha$ appears in $q, s, s^{\prime}$ are states of $\alpha$, and $v, v^{\prime} \in \operatorname{vars}(q)$

It is easy to see that the number of such atoms is bounded polynomially in $|\mathcal{T}|+|q|$, and thus, there are at most single-exponentially many distinct queries in $\operatorname{Rewrite}(q, \mathcal{T})$.

### 6.2 Query Evaluation

In Figure 16, we present a non-deterministic algorithm EvalQuery for C2RPQ answering in $\mathrm{DL}^{-L i t e}{ }_{\mathcal{R}}$ and $\mathcal{E L H}$. The algorithm starts by checking whether the input KB is satisfiable. If the check succeeds, then the algorithm guesses a query from $\operatorname{Rewrite}(q, \mathcal{T})$ and a variable assignment and calls the evaluation algorithm ${ }^{3}$ EvalAtom from Section 5 to check whether this assignment yields a match for this query in $\mathcal{I}_{\mathcal{K}}$. The following proposition establishes the correctness of EvalQuery.

Proposition 6.7. For every $C 2 R P Q ~ q$, DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H} K B K=(\mathcal{T}, \mathcal{A})$, and tuple of individuals $\vec{a}$ from $\operatorname{Ind}(\mathcal{A})$ of same arity as $q: \vec{a} \in \operatorname{cert}(q, \mathcal{K})$ if and only if there is some execution of $\operatorname{EvalQuery}(q, \mathcal{K}, \vec{a})$ that returns yes.
3. As we check KB satisfiability in the first step of EvalQuery, we may skip the satisfiability checks in the calls to EvalAtom.

## Algorithm EvalQuery

Input: $\mathrm{C} 2 \operatorname{RPQ} q\left(x_{1}, \ldots, x_{k}\right)$, DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H} \mathrm{KB} \mathcal{K}=\langle\mathcal{T}, \mathcal{A}\rangle$, tuple $\vec{a} \in \operatorname{Ind}(\mathcal{A})^{k}$

1. Test whether $\mathcal{K}$ is satisfiable, output yes if not.
2. Guess some $q^{\prime} \in \operatorname{Rewrite}(q, \mathcal{T})$ and an assignment $\vec{b}$ of individuals to quars $\left(q^{\prime}\right)$.
(a) Let $q^{\prime \prime}$ be the query obtained by substituting $\vec{a}$ for $\left(x_{1}, \ldots, x_{k}\right)$ and $\vec{b}$ for qvars $\left(q^{\prime}\right)$, then replacing each atom of the form $B(a)$ by the atom $\alpha_{B ?}(a)$ where $\alpha_{B}$ ? is an NFA with $L\left(\alpha_{B}\right.$ ) $=\{B ?\}$ consisting on an initial state $s_{0}^{B}$, a single final state $s_{f}^{B}$, and a single transition $\left(s_{0}^{B}, B\right.$ ?, $\left.s_{f}^{B}\right)$.
(b) For every atom $\alpha(a, b)$ in $q^{\prime \prime}$
check if $\operatorname{EvalAtom}(\alpha, \mathcal{K},(a, b))=$ yes
(c) If all checks succeed, return yes.

## 3. Return no.

Figure 16: Non-deterministic C2RPQ answering algorithm EvalQuery.

Proof. Let $\vec{x}$ be the set of answer variables of $q$.
To show the first direction, consider an execution of EvalQuery on input $(q, \mathcal{K}, \vec{a})$ that returns yes. If the algorithm returns yes in Step 1, then $\mathcal{K}$ is unsatisfiable, so we trivially have $\vec{a} \in \operatorname{cert}(q, \mathcal{K})$. Otherwise, in Step 2, the algorithm will guess a query $q^{\prime} \in \operatorname{Rewrite}(q, \mathcal{T})$ and assignment $\vec{b}$ for the quantified variables $\vec{y}$ of $q^{\prime}$. Let $q^{\prime \prime}$ be the query obtained by substituting $\vec{a}$ for $\vec{x}$ and $\vec{b}$ for $\vec{y}$, and writing all atoms in the form $\alpha(a, b)$ with $\alpha$ an NFA. In Step 2(c), for every atom $\alpha(a, b)$ in $q^{\prime \prime}$, we call EvalAtom on input $(\alpha, \mathcal{K},(a, b))$. Since the algorithm returns yes in Step 3, it must be the case that all of these calls return yes, and so by Proposition $5.8,(a, b) \in \operatorname{cert}(q, \mathcal{K})$ for every atom $\alpha(a, b)$ in $q^{\prime \prime}$. It follows that the mapping $\pi$ sending $\vec{x}$ to $\vec{a}$ and $\vec{y}$ to $\vec{b}$ defines a match for $q^{\prime}$ in $\mathcal{I}_{\mathcal{K}}$, so $\vec{a} \in \operatorname{cert}\left(q^{\prime}, \mathcal{K}\right)$. Applying Proposition 6.3, we obtain $\vec{a} \in \operatorname{cert}(q, \mathcal{K})$.

Next suppose that $\vec{a} \in \operatorname{cert}(q, \mathcal{K})$. If $\mathcal{K}$ is unsatisfiable, then the algorithm will return yes in Step 1. Otherwise, we know from Proposition 6.3 that there exists a query $q^{\prime} \in$ $\operatorname{Rewrite}(q, \mathcal{T})$ and a match $\pi$ for $q^{\prime}$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ that maps all variables to $\operatorname{Ind}(\mathcal{A})$ and is such that $\pi(\vec{x})=\vec{a}$. In Step 2, we choose the query $q^{\prime}$ and tuple $\pi(\vec{y})$, where $\vec{y}$ is the set of quantified variables in $q^{\prime}$. Let $q^{\prime \prime}$ be the query obtained by substituting $\vec{a}$ for $\vec{x}$ and $\vec{b}$ for $\vec{y}$, and writing all atoms in the form $\alpha(a, b)$ with $\alpha$ an NFA. Because of the match $\pi$ we know that $(a, b) \in \operatorname{cert}(q, \mathcal{K})$ for every atom $\alpha(a, b)$ in $q^{\prime \prime}$. It follows then by Proposition 5.8 that all of the calls to EvalAtom will return yes, and EvalQuery will return yes in Step 2(c).

By analyzing the complexity of the algorithm EvalQuery, we obtain the following upper bounds for C2RPQ answering, which match the lower bounds given in Section 4.

Theorem 6.8. C2RPQ answering is in

1. NL in data complexity for DL-Lite $\mathcal{R}_{\mathcal{R}}$ and DL-Lite $\mathrm{RDFS}^{\text {; }}$
2. P in data complexity for $\mathcal{E} \mathcal{L H}$;
3. NP in combined complexity for DL-Lite $e_{\text {RDS }}$;

## 4. PSpace in combined complexity for $D L-$ Lite $_{\mathcal{R}}$ and $\mathcal{E} \mathcal{L H}$.

Proof. For Statement (1), we consider the resources required to run EvalQuery on input $(\mathcal{T}, \mathcal{A}, q)$, where $\mathcal{T}$ is formulated in DL-Lite $\mathcal{R}_{\mathcal{R}}$. The consistency check in Step 1 can be performed in non-deterministic logarithmic space in $|\mathcal{A}|$ (Calvanese et al., 2007). If $\mathcal{T}$ and $q$ are treated as fixed, then computing $\operatorname{Rewrite}(q, \mathcal{T})$ requires only constant time (and space) in $|\mathcal{A}|$. It follows that the guessed query $q^{\prime}$ and tuple $\vec{b}$ and the set of atoms at the end of Step 2(a) can be stored using logarithmic space in $|\mathcal{A}|$. In Step 2(b), we call EvalAtom on each of the stored atoms. By Theorem 5.9, EvalAtom runs in non-deterministic logarithmic space in $|\mathcal{A}|$. Since $\mathrm{NL}^{\text {NL }}=\mathrm{NL}$, we obtain the desired NL upper bound in data complexity.

To show Statement 2, consider what happens if the input TBox $\mathcal{T}$ is formulated in $\mathcal{E L H}$. In this case, the consistency check in Step 1 takes constant time (since every $\mathcal{E L H} \mathrm{KB}$ is satisfiable), and by Theorem 5.9, EvalAtom runs in polynomial time in $|\mathcal{A}|$. We thus have a decision procedure for C 2 RPQ answering in $\mathcal{E L H}$ that runs in non-deterministic logarithmic space in $|\mathcal{A}|$ with access to a P oracle. Since $\mathrm{NL}^{\mathrm{P}}=\mathrm{P}$, this yields membership in P for data complexity.

To establish Statement (3), we first note that if $\mathcal{T}$ is a DL-LiterdFs TBox, then the query cannot be rewritten, i.e., $\operatorname{Rewrite}(q, \mathcal{T})=\{q\}$ (we cannot choose a $D$ as required in Step 5 of Algorithm OneStep because $\mathcal{T}$ does not entail any inclusions of the form $D \sqsubseteq \exists R$ ). Thus, in Step 2 of EvalQuery, we only need to guess a tuple $\vec{b}$ whose size is polynomial in $|\mathcal{A}|$ and $|q|$. The calls to EvalAtom in Step 2(b) run in polynomial time in the input (Theorem 5.2 ), so the overall procedure runs in NP.

