## On the Complexity of Evaluating Regular Path Queries over Linear Existential Rules

## - Source link

Meghyn Bienvenu, Michaël Thomazo
Institutions: University of Montpellier, French Institute for Research in Computer Science and Automation
Published on: 09 Sep 2016 - Web Reasoning and Rule Systems
Topics: Conjunctive query, EXPTIME and Ontology (information science)

Related papers:

- Answering Conjunctive Regular Path Queries over Guarded Existential Rules
- Regular path queries in lightweight description logics: complexity and algorithms
- Containment of Conjunctive Regular Path Queries with Inverse.
- Towards Efficient Reasoning Under Guarded-Based Disjunctive Existential Rules
- On the Complexity of Database Queries


# On the Complexity of Evaluating Regular Path Queries over Linear Existential Rules 

Meghyn Bienvenu, Michaël Thomazo

## - To cite this version:

Meghyn Bienvenu, Michaël Thomazo. On the Complexity of Evaluating Regular Path Queries over Linear Existential Rules. RR: Web Reasoning and Rule Systems, Sep 2016, Aberdeen, United Kingdom. pp.1-17, 10.1007/978-3-319-45276-0__1 . hal-01341787

## HAL Id: hal-01341787 <br> https://hal.inria.fr/hal-01341787

Submitted on 4 Jul 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# On the Complexity of Evaluating Regular Path Queries over Linear Existential Rules 

Meghyn Bienvenu ${ }^{1}$ and Michaël Thomazo ${ }^{2}$<br>${ }^{1}$ CNRS, Université de Montpellier, \& Inria, Montpellier, France<br>meghyn@lirmm.fr<br>${ }^{2}$ Inria, France<br>michael.thomazo@inria.fr


#### Abstract

In the setting of ontology-mediated query answering, a query is evaluated over a knowledge base consisting of a database instance and an ontology. While most work in the area focuses on conjunctive queries, navigational queries are gaining increasing attention. In this paper, we investigate the complexity of evaluating the standard form of navigational queries, namely two-way regular path queries, over knowledge bases whose ontology is expressed by means of linear existential rules. More specifically, we show how to extend an approach developed for DL-Lite ${ }_{\mathcal{R}}$ to obtain an exponential-time decision procedure for linear rules. We prove that this algorithm achieves optimal worst-case complexity by establishing a matching ExpTime lower bound.


## 1 Introduction

Ontology-mediated query answering (OMQA) has generated a lot of interest in the last years as a promising way of facilitating access to data (see [4] for a recent survey). In the OMQA approach, the ontology serves to define a conceptual view of an application domain, introducing a convenient vocabulary for query formulation and providing background knowledge that is exploited at query time to obtain the complete set of answers. So far, the vast majority of research on OMQA has considered user queries in the form of conjunctive queries (CQs), which are a standard query language for relational databases. However, in numerous application scenarios, data can naturally be seen as graphs, in which case so-called navigational queries are considered more suitable. The basic navigational query language is regular path queries (RPQs) [11], which allow one to find paths whose labels conform to a given regular language.

In recent years, the problem of answering navigational queries in the setting of OMQA has begun to be explored, first for ontologies formulated in highly expressive description logics (DLs) of the $\mathcal{Z}$ family [8,9,10], then for rich Horn DLs like Horn$\mathcal{S R O} \mathcal{I} \mathcal{Q}$ [18], and more recently, for lightweight DLs like DL-Lite $\mathcal{R}$ and $\mathcal{E} \mathcal{L}$ [19,5]. The latter DLs, which underlie the OWL 2 QL and EL profiles, are the most relevant for OMQA due to their favourable computational properties. In addition to plain RPQs, this line of work has also considered richer navigational languages like conjunctive RPQs (which extend both RPQs and CQs) and extensions with nesting and/or negation [3,6,15]. Although much work remains to be done in developing and implementing efficient algorithms, the complexity landscape for answering various forms of path queries
over DL knowledge bases is now rather well understood. The same cannot be said for ontologies formulated by means of decidable classes of existential rules (like linear and guarded rulesets), which constitute another important class of ontology languages [7,1]. A key feature that distinguishes existential rules from DLs is the possibility of using predicates of arity greater than two. Since regular path queries are defined only with respect to unary and binary predicates, one might wonder whether they make sense in higher arity settings. We argue however that unary and binary predicates form the backbone of real-world ontologies (irrespective of the choice of ontology language), and it is desirable to be able to use some higher-arity predicates without losing any expressivity in the query language.

In this paper, we take a step towards a better understanding of the combination of navigational query languages and existential rules by studying the complexity of answering two-way RPQs in the presence of linear rules, a well-studied class of existential rules that are a natural generalization of the DL-Lite description logics. After introducing the necessary background, we show how to adapt the RPQ algorithm for DL-Lite proposed in [5] to the setting of linear rules. Unfortunately, our adaptation incurs an exponential blow-up with respect to the maximum predicate arity. We can nevertheless show that the obtained algorithm is worst-case optimal, as RPQ answering is EXPTIME-complete in combined complexity.

## 2 Preliminaries

We adopt the notation of [13]. The notions of constants, function symbols and predicate symbols are standard. Each function or predicate symbol is associated with a nonnegative integer arity. Variables, terms, substitutions, atoms, first-order formulae, sentences, interpretations (i.e., structures), and models are defined as usual. By a slight abuse of notation, we often identify a conjunction with the set of its conjuncts. Furthermore, we often abbreviate a vector of terms $t_{1}, \ldots, t_{n}$ as $\mathbf{t}$, and define $|\mathbf{t}|=n$. By $\varphi \sigma$ we denote the result of applying a substitution $\sigma$ to $\varphi$. A term, atom, or formula is ground if it does not contain variables; a fact is a ground atom. A term $t^{\prime}$ is a subterm of a term $t$ if $t^{\prime}=t$ or $t=f(\mathbf{s})$ where $f$ is a function and $t^{\prime}$ is a subterm of some $s_{i} \in \mathbf{s}$. A term $s$ is contained in an atom $p(\mathbf{t})$ is $s \in \mathbf{t}$, and $s$ occurs in $p(\mathbf{t})$ if $s$ is a subterm of some term $t_{i} \in \mathbf{t}$; thus, if $s$ is contained in $p(\mathbf{t}), s$ occurs in $p(\mathbf{t})$, but the converse may not hold. A term $s$ is contained (resp. occurs) in a set of atoms $I$ if $s$ is contained (resp. occurs) in some atom in $I$. Let $T=\left\{t_{1}, \ldots, t_{n}\right\}$ be a set of terms. A term $t$ is generated by $T$ if (i) $t \in T$ or (ii) $t=f\left(x_{1}, \ldots, x_{k}\right)$ and all the $x_{k}$ are generated by $T$. An instance is a finite set of function-free facts. The terms appearing in an instance (resp. atom) are denoted by terms $(I)$ (resp. terms $(\alpha)$ ).

Existential Rules An existential rule (or just rule) takes the form:

$$
\forall \mathbf{x} \forall \mathbf{z} \cdot[\varphi(\mathbf{x}, \mathbf{z}) \rightarrow \exists \mathbf{y} \cdot \psi(\mathbf{x}, \mathbf{y})]
$$

where $\varphi(\mathbf{x}, \mathbf{z})$ and $\psi(\mathbf{x}, \mathbf{y})$ are non-empty conjunctions of function-free atoms, and tuples of variables $\mathbf{x}, \mathbf{y}$ and $\mathbf{z}$ are pairwise disjoint. We call $\varphi$ the body and $\psi$ the head of the rule. For brevity, quantifiers are often omitted.

We frequently use Skolemisation to interpret rules in Herbrand interpretations, which are defined as possibly infinite sets of facts. In particular, for each rule $\rho$ and each variable $y_{i} \in \mathbf{y}$, let $f_{\rho}^{i}$ be a function symbol globally unique for $\rho$ and $y_{i}$ of arity $|\mathbf{x}|$; furthermore, let $\theta_{s k}$ be the substitution such that $\theta_{s k}\left(y_{i}\right)=f_{\rho}^{i}(\mathbf{x})$ for each $y_{i} \in \mathbf{y}$. Then, the Skolemisation $\operatorname{sk}(\rho)$ of $\rho$ is the following rule: $\varphi(\mathbf{x}, \mathbf{z}) \rightarrow \psi(\mathbf{x}, \mathbf{y}) \theta_{s k}$.

A linear rule is an existential rule whose body is restricted to a single atom. For ease of presentation, we will consider only rules without any constants. As usual, we also assume that rules have only a single atom in the head. This can be done without loss of generality.

