Combining Existential Rules and Transitivity: Next Steps

Jean-François Baget
INRIA &
Univ. Montpellier II
Montpellier, France

Meghyn Bienvenu
CNRS &
Univ. Paris-Sud
Orsay, France

Marie-Laure Mugnier
INRIA &
Univ. Montpellier II
Montpellier, France

Swan Rocher
INRIA &
Univ. Montpellier II
Montpellier, France

Abstract

We consider existential rules (aka Datalog$\pm$) as a formalism for specifying ontologies. In recent years, many classes of existential rules have been exhibited for which conjunctive query (CQ) entailment is decidable. However, most of these classes cannot express transitivity of binary relations, a frequently used modelling construct. In this paper, we address the issue of whether transitivity can be safely combined with decidable classes of existential rules. First, we prove that transitivity is incompatible with one of the simplest decidable classes, namely aGRD (acyclic graph of rule dependencies), which clarifies the landscape of 'finite expansion sets' of rules. Second, we show that transitivity can be safely added to linear rules (a subclass of guarded rules, which generalizes the description logic DL-Lite$_R$) in the case of atomic CQs, and also for general CQs if we place a minor syntactic restriction on the rule set. This is shown by means of a novel query rewriting algorithm that is specially tailored to handle transitivity rules. Third, for the identified decidable cases, we pinpoint the combined and data complexities of query entailment.

1 Introduction

Ontology-based data access (OBDA) is a new paradigm in data management, which exploits the semantic information provided by ontologies when querying data. Briefly, the notion of a database is replaced by that of a knowledge base (KB), composed of a dataset and an ontology. *Existential rules*, aka Datalog$\pm$, have been proposed to represent ontological knowledge in this context [Calì *et al.*, 2009; Baget *et al.*, 2009; Baget *et al.*, 2011b; Krötzsch and Rudolph, 2011]. These rules are an extension of function-free first-order Horn rules (aka Datalog), that allows for existentially quantified variables in rule heads. The addition of existential quantification allows one to assert the existence of yet unknown entities and to reason about them, an essential feature of ontological languages, which is also at the core of description logics (DLs). Existential rules generalize the DLs most often considered in the OBDA setting, like the DL-Lite and $\mathcal{EL}$ families [Calvanese *et al.*, 2007; Baader, 2003; Lutz *et al.*, 2009] and Horn DLs [Krötzsch *et al.*, 2007].

The fundamental decision problem related to OBDA is the following: is a Boolean conjunctive query (CQ) entailed from a KB? This problem has long been known to be undecidable for general existential rules (this follows e.g., from [Beeri and Vardi, 1981]). Consequently, a significant amount of research has been devoted to the issue of finding decidable subclasses with a good expressivity / tractability tradeoff. It has been observed that most exhibited decidable classes fulfill one of the three following properties [Baget *et al.*, 2011a]: finiteness of a forward chaining mechanism known as the chase, which allows inferences to be materialized in the data (we call such rule sets *finite expansion sets*, *fes*); finiteness of query rewriting into a union of CQs, which allows to the rules to be compiled into the query (*finite unification sets*, *fus*); tree-like shape of the possibly infinite chase, which allows one to finitely encode the result (*bounded-treewidth sets*, *bts)*. The class of *guarded* rules [Calì *et al.*, 2008] is a well-known class satisfying the latter property.

Known decidable classes are able to express many useful properties of binary relations (e.g., inverses / symmetry) but most of them lack the ability to define a frequently required property, namely *transitivity*. This limits their applicability in key application areas like biology and medicine, for which transitivity of binary relations (especially the ubiquitous 'part of' relation) is an essential modelling construct. The importance of transitivity has long been acknowledged in the DL community [Horrocks and Sattler, 1999; Sattler, 2000], and many DLs support transitive binary relations. While adding transitivity to a DL often does not increase the complexity of CQ entailment (see [Eiter *et al.*, 2009] for some exceptions), it is known to complicate the design of query answering procedures [Glimm *et al.*, 2008; Eiter *et al.*, 2012], due to the fact that it destroys the tree structure of the chase upon which DL reasoning algorithms typically rely. In contrast to the extensive literature on transitivity in DLs, rather little is known about the compatibility of transitivity with decidable classes of existential rules. A notable exception is the recent result of [Gottlob *et al.*, 2013] on the incompatibility of transitivity with guarded rules, which holds even under strong syntactic restrictions (see Section 3).

In this paper, we investigate the issue of whether transitivity can be safely added to some well-known rule classes

and provide three main contributions. First, we show that adding transitivity to one of the simplest *fes* and *fus* classes (namely *aGRD*) makes atomic CQ entailment undecidable, which clarifies the issue for *fes* classes (Theorem 1). Second, we investigate the impact of adding transitivity to *linear* rules, a natural subclass of guarded rules which generalizes the well-known description logic DL-Lite$_R$. We introduce a query rewriting procedure that is sound and complete for all rule sets consisting of linear and transitivity rules (Theorem 2), and which is guaranteed to terminate for atomic CQs, and for arbitrary CQs if the rule set contains only unary and binary predicates or satisfies a certain safety condition, yielding decidability for these cases (Theorem 3). Third, based on a careful analysis of our algorithm, we establish tight bounds on the combined and data complexities of query entailment for the identified decidable cases. The obtained complexities are the lowest that could be expected, namely, NL in data complexity and PSPACE in combined.

2 Preliminaries

A *term* is a variable or a constant. An *atom* is of the form $p(t_1, \ldots, t_k)$ where p is a predicate of arity k, and the t_i are terms. We consider *(unions of) Boolean conjunctive queries ((U)CQs)*, which are (disjunctions of) existentially closed conjunctions of atoms. Note however that all results can be extended to non-Boolean queries. A CQ is often viewed as the *set* of atoms. An *atomic CQ* is a CQ consisting of a single atom. A *fact* is an atom without variables. A *fact base* is a finite set of facts.

An *existential rule* (hereafter abbreviated to *rule*) R is a formula $\forall\vec{x}\forall\vec{y}(B[\vec{x},\vec{y}] \rightarrow \exists\vec{z}\, H[\vec{x},\vec{z}])$ where B and H are conjunctions of atoms, resp. called the *body* and the *head* of R. The variables $\vec{z}$ (resp. $\vec{x}$), which occur only in H (resp. in B and in H) are called *existential* variables (resp. *frontier* variables). Hereafter, we omit quantifiers in rules and simply denote a rule by $B \rightarrow H$. For example, $p(x,y) \rightarrow p(x,z)$ stands for $\forall x\forall y(p(x,y) \rightarrow \exists z(p(x,z)))$. A *knowledge base* (KB) $\mathcal{K} = (\mathcal{F}, \mathcal{R})$ consists of a fact base $\mathcal{F}$ and a finite set of rules $\mathcal{R}$. The *(atomic) CQ entailment* problem consists in deciding whether $\mathcal{K} \models Q$, where $\mathcal{K}$ is a KB viewed as a first-order theory, Q is an (atomic) CQ, and $\models$ denotes standard logical entailment.

Query rewriting relies on a unification operation between the query and a rule head. Care must be taken when handling existential variables: when a term t of the query is unified with an existential variable in a rule head, all atoms in which t occurs must also be part of the unification, otherwise the result is unsound. Thus, instead of unifying one query atom at a time, we have to unify subsets ("pieces") of the query, hence the notion of a piece-unifier defined next. A partition P of a set of terms is said to be *admissible* if no class of P contains two constants; a substitution σ can be obtained from P by selecting an element e_i in each class C_i of P, with priority given to constants, and setting $\sigma(t) = e_i$ for all $t \in C_i$. A *piece-unifier* of a CQ Q with a rule $R = B \rightarrow H$ is a triple $\mu = (Q', H', P_\mu)$, where $Q' \subseteq Q$, $H' \subseteq H$ and P_μ is an admissible partition on the terms of $Q' \cup H'$ such that:

1. $\sigma(H') = \sigma(Q')$, where σ is any substitution obtained from P_μ;
2. if a class C_i in P_μ contains an existential variable, then the other terms in C_i are variables from Q' that do not occur in $(Q \setminus Q')$.

We say that Q' is a *piece* (and μ is a single-piece unifier) if there is no non-empty subset Q'' of Q' such that P_μ restricted to Q'' satisfies Condition 2. From now on, we consider only single-piece unifiers, which we simply call *unifiers*. The *(direct) rewriting* of Q with R w.r.t. μ is $\sigma(Q \setminus Q') \cup \sigma(B)$ where σ is a substitution obtained from P_μ. A *rewriting* of Q w.r.t. a set of rules $\mathcal{R}$ is a CQ obtained by a sequence $Q = Q_0, \ldots, Q_n$ ($n \geq 0$) where for all $i > 0$, Q_i is a direct rewriting of Q_{i-1} with a rule from $\mathcal{R}$. For any fact base $\mathcal{F}$, we have that $\mathcal{F}, \mathcal{R} \models Q$ iff there is a rewriting Q_n of Q w.r.t. $\mathcal{R}$ such that $\mathcal{F} \models Q_n$ [König *et al.*, 2013].

Example 1 *Consider the rule $R = h(x) \rightarrow p(x,y)$ and CQ $Q = q(u) \wedge p(u,v) \wedge p(w,v) \wedge r(w)$. If $p(u,v)$ is unified with $p(x,y)$, then v is unified with the existential variable y, hence $p(w,v)$ has to be part of the unifier. The triple $\mu = (\{p(u,v), p(w,v)\}, \{p(x,y)\}, \{\{x,u,w\}\{v,y\}\}$ is a unifier. The direct rewriting of Q associated with the substitution $\sigma = \{x \mapsto u, w \mapsto u, y \mapsto v\}$ is $h(u) \wedge q(u) \wedge r(u)$.*

We now define some important kinds of rule sets (see e.g., [Mugnier, 2011] for an overview). A model M of a KB $\mathcal{K}$ is called *universal* if for any CQ Q, M is a model of Q iff $\mathcal{K} \models Q$. A rule set $\mathcal{R}$ is a *finite expansion set (fes)* if any KB $(\mathcal{F}, \mathcal{R})$ has a finite universal model. It is a *bounded-treewidth set (bts)* if any KB $(\mathcal{F}, \mathcal{R})$ has a (possibly infinite) universal model of bounded treewidth. It is a *finite unification set (fus)* if, for any CQ Q, there is a finite set S of rewritings of Q w.r.t. $\mathcal{R}$ such that for any fact base $\mathcal{F}$, we have $\mathcal{F}, \mathcal{R} \models Q$ iff there is $Q' \in S$ such that $\mathcal{F} \models Q'$.

A *Datalog* rule has no existential variables, hence Datalog rule sets are *fes*. Other kinds of *fes* rules are considered in the next section. A rule $B \rightarrow H$ is *guarded* if there is an atom in B that contains all the variables occurring in B. Guarded rules are *bts*. A *linear* rule has a body composed of a single atom and does not contain any constant. Linear rules are guarded, hence *bts*, moreover they are *fus*.

As a special case of Datalog rules, we have *transitivity rules*, of the form $p(x,y) \wedge p(y,z) \rightarrow p(x,z)$, which are not *fus*. A predicate is called *transitive* if it appears in a transitivity rule. If $\mathcal{C}$ is a class of rule sets, *$\mathcal{C}$+trans* denotes the class obtained by adding transitivity rules to rule sets from $\mathcal{C}$.