For Statement (4), instead of computing the whole set $\operatorname{Rewrite}(q, \mathcal{T})$, which can contain exponentially many queries, we generate a single $q^{\prime} \in \operatorname{Rewrite}(q, \mathcal{T})$ non-deterministically. By Proposition 6.6, every query in $\operatorname{Rewrite}(q, \mathcal{T})$ can be generated after at most exponentially many steps, so we can use a polynomial-size counter to check when we have reached this limit. Since each rewritten query is of polynomial size (Proposition 6.6), and we keep only one query in memory at a time, the generation of a single query in $\operatorname{Rewrite}(q, \mathcal{T})$ requires only polynomial space. We can then proceed as for statement 3 , guessing a (polynomialsize) tuple $\vec{b}$ and performing a polynomial number of polynomial-time evaluation checks (Theorem 5.9). This yields a non-deterministic polynomial space procedure for deciding $\vec{a} \in \operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))$. Using the well-known fact that NPSpace $=$ PSpace, we obtain the desired PSpace upper bound.

### 6.3 Cases with Lower Complexity

Given the substantial jump in combined complexity - from NP to PSPace - when moving from CQs to C(2)RPQs, it is natural to look for interesting subcases that offer lower complexity. We pinpoint two such subcases, the first obtained by restricting the query language, and the second obtained by restricting the class of KBs.

Let us recall that 2 RPQs are single-atom C2RPQs that do not contain quantified variables. The following theorem shows that this restriction is inessential, as our complexity results for 2 RPQs hold also for single-atom queries with quantified variables. ${ }^{4}$

Theorem 6.9. The results in Theorem 5.9 hold also for single-atom C2RPQs.
Proof. Fix a DL-Lite $\mathcal{R}$ or $\mathcal{E} \mathcal{L H} \mathrm{KB} \mathcal{K}=(\mathcal{T}, \mathcal{A})$. There are six types of single-atom queries to consider: $\exists y . \alpha(a, y)$ (with $a$ an individual), $\exists y . \alpha(x, y)$ (with $x$ an answer variable), $\exists y . \alpha(y, a)$ (with $a$ an individual), $\exists y . \alpha(y, x)$ (with $x$ an answer variable), $\exists x, y . \alpha(x, y)$ (with $x \neq y$ ), and $\exists y . \alpha(y, y)$. For the first five types of queries, there are simple reductions to 2RPQs. When $q=\exists x, y \cdot \alpha(x, y)$ (with $x \neq y$ ), we can simply replace $q$ by the 2RPQ $q^{\prime}(x, y)=\alpha^{\prime}(x, y)$, where

$$
L\left(\alpha^{\prime}\right)=\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*} \cdot L(\alpha) \cdot\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*}
$$

The following claim establishes the correctness of this reduction.
Claim. $\mathcal{K} \models q$ if and only if $\operatorname{cert}\left(q^{\prime}, \mathcal{K}\right) \neq \emptyset$.
Proof of claim. First suppose that $\mathcal{K} \models q$. If $\mathcal{K}$ is unsatisfiable, then we trivially have $\operatorname{cert}\left(q^{\prime}, \mathcal{K}\right) \neq \emptyset$. Otherwise, there is a match $\pi$ for $q$ in $\mathcal{I}_{\mathcal{K}}$. This means that $(\pi(x), \pi(y)) \in$ $L(\alpha)^{\mathcal{I}_{\mathcal{K}}}$, and so there must exist some path $p_{0}$ from $\pi(x)$ to $\pi(y)$ with $\lambda\left(p_{0}\right) \in L(\alpha)$. Let $a$ be such that $\pi(x)$ is either equal to $a$ or begins by $a$, and let $p_{1}$ be any path from $a$ to $\pi(x)$. Since $\pi(y)$ is reachable from $\pi(x)$, it must also be reachable from $a$, and so we can find a path $p_{2}$ from $\pi(y)$ to $a$. We may choose $p_{1}$ and $p_{2}$ so that $\lambda\left(p_{1}\right)$ and $\lambda\left(p_{2}\right)$ belong to ( $\left.\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*}$. By combining the paths $p_{1}, p_{0}$ and $p_{2}$ (in that order), we obtain a path from $a$ to $a$ whose label belongs to $\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*} \cdot L(\alpha) \cdot\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*}$. It follows that $(a, a) \in \operatorname{cert}\left(q^{\prime}, \mathcal{K}\right)$.

Suppose next that $(a, b) \in \operatorname{cert}\left(q^{\prime}, \mathcal{K}\right)$ for two (not necessarily distinct) individuals $a, b$. If $\mathcal{K}$ is unsatisfiable or $\varepsilon \in L(\alpha)$, then we trivially have $\mathcal{K} \models q$. Otherwise, there is a path $p=$ $e_{0} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ in $\mathcal{I}_{\mathcal{K}}$ with $e_{0}=a, e_{n}=b$, and $\lambda(p) \in\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*} \cdot L(\alpha) \cdot\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*}$. Since $\varepsilon \notin L(\alpha)$, there exists $0<i<j \leq n$ such that $u_{i} \ldots u_{j-1} \in L(\alpha)$. By setting $\pi(x)=e_{i-1}$ and $\pi(y) \in e_{j}$, we obtain a match for $q$ in $\mathcal{I}_{\mathcal{K}}$. (end proof of claim)
The other four types of queries containing distinct terms can be handled similarly:

- For $\exists y . \alpha(x, y)$, we use the $2 \mathrm{RPQ} q^{\prime \prime}(x, y)=\alpha^{\prime \prime}(x, y)$, where $L\left(\alpha^{\prime \prime}\right)=L(\alpha) \cdot\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap\right.$ $\operatorname{sig}(\mathcal{T}))^{*}$. Arguing as in the claim, we can show that $b \in \operatorname{cert}(\exists y . \alpha(x, y), \mathcal{K})$ iff $(b, c) \in$ $\operatorname{cert}\left(q^{\prime \prime}(x, y), \mathcal{K}\right)$ for some $c \in \operatorname{Ind}(\mathcal{A})$.
- For $\exists y \cdot \alpha(a, y)$, we have $\mathcal{K} \models \exists y \cdot \alpha(a, y)$ iff $a \in \operatorname{cert}(\exists y \cdot \alpha(x, y), \mathcal{K})$, so we can reuse the 2 RPQ from the preceding point.
- For $\exists y . \alpha(y, x)$, we use the $2 \mathrm{RPQ} q^{\prime \prime \prime}(x, y)=\alpha^{\prime \prime \prime}(y, x)$ with $L\left(\alpha^{\prime \prime \prime}\right)=\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*}$. $L(\alpha)$. Using a similar argument to that used in the preceding claim, we can show that $b \in \operatorname{cert}(\exists y . \alpha(y, x), \mathcal{K})$ iff $(b, c) \in \operatorname{cert}\left(q^{\prime \prime \prime}(x, y), \mathcal{K}\right)$ for some $c \in \operatorname{Ind}(\mathcal{A})$.

4. In the preliminary version of this paper, we in fact used the more general notion of single-atom queries (possibly with quantified variables) as the definition of 2RPQs (Bienvenu et al., 2013).

- For $\exists y \cdot \alpha(y, a)$, we have $\mathcal{K} \models \exists y \cdot \alpha(y, a)$ iff $a \in \operatorname{cert}(\exists y \cdot \alpha(y, x), \mathcal{K})$, so we can reuse the 2 RPQ from the preceding point.

For queries of the form $\exists x . \alpha(x, x)$, the proof is more involved and passes by the definition of an alternative rewriting procedure for 2 RPQs, which is similar in spirit to the algorithm Rewrite but guaranteed to run in polynomial time. Details are given in the appendix.

More interestingly, we can adapt the techniques from the preceding proof in order to provide an NP upper bound for the class of C2RPQ s that do not contain any existential-join variables, i.e., existentially quantified variables that occur more than once in a query.

Theorem 6.10. C2RPQ answering is in NP in combined complexity for DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$ knowledge bases when restricted to C2RPQs without existential-join variables.

Proof. Consider a DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E L H} \mathrm{KB} \mathcal{K}=(\mathcal{T}, \mathcal{A})$ and a $\mathrm{C} 2 \mathrm{RPQ} q$ with answer variables $\vec{x}$ and existential variables $\vec{y}$ such that every variable in $\vec{y}$ occurs exactly once in $q$. We let $q^{\prime}$ be the C2RPQ obtained from $q$ as follows:

- Replace every atom $\alpha\left(y_{1}, y_{2}\right)$ such that $y_{1}, y_{2}$ are both existential variables by the atom $\alpha^{\prime}\left(y_{1}, y_{2}\right)$ where $L\left(\alpha^{\prime}\right)=\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*} \cdot L(\alpha) \cdot\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*}$.
- Replace every atom $\alpha(t, y)$ such that $y$ is an existential variable and $t$ is an individual or answer variable by the atom $\alpha^{\prime \prime}(t, y)$ where $L\left(\alpha^{\prime \prime}\right)=L(\alpha) \cdot\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*}$.
- Replace every atom $\alpha(y, t)$ such that $y$ is an existential variable and $t$ is an individual or answer variable by the atom $\alpha^{\prime \prime \prime}(y, t)$ where $L\left(\alpha^{\prime \prime \prime}\right)=\left(\mathrm{N}_{\mathrm{R}}^{ \pm} \cap \operatorname{sig}(\mathcal{T})\right)^{*} \cdot L(\alpha)$.
- Add $\vec{y}$ to the set of answer variables.