Skolem Chase The chase [16,14] (or canonical model) is a classical tool in OMQA. In this paper, we use the Skolem chase variant ([17]). Let $\rho=\varphi \rightarrow \psi$ be a Skolemised rule, and let $I$ be a set of facts. A set of facts $S$ is a consequence of $\rho$ on $I$ if a substitution $\sigma$ exists that maps the variables in $\rho$ to the terms occurring in $I$ (denoted by terms $(I)$ ) such that $\varphi \sigma \subseteq I$ and $S \subseteq \psi \sigma$. The result of applying $\rho$ to $I$, written $\rho(I)$, is the union of all consequences of $\rho$ on $I$. If $\Omega$ is a set of Skolemised rules, we set $\Omega(I)=\bigcup_{\rho \in \Omega} \rho(I)$. Let $I$ be a finite set of facts, let $\mathcal{R}$ be a set of rules, let $\mathcal{R}^{\prime}=\operatorname{sk}(\mathcal{R})$, and let $\mathcal{R}_{f}^{\prime \prime}$ and $\mathcal{R}_{n}^{\prime}$ be the subsets of $\mathcal{R}^{\prime}$ containing rules with and without function symbols, respectively. The chase sequence for $I$ and $\mathcal{R}$ is a sequence of sets of facts $I_{\mathcal{R}}^{0}, I_{\mathcal{R}}^{1}, \ldots$, where $I_{\mathcal{R}}^{0}=I$ and for each $i>0$, set $I_{\mathcal{R}}^{i}$ is defined as follows:

- if $\mathcal{R}_{n}^{\prime}\left(I_{\mathcal{R}}^{i-1}\right) \nsubseteq I_{\mathcal{R}}^{i-1}$, then $I_{\mathcal{R}}^{i}=I_{\mathcal{R}}^{i-1} \cup \mathcal{R}_{n}^{\prime}\left(I_{\mathcal{R}}^{i-1}\right)$
- otherwise $I_{\mathcal{R}}^{i}=I_{\mathcal{R}}^{i-1} \cup \mathcal{R}_{f}^{\prime}\left(I_{\mathcal{R}}^{i-1}\right)$

The chase of $I$ and $\mathcal{R}$, written chase $(I, \mathcal{R})$, is defined as $\bigcup_{i} I_{\mathcal{R}}^{i}$; note that chase $(I, \mathcal{R})$ can be infinite. However, the chase has a simple structure when linear rules are considered: each atom can be "chased" independently.

Property 1 (Decomposition of the chase). Let $\mathcal{R}$ be a set of linear rules and $I$ be an instance. It holds that:

$$
\operatorname{chase}(I, \mathcal{R})=\cup_{\alpha \in I} \operatorname{chase}(\{\alpha\}, \mathcal{R})
$$

Regular Languages A regular language can be represented either by a regular expression or by a non-deterministic finite automaton (NFA). Let $\Sigma$ be a finite set of symbols. A regular expression over $\Sigma$ is defined by the grammar: $\mathcal{E} \rightarrow \varepsilon|a| \mathcal{E} \cdot \mathcal{E}|\mathcal{E}+\mathcal{E}| \mathcal{E}^{*}$, where $a \in \Sigma$ and $\varepsilon$ denotes the empty word. We use $L(\mathcal{E})$ to denote the language defined by $\mathcal{E}$. An NFA over $\Sigma$ is a tuple $\mathbb{A}=\left(S, \Sigma, \delta, s_{0}, F\right)$, where $S$ is a finite set of states, $\delta \subseteq S \times \Sigma \times S$ is the transition relation, $s_{0} \in S$ is the initial state and $F \subseteq S$ is the set of final states. If $\mathbb{A}$ is an automaton and $s$ and $s^{\prime}$ are two states of $\mathbb{A}$, we denote by $\mathcal{L}_{\mathbb{A}}\left(s, s^{\prime}\right)$ the set of words $w$ for which there is path from $s$ to $s^{\prime}$ in $\mathbb{A}$ labeled by $w$.

Regular Path Queries Let $\mathcal{P}$ be a set of predicates. Let us define $\mathcal{P}_{2}^{ \pm}=\mathcal{P}_{2} \cup\left\{r^{-} \mid\right.$ $\left.r \in \mathcal{P}_{2}\right\}$ and $\mathcal{P}_{r}=\mathcal{P}_{2}^{ \pm} \cup \mathcal{P}_{1}$, where $\mathcal{P}_{i}(i \in\{1,2\})$ denotes the predicates of arity $i$. A two-way regular path query $\left(\mathrm{RPQ}^{3}\right)$ is a query of the form $q\left(x, x^{\prime}\right)=\mathcal{E}\left(x, x^{\prime}\right)$, where $\mathcal{E}$ is a regular expression defining a language over $\mathcal{P}_{r}$.

[^0]Given an interpretation $\mathcal{I}$, a path from $a_{0}$ to $a_{n}$ in $\mathcal{I}$ is a sequence $a_{0} r_{1} a_{1} r_{2} \ldots r_{n} a_{n}$ such that for any $i$ such that $1 \leq i \leq n, a_{i}$ is an element of the domain $\Delta^{\mathcal{I}}$ of $\mathcal{I}$, every $r_{i}$ is a symbol from $\mathcal{P}_{r}$ and:

- if $r_{i}=a \in \mathcal{P}_{1}$, then $a_{i}=a_{i-1} \in a^{\mathcal{I}}$;
- if $r_{i} \in \mathcal{P}_{2}$, then $\left(a_{i-1}, a_{i}\right) \in r_{i}^{\mathcal{I}}$;
- if $r_{i}=r^{-}$with $r \in \mathcal{P}_{2}$, then $\left(a_{i}, a_{i-1}\right) \in r^{\mathcal{I}}$.

The label $\lambda(p)$ of path $p=a_{0} r_{1} a_{1} r_{2} \ldots r_{n} a_{n}$ is the word $r_{1} r_{2} \ldots r_{n}$. For any language $L$ over $\mathcal{P}_{r}$, the semantics of $L$ with respect to an interpretation $\mathcal{I}$ is defined by:

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

A match for an RPQ $q\left(x, x^{\prime}\right)=\mathcal{E}\left(x, x^{\prime}\right)$ in an interpretation $\mathcal{I}$ is a mapping $\pi$ from the variables of $q$ to elements of $\Delta^{\mathcal{I}}$ such that $\left(\pi(x), \pi\left(x^{\prime}\right)\right) \in L(\mathcal{E})^{\mathcal{I}}$.

A certain answer to $q\left(x_{1}, x_{2}\right)$ with respect to $(I, \mathcal{R})$ is a pair of constants $\left(a_{1}, a_{2}\right)$ such that for every model $\mathcal{I}$ of $(I, \mathcal{R})$, there is a match $\pi$ for $q$ such that $\pi\left(x_{1}\right)=$ $a_{1}^{\mathcal{I}}$ and $\pi\left(x_{2}\right)=a_{2}^{\mathcal{I}}$. As matches are preserved under homomorphisms, it holds that $\left(a_{1}, a_{2}\right)$ is a certain answer to $q\left(x_{1}, x_{2}\right)$ w.r.t. $(I, \mathcal{R})$ if and only if there is a match for $\left(a_{1}^{\mathcal{I}}, a_{2}^{\mathcal{I}}\right)$ in $\mathcal{I}=$ chase $(I, \mathcal{R})$. The RPQ Answering problem asks, given an RPQ $q\left(x_{1}, x_{2}\right)$, an instance $I$, a set of existential rules $\mathcal{R}$, and two constants $\left(a_{1}, a_{2}\right) \in$ $\operatorname{terms}(\mathcal{I}) \times \operatorname{terms}(\mathcal{I})$, whether $\left(a_{1}, a_{2}\right)$ is a certain answer to $q\left(x_{1}, x_{2}\right)$.

Computational Complexity and Turing Machines We assume the reader to be familiar with standard complexity classes. In particular, we will consider P, NP, PSPACE, APSPACE (alternating PSPACE), and ExpTime. We recall that APSPACE = ExpTime.

To fix notations, we recall that an alternating Turing machine (TM) is given by a 5-tuple $\mathcal{M}=\left(Q, \Gamma, \delta, q_{0}, g\right)$ where:

- $Q$ is the finite set of states;
- $\Gamma$ is the finite tape alphabet;
- $\delta: Q \times \Gamma \rightarrow(Q \times \Gamma \times\{L, R\})^{2}$ is the transition function;
- $q_{0} \in Q$ is the initial state;
$-g: Q \rightarrow\{\wedge, \vee$, accept, reject $\}$ specifies the type of each state.
Note that without loss of generality, we consider TMs having the following properties:
- for every universal $(\wedge)$ or existential $(\vee)$ configuration, there exist exactly two applicable transitions;
- the machine directly accepts any configuration whose state $s$ is such that $g(s)=$ accept;
- the TM never tries to go to the left of the initial position.