3 Combining *fes / fus* and Transitivity

A large hierarchy of *fes* classes is known (see e.g., [Cuenca Grau *et al.*, 2013] for an overview). Beside Datalog, the simplest classes are *weakly-acyclic (wa)* sets, which prevent cyclic propagation of existential variables along predicate positions, and *aGRD* (acyclic Graph of Rule Dependencies) sets, which prevent cyclic dependencies between rules. Datalog is generalized by *wa*, while *wa* and *aGRD* are incomparable. Some classes generalize *wa* by a finer analysis of variable propagation (up to *super-weakly acyclic (swa)* sets). Most other *fes* classes generalize both *wa* and *aGRD*.

We show that *aGRD+trans* is undecidable even for atomic CQs. Since *aGRD* is both *fes* and *fus*, this negative result also transfers to *fes+trans* and *fus+trans.*

Theorem 1 *Atomic CQ entailment over* aGRD+*trans KBs is undecidable, even with a single transitivity rule.*

Proof: The proof is by reduction from atomic CQ entailment with general existential rules (which is known to be undecidable). Let $\mathcal{R}$ be a set of rules. We first translate $\mathcal{R}$ into an *aGRD* set of rules $\mathcal{R}^a$. We consider the following new predicates: p (which will be the transitive predicate) and, for each rule $R_i \in \mathcal{R}$, predicates a_i and b_i. Each rule $R_i = B_i \rightarrow H_i$ is translated into the two following rules:

- $R_i^1 = B_i \rightarrow a_i(\vec{x}, z_1) \wedge p(z_1, z_2) \wedge p(z_2, z_3) \wedge b_i(z_3)$
- $R_i^2 = a_i(\vec{x}, z_1) \wedge p(z_1, z_2) \wedge b_i(z_2) \rightarrow H_i$

where z_1,z_2 and z_3 are existential variables and $\vec{x}$ are the variables in B_i.

Let $\mathcal{R}^a = \{R_i^1, R_i^2 \mid R_i \in \mathcal{R}\}$, and let $GRD(\mathcal{R}^a)$ be the graph of rule dependencies of $\mathcal{R}^a$, defined as follows: the nodes of $GRD(\mathcal{R}^a)$ are in bijection with $\mathcal{R}^a$, and there is an edge from a node R_1 to a node R_2 if the rule R_2 depends on the rule R_1, i.e., if there is a piece-unifier of the body of R_2 (seen as a CQ) with the head of R_1.

We check that for any $R_i \in \mathcal{R}$, R_i^1 has no outgoing edge and R_i^2 has no incoming edge (indeed the z_j are existential variables). Hence, in $GRD(\mathcal{R}^a)$ all (directed) paths are of length less or equal to one. It follows that $GRD(\mathcal{R}^a)$ has no cycle, i.e., $\mathcal{R}^a$ is aGRD.

Let R^t be the rule stating that p is transitive. Let $\mathcal{R}' = \mathcal{R}^a \cup \{R^t\}$. The idea is that R^t allows to "connect" rules in $\mathcal{R}^a$ that correspond to the same rule in $\mathcal{R}$. For any fact base $\mathcal{F}$ (on the original vocabulary), for any sequence of rule applications from $\mathcal{F}$ using rules in $\mathcal{R}$, one can build a sequence of rule applications from $\mathcal{F}$ using rules from $\mathcal{R}'$, and reciprocally, such that both sequences produce the same fact base (restricted to atoms on the original vocabulary). Hence, for any $\mathcal{F}$ and Q (on the original vocabulary), we have that $\mathcal{F}, \mathcal{R} \models Q$ iff $\mathcal{F}, \mathcal{R}' \models Q$. □

Corollary 1 *Atomic CQ entailment over* fus+*trans or* fes+*trans KBs is undecidable.*

Most known *fes* classes that do not generalize *aGRD* range between Datalog and *swa* (inclusive). It can be easily checked that any *swa* set of rules remains *swa* when transitivity rules are added (and this is actually true for all known classes between Datalog and *swa*).

Proposition 1 *The classes* swa *and* swa+*trans coincide. Hence,* swa+*trans is decidable.*

Proof: It suffices to note that the addition of transitivity rules does not create new edges in the 'SWA position graph' from [Cuenca Grau *et al.*, 2013]. □

It follows that the effect of transitivity on the currently known *fes* landscape is now quite clear, which is not the case for *fus* classes. In the following, we focus on a well-known *fus* class, namely *linear* rules. We show by means of a query rewriting procedure that query entailment over *linear+trans* KBs is decidable in the case of atomic CQs, as well as for general CQs if we place a minor safety condition on the rule set. Such an outcome was not obvious in the light of existing results. Indeed, atomic CQ entailment over *guarded+trans* rules was recently shown undecidable, even when restricted to rule sets that belong to the two-variable fragment, use only unary and binary predicates, and contain only two transitive predicates [Gottlob *et al.*, 2013]. Moreover, inclusion dependencies (a subclass of linear rules) and functional dependencies (a kind of rule known to destroy tree structures, as do transitivity rules) are known to be incompatible [Chandra and Vardi, 1985].

4 Linear Rules and Transitivity

To obtain finite representations of sets of rewritings involving transitive predicates, we define a framework based on the notion of *pattern*.

4.1 Framework

To each transitive predicate we assign a *pattern name*. Each pattern name has an associated *pattern definition* $P := a_1 | \ldots | a_k$, where each a_i is an atom that contains the special variables $\#1$ and $\#2$. A *pattern* is either a *standard pattern* $P[t_1, t_2]$ or a *repeatable pattern* $P^+[t_1, t_2]$, where P is a pattern name and t_1 and t_2 are terms. A *union of patterned conjunctive queries* (UPCQ) is a pair $(\mathbb{Q}, \mathbb{P})$, where $\mathbb{Q}$ is a disjunction of conjunctions of atoms and patterns, and $\mathbb{P}$ is a set of pattern definitions that gives a unique definition to each pattern name occurring in $\mathbb{Q}$. A *patterned conjunctive query* (PCQ) $\mathcal{Q}$ is a UPCQ without disjunction. For the sake of simplicity, we will often denote a (U)PCQ by its first component $\mathbb{Q}$, leaving the pattern definitions implicit.

An *instantiation* T of a UPCQ $(\mathbb{Q}, \mathbb{P})$ is a node-labelled tree that satisfies the following conditions:

- the root of T is labelled by $\mathcal{Q} \in \mathbb{Q}$;
- the children of the root are labelled by the patterns and atoms occurring in $\mathcal{Q}$;
- each node that is labelled by a repeatable pattern $P^+[t_1, t_2]$ may be expanded into $k \geq 1$ children labelled respectively by $P[t_1, x_1]$, $P[x_1, x_2]$, …, $P[x_{k-1}, t_2]$, where the x_i are fresh variables;
- each node labelled by a standard pattern $P[t_1, t_2]$ may be expanded into a single child whose label is obtained from an atom a in the pattern definition of P in $\mathbb{P}$ by substituting $\#1$ (resp. $\#2$) by t_1 (resp. t_2), and freshly renaming the other variables.

For brevity, we will often refer to nodes in an instantiation using their labels.

The *instance* associated with an instantiation is the PCQ obtained by taking the conjunction of the labels of its leaves. An instance of a UPCQ is an instance associated with one of its instantiations. An instance is called *full* if it does not contain any pattern, and we denote by $full(\mathbb{Q}, \mathbb{P})$ the set of full instances of $(\mathbb{Q}, \mathbb{P})$.

Example 2 *Let* $(\mathcal{Q}, \mathbb{P})$ *be a PCQ, where* $\mathcal{Q} = P_1^+[a, z] \wedge P_2^+[z, b] \wedge s_1(a, b)$ *and* $\mathbb{P}$ *contains the pattern definitions:*

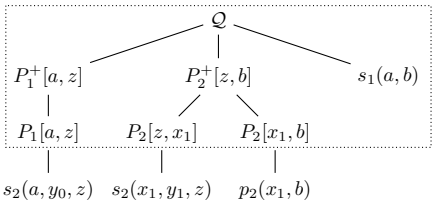

Figure 1: Instantiations of a PCQ

$P_1 := p_1(\#1,\#2)|s_2(\#1,y,\#2)$ *and* $P_2 := p_2(\#1,\#2)|$ $s_2(\#2,y,\#1)$.

Two instantiations of $\mathcal{Q}$ are displayed in Figure 1. The smaller instantiation (within the dotted lines) gives rise to the (non-full) instance $Q_1 = P_1[a,z] \wedge P_2[z,x_1] \wedge P_2[x_1,b] \wedge s_1(a,b)$. By expanding the three nodes labelled by patterns according to the definitions in $\mathbb{P}$, we may obtain the larger instantation (occupying the entire figure), whose associated instance $Q_2 = s_2(a,y_0,z) \wedge s_2(x_1,y_1,z) \wedge p_2(x_1,b) \wedge s_1(a,b)$ is a full instance for $(\mathcal{Q},\mathbb{P})$.

A UPCQ $(\mathbb{Q},\mathbb{P})$ can be translated into a set of Datalog rules $\Pi_{\mathbb{P}}$ and a UCQ $Q_{\mathbb{Q}}$ as follows. For each definition $P := a_1(\vec{t_1})|\ldots|a_k(\vec{t_k})$ in $\mathbb{P}$, we create the transitivity rule $p^+(x,y) \wedge p^+(y,z) \rightarrow p^+(x,z)$ and the rules $a_i(\vec{t_i}) \rightarrow p^+(\#1,\#2)$ $(1 \leq i \leq k)$. The UCQ $Q_{\mathbb{Q}}$ is obtained from $\mathbb{Q}$ by replacing each repeatable pattern $P^+[t_1,t_2]$ by the atom $p^+(t_1,t_2)$. Observe that $\Pi_{\mathbb{P}}$ is non-recursive except for the transitivity rules. The next proposition states that $(\Pi_{\mathbb{P}}, Q_{\mathbb{Q}})$ can be seen as a finite representation of the set of full instances of $(\mathbb{Q},\mathbb{P})$.

Proposition 2 *Let $\mathcal{F}$ be a fact base and $(\mathbb{Q},\mathbb{P})$ be a UPCQ. Then $\mathcal{F}, \Pi_{\mathbb{P}} \models Q_{\mathbb{Q}}$ iff $\mathcal{F} \models Q$ for some $Q \in full(\mathbb{Q},\mathbb{P})$.*

A *unifier* $\mu = (Q', H, P_u)$ of a *PCQ* is a unifier of one of its (possibly non-full) instances such that Q' is a set of (usual) atoms. We distinguish two types of unifiers (internal and external), defined next.

Let T be an instantiation, Q be its associated instance, and $\mu = (Q', H, P_u)$ be a unifier of Q. Assume T contains a repeatable pattern $P^+[t_1,t_2]$ that is expanded into $P[u_0,u_1]$, $\ldots, P[u_k,u_{k+1}]$, where $u_0 = t_1$ and $u_{k+1} = t_2$. We call $P[u_i,u_{i+1}]$ *relevant for* μ if it is expanded into an atom from Q'. Because we consider only single-piece unifiers (cf. Sec. 2), it follows that if such relevant patterns exist, they form a sequence $P[u_i,u_{i+1}]$,$P[u_{i+1},u_{i+2}]$,$\ldots, P[u_{j-1},u_j]$. Terms u_i and u_j are called *external* to $P^+[t_1,t_2]$ w.r.t. μ; the other terms occurring in the sequence are called *internal*. The unifier μ is said to be *internal* if all atoms from Q' are expanded from a single repeatable pattern, and no external terms are unified together or with an existential variable; otherwise μ is called *external*.