Clearly, it takes only polynomial time to construct the C2RPQ $q^{\prime}$. By exploiting the fact that every existential variable in $q$ occurs at most once, and applying similar reasoning to that used in the proof of Theorem 6.9, we can show that $\vec{a} \in \operatorname{cert}(q, \mathcal{K})$ iff $(\vec{a}, \vec{b}) \in \operatorname{cert}\left(q^{\prime}, \mathcal{K}\right)$ for some tuple of individuals $\vec{b}$ of the same arity as $\vec{y}$. To decide whether the latter holds, we non-deterministically guess a tuple $\vec{b}$ and let $\pi$ be the variable assignment that maps $\vec{x}$ to $\vec{a}$ and $\vec{y}$ to $\vec{b}$. We then use EvalAtom to verify that $\left.\left(\pi\left(t_{1}\right), \pi()_{2}\right)\right) \in \operatorname{cert}\left(\alpha\left(t_{1}, t_{2}\right), \mathcal{K}\right)$ for every atom $\alpha\left(t_{1}, t_{2}\right) \in q^{\prime}$, and we return yes if this is the case. Correctness of the described procedure follows from the correctness of EvalAtom (Proposition 5.8). Since EvalAtom can be implemented so as to run in polynomial time for both DL-Lite $\mathcal{R}$ and $\mathcal{E} \mathcal{L H}$ (Theorem 5.9), we obtain an NP procedure for answering C2RPQs without existential-join variables.

The preceding result can be extended a bit further to allow for simple chains of existential variables. Indeed, we remark that if a C2RPQ $q$ contains atoms $\alpha_{1}(x, y)$ and $\alpha_{2}(y, z)$ with $z$ an existential-join variable and $y$ an existential variable appearing only in these two atoms, then we can replace $\left\{\alpha_{1}(x, y), \alpha_{2}(y, z)\right\}$ by $\alpha_{3}(x, z)$ where $L\left(\alpha_{3}\right)=\left\{w_{1} w_{2} \mid w_{1} \in\right.$ $\left.L\left(\alpha_{1}\right), w_{2} \in L\left(\alpha_{2}\right)\right\}$ (such an NFA $\alpha_{3}$ is easily constructed from $\alpha_{1}, \alpha_{2}$ in polynomial time). By performing this and similar polynomial-time equivalence-preserving transformations, we can eliminate some existential-join variables and thereby enlarge the class of C2RPQs that can be handled using an NP procedure.

Finding further interesting classes of computationally well-behaved queries will likely prove difficult, given that the PSpace lower bound in Section 4 was shown to hold even under strong structural restrictions on C2RPQs. This suggests that it may be more fruitful to consider restrictions on knowledge bases. The next proposition identifies a natural restriction on knowledge bases that leads to an improved NP upper bound.

Theorem 6.11. C2RPQ answering is in NP in combined complexity for DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$ knowledge bases whose canonical models have finite domains.

Proof. Let $\mathcal{K}=(\mathcal{T}, \mathcal{A})$ be a satisfiable DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H}$ knowledge base whose canonical model $\mathcal{I}_{\mathcal{K}}$ contains only finitely many elements (note that if $\mathcal{K}$ is unsatisfiable, then it is trivial to perform query answering). It follows the construction of $\mathcal{I}_{\mathcal{K}}$ and the fact that $\Delta^{\mathcal{I}_{\mathcal{K}}}$ is finite that every element $a R_{1} C_{1} \ldots R_{n} C_{n} \in \Delta^{\mathcal{I}_{\mathcal{K}}}$ is such that $C_{i} \neq C_{j}$ for $i \neq j$, which implies in particular that $n \leq\left|\mathrm{TC}_{\mathcal{T}}\right|$ (indeed, if $C_{i}=C_{j}$ for $i<j$, then the domain of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ would contain the element $a R_{1} C_{1} \ldots R_{j} C_{j}\left(R_{i+1} C_{i+1} \ldots R_{j} C_{j}\right)^{m}$ for every $\left.m \geq 0\right)$. We rely on this property to devise a non-deterministic algorithm for deciding $\vec{a} \in \operatorname{cert}(q, \mathcal{K})$ that runs in polynomial time in the combined size of the inputs.

Without loss of generality, we may suppose that the input query $q$ contains only binary atoms of the form $\alpha\left(t, t^{\prime}\right)$ with $\alpha$ an NFA. In a first step, we guess a mapping $\pi$ from the terms in $q$ to sequences of the form $a R_{1} C_{1} \ldots R_{n} C_{n}$ where each $R_{i}$ is a role, each $C_{i}$ is a concept from $\mathrm{TC}_{\mathcal{T}}$, and $0 \leq n \leq\left|\mathrm{TC} \mathcal{T}_{\mathcal{T}}\right|$. In a second step, we verify that $\pi$ is indeed a match for $q$. First, we check that $\pi(b)=b$ for every individual $b$ in $q$, and that $\pi(\vec{x})=\vec{a}$ where $\vec{x}$ is the tuple of answer variables of $q$. Next, we check whether $\pi(z) \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ for every $z \in \operatorname{qvars}(q)$. This can be done with a polynomial number of (polynomial-time) entailment checks that verify conditions (A) and (B) of the definition of canonical models in Section 2.1.4. It only remains to check that all of the query atoms are satisfied under the mapping $\pi$. To this end, we construct a new ABox $\mathcal{A}_{\pi}$ as follows. Let $E_{\pi}$ be the set containing each $a R_{1} C_{1} \ldots R_{j} C_{j}$ with $1 \leq j \leq n$ such that there is some $a R_{1} C_{1} \ldots R_{j} C_{j} \ldots R_{n} C_{n}$ in the range of $\pi$. We introduce a fresh individual name $b_{e}$ for each $e \in E_{\pi}$, and let $\mathcal{A}_{\pi}$ be the ABox obtained by adding to $\mathcal{A}$ the following assertions:
$-R\left(a, b_{e}\right)$ if $R \in \mathrm{~N}_{\mathrm{R}}$ and $R^{-}\left(b_{e}, a\right)$ if $R^{-} \in \mathrm{N}_{\mathrm{R}}$ for each $e=a R C \in E_{\pi}$ where $a \in \mathrm{~N}_{\mathrm{I}}$
$-R\left(b_{e}, b_{e^{\prime}}\right)$ if $R \in \mathrm{~N}_{\mathrm{R}}$ and $R^{-}\left(b_{e^{\prime}}, b_{e}\right)$ if $R^{-} \in \mathrm{N}_{\mathrm{R}}$, for each $e^{\prime}=e R C \in E_{\pi}$ where $e \in E_{\pi}$

- $C\left(b_{e^{\prime}}\right)$ for each $e^{\prime}=e R C \in E_{\pi}$ where $e \in E_{\pi} \cup \mathrm{N}_{\mathrm{I}}$ and $C \in \mathrm{~N}_{\mathrm{C}}$

Define the mapping $\pi^{\prime}$ from terms in $q$ to individuals in $\mathcal{A}_{\pi}$ as follows: $\pi^{\prime}(t)=\pi(t)$ if $\pi(t) \in \operatorname{Ind}(\mathcal{A})$ and $\pi^{\prime}(t)=b_{\pi(t)}$ otherwise. For every atom $\alpha\left(t, t^{\prime}\right)$ in $q$, we call EvalAtom on input $\left(\alpha,\left(\mathcal{T}, \mathcal{A}_{\pi}\right),\left(\pi^{\prime}(t), \pi^{\prime}\left(t^{\prime}\right)\right)\right)$; by Theorem 5.9 , each of these calls needs only polynomial time. We output yes just in the case that every call to EvalAtom returns yes.

It should be clear that the algorithm just described runs in non-deterministic polynomial time. To show that the algorithm is sound, consider an execution of the algorithm that returns yes, and let $\pi$ be the mapping that was guessed. Since the algorithm returned yes, we know that for every atom $\alpha\left(t, t^{\prime}\right)$ in $q$, there is a successful execution of EvalAtom on input $\left(\alpha,\left(\mathcal{T}, \mathcal{A}_{\pi}\right),\left(\pi^{\prime}(t), \pi^{\prime}\left(t^{\prime}\right)\right)\right)$. Since EvalAtom is known to be correct (Theorem 6.7), this shows $\pi^{\prime}$ is a match for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}_{\pi}}$. It follows from the way we defined $\mathcal{A}_{\pi}$ and the construction of canonical models that $\mathcal{I}_{\mathcal{T}, \mathcal{A}_{\pi}}$ can be homomorphically embedded into $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. Moreover, we can choose the homomorphism $h$ such that $h\left(b_{e}\right)=e$ for each of the new
individuals $b_{e}$ in $\mathcal{A}_{\pi}$. Since matches of C2RPQs are preserved under homomorphisms, it follows that $\pi$ is a match for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\pi(\vec{x})=\vec{a}$, so $\vec{a} \in \operatorname{cert}(q,(\mathcal{T}, \mathcal{A}))$.

To show completeness, suppose that $\pi$ is a match for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. By construction, there is a homomorphism $h$ from $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ to $\mathcal{I}_{\mathcal{T}, \mathcal{A}_{\boldsymbol{\pi}}}$ that maps every individual $a \in \operatorname{Ind}(\mathcal{A})$ to itself and every $e \in E_{\pi}$ to the individual $b_{e}$. Using the fact that query matches are preserved under homomorphisms, we have that $\pi^{\prime}$ is a match for $q$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}_{\pi}}$. It follows that for every atom $\alpha\left(t, t^{\prime}\right)$ in $q$, there is an execution of EvalAtom on $\left(\alpha,\left(\mathcal{T}, \mathcal{A}_{\pi}\right),\left(\pi^{\prime}(t), \pi^{\prime}\left(t^{\prime}\right)\right)\right.$ ) that returns yes, and so the algorithm returns yes after guessing the mapping $\pi$.