We say $\mathcal{M}$ is polynomially space-bounded ( $\mathcal{M}$ is a PSPACE TM) if there exists a polynomial $p$ such that on input $x, \mathcal{M}$ visits only the first $p(|x|)$ tape cells. We assume w.l.o.g. that the alternating PSPACE TMs we consider terminate on every input.

## 3 Evaluating Regular Path Queries over Linear Rules

We consider the problem of computing the certain answers to a regular path query and show how to adapt the construction in [5] to the case of linear rules. There are two main ingredients in the original algorithm for DL-Lite:

- a path in the chase is guessed step by step, keeping in memory only the current constant of the instance and current state of the automaton;
- when a path goes through the Skolem part of the chase, these constants are not guessed, but the state in which the automaton is when the path returns to constants of the instance is guessed, thanks to a precomputed table.


### 3.1 Additional Challenges with Linear Rules

There are two main differences between DL-Lite and linear rules that need to be handled. First, in DL-Lite, it is enough to know the predicate of the atom in which an constant has been created during the chase and the position at which it appeared in that atom to determine all the atoms that contain that constant in the chase. This is not true if we consider general linear rules, as illustrated by the following example:

Example 1 (More complex types are needed). Let us consider the following rules:

$$
h(x, y, z) \rightarrow h(z, x, y) \quad h(x, x, y) \rightarrow q(y)
$$

and instance $I=\{h(a, b, b), h(c, d, e)\}$. Observe that while $a$ and $c$ occur in the same position of atoms with the same predicate, $q(a)$ is in chase $(I, \mathcal{R})$, while $q(c)$ is not.

Second, the following looping property is central to the algorithm from [5].
Definition 1 (Looping property). An ontology $\mathcal{R}$ fulfills the looping property if it holds that for any instance $I$, for any path $a_{0} r_{1} a_{1} \ldots r_{n} a_{n}$ in chase $(I, \mathcal{R})$ such that (i) $a_{i}$ and $a_{i+1}$ are Skolem terms, (ii) $a_{i}$ is a subterm of $a_{i+1}$, and (iii) $a_{1}$ and $a_{n}$ are original constants, there exists $k \geq i$ such that $a_{k}=a_{i}$.

Indeed, DL-Lite ${ }_{\mathcal{R}}$ fulfills the looping property (as do many other DLs). However, linear rules do not, as is witnessed by Example 2.

Example 2 (Failure of looping property). Consider the instance $I_{e}=\{t(a, b)\}$ and the ruleset $\mathcal{R}_{e}$ consisting of the following rules:

$$
\begin{array}{ll}
t(x, y) \rightarrow r(y, z) & q(x, y, z) \rightarrow p(y, z) \\
r(x, y) \rightarrow q(x, y, z) & q(x, y, z) \rightarrow p(z, x)
\end{array}
$$

The chase for $I_{e}$ and $\mathcal{R}_{e}$ contains the following atoms:

$$
r\left(b, f_{1}(b)\right) \quad q\left(b, f_{1}(b), f_{2}\left(b, f_{1}(b)\right)\right) \quad p\left(f_{1}(b), f_{2}\left(b, f_{1}(b)\right)\right) \quad p\left(f_{2}\left(b, f_{1}(b)\right), b\right)
$$

There is thus a path $b r f_{1}(b) p f_{2}\left(b, f_{1}(b)\right) p b$ going from the initial constant $b$ to $b$, that passes by $f_{1}(b)$ but does not return via $f_{1}(b)$.

### 3.2 Adapting the DL-Lite $\mathcal{R}_{\mathcal{R}}$ algorithm

To take care of the first difficulty, we utilize a finer notion of type, which has similar properties to the one used in [5].

Definition 2 (Type). A type is a pair $(r, \mathcal{P})$ where $r$ is a predicate of arity $k$ and $\mathcal{P}$ is a partition of $\{1, \ldots, k\}$.

With each atom, we can associate a type, representing the way terms are repeated in the atom.

Definition 3 (Type of an atom). Let $\alpha$ be an atom, whose arity is $k$. The type of $\alpha$ is the pair $(r, \mathcal{P})$ where $p$ is the predicate of $\alpha$ and $\mathcal{P}$ is the partition of $\{1, \ldots, k\}$ such that $i$ and $j$ belong to the same partition iff the $i^{\text {th }}$ and the $j^{\text {th }}$ arguments of $\alpha$ are equal.

Note that if two atoms $\alpha_{1}$ and $\alpha_{2}$ are of same type, there exists an injective substitution $\theta_{12}$ such that $\alpha_{2}=\alpha_{1} \theta_{12}$.

Property 2. Let $I$ be an instance, and $\mathcal{R}$ be a set of linear rules. Let $\alpha_{1}$ and $\alpha_{2}$ be two atoms of $I$ of same type and $\theta_{12}$ such that $\alpha_{2}=\alpha_{1} \theta_{12}$. Then for every atom $\beta$ such that $\beta \in \operatorname{chase}\left(\left\{\alpha_{1}\right\}, \mathcal{R}\right), \beta \theta_{12} \in \operatorname{chase}\left(\left\{\alpha_{2}\right\}, \mathcal{R}\right)$.

Let us define for any atom $\alpha \in \operatorname{chase}(I, \mathcal{R})$, the restriction of chase $(I, \mathcal{R})$ to $\alpha$, denoted chase $(I, \mathcal{R})_{\mid \alpha}$, as the subset of chase $(I, \mathcal{R})$ consisting of those atoms whose terms are generated by terms $(\alpha)$. Observe that by the preceding property, if type $(\alpha)=$ type $(\beta)$, then chase $(I, \mathcal{R})_{\mid \alpha}$ is isomorphic to chase $(\{\beta\}, \mathcal{R})$.

We can overcome the second difficulty by generalizing the Loop table introduced in [5], which keeps track of the paths that occur 'below' a given type. Intuitively, a type $T$ is in the cell indexed by $\left(s_{i}, j, s_{i}^{\prime}, j^{\prime}\right)$ if and only if below any atom of type $T$, there is a path going from the term in position $j$ to the term in position $j^{\prime}$ labeled by a word that takes $\mathbb{A}$ from state $s_{i}$ to state $s_{i}^{\prime}$.

Definition 4 (Loop). Let $\mathcal{R}$ be a set of linear rules and $\mathbb{A}$ be an NFA. A Loop table has cells indexed by tuples $\left(s_{i}, j, s_{i^{\prime}}, j^{\prime}\right)$ such that $s_{i}$ and $s_{i^{\prime}}$ are states of $\mathbb{A}$ and $j$ and $j^{\prime}$ are integers between 1 and $w$, where $w$ is the maximum arity appearing in the ruleset. Cells contain types. A Loop table is:

- sound if for every $T \in\left(s_{i}, j, s_{i^{\prime}}, j^{\prime}\right)$ it holds that for every atom $\alpha$ of type $T$ appearing in some chase $\left(\left\{\alpha^{\prime}\right\}, \mathcal{R}\right)$ (with the predicate of $\alpha^{\prime}$ appearing in $\mathcal{R}$ ), there is a path $p$ in the restriction of chase $(I, \mathcal{R})$ to $\alpha$ that goes from argument $j$ of $\alpha$ to argument $j^{\prime}$ of $\alpha$ such that $\lambda(p) \in \mathcal{L}_{\mathbb{A}}\left(s_{i}, s_{i^{\prime}}\right)$.
- complete if for every atom $\alpha$ of type $T$ (whose predicate appears in $\mathcal{R}$ ), if there is path $p$ from argument $j$ to argument $j^{\prime}$ of $\alpha$ in chase $(\{\alpha\}, \mathcal{R})$ such that $\lambda(p) \in$ $\mathcal{L}_{\mathbb{A}}\left(s_{i}, s_{i^{\prime}}\right)$, then $T \in\left(s_{i}, j, s_{i^{\prime}}, j^{\prime}\right)$.

It is direct from the definition that there exists a unique sound and complete Loop table, and in what follows, we use Loop to denote this table.