Example 3 *Consider Q_2 from Example 2 and the rules $R_1 = s_1(x',y') \rightarrow p_2(x',y')$ and $R_2 = s_1(x',y') \rightarrow s_2(x',y',z')$. The unifier of Q_2 with R_1 that unifies $p_2(x_1,b)$ with $p_2(x',y')$ is internal. The unifier of Q_2 with R_2 that unifies $\{s_2(a,y_0,z), s_2(x_1,y_1,z)\}$ with $s_2(x',y',z')$ is external because it involves two repeatable patterns.*

4.2 Overview of the Algorithm

Our query rewriting algorithm takes as input a CQ Q and a set of rules $\mathcal{R} = \mathcal{R}_L \cup \mathcal{R}_T$, with $\mathcal{R}_L$ a set of linear rules and $\mathcal{R}_T$ a set of transitivity rules, and produces a finite set of Datalog rules and a (possibly infinite) set of CQs. The main steps of the algorithm are outlined below.

Step 1 For each predicate p that appears in $\mathcal{R}_T$, create a pattern definition $P := p(\#1,\#2)$, where P is a fresh pattern name. Call the resulting set of definitions $\mathbb{P}_0$.

Step 2 Let $\mathcal{R}_L^+$ be the result of considering all of the rule bodies in $\mathcal{R}_L$ and replacing every body atom $p(t_1,t_2)$ such that p is a transitive predicate by the repeatable pattern $P^+[t_1,t_2]$.

Step 3 (Internal rewriting) Initialize $\mathbb{P}$ to $\mathbb{P}_0$ and repeat the following operation until fixpoint: select a pattern definition $P \in \mathbb{P}$ and a rule $R \in \mathcal{R}_L^+$ and compute the direct rewriting of $\mathbb{P}$ w.r.t. P and R.

Step 4 Replace in Q all atoms $p(t_1,t_2)$ such that p is a transitive predicate by the repeatable pattern $P^+[t_1,t_2]$, and denote the result by $\mathcal{Q}^+$.

Step 5 (External rewriting) Initialize $\mathbb{Q}$ to $\{\mathcal{Q}^+\}$ and repeat the following operation until fixpoint: choose $\mathcal{Q}_i \in \mathbb{Q}$, compute a direct rewriting of $\mathcal{Q}_i$ w.r.t. $\mathbb{P}$ and a rule from $\mathcal{R}_L^+$, and add the result to $\mathbb{Q}$ (except if it is isomorphic to some $\mathcal{Q}_j \in \mathbb{Q}$).

Step 6 Let $\Pi_{\mathbb{P}}$ be the Datalog translation of $\mathbb{P}$, and let $Q_{\mathbb{Q}}$ be the (possibly infinite) set of CQs obtained by replacing each repeatable pattern $P^+[t_1,t_2]$ in $\mathbb{Q}$ by $p^+(t_1,t_2)$.

The rewriting process in Step 3 is always guaranteed to terminate, and in Section 6, we propose a modification to Step 5 that ensures termination and formulate sufficient conditions that preserve completeness. When $\mathcal{Q}_{\mathbb{Q}}$ is finite (i.e., it is a UCQ), it can be evaluated over the fact base saturated by $\Pi_{\mathbb{P}}$, or alternatively, translated into a set of Datalog rules, which can be combined with $\Pi_{\mathbb{P}}$ and passed to a Datalog engine for evaluation. Observe that the construction of $\Pi_{\mathbb{P}}$ is query-independent and can be executed as a preprocessing step.

5 Rewriting Steps in Detail

A PCQ that contains a repeatable pattern has an infinite number of instances. Instead of considering all instances of a PCQ, we consider a finite set of 'instances of interest' for a given rule. Such instances will be used for both the internal and external rewriting steps.

Instances of interest Consider a PCQ $(\mathcal{Q},\mathbb{P})$ and a rule $R \in \mathcal{R}_L^+$ with head predicate p. The *instantiations of interest* of $(\mathcal{Q},\mathbb{P})$ w.r.t. R are constructed as follows. For each repeatable pattern $P_i^+[t_1,t_2]$ in $\mathcal{Q}$, let $a_1^i,\ldots,a_{n_i}^i$ be the atoms in the definition of P_i with predicate p. If $n_i > 0$, then expand $P_i^+[t_1,t_2]$ into k standard patterns, where $0 < k \leq min(arity(p), n_i) + 2$, and expand each of these standard patterns in turn into some a_ℓ^i. An *instance of interest* is the instance associated with an instantiation of interest.

Example 4 *Reconsider $\mathcal{Q}$, Q_2 and R_2 from Examples 2 and 3. Q_2 is not an instance of interest of $\mathcal{Q}$ w.r.t. R_2 since $P_2[x_1,b]$ is expanded into $p(\#1,\#2)$ whereas the head predicate of R_2 is s_2. If we expand $P_2[x_1,b]$ with $s_2(\#2,y,\#1)$*

instead, we obtain the instance of interest $Q_3 = s_2(a, y_0, z) \wedge s_2(x_1, y_1, z) \wedge s_2(b, y_2, x_1) \wedge s_1(a, b)$.

We next show that the set of unifiers computed on the instances of interest of a PCQ 'captures' the set of unifiers computed on all of its instances.

Proposition 3 *Let* $(\mathcal{Q}, \mathbb{P})$ *be a PCQ and* $R \in \mathcal{R}_L^+$. *For every instance* Q *of* $(\mathcal{Q}, \mathbb{P})$ *and unifier* μ *of* Q *with* R, *there exist an instance of interest* Q' *of* $(\mathcal{Q}, \mathbb{P})$ *w.r.t.* R *and a unifier* μ' *of* Q' *with* R *such that* μ' *is more general* [1] *than* μ.

5.1 Internal Rewriting

Rewriting w.r.t. internal unifiers is performed 'inside' a repeatable pattern, independently of the other patterns and atoms in the query. We will therefore handle this kind of rewriting in a query-independent manner by updating the pattern definitions.

To find all internal unifiers between instances under a repeatable pattern $P^+[t_1, t_2]$ and a rule head $H = p(\ldots)$, one may think that it is sufficient to consider each atom a_i in P's definition and check if there is an internal unifier of a_i with H. Indeed, this suffices when predicates are binary: in an internal unifier, t_1 and t_2 are unified with distinct variables, which cannot be existential; thus, the terms in H are frontier variables, and a piece must consist of a single atom. If the arity of p is greater than 2, the other variables can be existential, so it may be possible to unify a path of atoms from P's definition, but not a single such atom (see next example).

Example 5 *Let* $R = s(x, y) \rightarrow r(z_1, x, z_2, y)$ *and* $P := r(\#2, \#1, x_0, x_1) \mid\ r(\#1, x_2, \#2, x_3) \mid\ r(x_4, x_5, \#1, \#2)$. *There is no internal unifier of an atom in* P*'s definition with* $H = r(z_1, x, z_2, y)$. *However, if we expand* $P^+[t_1, t_2]$ *into a path* $P[t_1, y_0]P[y_0, y_1]P[y_1, t_2]$, *then expand the* i*th pattern of this path into the* i*th atom in* P*'s definition, the resulting instance can be unified with* H *by an internal unifier (with the partition* $\{\{z_1, y_0, x_4\}, \{x, t_1, x_2, x_5\}, \{z_2, x_0, y_1\}, \{y, x_1, x_3, t_2\}\}$*).*

Fortunately, we can bound the length of paths to be considered using both the arity of p and the number of atoms with predicate p in P's definition, allowing us to use instances of interest introduced earlier.

A *direct rewriting* $\mathbb{P}'$ of a set of pattern definitions $\mathbb{P}$ w.r.t. a pattern name P and a rule $R = B \rightarrow H \in \mathcal{R}_L^+$ is the set of pattern definitions obtained from $\mathbb{P}$ by updating P's definition as follows. We consider the PCQ $(\mathcal{Q} = P^+[x, y], \mathbb{P})$. We select an instance of interest Q of $\mathcal{Q}$ w.r.t. R, an internal unifier μ of Q with H, and a substitution σ associated with μ that preserves the external terms. Let B' be obtained from $\sigma(B)$ by substituting the first (resp. second) external term by $\#1$ (resp. $\#2$). If B' is an atom, we add it to P's definition. Otherwise, B' is a repeatable pattern of the form $S^+[\#1, \#2]$ or $S^+[\#2, \#1]$. Let f be a bijection on $\{\#1, \#2\}$: if B' is of the form $S^+[\#1, \#2]$, f is the identity, otherwise f permutes $\#1$ and $\#2$. For all s_i in the definition of S, we add $f(s_i)$ to P's definition.

Note that the addition of an atom to a pattern definition is up to *isomorphism* (with $\#1$ and $\#2$ treated as distinguished variables, i.e., $\#1$ and $\#2$ are mapped to themselves).

Example 6 *Reconsider* R, μ, *and the definition of* P *from Example 5. Performing a direct rewriting w.r.t.* P *using* R *and* μ *results in adding the atom* $s(\#1, \#2)$ *to* P*'s definition.*

Proposition 4 *Let* $(\mathcal{Q}, \mathbb{P})$ *be a PCQ where* $P^+[t_1, t_2]$ *occurs and* $R \in \mathcal{R}_L^+$. *For any instance* Q *of* $(\mathcal{Q}, \mathbb{P})$, *any classical direct rewriting* $\mathcal{Q}'$ *of* Q *with* R *w.r.t. to a unifier internal to* $P^+[t_1, t_2]$, *and any* $Q' \in full(\mathcal{Q}', \mathbb{P})$, *there exists a direct rewriting* $\mathbb{P}'$ *of* $\mathbb{P}$ *w.r.t.* P *and* R *such that* $(\mathcal{Q}, \mathbb{P}')$ *has a full instance that is isomorphic to* Q'.

5.2 External Rewriting

Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ, $R \in \mathcal{R}_L^+$, T be an instantiation of interest of $(\mathcal{Q}, \mathbb{P})$ w.r.t. R, Q be the instance associated with T, and $\mu = (Q', H, P)$ be an external unifier of Q with R. From this, several direct rewritings of $\mathcal{Q}$ w.r.t. $\mathbb{P}$ and R can be built. First, we mark all leaves in T that either have the root as parent or are labelled by an atom in Q', and we restrict T to branches leading to a marked leaf. Then, we consider each instantiation T_i that can be obtained from $\mathcal{Q}$ as follows. Replace each repeatable pattern $P^+[t_1, t_2]$ that has $k > 0$ children in T by one of the following:

(i) $P^+[t_1, x_1] \wedge X[x_1, x_2] \wedge P^+[x_2, t_2]$,

(ii) $P^+[t_1, x_1] \wedge X[x_1, t_2]$,

(iii) $X[t_1, x_2] \wedge P^+[x_2, t_2]$,

(iv) $X[t_1, t_2]$,

where $X[v_1, v_2]$ is a sequence $P[v_1, y_1], P[y_1, y_2], \ldots, P[y_{k-1}, v_2]$. Let $\mathcal{Q}_i$ be the instance associated with T_i.

If $P[x', y']$ in T has child $a(\vec{t})$, expand in T_i the corresponding $P[x, y]$ into $a(\rho(\vec{t}))$ where $\rho = \{x' \mapsto x, y' \mapsto y\}$. If $\mu' = (\rho(Q'), H, \rho(P))$ is still a unifier of $\mathcal{Q}_i$ with H, we say that $\mathcal{Q}_i$ is a *minimally-unifiable instance* of $\mathcal{Q}$ w.r.t. μ. In this case, $\mathcal{Q}'_i = \mu'(Q_i) \setminus \mu'(H) \cup \mu'(B)$ is a *direct rewriting* of $\mathcal{Q}$ w.r.t. $\mathbb{P}$ and R.