We point out that the class of knowledge bases considered in Theorem 6.11 is of practical relevance. Indeed, several important large-scale $\mathcal{E L H}$ terminologies, like the medical ontology SNOMED ${ }^{5}$, are acyclic terminologies (see, e.g., (Haase \& Lutz, 2008)), which are guaranteed to have finite canonical models. Moreover, it has been recently argued that real-world DL-Lite $\mathcal{R}_{\mathcal{R}}$ ontologies often yield canonical models of shallow depth (Kikot, Kontchakov, Podolskii, \& Zakharyaschev, 2013).

## 7. Beyond C2RPQs and Lightweight DLs

In this section, we discuss some of the implications of our results and give a brief overview of related results for similar settings.

### 7.1 Extensions of C2RPQs

We have argued in this paper that C2RPQs provide significantly more expressiveness than plain CQs as languages for querying ontologies, at a moderate computational cost. However, C2RPQs also have many limitations, and several application domains seem to call for even more expressive query languages. We now discuss some extensions to C2RPQs.

### 7.1.1 Complex Labels

In a preliminary version of the present work, we added to C2RPQs the ability to talk about combinations of concepts and roles that appear along a path (Bienvenu, Ortiz, \& Šimkus, 2012). In that language, which we called C2RPQs with complex labels, one can use, for example, the expression (sbLFT $\wedge$ sbLFSub)* to find paths between stations that are served by both low-floor tramway and subway. Complex labels provide a more concise and flexible syntax for many queries, and in fact they increase the expressiveness of C2RPQs. In particular in DLs that do not support role conjunction, like the ones considered here, with standard C2RPQs we can query for pairs of stations that are connected by two different means of transport, but we cannot require the existence of one route between them that is fully served by both. All the algorithms described in this paper can be extended straightforwardly to C2RPQs with complex labels, and the given complexity results apply also to that more expressive query language. However, the extension to complex labels causes a significant overhead in the notation and the technicalities of the algorithms. Hence, for the sake of readability, we decided not to include such an extension in this paper.
5. http://www.ihtsdo.org/snomed-ct/

### 7.1.2 RPQs with Nesting

Recent works in the database field advocate the extension of RPQs with nesting, allowing one to require that objects along a path satisfy complex conditions, in turn expressed through (nested) 2RPQs, in line with the XML query language XPath. RPQs with nesting were proposed as the basic component of the navigational language nSPARQL for RDF (Pérez, Arenas, \& Gutierrez, 2010) and they have received some attention in the setting of graph databases (Reutter, 2013; Barceló, Pérez, \& Reutter, 2012).

Building on the conference version of this work, we recently studied the query answering problem for nested (C)2RPQs in the presence of DL ontologies (Bienvenu, Calvanese, Ortiz, \& Šimkus, 2014). We establish tight complexity bounds in data and combined complexity for a variety of DLs, ranging from the lightweight DLs DL-Lite and $\mathcal{E L}$ considered in the present paper, up to highly expressive DLs. We show that adding nesting to (C)2RPQs does not increase the worst-case data complexity of query answering, but it leads to ExpTimehardness in combined complexity, even for (non-conjunctive) 2RPQs and the lightweight DLs DL-Lite or $\mathcal{E L}$. This contrasts sharply with the tractability result we obtained in this paper for the same setting but without nesting.

Other authors have considered nested navigational queries over DL KBs. Stefanoni, Motik, Krötzsch, and Rudolph (2014) studied the complexity of answering certain types of nested path queries over knowledge bases formulated in OWL 2 EL, which extends $\mathcal{E L H}$ with a number of constructs, most notably complex role inclusions. They establish PSpace membership for a query language that roughly corresponds to the nested (C)RPQs mentioned earlier extended with unary complex labels, and they show $P$ membership for the non-conjunctive fragment. These results demonstrate that nesting is computationally simpler when inverse roles are allowed neither in queries, nor in the ontology language. Kostylev, Reutter, and Vrgoc (2015) have recently investigated the complexity of the socalled DLXPath family of query languages over knowledge bases expressed in lightweight DLs, with a particular emphasis on the connection to propositional dynamic logic (PDL) and the effects of negation. The most expressive variant of DLXPath can be used to test Boolean conditions on nodes and edges (and thus fully captures the complex labels of the previous subsection), but unfortunately, answering such queries is undecidable even in the simplest of settings. Disallowing negation of binary relations restores decidability, but query answering remains coNP-hard in data complexity if negation of unary expressions is permitted. We finally note that Bourhis, Krötzsch, and Rudolph (2014) have recently explored several highly expressive extensions of RPQs with nesting for which query answering remains decidable in the presence of DL ontologies.

### 7.1.3 Path Variables and Path Relations

In the area of graph databases, it has been argued that C2RPQs are sometimes too weak since they can neither output the witnessing paths, nor talk about the relationships holding between different paths. This has motivated the recent extension of C2RPQs with path variables and relations among tuples of paths (Barceló et al., 2012; Barceló \& Muñoz, 2014). The introduction of path variables makes it possible to refer to the specific paths (or more precisely, the labels of those paths) that are used to witness the satisfaction of query atoms. By using a path variable in multiple atoms, one can enforce that paths with
the same label are used to connect different pairs of points. Moreover, path variables can appear as answer variables in the query, in which case a compact representation of the labels of witnessing paths is given as output (Barceló et al., 2012). A further extension allows queries to enforce that a tuple of path labels belongs to a given relation, either via regular relations such as prefix or equal length (Barceló et al., 2012), or using some common non-regular relations like subword and subsequence (Barceló \& Muñoz, 2014).

An in-depth study of these extensions in the presence of DL ontologies is ongoing work, but our preliminary results suggest that the addition of ontological knowledge makes things significantly harder. For instance, even in the presence of very simple DL-Lite ontologies, the labels of paths witnessing a query answer may form a non-regular language, as illustrated by the following example.

Example 7.1. Consider the DL-Lite KB $\mathcal{K}$ consisting of the TBox $\left\{A \sqsubseteq \exists R, \exists R^{-} \sqsubseteq \exists R\right\}$ and ABox $\{A(a)\}$, and let $q$ be the $2 \operatorname{RPQ} \mathcal{E}(x, y)$ with $\mathcal{E}=r^{*} \cdot\left(r^{-}\right)^{*}$. There are infinitely many paths witnessing that $(a, a)$ is an answer to $q$ in $\mathcal{I}_{\mathcal{K}}$, each obtained by taking $n$ steps away from $a$ via $r$, then $n$ steps back up to $a$ via $r^{-}$. It follows that the set of labels of the witnessing paths forms the non-regular language $\left\{r^{n} \cdot\left(r^{-}\right)^{n} \mid n \geq 0\right\}$.

The previous example crucially uses inverse roles, but non-regular (indeed, non-contextfree!) languages can also be enforced using shared path variables:

Example 7.2. Consider the $\mathcal{E} \mathcal{L}$ KB $\mathcal{K}$ consisting of the TBox $\{A \sqsubseteq \exists r . A, A \sqsubseteq \exists s . A\}$ and ABox $\{A(a)\}$. Let $q$ be the Boolean CRPQ $\exists x, y, z \cdot \mathcal{E}(x, y) \wedge \mathcal{E}(y, z) \wedge \mathcal{E}(x, z)$ with $\mathcal{E}=(r \cup s)^{*}$, and further suppose that we use path variables to require that the label of the path from $x$ to $y$ is the same as the label of the path from $y$ to $z$. Then every triple of paths $\left(p_{x y}, p_{y z}, p_{x z}\right)$ witnessing the satisfaction of $q$ is such that $\lambda\left(p_{x y}\right)=\lambda\left(p_{y z}\right)$ and $\lambda\left(p_{x z}\right)=\lambda\left(p_{x y}\right) \lambda\left(p_{y z}\right)$. It follows that the set of labels of witnessing paths for the third atom yields the language $\left\{w w \mid w \in L\left((r \cup s)^{*}\right)\right\}$, which is neither regular nor context-free.

While these examples may seem artificial, they highlight the difficulties that arise when combining path variables and ontologies. In particular, they demonstrate that we cannot use NFAs as a compact representation of path labels (as was the case for graph databases), and even if all path variables are existentially quantified, the sharing of path variables will likely require significant modification of our query answering algorithms. Interestingly, it seems that some of the techniques based on word equations with regular constraints that were recently explored for handling non-regular path relations (Barceló \& Muñoz, 2014) may be relevant in the presence of ontologies already for answering queries with (existentially quantified) path variables, and possibly also simple regular relations.

### 7.2 Other DLs

Our rewriting algorithm for C2RPQs is inspired by a technique first proposed for answering CQs in the DL Horn-SHIQ (Eiter et al., 2012), and it can be extended to all the constructs in that logic. In fact, a similar algorithm for nested C2RPQs was already developed for the DL $\mathcal{E L H} \mathcal{I}^{\perp}$ which subsumes both DL-Lite ${ }_{\mathcal{R}}$ and $\mathcal{E L H}$ (Bienvenu et al., 2014). $\mathcal{E L H} \mathcal{I}^{\perp}$ contains many of the constructors of Horn- $\mathcal{S H} \mathcal{I} \mathcal{Q}$, and they behave similarly in terms of computational complexity. We point out that transitive roles, which are often problematic
for query answering algorithms, are not a major issue in this setting. They can be easily accommodated in the $\mathcal{E L H} \mathcal{I}^{\perp}$ algorithm by combining the known techniques for axiomatizing in the TBox the transitivity of roles, and the machinery for handling the transitive closure constructor in the queries. The extension of our algorithm to $\mathcal{E L H} \mathcal{I}^{\perp}$ or Horn- $\mathcal{S H I Q}$ runs in polynomial time in the size of the data, and it is thus worst-case optimal in data complexity. Naturally, in combined complexity, it may require exponential time in some cases, but it always runs in single-exponential time, which again is worst-case optimal for these DLs. Moreover, we conjecture that if implemented smartly, this exponential behavior will rarely occur for real-world ontologies.