The table Loop can be constructed using Algorithm 1. Line 5 initializes the table by stating than one can go from a position to the same position without reading any word
(and thus not moving in the automaton). Lines 8 and 10 correspond to going through a single edge, reading its label either as an $r$ or an $r^{-}$, in the case where both terms are distinct. Lines 13 to 16 do the same thing when both arguments are equal. Line 19 deals with unary predicates. Finally, Lines 23 and 26 saturate the table through respectively transitive closure and propagation of paths from a child to its parent.

```
Algorithm 1: Creating the Loop table
    Data: A set of linear rules \(\mathcal{R}\)
    Result: A sound and complete Loop table
    /* Initialization step */
    foreach arity \(k\) do
        foreach type \(T\) of predicate of arity \(k\) do
            for \(j \in\{1, \ldots, k\}\) do
                for \(s_{i} \in Q(\mathbb{A})\) do
                    \(\operatorname{Loop}\left(s_{i}, j, s_{i}, j\right) \leftarrow \operatorname{Loop}\left(s_{i}, j, s_{j}, j\right) \cup\{T\} ;\)
    for type \(T\) based on \(r(x, y)\) do
        if \(s_{2} \in \delta\left(s_{1}, r\right)\) then
            \(\operatorname{Loop}\left(s_{1}, 1, s_{2}, 2\right) \leftarrow \operatorname{Loop}\left(s_{1}, 1, s_{2}, 2\right) \cup\{T\} ;\)
        if \(s_{2} \in \delta\left(s_{1}, r^{-}\right)\)then
            \(\operatorname{Loop}\left(s_{1}, 2, s_{2}, 1\right) \leftarrow \operatorname{Loop}\left(s_{1}, 2, s_{2}, 1\right) \cup\{T\} ;\)
    for type \(T\) based on \(r(x, x)\) do
        if \(s_{2} \in \delta\left(s_{1}, r\right) \cup \delta\left(s_{1}, r^{-}\right)\)then
            \(\operatorname{Loop}\left(s_{1}, 1, s_{2}, 1\right) \leftarrow \operatorname{Loop}\left(s_{1}, 1, s_{2}, 1\right) \cup\{T\} ;\)
            \(\operatorname{Loop}\left(s_{1}, 1, s_{2}, 2\right) \leftarrow \operatorname{Loop}\left(s_{1}, 1, s_{2}, 2\right) \cup\{T\} ;\)
            \(\operatorname{Loop}\left(s_{1}, 2, s_{2}, 1\right) \leftarrow \operatorname{Loop}\left(s_{1}, 2, s_{2}, 1\right) \cup\{T\}\);
            \(\operatorname{Lo\circ p}\left(s_{1}, 2, s_{2}, 2\right) \leftarrow \operatorname{Lo\circ p}\left(s_{1}, 2, s_{2}, 2\right) \cup\{T\} ;\)
    for type \(T\) based on \(a(x)\) do
        if \(s_{2} \in \delta\left(s_{1}, a\right)\) then
            \(\operatorname{Loop}\left(s_{1}, 1, s_{2}, 1\right) \leftarrow \operatorname{Loop}\left(s_{1}, 1, s_{2}, 1\right) \cup\{T\} ;\)
    /* Saturation step */
    while something added do
        for T a type do
            if \(T \in \operatorname{Loop}\left(s_{1}, j_{1}, s_{2}, j_{2}\right) \cap \operatorname{Loop}\left(s_{2}, j_{2}, s_{3}, j_{3}\right)\) then
                \(\operatorname{Lo\circ p}\left(s_{1}, j_{1}, s_{3}, j_{3}\right) \leftarrow \operatorname{Lo\circ p}\left(s_{1}, j_{1}, s_{3}, j_{3}\right) \cup\{T\} ;\)
        for \(\alpha \rightarrow \beta \in \mathcal{R}\), of respective types \(T_{\alpha}, T_{\beta}\) do
            if the same variable appears in \(\alpha\) at \(i_{\alpha}\) and \(\beta\) at \(i_{\beta}\) (resp. \(j_{\alpha}\) and \(j_{\beta}\) ),
                \(T_{\beta} \in \operatorname{Loop}\left(s_{1}, i_{\beta}, s_{2}, j_{\beta}\right)\) then
                \(\operatorname{Loop}\left(s_{1}, i_{\alpha}, s_{2}, j_{\alpha}\right) \leftarrow \operatorname{Loop}\left(s_{1}, i_{\alpha}, s_{2}, j_{\alpha}\right) \cup\left\{T_{\alpha}\right\} ;\)
```

Property 3. Let $\mathcal{R}$ be a set of linear rules, $I$ be an instance and $\alpha \in I$. The following are equivalent:

1. type $(\alpha) \in \operatorname{Loop}\left(s, i, s^{\prime}, j\right)$
2. there is a path $p=a_{0} r_{1} a_{1} \ldots r_{n} a_{n}$ in chase $(I, \mathcal{R})_{\mid \alpha}$ with $a_{0}$ appearing at position $i$ in $\alpha, a_{n}$ appearing at position $j$ in $\alpha$, and $\lambda(p) \in \mathcal{L}_{\mathbb{A}}\left(s, s^{\prime}\right)$.

Proof. $(\Rightarrow)$ We prove, by induction on the order of addition of types that whenever a type is added to a cell in $\operatorname{Loop}\left(s, i, s^{\prime}, j\right)$, the second condition is fulfilled as well. If type $(\alpha)$ is added to $\operatorname{Loop}\left(s_{i}, j, s_{i}, j\right)$ at Line 5, the empty word defines a trivial path from any position existing in $\alpha$ to itself, and takes the automaton from any state to itself. If type $(\alpha)$ is added to $\operatorname{Loop}\left(s_{1}, 1, s_{2}, 2\right)$ at Line $8, \alpha$ is a binary atom of the form $r\left(e_{1}, e_{2}\right)$, and there is indeed a path from $e_{1}$ to $e_{2}$ labeled $r$. Moreover, there is a transition in $\mathbb{A}$ from $s_{1}$ to $s_{2}$ labeled by $r$, which concludes this case. The reasoning is similar for types added via Line 10 and Lines 13 to 16. If type $(\alpha)$ is added at Line 23, it must have already been added to $\operatorname{Loop}\left(s_{1}, j_{1}, s_{2}, j_{2}\right)$ and $\operatorname{Loop}\left(s_{2}, j_{2}, s_{3}, j_{3}\right)$. By the induction assumption, there is a word $w_{1}\left(\right.$ resp. $\left.w_{2}\right)$ in $\mathcal{L}_{\mathbb{A}}\left(s_{1}, s_{2}\right)\left(\right.$ resp. $\left.\mathcal{L}_{\mathbb{A}}\left(s_{2}, s_{3}\right)\right)$ that labels a path from the position $j_{1}$ (resp. $j_{2}$ ) of an atom $\alpha$ of type $T$ to the position $j_{2}$ (resp. $j_{3}$ ). Thus $w_{1} \cdot w_{2}$ labels a path from position $j_{1}$ in $\alpha$ to position $j_{3}$ in $\alpha$ and belongs to $\mathcal{L}_{\mathbb{A}}\left(s_{1}, s_{3}\right)$. Finally, let us assume that type $(\alpha)$ is added to $\operatorname{Loop}\left(s_{1}, i_{\alpha^{\prime}}, s_{2}, j_{\alpha^{\prime}}\right)$ at Line 26. By assumption, there is a rule $\alpha^{\prime} \rightarrow \beta^{\prime}$ in $\mathcal{R}$ such that $\alpha$ and $\alpha^{\prime}$ have the same type, type $\left(\beta^{\prime}\right)$ is in $\operatorname{Loop}\left(s_{1}, i_{\beta^{\prime}}, s_{2}, j_{\beta^{\prime}}\right)$, and the same variable appears at position $i_{\alpha^{\prime}}$ (resp. $j_{\alpha^{\prime}}$ ) in $\alpha^{\prime}$ and $i_{\beta^{\prime}}$ (res. $j_{\beta^{\prime}}$ ) in $\beta^{\prime}$. By the induction assumption, there is a word $w \in \mathcal{L}_{\mathbb{A}}\left(s_{1}, s_{2}\right)$ that labels a path from $i_{\beta^{\prime}}$ to $j_{\beta^{\prime}}$. Now, let us observe that any two terms that are at positions $i_{\alpha^{\prime}}$ and $j_{\alpha^{\prime}}$ of the same atom of type type $\left(\alpha^{\prime}\right)$ are also at position $i_{\beta^{\prime}}$ and $j_{\beta^{\prime}}$ of an atom of type type $\left(\beta^{\prime}\right)$ in chase $(D, \mathcal{R})_{\mid \alpha}$ because it is a model of $\alpha^{\prime} \rightarrow \beta^{\prime}$. Thus, $w$ is also the label of a path from the term at position $i_{\alpha}^{\prime}$ to the term at position $j_{\alpha}^{\prime}$, which concludes the proof.
$(\Leftarrow)$ We suppose that the second statement holds and reason by induction on the length $n$ of the path $p=a_{0} r_{1} a_{1} \ldots r_{n} a_{n}$.
Base case, path of length 0: both states and database constants are thus equal, and the type is added by the initialization in Line 5.
Base case, path of length 1: $\alpha^{\prime}=r_{1}\left(a_{0}, a_{1}\right)$ belongs to chase $(I, \mathcal{R})_{\mid \alpha}$, and $r_{1} \in$ $\mathcal{L}_{\mathbb{A}}\left(s, s^{\prime}\right)$. If $a_{0} \neq a_{1}$, then type $\left(\alpha^{\prime}\right)$ is added to the cells $\left(s, 1, s^{\prime}, 2\right)$ and $\left(s, 1, s^{\prime}, 2\right)$ in Lines 8 and 10. If $a_{0}=a_{1}$, then type $\left(\alpha^{\prime}\right)$ is added to the four cells $\left(s, i^{\prime}, s^{\prime}, j^{\prime}\right)$ with $i^{\prime}, j^{\prime} \in\{1,2\}$ (Lines 13-16). As $\alpha^{\prime}$ belongs to chase $(I, \mathcal{R})_{\mid \alpha}$, there exists a finite sequence of atoms $\alpha=\alpha_{0}, \ldots, \alpha_{m}=\alpha^{\prime}$ such that $\alpha_{i+1}$ belongs to $\rho_{i}\left(\alpha_{i}\right)$ for some rule $\rho_{i} \in \mathcal{R}$. By using $m$ applications of Line 26, we obtain type $(\alpha) \in \operatorname{Loop}\left(s, i, s^{\prime}, j\right)$. Induction step: let us assume that the result holds for any path of length up to $n-1, n \geq$ 2 , and consider the path $p=a_{0} r_{1} a_{1} \ldots r_{n} a_{n}$. First consider the case in which $a_{k}$ is contained in $\alpha$ for some $1 \leq k<n$, and let $l$ be a position of $a_{k}$ in $\alpha$. There exists a path from $a_{0}$ to $a_{k}$ of length strictly smaller than $n$, and similarly from $a_{k}$ to $a_{n}$. By the induction assumption, type $(\alpha)$ is in both Loop $\left(s, i, s^{\prime \prime}, l\right)$ and Loop $\left(s^{\prime \prime}, l, s^{\prime}, j\right)$ for some state $s^{\prime \prime}$. An application of Line 23 yields type $(\alpha) \in \operatorname{Loop}\left(s, i, s^{\prime}, j\right)$. Next suppose there is no $a_{k}(1 \leq k<n)$ that occurs in $\alpha$, and let $\beta$ be the atom in which $a_{1}$ is created (at position $k^{\prime}$ ). This atom is well defined as we consider rules with atomic head. We know that $a_{0}$ (resp. $a_{n}$ ) must occur in $\beta$, let us say at position $i^{\prime}$ (resp. $j^{\prime}$ ). Indeed, if it was not the case, $\alpha$ should contain a term among $a_{1}, \ldots, a_{n-1}$ which contradicts our earlier assumption. By the induction hypothesis, type $(\beta)$ belongs to $\operatorname{Loop}\left(s, i^{\prime}, s^{\prime \prime}, k^{\prime}\right)$ and to $\operatorname{Loop}\left(s^{\prime \prime}, k^{\prime}, s^{\prime}, j^{\prime}\right)$ for some state $s^{\prime \prime}$. Hence, by Line 23,
type $(\beta)$ is in the cell $\operatorname{Loop}\left(s, i^{\prime}, s^{\prime}, j^{\prime}\right)$. By (repeated) application of Line 26, type $(\alpha)$ is in the cell Loop $\left(s, i, s^{\prime}, j\right)$, which concludes the proof.