Example 7 *Reconsider* Q_3 *and* R_2, *and let* $\mu = (\{s_2(a, y_0, z),\ s_2(x_1, y_1, z)\},\ H_2,\ \{\{a, x_1, x'\},\ \{y_0, y_1, y'\},\ \{z, z'\}\})$. *First, we consider the instantiation that generated* Q_3, *and we remove the node labelled* $P_2[x_1, b]$ *and its child* $s_2(b, y_2, x_1)$, *since the latter atom is not involved in* μ. *Next will replace the repeatable pattern* $P_1^+[a, z]$ *(resp.* $P_2^+[z, b]$*) using one of the four cases detailed above, and we check whether* μ' *(obtained from* μ*) is still a unifier. We obtain in this manner the following minimally-unifiable instances:* $\mathcal{Q}_1 = P_1^+[a, x_2] \wedge\ s_2(x_2, y_0, z) \wedge\ s_2(x_1, y_1, z) \wedge\ P_2^+[x_1, b] \wedge\ s_1(a, b)$, *and* $\mathcal{Q}_2 = s_2(a, y_0, z)\ \wedge\ s_2(x_1, y_1, z)\ \wedge\ P_2^+[x_1, b]\ \wedge\ s_1(a, b)$. *Finally, we rewrite* $\mathcal{Q}_1$ *and* $\mathcal{Q}_2$ *into:* $\mathcal{Q}'_1 = P_1^+[a, x'] \wedge s_1(x', y') \wedge P_2^+[x', b] \wedge s_1(a, b)$ *and* $\mathcal{Q}'_2 = s_1(a, y') \wedge P_2^+[a, b] \wedge s_1(a, b)$.

[1] Consider unifiers $\mu = (Q, H, P_\mu)$ and $\mu' = (Q', H, P_{\mu'})$, and let σ (resp. σ') be a substitution associated with P_μ (resp. $P_{\mu'}$). We say that μ' is *more general* than μ if there is a substitution h from $\sigma'(Q')$ to $\sigma(Q)$ such that $h(\sigma'(Q')) \subseteq \sigma(Q)$ (i.e., h is a homomorphism from $\sigma'(Q')$ to $\sigma(Q)$), and for all terms x and y in $Q' \cup H$, if $\sigma'(x) = \sigma'(y)$ then $\sigma(h(x)) = \sigma(h(y))$.

Proposition 5 *Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ and $R \in \mathcal{R}_L^+$. For every $Q \in full(\mathcal{Q}, \mathbb{P})$ and every classical direct rewriting Q' of Q with R w.r.t. an external unifier, there is a direct rewriting $\mathcal{Q}'$ of $\mathcal{Q}$ w.r.t. $\mathbb{P}$ and R that has an instance isomorphic to Q'.*

6 Termination and Correctness

To establish the correctness of the query rewriting algorithm, we utilize Propositions 2, 4 and 5.

Theorem 2 *Let Q be a CQ, $(\mathcal{F}, \mathcal{R})$ be a linear+trans KB, and ($\Pi_\mathbb{P}$,$Q_\mathbb{Q}$) be the (possibly infinite) output of the algorithm. Then: $\mathcal{F}, \mathcal{R} \models Q$ iff $\mathcal{F}, \Pi_\mathbb{P} \models Q'$ for some $Q' \in Q_\mathbb{Q}$.*

Regarding termination, we observe that Step 3 (internal rewriting) must halt since every direct rewriting step adds a new atom (using a predicate from $\mathcal{R}_L^+$) to a pattern definition, and there are finitely many such atoms, up to isomorphism.

By contrast, Step 5 (external rewriting) need not halt, as the rewritings may grow unboundedly in size. Thus, to ensure termination, we will modify Step 5 to exclude direct rewritings that increase rewriting size. Specifically, we identify the following 'problematic' minimally-unifiable instances:

- Q' is composed of atoms expanded from a single pattern $P^+[t_1, t_2]$, $\mu'(t_1) = \mu'(t_2)$, and $P^+[t_1, t_2]$ is replaced as in case (i), (ii) or (iii).
- Q' is obtained from the expansion of repeatable patterns, a term t of Q is unified with an existential variable of the head of the rule, t appears only in repeatable patterns of form $P_i^+[t_i, t]$ (resp. $P_i^+[t, t_i]$), and all these repeatable patterns are rewritten as in case *(ii)* $P_i^+[t_i, t_i'] \wedge X[t_i', t]$ (resp. as in case *(iii)* $X[t, t_i'] \wedge P_i^+[t_i', t_i]$).

We will call a direct rewriting *excluded* if it is based on such a minimally-unifiable instance; otherwise, it is *non-excluded*.

Example 8 *The rewriting $\mathcal{Q}_1'$ from Example 7 is excluded because it is obtained from the minimally-unifiable instance Q_1 in which the repeatable patterns $P_1^+[a, z]$ is expanded as in case (ii) and $P_2^+[z, b]$ as in case (iii), and z is unified with the existential variable z'.*

Proposition 6 *Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ and $R \in \mathcal{R}_L^+$. If $\mathcal{Q}'$ is a non-excluded direct rewriting of $\mathcal{Q}$ with R, then $|\mathcal{Q}'| \leq |\mathcal{Q}|$.*

Let us consider the 'modified query rewriting algorithm' that is obtained by only performing non-excluded direct rewritings in Step 5. This modification ensures termination but may comprise completeness. However, we can show that the modified algorithm is complete in the following key cases: when the CQ is atomic, when there is no specialization of a transitive predicate, or when all predicates have arity at most two. By further analyzing the latter case, we can formulate a safety condition, defined next, that guarantees completeness for a much wider class of rule sets.

Safe rule sets We begin by defining a specialization relationship between predicates. A predicate q is a *direct specialization* of a binary predicate p on positions $\{\vec{i}, \vec{j}\}$ ($\vec{i} \neq \emptyset, \vec{j} \neq \emptyset$) if there is a rule of the form $q(\vec{u}) \rightarrow p(x, y)$ such that $\vec{i}$ (resp. $\vec{j}$) contains those positions of $\vec{u}$ that contain the term x (resp. y). It is a *specialization* of p on positions $\{\vec{i}, \vec{j}\}$ if (a) it is a direct specialization of p on positions $\{\vec{i}, \vec{j}\}$, or (b) there is a rule of the form $q(\vec{u}) \rightarrow r(\vec{v})$ such that $r(\vec{v})$ is a specialization of p on positions $\{\vec{k}, \vec{l}\}$ and the terms occurring in positions $\{\vec{k}, \vec{l}\}$ of $\vec{v}$ occur in positions $\{\vec{i}, \vec{j}\}$ of $\vec{u}$ with $\vec{i} \neq \emptyset$ and $\vec{j} \neq \emptyset$. We say that q is a *pseudo-transitive predicate* if it is a specialization of at least one transitive predicate.

We call a linear+trans rule set *safe* if it satisfies the following *safety condition*: for every pseudo-transitive predicate q, there exists a pair of positions $\{i, j\}$ with $i \neq j$ such that for all transitive predicates p of which q is a specialization on positions $\{\vec{i}, \vec{j}\}$, either $i \in \vec{i}$ and $j \in \vec{j}$, or $i \in \vec{j}$ and $j \in \vec{i}$.

Note that if we consider binary predicates, the safety condition is always fulfilled. Then, specializations correspond exactly to the subroles considered in DLs.

Example 9 *Let $R_1 = s_1(x, x, y) \rightarrow p_1(x, y)$, $R_2 = s_2(x, y, z) \rightarrow p_2(x, y)$, $R_3 = s_1(x, y, z) \rightarrow s_2(z, x, y)$, and p_1 and p_2 be two transitive predicates.*

The following specializations have to be considered: s_1 is a direct specialization of p_1 on positions $\{\{1, 2\}, \{3\}\}$, s_2 is a direct specialization of p_2 on positions $\{\{1\}, \{2\}\}$, s_1 is a specialization of p_2 on positions $\{\{3\}, \{1\}\}$. We then have two pseudo-transitive predicates: s_1 and s_2. By choosing the pair $\{1, 3\}$ for s_1 and $\{1, 2\}$ for s_2, we observe that $\{R_1, R_2, R_3\}$ satisfies the safety condition.

If we replace R_3 by $R_4 = s_1(x, y, z) \rightarrow s_2(x, y, z)$, s_1 is a specialization of p_2 on positions $\{\{1\}, \{2\}\}$, and $\{R_1, R_2, R_4\}$ is not safe.

Theorem 3 *The modified query rewriting algorithm halts. Moreover, Theorem 2 (soundness and completeness) holds for the modified algorithm if either the input CQ is atomic, or the input rule set is safe.*

7 Complexity

A careful analysis of our query rewriting algorithm allows us to pinpoint the worst-case complexity of atomic CQ entailment over linear+trans KBs, and general CQ entailment over safe linear+trans KBs. As usual, we consider two complexity measures: *combined complexity* (measured in terms of the size of the whole input), and *data complexity* (measured in terms of the size of the fact base). The latter is often considered more relevant since the fact base is typically significantly larger than the rest of the input.

With regards to data complexity, we show completeness for NL (non-deterministic logarithmic space), which is the same complexity as in the presence of transitivity rules alone.

Theorem 4 *Both (i) atomic CQ entailment over linear+trans KBs and (ii) CQ entailment over safe linear+trans KBs are* NL*-complete in data complexity.*

Regarding combined complexity, we show that the addition of transitivity rules does not increase the complexity of query entailment for the two considered cases.

Theorem 5 *Both (i) atomic CQ entailment over linear+trans KBs and (ii) CQ entailment over safe linear+trans KBs are* PSPACE*-complete in combined complexity.*

8 Conclusion

In this paper, we made some steps towards a better understanding of the interaction between transitivity and decidable classes of existential rules. We obtained an undecidability result for aGRD+*trans*, hence for *fes*+*trans* and *fus*+*trans*. More positively, we established decidability (with the lowest possible complexity) of atomic CQ entailment over linear+*trans* KBs and general CQ entailment for safe linear+*trans* rule sets. The safety condition was introduced to ensure termination of the rewriting mechanism when predicates of arity more than two are considered (rule sets which use only unary and binary predicates are trivially safe). We believe the condition can be removed with a much more involved termination proof.

In future work, we plan to explore the effect of transitivity on decidable rule classes that are incomparable with linear rules, namely frontier-one rules, a *bts* class that has close connections to Horn DLs, and two *fus* classes: domain-restricted and sticky rule sets [Baget *et al.*, 2011a; Calì *et al.*, 2010].

Acknowledgements

This work was supported by ANR project PAGODA (contract ANR 12 JS02 007 01).

References

[Baader, 2003] F. Baader. Terminological cycles in a description logic with existential restrictions. In *Proc. of IJCAI*, pages 325–330, 2003.

[Baget *et al.*, 2009] J.-F. Baget, M. Leclère, M.-L. Mugnier, and E. Salvat. Extending decidable cases for rules with existential variables. In *Proc. of IJCAI*, pages 677–682, 2009.

[Baget *et al.*, 2011a] J.-F. Baget, M. Leclère, M.-L. Mugnier, and E. Salvat. On rules with existential variables: Walking the decidability line. *Art. Intell. (AIJ)*, 175(9-10):1620–1654, 2011.

[Baget *et al.*, 2011b] J.-F. Baget, M.-L. Mugnier, S. Rudolph, and M. Thomazo. Walking the complexity lines for generalized guarded existential rules. In *Proc. of IJCAI*, pages 712–717, 2011.

[Beeri and Vardi, 1981] C. Beeri and M. Y. Vardi. The implication problem for data dependencies. In *Proc. of ICALP*, volume 115 of *LNCS*, pages 73–85. Springer, 1981.

[Calì *et al.*, 2008] A. Calì, G. Gottlob, and M. Kifer. Taming the infinite chase: Query answering under expressive relational constraints. In *Proc. of KR*, pages 70–80, 2008.