### 7.3 OWL 2 Profiles and Query Answering for the Semantic Web

The Web Ontology Language is a family of languages for specifying ontologies, endorsed as a standard by the W3C. The current version of the standard, called OWL 2 (OWL Working Group, 2009), features three profiles (Motik et al., 2012) or sublanguages that restrict the expressivity in such a way that logical inference over ontologies can be achieved by efficient algorithms. The three profiles provide different modeling capabilities, making them suitable for different applications: the EL profile is the preferred language for life science ontologies, the QL profile is geared towards applications that enrich relational data with ontological information, and the RL profile is used mostly for reasoning with Web data. For more information about the profiles, their modeling capabilities and supported inference services, we refer the reader to the introductory text by Krötzsch and references therein (Krötzsch, 2012). The $\mathcal{E L}$ and DL-Lite families of lightweight description logics studied in the present paper provide the logical underpinnings of the EL and QL profiles (the third profile, RL, is based upon Datalog). As a consequence, our results are immediately relevant to the problem of answering regular path queries over OWL 2 knowledge bases formulated using the QL and EL profiles. In particular, our algorithms can be adapted for querying datasets, such as RDF triplestores, enriched with ontological knowledge expressed in the fragments of the profiles that correspond to DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$.

## 8. Conclusion and Future Work

In this paper, we have provided algorithms and tight complexity bounds for answering various forms of regular path queries over knowledge bases formulated in lightweight DLs from the DL-Lite and $\mathcal{E L}$ families. Our results demonstrate that the query answering problem for these richer query languages is often not much harder than for the CQs and IQs typically considered. Indeed, for both DL-Lite $\mathcal{R}_{\mathcal{R}}$ and $\mathcal{E L H}$, query answering remains tractable in data complexity and in PSPACE in combined complexity for the highly expressive class of C2RPQs, and for 2RPQs, we even retain tractability in combined complexity. This computational price does not seem too high, particularly if we consider that the rich navigational features of these queries can partially compensate for the limited expressiveness of lightweight ontology languages. We thus believe that C2RPQs constitute a promising language for ontology-mediated query answering.

An important challenge for future work is to implement and experimentally evaluate the developed algorithms. Although C2RPQs is the natural query language to aim at, we believe that it makes more sense to start with a prototype implementation of the algorithm
for (2)RPQs. Indeed, not only does the algorithm for (2)RPQs have a significantly lower worst-case complexity, but it is also considerably simpler than the rewriting-based approach for C2RPQs, and we are confident that it can be easily translated into a practical procedure. By contrast, the rewriting algorithm used to establish the PSPACE upper bound for C2RPQs involves a considerable amount of non-determinism, and a naïve implementation can be expected to perform poorly. We nonetheless believe that the proposed rewriting approach, when suitably modified to avoid unnecessary non-deterministic guesses, provides a good basis for the development of practical methods for C2RPQ answering. Finally, identifying further restrictions on queries and ontologies that lead to lower combined complexity is another interesting problem for future study.

## Acknowledgments

The authors would like to thank the anonymous reviewers for their very careful reading of the paper and their many helpful comments. This work was supported by the French National Research Agency (ANR) project PAGODA 12-JS02-007-01, the Austrian Science Fund (FWF) project T515, the FWF project P25518 and the Vienna Science and Technology Fund (WWTF) project ICT12-015.

## Appendix A. Proof of Theorem 6.9

To complete the proof of Theorem 6.9, we must show how to handle queries of the form $\exists x . \alpha(x, x)$. As there does not seem to be a simple reduction to 2RPQs for queries of this form, we propose instead an approach based upon query rewriting.

To define our new query rewriting algorithm and prove its correctness, we will require the following notion. Given two concepts $C, D$ from $\mathrm{BC}_{\mathcal{T}}$, we say that $C$ causes $D$ w.r.t. a TBox $\mathcal{T}$, denoted $C \rightsquigarrow \mathcal{T} D$, if for every ABox $\mathcal{A}$ and any model $\mathcal{I}$ of $\mathcal{T}$ and $\mathcal{A}$, we have that $C^{\mathcal{I}} \neq \emptyset$ implies $D^{\mathcal{I}} \neq \emptyset$. It is not difficult to see that checking $C \rightsquigarrow \mathcal{T} D$ is feasible in polynomial time:

Lemma A.1. The following problem is in P : given a DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E L H}$ TBox $\mathcal{T}$ and concepts $C, D \in \mathrm{BC}_{\mathcal{T}}$, decide whether $C \rightsquigarrow \mathcal{T} D$.

Proof. We start with the case where $\mathcal{T}$ is a DL-Lite $\mathcal{R}_{\mathcal{R}}$ TBox. It is easy to see that $C \rightsquigarrow_{\mathcal{T}} D$ iff $\mathcal{T} \models C \sqsubseteq D$ or there exists a sequence of role names $R_{1}, \ldots, R_{n}$ such that
$-\mathcal{T} \models C \sqsubseteq \exists R_{1}$,

- $\mathcal{T} \models \exists R_{n}^{-} \sqsubseteq D$, and
- for $1 \leq i<n, \mathcal{T} \models \exists R_{i}^{-} \sqsubseteq \exists R_{i+1}$.

We then observe that the existence of a sequence $R_{1}, \ldots, R_{n}$ satisfying these conditions can be decided in polynomial time by (i) initializing a set Reach with all roles $R_{1}$ such that $\mathcal{T} \models C \sqsubseteq \exists R_{1}$, (ii) saturating Reach by adding role $S$ to Reach whenever $\mathcal{T} \models \exists U^{-} \sqsubseteq \exists S$ for some $U \in$ Reach, and (iii) checking whether there is some $U \in \operatorname{Reach}$ such that $\mathcal{T} \models$ $\exists U^{-} \sqsubseteq D$. Since TBox reasoning is tractable for DL-Lite $\mathcal{R}_{\mathcal{R}}$, the procedure we have just described can be performed in polynomial time.

Assume $\mathcal{T}$ is an $\mathcal{E L H}$ TBox. We have that $C \rightsquigarrow \mathcal{T} D$ iff $\mathcal{T} \models C \sqsubseteq D$ or there exists a sequence of concepts $\exists r_{1} . A_{1}, \ldots, \exists r_{n} . A_{n}$, where $\left\{A_{1}, \ldots, A_{n}\right\} \subseteq \mathrm{N}_{\mathrm{C}}$, such that
$-\mathcal{T} \models C \sqsubseteq \exists r_{1} . A_{1}$,

- $\mathcal{T} \models A_{n} \sqsubseteq D$, and
- for $1 \leq i<n, \mathcal{T} \models A_{i} \sqsubseteq \exists r_{i+1} . A_{i+1}$.

The existence of a sequence $\exists r_{1} . A_{1}, \ldots, \exists r_{n} . A_{n}$ satisfying the above conditions can be decided in polynomial time, using a saturation procedure analogus to the one used for DLLite $_{\mathcal{R}}$ (recall that TBox reasoning is tractable for $\mathcal{E L H}$ ).

In Figure 17, we present a deterministic query rewriting algorithm PolyRewrite that takes as input an NFA $\alpha$ and a DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H}$ TBox $\mathcal{T}$ and outputs a set of queries, which will be denoted PolyRewrite $(\alpha, \mathcal{T})$. As was the case in Section 6, the purpose of query rewriting is to ensure that we only need to consider query matches that map answer variables to ABox individuals. Since we are only interested in queries of the form $\exists x . \alpha(x, x)$, it turns out that aside from the trivial rewriting $\alpha(x, x)$, it is sufficient to consider rewritings of the forms $C(x)$ or $C(x) \wedge \alpha_{s_{1}, s_{2}}(x, x)$ in which $C$ is a basic concept and $s_{1}$ and $s_{2}$ are states in $\alpha$. In Step 1, we initialize Frontier with all tuples ( $C, s_{0}, s_{f}$ ) such that $C$ is a basic concept, $s_{0}$ is the initial state of $\alpha$, and $s_{f}$ is a final state. Then, at each iteration of the while loop, we remove a tuple ( $C, s_{1}, s_{2}$ ) from Frontier and add it to Visited to record that it has already been examined. If $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$, then the corresponding query $C(x) \wedge \alpha_{s_{1}, s_{2}}(x, x)$ is equivalent (under $\mathcal{T}$ ) to the simpler query $C(x)$, and so we add to $Q$ all queries $D(x)$ that ensure that $\exists x . C(x)$ holds. If $C \notin \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$, then we add the corresponding query $C(x) \wedge \alpha_{s_{1}, s_{2}}(x, x)$ to $Q$. We next add to Frontier all those unvisited tuples ( $D, s_{5}, s_{6}$ ) such that a match for $D(x) \wedge \alpha_{s_{5}, s_{6}}(x, x)$ that maps $x$ to $e$ implies the existence of a match for $C(x) \wedge \alpha_{s_{1}, s_{2}}(x, x)$ that maps $x$ to a child of $e$ in the anonymous part. This operation intuitively 'moves' the query match one step closer to the ABox and can be viewed as the analogue of Steps 6 and 7 in the algorithm Rewrite from Section 6.

A simple inspection of the algorithm PolyRewrite reveals that each tuple in $\mathrm{BC}_{\mathcal{T}} \times S \times S$ is examined at most once, and so there can be at most $\left|\mathrm{BC}_{\mathcal{T}}\right| \cdot|S|^{2}$ iterations of the while loop in Step 2. Since we know that all of the loop, entailment, and causation checks can be carried out in polynomial time, we obtain the following:

Lemma A.2. The algorithm PolyRewrite runs in polynomial time in $|\alpha|$ and $|\mathcal{T}|$, and hence PolyRewrite $(\alpha, \mathcal{T})$ contains a polynomial number of queries.