Property 4. Algorithm 1 runs in exponential time, and in polynomial time if the predicate arity is bounded.

Proof. There are polynomially many cells in the table, each of which can contain at most all types. The number $n_{t}$ of distinct types is single exponential (and polynomial for bounded-arity predicates). The first for loop runs in $\mathcal{O}\left(n_{t}\right)$, the next two run in polynomial time, and the while loop is performed at most $n_{t}$ times.

The remainder of the decision procedure is very close to the original algorithm for DL-Lite $_{\mathcal{R}}$, but we recall it here (Algorithm 2) in the interest of self-containment. The idea is as follows: starting from a constant $a$ and the initial state of $\mathbb{A}$, we guess the next constant in $I$ on a path from $a$ to $b$ and the state of $\mathbb{A}$ after taking this step (Line 7). We then check that this choice is valid, i.e., there is indeed a path from $a$ to the guessed constant which takes the automaton from the initial state to the current guessed state. This can be done either by a checking that a corresponding unary or binary atom is entailed (Lines 9 and 10), or by checking that a path going through the Skolem part of the chase allows us to reach the next constant in the required state, using the Loop table (Lines 12 to 14). We repeat this procedure until we reach the constant $b$ in a final state, or hit the maximal path length. Note that at Line 12, $\alpha$ is uniquely defined if it exists (it may not exist e.g., if $c$ and $d$ are different but are at positions that should have identical terms according to $T$ ).

The following property will be used to establish correctness of the algorithm.
Property 5. At the beginning of each iteration of the while loop of Algorithm 2, it holds that there is a path from $a$ to the first element of current that takes the NFA $\mathbb{A}$ from the initial state $s_{0}$ to the state in the second argument of current.

Proof. At the beginning of the first iteration of the while loop, current is equal to $\left(a, s_{0}\right)$. Thus, the path $a$, whose label is $\varepsilon$, goes from $a$ to $a$ and $\varepsilon \in \mathcal{L}_{\mathbb{A}}\left(s_{0}, s_{0}\right)$.

Let $\left(a_{i}, s_{i}\right)$ be the content of current at the beginning of the $i^{\text {th }}$ iteration of the while loop. Let $w_{i}$ be the label of a path from $a_{0}$ to $a_{i}$ such that $w_{i} \in \mathcal{L}_{\mathbb{A}}\left(s_{0}, s_{i}\right)$. If there is an $(i+1)^{\text {th }}$ iteration, either $\left(s, \sigma, s^{\prime}\right)$ or $\left(T, i_{c}, i_{d}\right)$ has been guessed, and the corresponding check was successful. Let us consider each case:

- if $\left(s, \sigma, s^{\prime}\right)$ has been guessed and checked, we have two cases:
- $\sigma \in \mathcal{P}_{2}^{ \pm}$, and there is a path from $a_{i}$ to $a_{i+1}$ in chase $(I, \mathcal{R})$ labeled by $\sigma$. Moreover, $\sigma$ labels an edge from $s$ to $s^{\prime}$ in $\mathbb{A}$. We can thus define $w_{i+1}=w_{i} . \sigma$ - $\sigma=A$, and $I, \mathcal{R} \models A(c)$. As $c=d$, we can again define $w_{i+1}=w_{i} . \sigma$
- if $\left(T, i_{c}, i_{d}\right)$ has been guessed, it means that $T$ belongs to $\operatorname{Lo\circ p}\left(s_{i}, i_{c}, s_{i+1}, i_{d}\right)$. By the definition of Loop, there is a path $p$ (in the Skolem part) from any term at position $i_{c}$ of an atom of type $T$ to the position $i_{d}$ of an atom of type $T$ such that $\lambda(p) \in \mathcal{L}_{\mathbb{A}}\left(s, s^{\prime}\right)$. Let $\alpha$ be as defined Line 12 . As $I, \mathcal{R} \models \alpha$, where type $(\alpha)=T$, $a_{i}$ appears at position $i_{c}$ of $\alpha$, and $a_{i+1}$ appears at position $i_{d}$ of $\alpha$, there is such a path from $a_{i}$ to $a_{i+1}$. We can thus set $w_{i+1}=w_{i} . p$.

```
Algorithm 2: RPQ answering over linear rules
    Input: An NFA \(\mathbb{A}\), an instance \(I\), a set of linear rules \(\mathcal{R},(a, b) \in \operatorname{terms}(I) \times \operatorname{terms}(I)\)
    Output: Yes if and only if \((a, b)\) is a certain answer to the query \(q\) defined by \(\mathbb{A}\)
    if \((I, \mathcal{R})\) is not satisfiable then
        return Yes
    current \(=\left(a, s_{0}\right)\);
    count \(=0, \max =|\mathbb{A}| \times|I| ;\)
    while count < max and current \(\notin\left\{\left(b, s_{f}\right) \mid s_{f} \in F\right\}\) do
        Define ( \(c, s\) ) = current;
        Guess \(\left(d, s^{\prime}\right)\) together with \(\left(s, \sigma, s^{\prime}\right) \in \delta\) or \(T, i_{c}, i_{d}\) such that
            \(T \in \operatorname{Loop}\left(s, i_{c}, s^{\prime}, i_{d}\right)\);
        if \(\left(s, \sigma, s^{\prime}\right)\) was guessed then
            if \(\sigma \in \mathcal{P}_{2}^{ \pm} \wedge(I, \mathcal{R} \not \vDash \sigma(c, d))\) then return \(N o\);
            if \(\sigma=A \wedge(c \neq d \vee I, \mathcal{R} \not \vDash A(c))\) then return \(N o\);
        if \(T, i_{c}, i_{d}\) was guessed then
            Let \(\alpha\) be of type \(T\) such that \(c\) is at position \(i_{c}\) and \(d\) is at position \(i_{d}\); other terms
                are set to fresh variables
                if \(\alpha\) does not exist then return No;
                if \(I, \mathcal{R} \notin \alpha\) then return \(N o\);
        current \(=\left(d, s^{\prime}\right)\), count \(=\) count +1 ;
    if current \(=\left(b, s_{f}\right)\) for some \(s_{f} \in F\) then return Yes else return \(N o\);
```

Property 6. There is an execution of Algorithm 2 that outputs Yes iff the RPQ given by $\mathbb{A}$ is entailed from $(I, \mathcal{R})$.