[Calì *et al.*, 2009] A. Calì, G. Gottlob, and T. Lukasiewicz. A general Datalog-based framework for tractable query answering over ontologies. In *Proc. of PODS*, pages 77–86, 2009.

[Calì *et al.*, 2010] A. Calì, G. Gottlob, and A. Pieris. Query answering under non-guarded rules in Datalog+/-. In *Proc. of RR*, pages 1–17, 2010.

[Calvanese *et al.*, 2007] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *J. Autom. Reasoning (JAR)*, 39(3):385–429, 2007.

[Chandra and Vardi, 1985] A. K. Chandra and M. Y. Vardi. The implication problem for functional and inclusion dependencies is undecidable. *SIAM J. Comput.*, 14(3):671–677, 1985.

[Cuenca Grau *et al.*, 2013] B. Cuenca Grau, I. Horrocks, M. Krötzsch, C. Kupke, D. Magka, B. Motik, and Z. Wang. Acyclicity notions for existential rules and their application to query answering in ontologies. *J. Art. Intell. Res. (JAIR)*, 47:741–808, 2013.

[Eiter *et al.*, 2009] T. Eiter, C. Lutz, M. Ortiz, and M. Simkus. Query answering in description logics with transitive roles. In *Proc. of IJCAI*, pages 759–764, 2009.

[Eiter *et al.*, 2012] T. Eiter, M. Ortiz, M. Simkus, T.-K. Tran, and G. Xiao. Query rewriting for Horn-$\mathcal{SHIQ}$ plus rules. In *Proc. of AAAI*, 2012.

[Glimm *et al.*, 2008] B. Glimm, C. Lutz, I. Horrocks, and U. Sattler. Conjunctive query answering for the description logic $\mathcal{SHIQ}$. *J. Artif. Intell. Res. (JAIR)*, 31:157–204, 2008.

[Gottlob *et al.*, 2013] G. Gottlob, A. Pieris, and L. Tendera. Querying the guarded fragment with transitivity. In *Proc. of ICALP*, volume 7966 of *LNCS*, pages 287–298. Springer, 2013.

[Horrocks and Sattler, 1999] I. Horrocks and U. Sattler. A description logic with transitive and inverse roles and role hierarchies. *J. Log. Comput.*, 9(3):385–410, 1999.

[König *et al.*, 2013] M. König, M. Leclère, M.-L. Mugnier, and M. Thomazo. On the exploration of the query rewriting space with existential rules. In *Proc. of RR*, pages 123–137, 2013.

[Krötzsch and Rudolph, 2011] M. Krötzsch and S. Rudolph. Extending decidable existential rules by joining acyclicity and guardedness. In *Proc. of IJCAI*, pages 963–968, 2011.

[Krötzsch *et al.*, 2007] M. Krötzsch, S. Rudolph, and P. Hitzler. Complexity boundaries for Horn description logics. In *Proc. of AAAI*, pages 452–457. AAAI Press, 2007.

[Lutz *et al.*, 2009] C. Lutz, D. Toman, and F. Wolter. Conjunctive query answering in the description logic $\mathcal{EL}$ using a relational database system. In *Proc. of IJCAI*, pages 2070–2075, 2009.

[Mugnier, 2011] M.-L. Mugnier. Ontological query answering with existential rules. In *Proc. of RR*, pages 2–23, 2011.

[Sattler, 2000] U. Sattler. Description logics for the representation of aggregated objects. In *Proc. of ECAI*, pages 239–243, 2000.

Appendix

Notations In the following we use several notations that were not included in the main paper for space restrictions. Let Q_1 and Q_2 be two CQs, if there is a homomorphism from Q_1 to Q_2, we say that Q_1 is *more general* than Q_2, and we note $Q_1 \geq Q_2$. This definition is naturally extended to (U)PCQs: let $(\mathcal{Q}_1, \mathbb{P}_1)$ and $(\mathcal{Q}_2, \mathbb{P}_2)$ be two (U)PCQs, if for any full instance Q_2 of $(\mathcal{Q}_2, \mathbb{P}_2)$ there is a full instance Q_1 of $(\mathcal{Q}_1, \mathbb{P}_1)$ such that $Q_1 \geq Q_2$, we say that $(\mathcal{Q}_1, \mathbb{P}_1)$ is *more general* than $(\mathcal{Q}_2, \mathbb{P}_2)$ and we note $(\mathcal{Q}_1, \mathbb{P}_1) \geq (\mathcal{Q}_2, \mathbb{P}_2)$.

We recall from the body of the paper that a unifier $\mu' = (Q', H, P_{\mu'})$ is more general than $\mu = (Q, H, P_\mu)$ if there is a substitution h from $\sigma'(Q')$ to $\sigma(Q)$ such that $h(\sigma'(Q')) \subseteq \sigma(Q)$ (i.e., h is a homomorphism from $\sigma'(Q')$ to $\sigma(Q)$), and for all terms x and y in $Q' \cup H$, if $\sigma_{\mu'}(x) = \sigma_{\mu'}(y)$ then $\sigma_\mu(h(x)) = \sigma_\mu(h(y))$, where σ_μ and $\sigma_{\mu'}$ are substitutions associated respectively with P_μ and $P_{\mu'}$. In what follows, we will write $\mu' \geq \mu$ to indicate that μ' is more general than μ.

Proposition 2 *Let $\mathcal{F}$ be a fact base and $(\mathbb{Q}, \mathbb{P})$ be a UPCQ. Then $\mathcal{F}, \Pi_\mathbb{P} \models Q_\mathbb{Q}$ iff $\mathcal{F} \models Q$ for some $Q \in full(\mathbb{Q}, \mathbb{P})$.*

Proof: We successively prove the two directions.

($\Rightarrow$) Let T be an instantiation of $(\mathbb{Q}, \mathbb{P})$ such that there exists a homomorphism π from its associated full instance to the fact base $\mathcal{F}$. Let us consider a node of T labelled by a standard pattern atom $P[t_1, t_2]$. The label $r(\rho(\vec{t}))$ of its child was obtained by choosing an atom $r(\vec{t})$ in the pattern definition of P. Thus our Datalog program $\Pi_\mathbb{P}$ contains the rule $r(\vec{t}) \rightarrow p^+(\#1, \#2)$ where $\#1, \#2 \in \vec{t}$. Applying the rule according to the homomorphism $\pi \circ \rho$, we can add the atom $p^+(t_1, t_2)$ to $\mathcal{F}$. Let us repeat this procedure for every node of T labelled by a standard pattern atom. Consider next a repeatable pattern atom labelled $P^+[t, t']$ whose children are respectively labelled $P[t = t_1, t_2], P[t_2, t_3], \ldots, P[t_{k-1}, t_k = t']$. According to the rule applications already described, $\mathcal{F}$ now contains the atoms $p^+(t = t_1, t_2), p^+(t_2, t_3), \ldots, p^+(t_{k-1}, t_k = t')$. Then, by successive applications of the rule in $\Pi_\mathbb{P}$ expressing the transitivity of p^+, we finally add to $\mathcal{F}$ the atom $p^+(t, t')$. Repeat this procedure for every node of T labelled by a repeatable pattern atom. Now the root of T is labelled by some $\mathcal{Q} \in \mathbb{Q}$. The UCQ $Q_\mathbb{Q}$ contains a CQ Q that was obtained from $\mathcal{Q}$ by replacing each repeatable pattern $P^+[t_1, t_2]$ by $p^+(t_1, t_2)$. Observe that the restriction of π to the terms of Q is a homomorphism from Q to the fact based obtained from $\mathcal{F}$ by previous rule applications.

($\Leftarrow$) Conversely, let us consider a fact base $\mathcal{F}'$ obtained by saturating the initial fact base $\mathcal{F}$ with the rules of $\Pi_\mathbb{P}$, and a homomorphism π from some CQ Q' in the UCQ $Q_\mathbb{Q}$ to $\mathcal{F}'$. Let us now build an instantiation T whose full instance can be mapped to $\mathcal{F}$ thanks to a homomorphism π'. The root node of T is labelled by the UPCQ $\mathcal{Q}$ in $\mathbb{Q}$ from which Q' was obtained. Its children are labelled by the atoms and pattern atoms of $\mathcal{Q}$. Now we define the restriction of π' to the terms of $\mathcal{Q}$ as π. Let us now consider a child of the root labelled by a repeatable pattern atom $P^+[t, t']$. It follows that $p^+(\pi(t), \pi(t'))$ is an atom of $\mathcal{F}'$. Since this atom is not in the initial fact base, it means that it has been obtained by a (possibly empty) sequence of applications of the rule expressing the transitivity of p^+ on a p^+-path $\pi(t) = t_1, \ldots, t_k = \pi(t')$ such that no atom $p^+(t_i, t_{i+1})$ in $\mathcal{F}'$ has been obtained by a transitivity rule. Then the node labelled $P^+[t, t']$ has $k + 1$ children respectively labelled $P[t = x_1, x_2], \ldots, P[x_{k-1}, x_k = t']$. For the fresh variables $x_2, \ldots, x_{k-1}$, we define $\pi'(x_i) = t_i$. Repeat this procedure for every repeatable pattern atom in T. Let us next consider a node of T labelled by a standard pattern atom $P[x, x']$. Since that node was obtained in the previous phase, we know that the atom $p^+(\pi'(x), \pi'(x'))$ is in $\mathcal{F}'$, and that it was not obtained from the application of a transitivity rule. Thus, the Datalog rule used to produce that atom is necessarily a rule obtained from the definition of the pattern P. Let $r(\vec{t}) \rightarrow p^+(\#1, \#2)$, where $\#1, \#2 \in \vec{t}$, be that rule. According to that pattern definition, we can add to the node labelled $P[x, x']$ a child labelled $r(\rho(\vec{t}))$. Since the Datalog rule was applied according to a homomorphism π'', we define, for every fresh variable $\rho(t)$, $\pi'(\rho(t)) = \pi''(t)$. Do the same for every standard pattern atom of T. The instance associated with T is full, and π' is a homomorphism from this full instance to the initial fact base. □

Proposition 3 *Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ and $R \in \mathcal{R}_L^+$. For every instance Q of $(\mathcal{Q}, \mathbb{P})$ and unifier μ of Q with R, there exist an instance of interest Q' of $(\mathcal{Q}, \mathbb{P})$ w.r.t. R and a unifier μ' of Q' with R such that μ' is more general than μ.*

Proof: Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ, $R \in \mathcal{R}_L^+$ be a rule with head $p(\vec{t})$, Q be an instance of $(\mathcal{Q}, \mathbb{P})$, and $\mu = (Q_2, p(\vec{t}), P_u)$ be a unifier of Q with R.

Consider a repeatable pattern $P^+[t_1, t_2]$ from which some atoms in Q_2 are expanded. If $P^+[t_1, t_2]$ is expanded into $k \leq arity(p) + 2$ standard patterns in Q[2], then there is an instance of interest Q' that expands $P^+[t_1, t_2]$ into exactly k standard patterns. Thus, there is an isomorphism π between the atoms expanded under $P^+[t_1, t_2]$ in Q and in Q'. Assume instead that $P^+[t_1, t_2]$ is expanded in Q into $k > arity(p) + 2$ standard patterns. We denote by σ a substitution associated with P_u, and by $P[t_1 = x_0, x_1], P[x_1, x_2], \ldots, P[x_{k-2}, x_{k-1}], P[x_{k-1}, x_k = t_2]$ the sequence of standard patterns expanded from $P^+[t_1, t_2]$ in Q, and let x_s and x_e ($s < e$) be the external terms of $P^+[t_1, t_2]$ w.r.t. μ. The unifier is single-piece (cf. Section 2), thus, for every $0 < i < k$, $\sigma(x_i) = \sigma(z_i)$ for some existential variable z_i from the head of R.