The next two lemmas establish the correctness of the rewriting procedure.
Lemma A.3. If $\mathcal{T}, \mathcal{A} \models \exists x . \alpha(x, x)$, then there exists a query $q(x) \in \operatorname{PolyRewrite}(\alpha, \mathcal{T})$ which has a match $\pi$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\pi(x) \in \operatorname{Ind}(\mathcal{A})$.

Proof. Suppose that $\mathcal{T}, \mathcal{A} \models \exists x . \alpha(x, x)$, and let $\pi$ be a match for $\exists x . \alpha(x, x)$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. If $\pi(x) \in \operatorname{Ind}(\mathcal{A})$, then the statement trivially holds since $\alpha(x, x)$ is added to $Q$ in Step 1 . Thus, suppose that $\pi(x) \notin \operatorname{Ind}(\mathcal{A})$. We start by proving the following claim, which captures how query matches in the anonymous part of the canonical model can be moved up closer to individuals:

Claim: Suppose that $\left(C, s_{1}, s_{2}\right)$ is added to Frontier at some point during the execution of PolyRewrite on input $(\alpha, \mathcal{T})$. Further suppose that $C \notin \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$, and there is a match $\tau$ for $\alpha_{s_{1}, s_{2}}(x, x)$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ such that $\tau(x)=d R C$ for some $d \in \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. Then there is a tuple

Algorithm PolyRewrite $(\alpha, \mathcal{T})$
Input: NFA $\alpha=\left(S, \Sigma, \delta, s_{0}, F\right)$, DL-Lite $\mathcal{R}_{\mathcal{R}}$ or $\mathcal{E} \mathcal{L H}$ TBox $\mathcal{T}$

1. Set $Q=\{\alpha(x, x)\}$, Visited $=\emptyset$, and Frontier $=\left\{\left(C, s_{0}, s_{f}\right) \mid C \in \mathrm{BC}_{\mathcal{T}}, s_{f} \in F\right\}$.
2. While Frontier $\neq \emptyset$
(a) Move a tuple ( $C, s_{1}, s_{2}$ ) from Frontier to Visited ${ }^{a}$.
(b) If $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$, then

- For every $D \in \mathrm{BC}_{\mathcal{T}}$ such that $D \rightsquigarrow \mathcal{T} C$, add $D(x)$ to $Q$.
(c) If $C \notin \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$, then
- Add $C(x) \wedge \alpha_{s_{1}, s_{2}}(x, x)$ to $Q$.
- For every tuple $\left(D, R, s_{3}, s_{4}, s_{5}, s_{6}\right) \in \mathrm{BC}_{\mathcal{T}} \times \mathrm{N}_{\mathrm{R}}^{ \pm} \times S^{4}$ such that
$-C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{3}\right]$ and $C \in \operatorname{Loop}_{\alpha}\left[s_{4}, s_{2}\right]$,
$-\mathcal{T} \models D \sqsubseteq \exists R$ and $C=\exists R^{-}\left[\right.$DL-Lite $\left._{R}\right]$ or $\mathcal{T} \models D \sqsubseteq \exists R . C[\mathcal{E L H}]$,
- there exist roles $R^{\prime}, R^{\prime \prime}$ with $\mathcal{T} \models R^{-} \sqsubseteq R^{\prime}, \mathcal{T} \models R \sqsubseteq R^{\prime \prime},\left(s_{3}, R^{\prime}, s_{5}\right) \in \delta$, and $\left(s_{6}, R^{\prime \prime}, s_{4}\right) \in \delta$, and
- $\left(D, s_{5}, s_{6}\right) \notin($ Frontier $\cup$ Visited $)$, add $\left(D, s_{5}, s_{6}\right)$ to Frontier.

3. Output $Q$.
a. We choose the least tuple in Frontier according to some arbitrary lexicographic ordering on $\mathrm{BC} C_{\mathcal{T}} \times S \times S$. We note however that the particular choice of a tuple does not affect the output of the procedure.

Figure 17: Query rewriting algorithm PolyRewrite.
$\left(D, s_{5}, s_{6}\right)$ that is added to Frontier at some point such that the query $D(x) \wedge \alpha_{s_{5}, s_{6}}(x, x)$ has a match $\tau^{\prime}$ such that $\tau^{\prime}(x)=d, \mathcal{T} \models D \sqsubseteq \exists R . C$, and either $d \in \operatorname{Ind}(\mathcal{A})$ or $\operatorname{Tail}(d)=D$.

Proof of claim. Suppose that ( $C, s_{1}, s_{2}$ ) and $\tau$ satisfy the conditions of the claim. Since $\tau(x)=d R C$, it follows from the definition of canonical models that there is a concept $D \in \mathrm{BC}_{\mathcal{T}}$ such that $d \in D^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$ and either $\mathcal{T} \models D \sqsubseteq \exists R$ and $C=\exists R^{-}$(if $\mathcal{T}$ is formulated in DL-Lite ${ }_{R}$ ), or $\mathcal{T} \models D \sqsubseteq \exists R$. $C$ (for $\mathcal{T}$ an $\mathcal{E} \mathcal{L H}$ TBox). Moreover, we may choose $D$ such that $D=\operatorname{Tail}(d)$ if $d \notin \operatorname{Ind}(\mathcal{A})$. We also know that $\tau$ is a match for $\alpha_{s_{1}, s_{2}}(x, x)$, so there exists a path $p=e_{0} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ with $e_{0}=e_{n}=\tau(x)$ whose label $\lambda(p)$ belongs to $L\left(\alpha_{s_{1}, s_{2}}\right)$. Since $\lambda(p) \in L\left(\alpha_{s_{1}, s_{2}}\right)$, we can find a sequence of states $s_{0}^{\prime} s_{1}^{\prime} \ldots s_{n}^{\prime}$ with $s_{0}^{\prime}=s_{1}$ and $s_{n}^{\prime}=s_{2}$ such that for every $1 \leq i \leq n,\left(s_{i-1}^{\prime}, u_{i}, s_{i}^{\prime}\right) \in \delta$. Because $C \notin \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$, we know that the match $\tau$ is not fully contained within $\left.\mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{\pi(x)}$. Thus, there must be at least one occurrence of the parent $d$ of $\pi(x)$ in the path $p$. Let $e_{j}$ and $e_{k}$ respectively be the first and last occurrences of $d$ in $p$ (if there is a single occurrence of $d$, then $j=k$ ). Observe that $e_{j-1}=e_{k+1}=\pi(x)$. Set $s_{3}=s_{j-1}^{\prime}, s_{4}=s_{k+1}^{\prime}, s_{5}=s_{j}^{\prime}$, and $s_{6}=s_{k}^{\prime}$. The paths $e_{0} u_{1} e_{1} \ldots u_{j-1} e_{j-1}$ and $e_{k+1} u_{k+2} e_{k+2} \ldots u_{n} e_{n}$ witness that $(\pi(x), \pi(x)) \in L\left(\alpha_{s_{1}, s_{3}}\right)^{\mathcal{I} \mathcal{T}, \mathcal{A}}$ and $(\pi(x), \pi(x)) \in L\left(\alpha_{s_{4}, s_{2}}\right)^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. As the paths $e_{0} \ldots e_{j-1}$ and $e_{k+1} \ldots e_{n}$ begin and end at $\pi(x)$
and are fully contained within $\left.\mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{\pi(x)}$, we obtain $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{3}\right]$ and $C \in \operatorname{Loop}_{\alpha}\left[s_{4}, s_{2}\right]$. Finally, we know from the definiton of paths and the construction of the canonical model that there must exist roles $R^{\prime}, R^{\prime \prime}$ with $\mathcal{T} \models R^{-} \sqsubseteq R^{\prime}, \mathcal{T} \models R \sqsubseteq R^{\prime \prime},\left(s_{3}, R^{\prime}, s_{5}\right) \in \delta$, and $\left(s_{6}, R^{\prime \prime}, s_{4}\right) \in \delta$.

By assumption, the tuple ( $C, s_{1}, s_{2}$ ) is added to Frontier at some point, and it will eventually be selected during Step 2. Moreover, since $C \notin \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$, we will enter 2(c) when examining ( $C, s_{1}, s_{2}$ ). From what we have shown above, we know that the tuple ( $D, R, s_{3}, s_{4}, s_{5}, s_{6}$ ) satisfies the first three requirements of the for-loop in Step 2(c). If the fourth requirement also holds, then this means that ( $D, s_{5}, s_{6}$ ) is added to Frontier, and if it fails, then this is because this triple has already been added to Frontier earlier in the execution of the algorithm. To complete the proof of the claim, we remark that the path $e_{j} u_{j+1} \ldots u_{k} e_{k}$ witnesses that $(d, d) \in L\left(\alpha_{s_{5}, s_{6}}\right)^{\mathcal{I}_{\mathcal{T}}, \mathcal{A}}$. Moreover, we have seen that $d \in D^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. It follows that by setting $\tau^{\prime}(x)=d$, we obtain a match for the query $D(x) \wedge \alpha_{s_{5}, s_{6}}(x, x)$ with the required properties. (end proof of claim)