Proof. $(\Rightarrow)$ If the algorithm outputs Yes, the while loop has been exited with current equal to $\left(b, s_{f}\right)$, with $s_{f}$ a final state of $\mathbb{A}$. By Property 5 , this means that there is a path from $a$ to $b$ whose label takes $\mathbb{A}$ from $s_{0}$ to $s_{f}$, hence is accepted by $\mathbb{A}$. This show that whenever Algorithm 2 accepts, $(a, b)$ is a certain answer to the RPQ given by $\mathbb{A}$.
$(\Leftarrow)$ If $(a, b)$ is a certain answer to the RPQ based upon $\mathbb{A}$, then there is path of minimal length $p=a_{0}^{\prime} r_{1} a_{1}^{\prime} \ldots r_{n} a_{n}^{\prime}$ from $a=a_{0}^{\prime}$ to $b=a_{n}^{\prime}$ in chase $(I, \mathcal{R})$ such that $\lambda(p)=r_{1} \ldots r_{n} \in \mathcal{L}_{\mathbb{A}}\left(s_{0}, s_{f}\right)$ for some final state $s_{f}$. Let $s_{0}^{\prime} s_{1}^{\prime} \ldots s_{n}^{\prime}$ be a sequence of states of $\mathbb{A}$ such that $s_{n}^{\prime}$ is a final state of $\mathbb{A}$ and for every $1 \leq i \leq n,\left(s_{i-1}, r_{i}, s_{i}\right) \in \delta$. Since $p$ is of minimal length, there is no pair $(i, j)$ with $i \neq j$ such that $\left(a_{i}, s_{i}\right)=$ $\left(a_{j}, s_{j}\right)$. Let us consider the sequence $p^{\prime}=\left(\left(a_{i}, s_{i}\right)\right)_{i}$ such that:

- for any $i, a_{i}$ is the $i^{\text {th }}$ constant, say $a_{k_{i}}^{\prime}$, in $p$ belonging to terms $(I)$;
- for any $i, s_{i}=s_{k_{i}}^{\prime}$.

Moreover, for any $i$, if $k_{i+1}=k_{i}+1$, we define aux ${ }_{i}=\left(s_{i}, r_{i+1}, s_{i+1}\right)$. Otherwise, let aux ${ }_{i}=\left(\operatorname{type}(\alpha), i_{c}, i_{d}\right)$,where:

- $\alpha$ is such that $\alpha \in I$ and type $(\alpha) \in \operatorname{Loop}\left(s_{i}, i_{c}, s_{i+1}, i_{d}\right)$;
- $a_{k_{i}}$ appears at position $i_{c}$ of $\alpha$ and $a_{k_{i+1}}$ appears at position $i_{d}$ of $\alpha$.

In the second case, it is possible to define aux $_{i}$ in such a way, as the path $p_{s}=$ $a_{k_{i}}^{\prime} r_{k_{i}+1} \ldots a_{k_{i+1}}^{\prime}$ goes from $a_{k_{i}}$ to $a_{k_{i+1}}$ and belongs to $\mathcal{L}_{\mathbb{A}}\left(s_{i}, s_{i+1}\right)$ by definition
of $s_{i}$. We show that the sequence of guesses $\left(a_{i}, s_{i}\right.$, aux $\left._{i}\right)$ leads Algorithm 2 to accept. Since $p$ is minimal, the length of $p^{\prime}$ is less than $|\mathbb{A}| \times|I|$. Moreover, $a_{n}=b$ and $s_{f}$ is a final state. Thus, the only way for Algorithm 2 to reject with this sequence of guesses is to reject during checks, i.e., one of the checks performed at Lines 9, 10, 12 or 14 fails. Let $\left(a_{i}, s_{i}\right.$, aux $\left._{i}\right)$ be the guess at one of the steps. If aux ${ }_{i}$ is of the form $\left(s_{i}, r_{i+1}, s_{i+1}\right)$, then $a_{k_{i}}$ and $a_{k_{i+1}}$ are consecutive elements in $p$, and there is an atom $r_{i+1}\left(a_{k_{i}}, a_{k_{i+1}}\right)$ in chase $(I, \mathcal{R})$. Thus, $r_{i+1}\left(a_{k_{i}}, a_{k_{i+1}}\right)$ is entailed by $I$ and $\mathcal{R}$, and the check at Line 9 or 10 (depending on $r_{i+1}$ being a binary or unary atom) is successful. If aux $x_{i}$ is of the form $\left(\operatorname{type}(\alpha), i_{c}, i_{d}\right)$, then there is $\alpha \in I$ such that type $(\alpha) \in \operatorname{Loop}\left(s_{i}, i_{c}, s_{i+1}, i_{d}\right)$, and with $a_{k_{i}}$ (resp. $a_{k_{i+1}}$ ) appearing at position $i_{c}$ (resp. $i_{d}$ ) of $\alpha$. The atom $\alpha$ fulfills the conditions of Lines 12 and 14. Thus the defined sequence never triggers a rejection from Algorithm 2, which concludes the proof.

Theorem 1. RPQ Answering in the presence of linear existential rules is:

- in NL in data complexity
- in PTIME in combined complexity with bounded arity
- in EXPTIME in combined complexity with unbounded arity

Proof. Algorithm 2 is a non-deterministic algorithm that needs to keep in memory the current state, the current constant, and the number of iterations done so far. It performs two types of operations: entailment checks and accessing the contents of the Loop table (more precisely, deciding whether $T \in \operatorname{Lo\circ p}\left(s, i_{c}, s^{\prime}, i_{d}\right)$ ). Hence, it can be seen as an NL algorithm making oracle calls whenever an entailment check is performed or a cell of Loop is retrieved. Entailment checks are in NL in data complexity, and Loop is independent from the data: the overall algorithm thus runs in NL in data complexity. In combined complexity with bounded arity, entailment checks can be performed in PTime, while Loop can be computed in polynomial time: the overall algorithm is thus in PTime with bounded arity. In the unbounded arity case, the entailment checks can be performed in PSpace, while the Loop table can be computed in ExpTime: the algorithm thus runs in EXPTIME.

## 4 Lower Bound

It is already known that the data complexity (resp. combined complexity) of RPQs under linear rules (resp. linear rules with bounded arity) is NL-hard (resp. PTime-hard) [5], which matches the upper bounds obtained in the preceding section. We thus focus on providing a matching EXPTIME lower bound for the combined complexity of evaluating RPQs under linear rules of unbounded arity. The proof is done by simulating an alternating PSPACE TM. It is already known that PSPACE TMs can be simulated by means of linear rules [12]. In the following, we explain how to adapt this construction to simulate alternating TMs. Note that in this section, we will use rules with multiple atoms in the head: this is done to simplify the presentation, and a classical transformation allows us to get the same lower bound for rules with atomic heads.

The intuition is as follows: the construction in [12] represents the configuration of a $\mathrm{TM} \mathcal{M}$ by a single atom of polynomial arity. The initial configuration can thus be
represented by an instance $I_{\mathcal{M}}$ containing a single atom. Then, for each transition of the TM, polynomially many linear rules are created, each one representing the action of the transition on a cell at a given position. All these rules are part of $\mathcal{R}_{\mathcal{M}}$. The initial configuration of the TM is accepted if and only if an atom encoding a configuration having an accepting state is entailed by $I_{\mathcal{M}}$ and $\mathcal{R}_{\mathcal{M}}$.