We construct an instance Q' and function π as follows. Starting from $\mathcal{Q}$, we expand every repeatable pattern $P^+[t_1, t_2]$ that is relevant for μ into $e - s$ standard patterns (where e and s are defined as above and depend on the particular pattern):

$$P[t_1 = x'_s, x'_{s+1}], P[x'_{s+1}, x'_{s+2}], \ldots, P[x'_{e-1}, x'_e = t_2].$$

Then for every $s \leq i \leq e$, we set $\pi(x'_i) = x_i$, and we expand $P[x'_i, x'_{i+1}]$ into the atom a'_i that is obtained from the the atom a_i expanded under $P[x_i, x_{i+1}]$ in Q by replacing every x_j

[2]Strictly speaking, we mean the instantiation underlying Q, but to simplify the notation, here and later in the appendix, we will often refer to instances, leaving the instantiation implicit.

with x'_j. If $e - s \leq arity(p) + 2$, we are done. Otherwise, we will need to remove some patterns in order to satisfy the definition of instances of interest. To this end, we define a sequence $s < i_1 < j_1 < \ldots < i_m < j_m < e$ of indices as follows:

- We call $i < j$, with $s < i < j < e$, a *matching pair* if a_i and in a_j, x'_{i+1} and x'_j occur at the same position of p (hence, $\sigma(x_{i+1}) = \sigma(x_j)$);
- We say that a matching pair $i < j$ is *maximal w.r.t. index* ℓ if the following conditions hold:
  - $i \geq \ell$,
  - there is no matching pair $i' < j'$ with $\ell \leq i' < i$
  - there is no matching pair $i < j'$ with $j' > j$
- We let $i_1 < j_1$ be the matching pair that is maximal w.r.t. index $s + 1$
- If $i_k < j_k$ is already defined, then we let $i_{k+1} < j_{k+1}$ be the matching pair that is maximal w.r.t. index j_k, if such a pair exists (otherwise, $i_k < j_k$ is the final pair in the sequence).

Now remove from Q' all of the patterns $P[x'_\ell, x'_{\ell+1}]$ such that $i_g < \ell < j_g$ for some $1 \leq g \leq m$, as well as the atoms that are expanded from such patterns. We claim that there are now at most $arity(p)+2$ patterns $P[x'_\ell, x'_{\ell+1}]$ below $P^+[t_1, t_2]$ in Q'. Indeed, if this were not the case, we could find a matching pair $i < j$ among the remaining patterns. Since $i_1 < j_1$ is maximal w.r.t. index $s + 1$, and there are no further matching pairs starting from j_m, we know that $i \geq i_1$ and $i < j_m$. Moreover, since a_i is still present in Q', it must be the case that $j_g < i < i_{g+1}$ for some $1 \leq g < m$. But this contradicts the fact that $i_{g+1} < j_{g+1}$ is maximal w.r.t. j_g.

In order for the different remaining patterns to form a sequence, we will need to perform a renaming of terms. If there are n patterns left under $P^+[t_1, t_2]$, then we will rename these patterns from left to right by:

$$P[t_1 = x''_0, x''_1], P[x''_1, x''_2], \ldots, P[x''_{n-1}, x''_n = t_2].$$

and will rename the atoms underneath these patterns accordingly. We will also update π by setting $\pi(x''_i) = \pi(x'_j)$ if x'_j was renamed into x''_i and there is no $x'_{j'}$ with $j' < j$ that was also renamed into x''_i.

Let Q' be the instance obtained in this manner. We note that by construction, it is an instance of interest of $(\mathcal{Q}, \mathbb{P})$ w.r.t. R, as we only expand patterns into atoms that use the predicate p from the rule head, and the number of patterns generated from any repeatable pattern is at most $arity(p)+2$.

Regarding π, note that a term may be shared among several repeatable patterns that are relevant for μ. However, we can show that if a term is shared by multiple relevant patterns, then the (partial) mapping associated with those patterns will agree on the shared term, i.e. π is well defined. First note if a term is shared by two repeatable patterns, then it must appear as one of the distinguished terms (t_1, t_2) in both patterns. Moreover, by tracing the above construction, we find that π is the identity on such terms.

To complete the definition of π, we extend it to all of the terms of Q' by letting π be the identity on all terms that do not occur underneath a developed repeatable pattern (i.e., terms that appear in a repeatable pattern that is not expanded, or in one of the standard atom of $\mathcal{Q}$). Observe that π is an injective function, so its inverse π^{-1} is well-defined.

Now let Q'_2 consist of all atoms in $Q' \cap Q_2$ that are not expanded from a repeatable pattern (i.e., they are standard atoms from $\mathcal{Q}$) as well as all atoms in Q' that lie under a repeatable pattern.

Note that by construction every term t in Q'_2 is such that $\pi(t)$ appears in Q_2. We can thus define a partition P'_u of the terms in $Q'_2 \cup H$ by taking every class C in P_u and replacing every term t from Q_2 by $\pi^{-1}(t)$, if such a term exists, and otherwise deleting t; terms from $p(\vec{t})$ are left untouched. Moreover, by the injectivity of π, every term appears in at most one class, i.e., P'_u is indeed a partition.

We aim to show that $\mu' = (Q'_2, p(\vec{t}), P'_u)$ is the desired unifier, We first show that μ' is a unifier of Q' with R. In what follows, it will prove convenient to extend π to the terms in the head atom $p(\vec{t})$, by letting π be the identity on such terms. We will let σ be a substitution associated with μ, and let σ' be the corresponding substitution for μ' defined by setting $\sigma'(t) = \sigma(\pi(t))$.

- P'_u is admissible: since π is the identity on constants, if a class in P'_u contains two constants c, d, then the corresponding class in P_u must also contain c, d (a contradiction).
- $\sigma'(p(\vec{t})) = \sigma'(Q'_2)$: since $\sigma'(p(\vec{t})) = \sigma(p(\vec{t}))$ (due to our choice of σ'), it suffices to prove that $\sigma'(Q'_2) \subseteq \sigma(Q_2)$. First take some atom α that belongs to $Q'_2 \cap Q_2$. Then we have $\pi(\alpha) = \alpha$, so $\sigma'(\alpha) = \sigma(\pi(\alpha)) \in \sigma(Q_2)$. Next consider the case of an atom α that belongs to Q_2 but not Q'_2. Then α must lie below a repeatable pattern $P^+[t_1, t_2]$ that is expanded into $k > arity(p) + 2$ standard patterns $P[t_1 = x_0, x_1], P[x_1, x_2], \ldots, P[x_{k-2}, x_{k-1}], P[x_{k-1}, x_k = t_2]$ in Q. In this case, $P^+[t_1, t_2]$ is expanded in Q' into

  $$P[t_1 = x'_s, x'_{s+1}], P[x'_{s+1}, x'_{s+2}], \ldots, P[x'_{e-1}, x'_e = t_2],$$

  and each $P[x'_i, x'_{i+1}]$ is expanded into a'_i. If $e - s \leq arity(p) + 2$, then the atoms a'_i all belong to Q'_2. If we have $\alpha = a'_i$, then we have $\pi(a'_i) = a_i$, hence $\sigma'(\alpha) = \sigma(\pi(a'_i)) = \sigma(a_i) \in \sigma(Q_2)$. The final possibility is that $e - s > arity(p) + 2$, in which case some of the patterns will be removed and the remaining patterns will be renamed (as will be their corresponding atoms). Suppose that α is the atom a''_h below the pattern $P[x''_\ell, x''_{\ell+1}]$, which was obtained from renaming the pattern $P[x_h, x_{h+1}]$. We claim that $\sigma'(\alpha) = \sigma(\pi(\alpha)) = \sigma(a_h)$, hence $\sigma'(\alpha) \in \sigma(Q_2)$. By examining the way renaming is performed, there are two situations that can occur:
  - $\pi(x''_\ell) = x_h$ and $\pi(x''_{\ell+1}) = x_{h+1}$: in this case, $\pi(a''_h) = a_h$, hence $\sigma'(\alpha) = \sigma(a_h)$.
  - $\pi(x''_\ell) \neq x_h$: in this case, there must exist a matching pair $i_g < j_g$ such that $h = j_g$, $\pi(x''_\ell) = x_{i_g+1}$, and $\pi(x''_{\ell+1}) = x_{h+1}$. From the definition of matching pairs, we know that $\sigma(x_{i_g+1}) = \sigma(x_{j_g})$.

It follows that $\sigma'(x''_\ell) = \sigma(\pi(x''_\ell)) = \sigma(x_{i_g+1}) = \sigma(x_h)$ and $\sigma'(x''_{\ell+1}) = \sigma(\pi(x''_{\ell+1})) = \sigma(x_{h+1})$. We can thus conclude that $\sigma'(\alpha) = \sigma(a_h)$.

- for a contradiction, suppose the class C' in P'_u contains an existential variable z from H and either a constant or a variable that occurs in $Q' \setminus Q'_2$. If it contains a constant c, then the corresponding class C in P_u will contain both z and c, i.e., C is not a valid class. Next suppose that C' contains a variable x that occurs in $Q' \setminus Q'_2$, which means that the corresponding class C in P_u contains $\pi(x)$. Since x that occurs in $Q' \setminus Q'_2$, it must either appear in a standard atom of Q' that does not appear under any repeatable pattern or in a repeatable pattern that is not developed in Q'. In the former case, the same atom appears in $Q \setminus Q_2$, and in the latter case, since Q is full, there is an atom in Q that is developed from the repeatable pattern and contains $\pi(x)$, but which does not participate in Q_2. In both cases, we obtain a contradiction.

Finally, we show that μ' is more general than μ:

- $\sigma'(Q'_2)) \subseteq \sigma(Q_2)$: proven above.
- if $\sigma'(u_1) = \sigma'(u_2)$, then u_1, u_2 belong to the same class in P'_u, and so $\pi(u_1)$ and $\pi(u_2)$ must belong to the same class in P_u.

We have thus shown that Q' is an instance of interest of $(\mathcal{Q}, \mathbb{P})$ w.r.t. R such that there is a unifier μ' of Q' with R with $\mu' \geq \mu$. □

Proposition 4 *Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ where $P^+[t_1, t_2]$ occurs and $R \in \mathcal{R}_L^+$. For any instance Q of $(\mathcal{Q}, \mathbb{P})$, any classical direct rewriting $\mathcal{Q}'$ of Q with R w.r.t. to a unifier internal to $P^+[t_1, t_2]$, and any $Q' \in full(\mathcal{Q}', \mathbb{P})$, there exists a direct rewriting $\mathbb{P}'$ of $\mathbb{P}$ w.r.t. P and R such that $(\mathcal{Q}, \mathbb{P}')$ has a full instance that is isomorphic to Q'.*

Proof: Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ where $P^+[t_1, t_2]$ occurs, $R = (B \to H) \in R_L^+$, Q be an instance of $(\mathcal{Q}, \mathbb{P})$, $\mu = (Q_2, H, P_u)$ be a unifier internal to $P^+[t_1, t_2]$ of Q with R, $\mathcal{Q}'$ be the classical direct rewriting of Q with R w.r.t. μ, and Q' be a full instance of $(\mathcal{Q}', \mathbb{P})$.