Observe that if $\tau^{\prime}$ as described in the claim exists, then $\tau^{\prime}(x)$ is the parent of $\tau(x)$ in the canonical model $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. We can now finalize the proof. Indeed, since $\pi(x) \notin \operatorname{Ind}(\mathcal{A})$, we have $\pi(x)=d R C$ for some $R \in \mathrm{~N}_{\mathrm{R}}^{ \pm}$and $C \in \mathrm{BC}_{\mathcal{T}}$. As $\pi$ is a match for $\alpha(x, x)$, it must also be a match for $\alpha_{s_{0}, s_{f}}(x, x)$ for some $s_{f} \in F$. In Step 1, the tuple ( $C, s_{0}, s_{f}$ ) will be added to Frontier. Repeated applications of the claim either yield a query $q(x) \in \operatorname{PolyRewrite}(\alpha, \mathcal{T})$ having a match $\tau$ mapping $x$ to the ABox, or result in the insertion of a tuple ( $D, s_{1}, s_{2}$ ) into Frontier such that $D(x) \wedge \alpha_{s_{1}, s_{2}}(x, x)$ has a match in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ and $D \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{2}\right]$. In the latter case, let $\tau$ be a match for $D(x) \wedge \alpha_{s_{1}, s_{2}}(x, x)$, and let $a \in \operatorname{Ind}(\mathcal{A})$ be such that $\left.\tau(x) \in \mathcal{I}_{\mathcal{T}, \mathcal{A}}\right|_{a}$. Then it follows from the definition of canonical models that there is some $E \in \mathrm{BC}_{\mathcal{T}}$ such that $E \rightsquigarrow \mathcal{T} D$ and $a \in E^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$. Thus, there is a match for $E(x) \in \operatorname{PolyRewrite}(\alpha, \mathcal{T})$ that maps $x$ to $a \in \operatorname{Ind}(\mathcal{A})$.

Lemma A.4. If $q(x) \in \operatorname{PolyRewrite}(\alpha, \mathcal{T})$ and $\pi$ is a match for $q(x)$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ with $\pi(x) \in$ $\operatorname{Ind}(\mathcal{A})$, then $\mathcal{T}, \mathcal{A} \models \exists x . \alpha(x, x)$.

Proof. To simplify presentation we introduce the notion of containment of Boolean queries w.r.t. a TBox. Given a TBox $\mathcal{T}$ and two Boolean queries $q_{1}, q_{2}$, we write $q_{1} \subseteq_{\mathcal{T}} q_{2}$ if for every ABox $\mathcal{A}$ we have that $\mathcal{T}, \mathcal{A} \models q_{1}$ implies $\mathcal{T}, \mathcal{A} \models q_{2}$. With a slight abuse of notation we allow $q_{1}$ and $q_{2}$ to contain atoms of the form $\exists R(x)$. Each occurrence of $\exists R(x)$ in a query $q$ is a shorthand for $R(x, y)$ where $y$ is a variable that occurs once in $q$.

We start by establishing the following claim:
Claim: If $\left(C, s_{1}, s_{2}\right)$ is added to Frontier at some point during the execution of PolyRewrite on input $(\alpha, \mathcal{T})$, then $\exists x . C(x) \wedge \alpha_{s_{1}, s_{2}}(x, x) \subseteq \mathcal{T} \exists x . \alpha(x, x)$.
Proof of claim. The proof is by induction on the precedence relation obtained by setting $\left(C, s, s^{\prime}\right) \prec\left(D, s^{\prime \prime}, s^{\prime \prime \prime}\right)$ if the tuple $\left(D, s^{\prime \prime}, s^{\prime \prime \prime}\right)$ is added to Frontier during the examination of tuple ( $C, s, s^{\prime}$ ). For the base case, we have the tuples ( $C, s_{0}, s_{f}$ ) which are inserted in Step 1. Every such tuple has the form $\left(C, s_{0}, s_{f}\right)$, where $s_{f} \in F$, and thus we trivially have $\exists x . C(x) \wedge \alpha_{s_{0}, s_{f}}(x, x) \subseteq_{\mathcal{T}} \exists x . \alpha(x, x)$. Next suppose that we have already shown the property for $\left(C, s_{1}, s_{2}\right)$, and let $\left(D, s_{5}, s_{6}\right)$ be such that $\left(C, s_{1}, s_{2}\right) \prec\left(D, s_{5}, s_{6}\right)$. Further suppose that there is a match $\pi$ for $D(x) \wedge \alpha_{s_{5}, s_{6}}(x, x)$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. Then we can find a path $p_{0}=e_{0} u_{1} e_{1} u_{2} \ldots u_{n} e_{n}$ with $e_{0}=e_{n}=\pi(x)$ such that $\lambda(p) \in L\left(\alpha_{s_{5}, s_{6}}\right)$. It follows that
there must exist a sequence of states $s_{0}^{\prime} s_{1}^{\prime} \ldots s_{n}^{\prime}$ with $s_{0}^{\prime}=s_{5}$ and $s_{n}^{\prime}=s_{6}$ such that for every $1 \leq i \leq n,\left(s_{i-1}^{\prime}, u_{i}, s_{i}^{\prime}\right) \in \delta$. Because $\left(C, s_{1}, s_{2}\right) \prec\left(D, s_{5}, s_{6}\right)$, there must exist a tuple $\left(D, R, s_{3}, s_{4}, s_{5}, s_{6}\right)$ such that $\left(D, s_{5}, s_{6}\right)$ was added to Frontier when examining ( $D, R, s_{3}, s_{4}, s_{5}, s_{6}$ ). We know that this tuple must have satisfied the four conditions, and so we must have

- $\mathcal{T} \models D \sqsubseteq \exists R$ and $C=\exists R^{-}\left[\right.$DL-Lite $\left._{R}\right]$ or $\mathcal{T} \models D \sqsubseteq \exists R . C[\mathcal{E L H}]$,
- $C \in \operatorname{Loop}_{\alpha}\left[s_{1}, s_{3}\right]$ and $C \in \operatorname{Loop}_{\alpha}\left[s_{4}, s_{2}\right]$, and
- there exist roles $R^{\prime}, R^{\prime \prime}$ with $\mathcal{T} \models R^{-} \sqsubseteq R^{\prime}, \mathcal{T} \models R \sqsubseteq R^{\prime \prime},\left(s_{3}, R^{\prime}, s_{5}\right) \in \delta$, and $\left(s_{6}, R^{\prime \prime}, s_{4}\right) \in \delta$.

By the first point, the element $\pi(x) R C$ belongs to the canonical model, and by the second point, we can find paths $p_{1}=e_{0}^{\prime} \ldots u_{m}^{\prime} e_{m}^{\prime}$ and $p_{2}=e_{0}^{\prime \prime} \ldots u_{\ell}^{\prime \prime} e_{\ell}^{\prime \prime}$ such that $e_{0}^{\prime}=e_{0}^{\prime \prime}=e_{m}^{\prime}=$ $e_{\ell}^{\prime \prime}=\pi(x) R C, \lambda\left(p_{1}\right) \in L\left(\alpha_{s_{1}, s_{3}}\right)$, and $\lambda\left(p_{2}\right) \in L\left(\alpha_{s_{4}, s_{2}}\right)$. Using this and the third point, we can show that the path $p^{*}=p_{1} R^{\prime} p_{0} R^{\prime \prime} p_{2}$ is such that $\lambda\left(p^{*}\right) \in L\left(\alpha_{s_{1}, s_{2}}\right)$. Since $p^{*}$ begins and ends at $\pi(x) R C$ and $\pi(x) R C \in C^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}}$, it follows that $\mathcal{T}, \mathcal{A} \vDash \exists x . C(x) \wedge \alpha_{s_{1}, s_{2}}(x, x)$. By the induction hypothesis, $\exists x . C(x) \wedge \alpha_{s_{1}, s_{2}}(x, x) \subseteq_{\mathcal{T}} \exists x . \alpha(x, x)$, so we must also have $\mathcal{T}, \mathcal{A} \models \exists x . \alpha(x, x)$. This establishes the desired containment $\exists x . D(x) \wedge \alpha_{s_{5}, s_{6}}(x, x) \subseteq_{\mathcal{T}}$ $\exists x \cdot \alpha(x, x)$. (end proof of claim)

Now suppose that $\pi$ is a match for $q(x) \in \operatorname{PolyRewrite}(\alpha, \mathcal{T})$ in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ such that $\pi(x) \in$ $\operatorname{Ind}(\mathcal{A})$. There are three possibilities. The first is that $q(x)=\alpha(x, x)$, in which case we trivially have $\mathcal{T}, \mathcal{A} \models \exists x . \alpha(x, x)$. The next possibility is that $q(x)=\exists x . C(x) \wedge \alpha_{s, s^{\prime}}(x, x)$, in which case we can apply the preceding claim to show that $\mathcal{T}, \mathcal{A} \models \exists x . \alpha(x, x)$. The final possibility is that $q(x)=D(x)$, in which case there must have been some $\left(C, s, s^{\prime}\right) \in$ Frontier such that $C \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]$ and $\exists x . D(x) \subseteq_{\mathcal{T}} \exists x . C(x)$. In this case, we have $\exists x . D(x) \subseteq_{\mathcal{T}} \exists x . C(x), \exists x . C(x) \subseteq_{\mathcal{T}} \exists x . C(x) \wedge \alpha_{s, s^{\prime}}(x, x)$ (since $\left.C \in \operatorname{Loop}_{\alpha}\left[s, s^{\prime}\right]\right)$, and $\exists x . C(x) \wedge \alpha_{s, s^{\prime}}(x, x) \subseteq_{\mathcal{T}} \exists x . \alpha(x, x)$ (by the claim). Putting these statements together, we obtain $\exists x . D(x) \subseteq_{\mathcal{T}} \exists x . \alpha(x, x)$, which yields $\mathcal{T}, \mathcal{A} \models \exists x$. $\alpha(x, x)$.