We modify this construction in the following way to deal with alternating Turing machines: to each atom, we add two positions, that will act as "input" and "output" positions. Moreover, we will maintain the following property: there is a path, whose edges are all labeled by the same predicate $p$, from the input position of $\alpha$ to the output position of $\alpha$ entailed by chase $\left(I_{\{\alpha\}}, \mathcal{R}_{\mathcal{M}}\right)$ if and only if the configuration represented by $\alpha$ is accepted by $\mathcal{M}$. This is true in the following cases:

- the state of the current configuration is accepting. It is then enough to add a $p$-edge from $i_{c}$ to $o_{c}$; this is possible as the Turing machine is assumed to never leave an accepting state;
- the current state is existential and one of the two successor configurations is accepting: we thus add $p$-edges from the input of the current configuration to the input of the two children, and from the output of the two children to the output of the current configuration;
- the current state is universal, and both successor configurations are accepting: we thus add $p$-edges from the input of the current configuration to the input of the first successor configuration, then from the output of that configuration to the input of the other successor, and lastly from the output of the second successor to the output of the current configuration.

We now formalize the construction sketched above, staying as close as possible to the notations in [12].

Turing Machine Given an alternating PSPACE TM and an input $x$, we can represent a configuration $c$ reached during the computation by storing the content of the first $p(|x|)$ cells, as well as the position of the head of the tape and the current state of the TM. Adding input and output positions, this can be encoded by a predicate conf of arity $2 p(|x|)+3$ :

$$
\operatorname{conf}\left(i_{c}, \text { state }, \operatorname{cell}_{1}, \operatorname{cur}_{1}, \operatorname{cell}_{2}, \operatorname{cur}_{2}, \ldots, \operatorname{cell}_{p(|x|)}, \operatorname{cur}_{p(|x|)}, o_{c}\right)
$$

where state contains the state identifier, cell $_{i}$ represents the content of the $i^{\text {th }}$ cell, $c u r_{i}$ is equal to 1 if the head of the Turing machine is on cell $i$ and 0 otherwise, and $i_{c}$ and $o_{c}$ are the input and output terms of this atom. We say that the above atom represents configuration $c$. Given an atom $\alpha$, the term at its input (resp. output) position is denoted by $i(\alpha)$ (resp. $o(\alpha)$ ). We denote by $I_{\mathcal{M}, x}$ the instance containing a single atom representing the initial configuration of $\mathcal{M}$ on input $x$.

For every state $q_{f}$ with $g\left(q_{f}\right)=$ accept, we create the following rule:

$$
\begin{equation*}
\operatorname{conf}\left(i_{c}, q_{f}, \ldots, o_{c}\right) \rightarrow p\left(i_{c}, o_{c}\right) \tag{1}
\end{equation*}
$$

For each transition $\delta(q, \gamma)=\left\{\left(q^{\prime}, \gamma^{\prime}, L\right),\left(q^{\prime \prime}, \gamma^{\prime \prime}, L\right)\right\}$ such that $g(q)=\vee$, we create the rule

$$
\begin{align*}
& \operatorname{conf}\left(i_{c}, q, \operatorname{cell}_{1}, \text { cur }_{1}, \ldots, \operatorname{cell}_{i-1}, 0, \gamma, 1, \ldots, o_{c}\right) \rightarrow \\
& \exists i_{c^{\prime}}, o_{c^{\prime}}, i_{c^{\prime \prime}}, o_{c^{\prime \prime}} \operatorname{conf}\left(i_{c^{\prime}}, q^{\prime}, \operatorname{cell}_{1}, \operatorname{cur}_{1}, \ldots, \operatorname{cell}_{i-1}, 1, \gamma^{\prime}, 0, \ldots, o_{c^{\prime}}\right) \\
& \operatorname{conf}\left(i_{c^{\prime \prime}}, q^{\prime \prime}, \operatorname{cell}_{1}, \operatorname{cur}_{1}, \ldots, \operatorname{cell}_{i-1}, 1, \gamma^{\prime}, 0, \ldots, o_{c^{\prime \prime}}\right) \\
& \quad p\left(i_{c}, i_{c^{\prime}}\right), p\left(o_{c^{\prime}}, o_{c}\right), p\left(i_{c}, i_{c^{\prime \prime}}\right), p\left(o_{c^{\prime \prime}}, o_{c}\right) . \tag{2}
\end{align*}
$$

for each position $i$ on the tape, and similarly when the head is moving to the right.
When $g(q)=\wedge$, we associate with each transition $\delta(q, \gamma)=\left\{\left(q^{\prime}, \gamma^{\prime}, L\right),\left(q^{\prime \prime}, \gamma^{\prime \prime}, L\right)\right\}$ the following rule:

$$
\begin{align*}
& \operatorname{conf}\left(i_{c}, q, \operatorname{cell}_{1}, \operatorname{cur}_{1}, \ldots, \operatorname{cell}_{i}, 0, \gamma, 1, \ldots, o_{c}\right) \rightarrow \\
& \quad \exists i_{c^{\prime}}, o_{c^{\prime}}, i_{c^{\prime \prime}}, o_{c^{\prime \prime}} \operatorname{conf}\left(i_{c^{\prime}}, q^{\prime}, \operatorname{cell}_{1}, \operatorname{cur}_{1}, \ldots, \operatorname{cell}_{i}, 1, \gamma^{\prime}, 0, \ldots, o_{c^{\prime}}\right) \\
& \qquad \operatorname{conf}\left(i_{c^{\prime \prime}}, q^{\prime \prime}, \operatorname{cell}_{1}, \operatorname{cur}_{1}, \ldots, \operatorname{cell}_{i}, 1, \gamma^{\prime \prime}, 0, \ldots, o_{c^{\prime \prime}}\right) \\
& \quad p\left(i_{c}, i_{c^{\prime}}\right), p\left(o_{c^{\prime}}, i_{c^{\prime \prime}}\right), p\left(o_{c^{\prime \prime}}, o_{c}\right) . \tag{3}
\end{align*}
$$

Figure 1 illustrates the functioning of rules of types (2) and (3). We denote by $\mathcal{R}_{\mathcal{M}, x}$ the set containing all the rules defined above ${ }^{4}$. The above rules (where input and output positions are removed) simulate the run of a PSPACE TM [12].

The following property formalizes the reduction and establishes its correctness.


Fig. 1. Existential (left) and universal (right) gadgets

Property 7. Let $\mathcal{M}$ be an alternating PSPACE Turing machine, and let $\alpha$ be an atom of chase $\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$ representing a configuration $c(\alpha)$. Then $c(\alpha)$ is an accepting configuration of $\mathcal{M}$ if and only if there is a path in chase $\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$ from $i(\alpha)$ to $o(\alpha)$ whose label belongs to $p^{*}$.

Proof. $(\Leftarrow)$ Let $\alpha \in \operatorname{chase}\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$ represent a configuration $c(\alpha)$, and let $C_{\alpha}$ be the restriction of chase $\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$ to $\alpha$. We show by induction on the number of atoms of $C_{\alpha}$ that the required path exists. Note that the induction is well-founded as the Skolem chase is finite (recall that the considered Turing machines terminate).