Since μ is internal to $P^+[t_1, t_2]$, all atoms in Q_2 are expanded from $P^+[t_1, t_2]$ in Q, and do not unify t_1 with t_2, nor t_1 (resp. t_2) with an existential variable from H. We denote by $P[t_1 = x_0, x_1], P[x_1, x_2], \ldots, P[x_{k-1}, x_k = t_2]$ the sequence of standard patterns expanded under $P^+[t_1, t_2]$ in Q, x_s and x_e ($s < e$) the external terms of $P^+[t_1, t_2]$ w.r.t. μ, and a_i the atom expanded under $P[x_i, x_{i+1}]$. From Prop. 3, there is a unifier μ' of an instance of interest Q_3 of $\mathcal{Q}$ with R with $\mu' \geq \mu$. Since x_s and x_e are not unified with existential variables, let Q_4 be the CQ obtained from Q_3 by removing all atoms and patterns that are not relevant for μ'. Obviously Q_4 is an instance of interest of a PCQ of form $P^+[t_1, t_2]$. Let $\mathbb{P}'$ be the direct rewriting of $\mathbb{P}$ w.r.t. μ', obtained from Q_4.

Let $A_l = \{P[x_i, x_{i+1}] \mid 0 \leq i < s\}$, $A_m = \{P[x_i, x_{i+1}] \mid s \leq i < e\}$, $A_r = \{P[x_i, x_{i+1}] \mid e \leq i < k\}$, and $A = A_l \cup A_m \cup A_r$. Further let A'_l (resp. A'_m, A'_r, A') be the set of atoms expanded under A_l (resp. A_m, A_r, A) in Q.

Initialize Q'' to $Q \setminus A' \cup \{P^+[t_1, t_2]\}$. One can see that Q'' is an instance of both $(\mathcal{Q}, \mathbb{P})$ and $(\mathcal{Q}, \mathbb{P}')$. If B is a not a repeatable pattern, let $\ell = 1$, otherwise let S be the repeatable pattern in B, and $S[x'_0, x'_1], \ldots, S[x'_{\ell-1}, x'_\ell]$ be the sequence expanded from $S^+[x'_0, x'_\ell]$ in Q'. We denote by a'_i the atom expanded under $S[x'_i, x'_{i+1}]$. Then expand $P^+[t_1, t_2]$ in Q'' into $k' = |A_l| + |A_r| + \ell$ standard patterns: $P[t_1 = x''_0, x''_1], \ldots, P[x''_{k-1}, x''_{k'} = t_2]$. Let π be the function defined as follows:

- for all $0 \leq i \leq s$, $\pi(x''_i) = x_i$;
- for all $s < i < s + \ell$, $\pi(x''_i) = x'_{i-s}$;
- for all $s + \ell \leq i \leq k'$, $\pi(x''_i) = x_{i-\ell+(e-s)}$.

Note that π is injective, so its inverse exists. Expand all $P[x''_i, x''_{i+1}]$ with $0 \leq i < s$ or $s + \ell \leq i < k$ (resp. $s \leq i < s + \ell$) into $\pi^{-1}(a_i)$ (resp. $\pi^{-1}(a'_i)$). Finally, for all terms u in Q'' for which π is not defined (i.e., those terms appearing in atoms that were not expanded from the pattern $P^+[t_1, t_2]$), we set $\pi(u) = u$.

By construction, Q'' is still an instance of $(\mathcal{Q}, \mathbb{P}')$ and π is an isomorphism between Q' and Q''. □

Proposition 5 *Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ and $R \in \mathcal{R}_L^+$. For every $Q \in full(\mathcal{Q}, \mathbb{P})$ and every classical direct rewriting Q' of Q with R w.r.t. an external unifier, there is a direct rewriting $\mathcal{Q}'$ of $\mathcal{Q}$ w.r.t. $\mathbb{P}$ and R that has an instance isomorphic to Q'.*

Proof: Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ, $R = (B \to H) \in R_L^+$, $Q \in full(\mathcal{Q}, \mathbb{P})$, $\mu = (Q_u, H, P_u)$ be an external unifier of Q with R, and Q' be the classical direct rewriting of Q with R w.r.t. μ.

From Proposition 3, there is an instance of interest Q_2 of $(\mathcal{Q}, \mathbb{P})$ such that there is a unifier $\mu' = (Q_{u'}, H, P_{u'}) \geq \mu$ of Q_2 with R. We denote by σ (resp. σ') a substitution associated with μ (resp. μ').

For any repeatable pattern $P^+[t_1, t_2]$ in $\mathcal{Q}$, build A, A_l, A_m and A_r as in the proof of Proposition 4 using the instance Q_2 and unifier μ'. Assume t_1 (or t_2) is unified with an existential variable, then from the condition on external unifiers, either A_l or A_r is empty. Consider the minimally-unifiable instance $\mathcal{Q}_M$ of $\mathcal{Q}$ w.r.t. μ' that replaces $P^+[t_1, t_2]$ by: (i) A_m if $A_l = A_r = \emptyset$; (ii) $P^+[t_1, x_s], A_m$ if $A_r = \emptyset$ and $A_l \neq \emptyset$; or (iii) $A_m, P^+[x_e, t_2]$ if $A_l = \emptyset$ and $A_r \neq \emptyset$. In case (ii) (resp. (iii)), since all atoms in A_l (resp. A_r) are not involved in μ', x_s (resp. x_e) is not unified with an existential variable (or the piece condition on unifiers would not be satisfied). Therefore, μ' is a unifier of $\mathcal{Q}_M$ with R. We let $\mathcal{Q}'$ be the direct rewriting of $\mathcal{Q}_M$ w.r.t. μ' and R.

Note that each repeatable pattern $P^+[t_1, t_2]$ in Q' expands into $A_l \wedge \sigma(B) \wedge A_r$, and in $\mathcal{Q}'$ there is a $P^+[t_1, x_s]$ (resp. $P^+[x_e, t_2]$) iff A_l (resp. A_r) is not empty. Thus consider Q'' obtained from $\mathcal{Q}'$ by expanding $P^+[t_1, x_s]$ (resp. $P^+[x_e, t_2]$) into k standard patterns where $k = |A_l|$ (resp. $k = |A_r|$), and choose the same atoms as in A'_l (resp. A'_r). Since $\mu' \geq \mu$, there is an homomorphism π from $\sigma'(Q_{u'})$ to $\sigma(Q_u)$. Note that if we restrict π to terms in $\sigma'(B)$, π is an isomorphism. Furthermore, we can extend π to Q'' in the same way as we did in the previous proof. Thus Q'' is isomorphic to Q'. □

Proposition 6 *Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ and $R \in \mathcal{R}_L^+$. If $\mathcal{Q}'$ is a non-excluded direct rewriting of $\mathcal{Q}$ with R, then $|\mathcal{Q}'| \leq |\mathcal{Q}|$.*

Proof: Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ, $R = (B \rightarrow H) \in R_L^+$, Q be a non-excluded minimally-unifiable instance of $(\mathcal{Q}, \mathbb{P})$, $\mu = (Q', H, P_u)$ be an external unifier of Q with R, and σ the substitution induced by P_u.

Note that all repeatable patterns $P^+[t_1, t_2]$ are at most replaced by the sequence S needed by the unifier (i.e., $S \subseteq Q'$), plus a single repeatable pattern $P^+[t_1, x_1]$ (or $P^+[x_k, t_2]$). Indeed, the only situation that would lead us to introduce more than one more repeatable pattern (i.e., as in External Rewriting case (i)) is when either t_1 or t_2 is unified with an existential variable. However, if t_1 (or t_2) is unified with an existential variable, because of the piece condition on unifiers, no unifier of $P^+[t_1, x_1] \wedge S \wedge P^+[x_k, t_2]$ can be found.

Since $|B| = 1$, we have to show that all atoms that were introduced when replacing a repeatable pattern are erased by the direct rewriting of Q w.r.t. μ.

If Q' consists of at least one atom that is not expanded from a pattern, the direct rewriting of Q w.r.t. μ erases this atom.

Next assume Q' consists only of atoms expanded from repeatable patterns. If $Q' = \{P^+[t_1, t_2]\}$ and neither t_1 nor t_2 is unified with an existential variable, then $\sigma(t_1) = \sigma(t_2)$, so the only non-excluded minimally-unifiable instance of $\mathcal{Q}$ w.r.t. μ replaces $P^+[t_1, t_2]$ only by the sequence S needed by the unifier (see the first condition on non-excluded minimally-unifiable instances). Thus, the direct rewriting erases $P^+[t_1, t_2]$.

Otherwise, we know that at least one $P^+[t_1, t_2]$ from $\mathcal{Q}$ is replaced by the sequence S involved in the unifier (see the second condition on non-excluded minimally-unifiable instances), thus there is at least one $P^+[t_1, t_2]$ erased by the direct rewriting. □

We will break the proof of Theorem 2 into the following five lemmas.

Lemma 1 *Let Q be a CQ, $R \in R^T$, $\mathbb{P}_0$ be the initial set of pattern definitions relative to R^T (see Step 1 of the algorithm overview), and $\mathcal{Q}^+$ be obtained from Q by replacing all atoms $p(t_1, t_2)$ such that p is a transitive predicate by $P^+[t_1, t_2]$. If there is a classical direct rewriting Q' of Q with R, then there is a full instance Q'' of $(\mathcal{Q}^+, \mathbb{P}_0)$ that is isomorphic to Q'.*

Proof: Let $p(t_1, t_2)$ be the atom of Q that is rewritten to obtain Q'. Since p is a transitive predicate, it occurs in a pattern definition P_0 in $\mathbb{P}$, and $\mathcal{Q}^+$ contains the atom $P^+[t_1, t_2]$. In Q', $p(t_1, t_2)$ is rewritten into $p(t_1, x_1) \wedge p(x_1, t_2)$. Let Q'' be the full instance of $(\mathcal{Q}^+, \mathbb{P}_0)$ that expands all repeatable patterns but $P^+[t_1, t_2]$ into a single standard pattern, expands $P^+[t_1, t_2]$ into two standard patterns $P[t_1, x_1']$, $P[x_1', t_2]$, and then further expands the standard patterns using the unique atom in each of the pattern definitions. It is clear that Q'' is isomorphic to Q' (simply map x_1' to x_1 and all other terms to themselves). □

Lemma 2 *Let $\mathbb{P}$ be a set of pattern definitions, $\mathbb{P}_0 \subseteq \mathbb{P}$ be the initial set of patterns definitions built from the set $\mathcal{R}_T$ of transitivity rules, $(\mathcal{Q}, \mathbb{P})$ be a PCQ that does not contain any standard atom using a transitive predicate, Q be a full instance of $(\mathcal{Q}, \mathbb{P})$, and $\mathcal{Q}^+$ be obtained from Q by replacing all atoms $p(t_1, t_2)$ with p transitive by $P^+[t_1, t_2]$.*

Then, for every full instance Q' of $(\mathcal{Q}^+, \mathbb{P}_0)$, there is a full instance Q'' of $(\mathcal{Q}, \mathbb{P})$ such that Q'' is isomorphic to Q'.

Proof: We build the instance Q'' as follows. Initialize Q'' to the atoms and repeatable patterns occurring in $\mathcal{Q}$. Next, for all repeatable patterns $P_i^+[t_1, t_2]$ in the instantiation underlying Q consider each of the atom that is expanded from a child of $P_i^+[t_1, t_2]$ in turn, working from left to right. If the atom $p(\vec{t})$ under $P_i[u, v]$ is being considered, then do the following:

- if p is not a transitive predicate, then add a single child $P_i[u, v]$ to $P_i^+[t_1, t_2]$, and expand it into $p(\vec{t})$.
- if p is transitive, then $p(\vec{t})$ has been replaced in $\mathcal{Q}^+$ by $P^+[\vec{t}]$. We also know that $\vec{t}$ consists of the terms u, v from $P_i[u, v]$. We suppose that $p(\vec{t}) = p(u, v)$ (hence $P^+[\vec{t}] = P^+[u, v]$); a similar argument can be used if the positions are reversed. Let $P[u = x_0, x_1], \ldots, P[x_{k-1}, x_k = v]$ be the children of $P^+[u, v]$ in Q', and a_ℓ be the atom expanded under $P[x_\ell, x_{\ell+1}]$ $(0 \leq \ell < k)$. In place of the child $P_i^+[u, v]$ in $\mathcal{Q}^+$, we will add k children to $P_i^+[t_1, t_2]$ in Q'': $P_i[u = x_0, x_1], \ldots, P_i[x_{k-1}, x_k = v]$, and expand $P_i[x_j, x_{j+1}]$ into a_j. Note that we may assume that the terms x_i $(0 < i < k)$ are fresh, i.e., they do not already appear in Q''.