To complete the argument, we observe that the queries output by PolyRewrite are either 2RPQs or take the form $C(x)$ or $C(x) \wedge \alpha_{s, s^{\prime}}(x, x)$, and queries of the latter forms are trivially transformed into 2RPQs. It follows that we can answer a query of the form $\exists x . \alpha(x, x)$ by (i) computing the set PolyRewrite $(\alpha, \mathcal{T})$, and (ii) using EvalAtom to check, for each 2 RPQ $\alpha^{\prime}(x, x)$ obtained from PolyRewrite $(\alpha, \mathcal{T})$, whether $\operatorname{cert}\left(\alpha^{\prime}(x, x),(\mathcal{T}, \mathcal{A})\right) \neq \emptyset$. Correctness of this procedure follows from Proposition 5.8 and Lemmas A. 3 and A.4, and the complexity bounds then follow from Lemma A. 2 and Theorem 5.9.

## References

Abiteboul, S., Hull, R., \& Vianu, V. (1995). Foundations of Databases. Addison-Wesley.
Arora, S., \& Barak, B. (2009). Computational Complexity - A Modern Approach. Cambridge University Press.
Artale, A., Calvanese, D., Kontchakov, R., \& Zakharyaschev, M. (2009). The DL-Lite family and relations. Journal of Artificial Intelligence Research (JAIR), 36, 1-69.

Bienvenu, Ortiz, \& Šimkus

Baader, F., Brandt, S., \& Lutz, C. (2005). Pushing the $\mathcal{E L}$ envelope. In Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI 2005).

Bala, S. (2002). Intersection of regular languages and star hierarchy. In Proceedings of the Twenty-Ninth International Colloquium on Automata, Languages and Programming (ICALP 2002).
Barceló, P. (2013). Querying graph databases. In Proceedings of the Thirty-Second Symposium on Principles of Database Systems (PODS 2013).

Barceló, P., Libkin, L., Lin, A. W., \& Wood, P. T. (2012). Expressive languages for path queries over graph-structured data. ACM Transactions on Database Systems (TODS), $37(4), 31$.

Barceló, P., \& Muñoz, P. (2014). Graph logics with rational relations: The role of word combinatorics. In Proceedings of the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2014).

Barceló, P., Pérez, J., \& Reutter, J. L. (2012). Relative expressiveness of nested regular expressions. In Proceedings of the Sixth Alberto Mendelzon International Workshop on Foundations of Data Management (AMW 2012).

Berglund, A., Boag, S., Chamberlin, D., Fernández, M. F., Kay, M., Robie, J., \& Siméon, J. (2007). XML Path Language (XPath) 2.0. W3C Recommendation. Available at http://www.w3.org/TR/xpath20/.
Bienvenu, M., Calvanese, D., Ortiz, M., \& Šimkus, M. (2014). Nested regular path queries in description logics. In Proceedings of the Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning (KR 2014).

Bienvenu, M., Ortiz, M., \& Šimkus, M. (2012). Answering expressive path queries over lightweight DL knowledge bases. In Proceedings of the Twenty-Fifth International Workshop on Description Logics (DL 2012).

Bienvenu, M., Ortiz, M., \& Šimkus, M. (2013). Conjunctive regular path queries in lightweight description logics. In Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI 2013).
Bourhis, P., Krötzsch, M., \& Rudolph, S. (2014). How to best nest regular path queries. In Proceedings of the Twenty-Seventh International Workshop on Description Logics (DL 2014).

Brickley, D., \& Guha, R. (2014). RDF Schema 1.1. W3C Recommendation. Available at http://www.w3.org/TR/rdf-schema/.
Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., \& Rosati, R. (2006). Data complexity of query answering in description logics. In Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2006).
Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., \& Rosati, R. (2007). Tractable reasoning and efficient query answering in description logics: The DL-Lite family. Journal of Automated Reasoning, 39(3), 385-429.

Calvanese, D., De Giacomo, G., \& Lenzerini, M. (1998). On the decidability of query containment under constraints. In Proceedings of the Seventeenth Symposium on Principles of Database Systems (PODS 1998).
Calvanese, D., Eiter, T., \& Ortiz, M. (2007). Answering regular path queries in expressive description logics: An automata-theoretic approach. In Proceedings of the TwentySecond AAAI Conference on Artificial Intelligence (AAAI 2007).
Calvanese, D., Eiter, T., \& Ortiz, M. (2009). Regular path queries in expressive description logics with nominals. In Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI 2009).

Calvanese, D., Eiter, T., \& Ortiz, M. (2014). Answering regular path queries in expressive description logics via alternating tree-automata. Information and Computation, 237, 12-55.

Consens, M. P., \& Mendelzon, A. O. (1990). GraphLog: A visual formalism for real life recursion. In Proceedings of the Ninth Symposium on Principles of Database Systems (PODS 1990).
Ehrenfeucht, A., \& Zeiger, P. (1974). Complexity measures for regular expressions. In Proceedings of the Sixth Annual ACM Symposium on Theory of Computing (STOC 1974).

Eiter, T., Ortiz, M., Šimkus, M., Tran, T., \& Xiao, G. (2012). Query rewriting for Horn$\mathcal{S H I Q}$ plus rules. In Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012).
Florescu, D., Levy, A., \& Suciu, D. (1998). Query containment for conjunctive queries with regular expressions. In Proceedings of the Seventeenth Symposium on Principles of Database Systems (PODS 1998).
Haase, C., \& Lutz, C. (2008). Complexity of subsumption in the $\mathcal{E L}$ family of description logics: Acyclic and cyclic TBoxes. In Proceedings of the Eighteenth European Conference on Artificial Intelligence (ECAI 2008).

Harris, S., \& Seaborne, A. (2013). SPARQL 1.1 Query Language. W3C Recommendation. Available at http://www.w3.org/TR/sparql11-query/.

Kikot, S., Kontchakov, R., Podolskii, V. V., \& Zakharyaschev, M. (2013). Query rewriting over shallow ontologies. In Proceedings of the Twenty-Sixth International Workshop on Description Logics (DL 2013).
Kostylev, E. V., Reutter, J. L., \& Vrgoc, D. (2015). XPath for DL ontologies. In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015).
Kozen, D. (1977). Lower bounds for natural proof systems. In Proceedings of the Eighteenth Annual Symposium on Foundations of Computer Science (SFCS 1977).
Krisnadhi, A., \& Lutz, C. (2007). Data complexity in the $\mathcal{E L}$ family of DLs. In Proceedings of the Twentieth International Workshop on Description Logics (DL 2007).

Krötzsch, M. (2012). OWL 2 Profiles: An introduction to lightweight ontology languages. In Proceedings of the Eighth Reasoning Web Summer School (RW 2012).

Bienvenu, Ortiz, \& Šimkus

Krötzsch, M., \& Rudolph, S. (2007). Conjunctive queries for $\mathcal{E} \mathcal{L}$ with composition of roles. In Proceedings of the Twentieth International Workshop on Description Logics (DL 2007).

Levy, A. Y., \& Rousset, M. (1996). The limits on combining recursive Horn rules with description logics. In Proceedings of the Thirteenth National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference (AAAI 96).
Lutz, C. (2008). The complexity of conjunctive query answering in expressive description logics. In Proceedings of the Fourth Joint Conference on Automated Reasoning (IJCAR 2008).

Motik, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokoue, A., \& Lutz, C. (2012). OWL 2 Web Ontology Language Profiles. W3C Recommendation. Available at http://www. w3.org/TR/owl2-profiles/.

Ortiz, M. (2013). Ontology based query answering: The story so far. In Proceedings of the Seventh Alberto Mendelzon International Workshop on Foundations of Data Management (AMW 2013).
Ortiz, M., Rudolph, S., \& Šimkus, M. (2011). Query answering in the Horn fragments of the description logics $\mathcal{S H O I Q}$ and $\mathcal{S R O I Q}$. In Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011).
Ortiz, M., \& Šimkus, M. (2012). Reasoning and query answering in description logics. In Proceedings of the Eighth Reasoning Web Summer School (RW 2012).
Ortiz, M., \& Šimkus, M. (2014). Revisiting the hardness of query answering in expressive description logics. In Proceedings of the Eighth International Conference on Web Reasoning and Rule Systems (RR 2014).
OWL Working Group, W. (2009). OWL 2 Web Ontology Language: Document Overview. W3C Recommendation. Available at http://www.w3.org/TR/ow12-overview/.
Pérez, J., Arenas, M., \& Gutierrez, C. (2010). nSPARQL: A navigational language for RDF. Journal of Web Semantics, 8(4), 255-270.
Reutter, J. L. (2013). Containment of nested regular expressions. CoRR, abs/1304.263\%.
Rosati, R. (2007). On conjunctive query answering in $\mathcal{E L}$. In Proceedings of the Twentieth International Workshop on Description Logics (DL 2007).
Stefanoni, G., Motik, B., Krötzsch, M., \& Rudolph, S. (2014). The complexity of answering conjunctive and navigational queries over OWL 2 EL knowledge bases. Journal of Artificial Intelligence Research (JAIR), 51, 645-705.
Thompson, K. (1968). Regular expression search algorithm. Communications of the ACM, 11(6), 419-422.


[^0]:    1. The $\mathcal{Z}$ symbol was introduced as an abbreviation for $\mathcal{A} \mathcal{L C} b_{\text {reg }}^{\text {Self }}$ (Calvanese et al., 2009).
[^1]:    2. Recall that throughout the paper, we assume that $\mathcal{E L H}$ KBs are in normal form, and for this reason, we only need to consider existential concepts of the form $\exists r . A$ with $A \in \mathrm{~N}_{\mathcal{C}} \cup\{\top\}$.