[^1]- If $C_{\alpha}$ contains one atom, then there can be no path in chase $\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$ witnessing $p^{*}(i(\alpha), o(\alpha))$. Suppose then that $C_{\alpha}$ contains two atoms. In this case, the only atom in $C_{\alpha}$ other than $\alpha$ must be $p(i(\alpha), o(\alpha))$. The only way to derive such an atom is to apply a rule of the form (1), which is applied if and only if $c(\alpha)$ is in an accepting state, hence $c(\alpha)$ is an accepting configuration of $\mathcal{M}$.
- Next assume that the result holds for any atom $\alpha$ such that $C_{\alpha}$ has less than $n$ atoms, and let $\alpha$ be an atom such that $C_{\alpha}$ contains $n$ atoms. We distinguish two cases:
- Case 1: the state of $c(\alpha)$ is existential. Then, since the rules of type (2) must be satisfied, $C_{\alpha}$ contains atoms $\alpha_{1}$ and $\alpha_{2}$ representing the successor configurations of $c(\alpha)$. The existence of a path from $i(\alpha)$ to $o(\alpha)$ implies that there is either a path from $i\left(\alpha_{1}\right)$ to $o\left(\alpha_{1}\right)$ or a path from $i\left(\alpha_{2}\right)$ to $o\left(\alpha_{2}\right)$. To see why, observe that every $p$-atom involving $i(\alpha)$ or $o(\alpha)$ is added either by the same rule application as created $\alpha$ or by a rule of type (2) applied to $\alpha$. Only atoms of the second kind (refer to Fig. 1, left) can belong to a shortest path from $i(\alpha)$ to $o(\alpha)$, as atoms of the first kind have $i(\alpha)$ (resp. $o(\alpha))$ as second (resp. first) argument. If we have a path from $i\left(\alpha_{1}\right)$ to $o\left(\alpha_{1}\right)$, then we can apply the induction assumption to $\alpha_{1}$ to get that $c\left(\alpha_{1}\right)$ is an accepting configuration, which implies that $c(\alpha)$ is also accepting. We can proceed analogously if we have path from $i\left(\alpha_{2}\right)$ to $o\left(\alpha_{2}\right)$.
- Case 2: the state of $c(\alpha)$ is universal. As the rules of type (3) must be satisfied, the existence of a path from $i(\alpha)$ to $o(\alpha)$ implies the existence of a path from $i\left(\alpha_{1}\right)$ to $o\left(\alpha_{1}\right)$ and a path from $i\left(\alpha_{2}\right)$ to $o\left(\alpha_{2}\right)$, where $\alpha_{1}$ and $\alpha_{2}$ represent the successor configurations of $c(\alpha)$ (refer to Fig. 1, right). By the induction assumption, $c\left(\alpha_{1}\right)$ and $c\left(\alpha_{2}\right)$ are both accepting configurations, which means that $c(\alpha)$ is also accepting.
$(\Rightarrow)$ We prove the other direction by induction on the number of transitions that need to be performed to prove that $c(\alpha)$ is accepted by $\mathcal{M}$.
- If no transitions are required, this means that $c(\alpha)$ is in an accepting state. Thus, Rule (1) is applicable, and $p(i(\alpha), o(\alpha))$ is present in chase $\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$.
- Assume the result holds up to $n$ required transitions. We distinguish two cases:
- Case 1: the state of $c(\alpha)$ is existential. As $c(\alpha)$ is accepting, this means that one of its two successor configurations, say $c\left(\alpha_{1}\right)$, is accepting. Moreover, the number of transitions required to accept $c\left(\alpha_{1}\right)$ is strictly smaller than for $c(\alpha)$. By the induction assumption, $p^{*}\left(i\left(\alpha_{1}\right), o\left(\alpha_{1}\right)\right)$ is present in chase $\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$. As $p\left(i(\alpha), i\left(\alpha_{1}\right)\right)$ and $p\left(o\left(\alpha_{1}\right), o(\alpha)\right)$ are also present (since the rules of the form (2) generate them), this proves that $p^{*}(i(\alpha), o(\alpha))$ is present as well.
- Case 2: the state of $c(\alpha)$ is universal. As $c(\alpha)$ is accepting, this means that its two successor configuration are also accepting. By the induction assumption, this means that $p^{*}\left(i\left(\alpha_{1}\right), o\left(\alpha_{1}\right)\right)$ and $p^{*}\left(i\left(\alpha_{2}\right), o\left(\alpha_{2}\right)\right)$ are present in chase $\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$. As the rules of the form (3) also generate $p\left(i(\alpha), i\left(\alpha_{1}\right)\right)$, $p\left(o\left(\alpha_{1}\right), i\left(\alpha_{2}\right)\right)$, and $p\left(o\left(\alpha_{2}\right), o(\alpha)\right)$, this proves that $p^{*}(i(\alpha), o(\alpha))$ is present in chase $\left(I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x}\right)$.

Now let $\mathcal{M}$ be an alternating PSpace Turing machine, $x$ be an input to $\mathcal{M}$, and $\alpha$ be the unique atom in $I_{\mathcal{M}, x}$. Then by Property $7, c(\alpha)$ is an accepting configuration
of $\mathcal{M}$ if and only if $I_{\mathcal{M}, x}, \mathcal{R}_{\mathcal{M}, x} \models p^{*}(i(\alpha), o(\alpha))$. This, together with known results, yields the following lower bounds:

Theorem 2. RPQ Answering in the presence of linear existential rules is NL-hard in data complexity, PTIME-hard in combined complexity with bounded arity and EXPTIMEhard in combined complexity without arity bound, even for a fixed $R P Q$.

Note that the preceding reduction can be easily adapted to show that atomic query answering under rulesets containing linear rules and transitivity rules is ExPTime-hard. Assuming ExpTime $\neq$ PSpace, this result is in contradiction with Theorem 5 in [2], which purports to show a PSPACE upper bound. Indeed, after reexamining the proofs, the authors of the latter work have identified the flaw, which occurs in the analysis of the combined complexity of their rewriting-based decision procedure. It turns out that the procedure runs in exponential time, rather than in polynomial space (the NL upper bound in data complexity remains valid). Combining our lower bound with their procedure shows that the problem is ExpTIME-complete in combined complexity.

## 5 Conclusion and Future Work

In this paper, we have investigated the complexity of evaluating regular path queries under linear existential rules. We have shown that it is NL-complete in data complexity, PTIME-complete in combined complexity when the predicate arity is bounded, and EX-PTIME-complete otherwise. This behavior is somewhat surprising with respect to prior work: indeed, for DL-Lite ${ }_{\mathcal{R}}$, the combined complexity of RPQ answering is lower than for CQs, whereas we observe just the opposite in the linear case (recall CQ answering is PSPACE-complete under linear rules). The upper bound was shown by adapting an existing decision procedure for DL-Lite, using a refined definition of type. The lower bound builds upon a PSPACE-hardness result for CQ answering under linear rules.

There are two natural ways to extend the present work: either investigate more expressive forms of path queries (with conjunction and/or nesting) over linear rules, or consider the effect of moving to more expressive decidable classes of existential rules.

Acknowledgements. This work was supported by the ANR project 12 JS02 00701.

## References

1. Baget, J., Leclère, M., Mugnier, M., Salvat, E.: Extending decidable cases for rules with existential variables. In: Proc. of IJCAI. pp. 677-682 (2009)
2. Baget, J., Bienvenu, M., Mugnier, M., Rocher, S.: Combining existential rules and transitivity: Next steps. In: Proc. of IJCAI. pp. 2720-2726 (2015)
3. Bienvenu, M., Calvanese, D., Ortiz, M., M.Šimkus: Nested regular path queries in description logics. In: Proc. of KR (2014)
4. Bienvenu, M., Ortiz, M.: Ontology-mediated query answering with data-tractable description logics. In: Reasoning Web. pp. 218-307 (2015)
5. Bienvenu, M., Ortiz, M., Simkus, M.: Regular path queries in lightweight description logics: Complexity and algorithms. J. Artif. Intell. Res. (JAIR) 53, 315-374 (2015)
6. Bourhis, P., Krötzsch, M., Rudolph, S.: How to best nest regular path queries. In: Proc. of DL. pp. 404-415 (2014)
7. Calì, A., Gottlob, G., Kifer, M.: Taming the infinite chase: Query answering under expressive relational constraints. In: Proc. of KR. pp. 70-80 (2008)
8. Calvanese, D., Eiter, T., Ortiz, M.: Answering regular path queries in expressive description logics: An automata-theoretic approach. In: Proc. of AAAI. pp. 391-396 (2007)
9. Calvanese, D., Eiter, T., Ortiz, M.: Regular path queries in expressive description logics with nominals. In: Proc. of IJCAI. pp. 714-720 (2009)
10. Calvanese, D., Eiter, T., Ortiz, M.: Answering regular path queries in expressive description logics via alternating tree-automata. Inf. Comput. 237, 12-55 (2014)
11. Florescu, D., Levy, A., Suciu, D.: Query containment for conjunctive queries with regular expressions. In: Proc. of PODS (1998)
12. Gottlob, G., Papadimitriou, C.H.: On the complexity of single-rule datalog queries. Inf. Comput. 183(1), 104-122 (2003)
13. Grau, B.C., Horrocks, I., Krötzsch, M., Kupke, C., Magka, D., Motik, B., Wang, Z.: Acyclicity notions for existential rules and their application to query answering in ontologies. J. Artif. Intell. Res. (JAIR) 47, 741-808 (2013)
14. Johnson, D.S., Klug, A.C.: Testing containment of conjunctive queries under functional and inclusion dependencies. J. Comput. Syst. Sci. 28(1), 167-189 (1984)
15. Kostylev, E.V., Reutter, J.L., Vrgoc, D.: XPath for DL ontologies. In: Proc. of AAAI (2015)
16. Maier, D., Mendelzon, A.O., Sagiv, Y.: Testing implications of data dependencies. ACM Trans. Database Syst. 4(4), 455-469 (1979)
17. Marnette, B.: Generalized schema-mappings: from termination to tractability. In: Proc. of PODS. pp. 13-22 (2009)
18. Ortiz, M., Rudolph, S., Šimkus, M.: Query answering in the Horn fragments of the description logics $\mathcal{S H O} \mathcal{I}$ and $\mathcal{S R O I Q}$. In: Proc. of IJCAI (2011)
19. Stefanoni, G., Motik, B., Krötzsch, M., Rudolph, S.: The complexity of answering conjunctive and navigational queries over OWL 2 EL knowledge bases. J. of Art. Intell. Res. (JAIR) 51, 645-705 (2014)

[^0]:    ${ }^{3}$ As we only consider the two-way variant, we will use the abbreviation RPQ instead of the more traditional 2RPQ.

[^1]:    ${ }^{4}$ Note that $x$ is required to determine the arity of conf.