It can be verified that the resulting full instance Q'' is isomorphic to Q'. Indeed, all atoms in Q' that are also in Q are present in Q''. All other atoms belong to a sequence of transitive atoms, which we have reproduced (modulo renaming of variables) in Q''. □

Lemma 3 *Let Q be a CQ and $\mathcal{R}$ be a set of linear+trans rules, and let $(\mathbb{Q}, \mathbb{P})$ be the output of the algorithm. For any Q' obtained from a sequence of classical direct rewritings of Q with $\mathcal{R}$, there is a PCQ $(\mathcal{Q}, \mathbb{P})$ with $\mathcal{Q} \in \mathbb{Q}$ and a full instance Q'' of $(\mathcal{Q}, \mathbb{P})$ s.t. Q'' is isomorphic to Q'.*

Proof: Let $Q = Q_0, \mu_1, Q_1, \mu_2, Q_2, \ldots, \mu_k, Q_k = Q'$ be a sequence of classical direct rewritings from Q to Q', and let $R_1, \ldots, R_k$ be the associated sequence of rules from $\mathcal{R}$.

We show the desired property, by induction on $0 \leq i \leq k$. For the base case ($i = 0$), we can set $\mathcal{Q}_0 = Q^+$, since $Q_0 = Q$ is clearly a full instance of $(\mathcal{Q}_0, \mathbb{P})$.

For the induction step, suppose that we have $\mathcal{Q}_{i-1} \in \mathbb{Q}$ and a full instance Q_{i-1}'' of $(\mathcal{Q}_{i-1}, \mathbb{P})$ that is isomorphic to the CQ Q_{i-1}. There are two cases to consider, depending on the type of the rule R_i.

If R_i is a transitivity rule, then from Lemma 1, we know that $\mathcal{Q}_{i-1}^+$ (obtained from Q_{i-1} by replacing every transitive predicate p by pattern P^+) is such that there is a full instance Q_{i-1}^+ of $(\mathcal{Q}_{i-1}^+, \mathbb{P})$ that is isomorphic to Q_i. Furthermore, we know that $\mathcal{Q}_{i-1}$ cannot contain any standard atoms with transitive predicates, since every PCQ produced in Step 5 contains patterns for the transitive predicates. Thus, we may apply Lemma 2 and infer that Q_{i-1}^+ is isomorphic to some full instance Q_{i-1}'' of $(\mathcal{Q}_{i-1}, \mathbb{P})$. Therefore, Q_{i-1}'' is isomorphic to Q_i.

If R_i is not a transitive rule, since Q_{i-1} is isomorphic to some full instance Q''_{i-1} of $(\mathcal{Q}_{i-1}, \mathbb{P})$, let μ''_i be the unifier of Q''_{i-1} with R_i obtained from μ and the isomorphism between Q_{i-1} and Q''_{i-1}. If μ''_i is internal to some repeatable pattern, then from Proposition 4, we know that there is an instance Q'_i of $(\mathcal{Q}_{i-1}, \mathbb{P})$ that is isomorphic to Q_i. Otherwise, from Proposition 5, there exists μ'_i and a direct rewriting $\mathcal{Q}_i$ of $\mathcal{Q}_{i-1}$ with μ'_i such that there is an instance Q'_i of $(\mathcal{Q}_i, \mathbb{P})$ that is isomorphic to Q_i.

We have thus completed the inductive argument and can conclude that there is a PCQ $(\mathcal{Q}, \mathbb{P})$ with $\mathcal{Q} \in \mathbb{Q}$ and a full instance Q'' of $(\mathcal{Q}, \mathbb{P})$ s.t. Q'' is isomorphic to $Q' = Q_k$. □

Lemma 4 *Let Q be a CQ, $(\mathcal{F}, \mathcal{R})$ be a linear+trans KB, and $(\Pi_{\mathbb{P}}, Q_{\mathbb{Q}})$ be the output of the algorithm. If $\mathcal{F}, \mathcal{R} \models Q$ then $\mathcal{F}, \Pi_{\mathbb{P}} \models Q'$ for some $Q' \in Q_{\mathbb{Q}}$.*

Proof: Since $\mathcal{F}, \mathcal{R} \models Q$, there is a (finite) classical rewriting Q' of Q with $\mathcal{R}$ such that $\mathcal{F} \models Q'$. From Proposition 3, there is there is a PCQ $(\mathcal{Q}, \mathbb{P})$ with $\mathcal{Q} \in \mathbb{Q}$ and a full instance Q'' of $(\mathcal{Q}, \mathbb{P})$ s.t. Q'' is isomorphic to Q'. Therefore, $\mathcal{F} \models Q''$. We conclude by Proposition 2. □

Lemma 5 *Let Q be a CQ, $(\mathcal{F}, \mathcal{R})$ be a linear+trans KB, and $(\Pi_{\mathbb{P}}, Q_{\mathbb{Q}})$ be the output of the algorithm. If $\mathcal{F}, \Pi_{\mathbb{P}} \models Q_{\mathbb{Q}}$ then $\mathcal{F}, \mathcal{R} \models Q$.*

Proof: Let $\mathbb{P}$ be the set of pattern definitions computed in Step 3 of the algorithm, and $\Pi_{\mathbb{P}}$ the corresponding set of Datalog rules. Consider the CQ Q^{++} obtained from Q by replacing every atom $p(t_1, t_2)$ such that p is transitive by the atom $p^+(t_1, t_2)$. The following claim establishes the soundness of the internal rewriting mechanism in Step 3:

Claim 1 *If $\mathcal{F}, \Pi_{\mathbb{P}} \models Q^{++}$, then $\mathcal{F}, \mathcal{R} \models Q$.*

Proof of claim. Let $\mathbb{P}_0, \mathbb{P}_1, \ldots, \mathbb{P}_k = \mathbb{P}$ be the sequence of sets of pattern definitions that led to $\mathbb{P}$ in Step 3, with $\mathbb{P}_{i+1}$ being obtained from $\mathbb{P}_i$ by a single direct (internal) rewriting step. We prove by induction two distinct properties expressed at rank $0 \leq j \leq k$:

P1 every rule in $\Pi_{\mathbb{P}_j}$ is a semantic consequence of $\Pi_{\mathbb{P}_0} \cup \mathcal{R}$.

P2 for every fact base $\mathcal{F}'$ and CQ Q' (over the original vocabulary):

$$\mathcal{F}', \Pi_{\mathbb{P}_j} \models (Q')^{++} \quad \Rightarrow \quad \mathcal{F}', \mathcal{R} \models Q'$$

In the second property, $(Q')^{++}$ denotes the CQ obtained by replacing every atom $p(t_1, t_2)$ such that p is transitive by the atom $p^+(t_1, t_2)$. Observe that **P2** at rank k yields the claim: we simply take $\mathcal{F}' = \mathcal{F}$ and $Q' = Q$.

Base case ($i = 0$): property **P1** is obviously verified. For property **P2**, we note that $\mathbb{P}_0$ consists of the following rules for every transitive predicate p: the transitivity rule $p^+(x, y) \wedge p^+(y, z) \rightarrow p^+(x, z)$ and the initialization rule $p(x, y) \rightarrow p^+(x, y)$. Clearly, if $\mathcal{F}, \Pi_{\mathbb{P}_0} \models (Q')^{++}$, then we have $\mathcal{F}, \mathcal{R} \models Q'$, since if we can derive $p^+(a, b)$ using $\mathcal{F}, \Pi_{\mathbb{P}_0}$, then we can also derive $p(a, b)$ from $\mathcal{F}, \mathcal{R}$ using the transitivity rule for p in $\mathcal{R}$.

*Induction step for **P1**:* we assume property **P1** holds for some rank $0 \leq i < k$ and show that it holds also for $i + 1$.

Suppose that $\mathbb{P}_{i+1}$ is obtained from $\mathbb{P}_i$ by a single direct rewriting step w.r.t. pattern name P and the rule $R = B \rightarrow H \in \mathcal{R}^+_L$. Let $\mathcal{Q} = P^+[x, y]$, Q be the considered instance of interest of $\mathcal{Q}$ w.r.t. R, $\mu = (Q', H, P_u)$ be the considered internal unifier of Q with H, and σ be the considered substitution associated with μ that preserves the external terms. Finally, let B' be obtained from $\sigma(B)$ by substituting the first (resp. second) external term by $\#1$ (resp. $\#2$).

Since we know that μ is an internal unifier, the external terms of Q' cannot be unified together or with an existential variable. Thus by considering Q'' and $P_{u'}$ obtained from Q' and P_u by substituting the first (resp. second) external term by $\#1$ (resp. $\#2$), it is clear that $\mu' = (Q'', H, P_{u'})$ is a unifier of Q'' with R such that $\sigma'(B) = B'$, where σ' is the substitution associated with μ' that preserves the special terms $\#1$ and $\#2$.

We consider two cases depending on the nature of B'.

Case 1: The first possibility is that B' is an atom (as opposed to a repeatable pattern), in which case we add the following rule to $\Pi_{\mathbb{P}_i}$: $B' \rightarrow p^+(\#1, \#2)$.

Let $a_1, \ldots, a_k$ be the atoms of Q'', and let a'_j be the atom in P's definition from which a_j is obtained. Since there is a rewriting of $\{a_j \mid 0 < j \leq k\}$ with R into B' (using the unifier μ'), and the rule R appears in the original set of rules $\mathcal{R}$, it follows that

$$\mathcal{R} \models B' \rightarrow a_1 \wedge \ldots \wedge a_k$$

From the induction hypothesis, we know that the rules $a'_j \rightarrow p^+(\#1, \#2)$ $(0 < j \leq k)$ are entailed by $\Pi_{\mathbb{P}_0}, \mathcal{R}$. We also know that for all $1 < j \leq k$, the atoms a_{j-1} and a_j share a variable corresponding respectively to $\#2$ in a'_{j-1} and to $\#1$ in a'_j. Thus, by applying the rules $a'_j \rightarrow p^+(\#1, \#2)$ $(0 < j \leq k)$ to the conjunction $a_1 \wedge \ldots \wedge a_k$, we obtain $p^+(\#1, x_1) \wedge p^+(x_1, x_2) \wedge \ldots \wedge p^+(x_{k-1}, \#2)$. Hence:

$$\Pi_{\mathbb{P}_0}, \mathcal{R} \models \bigwedge_{j=0}^{k} a_j \rightarrow p^+(\#1, x_1) \wedge p^+(x_1, x_2) \wedge \ldots \wedge p^+(x_{k-1}, \#2)$$

Since $\Pi_{\mathbb{P}_0}$ contains a transitivity rule for p^+, we can further infer that

$$\Pi_{\mathbb{P}_0} \models p^+(x_1, x_2) \wedge \ldots \wedge p^+(x_{k-1}, \#2) \rightarrow p^+(\#1, \#2)$$

By chaining together the preceding entailments, we obtain $\Pi_{\mathbb{P}_0}, \mathcal{R} \models (B' \rightarrow p^+(\#1, \#2))$, as desired.

Case 2: The other possibility is that B' is a repeatable pattern of the form $S^+[\#1, \#2]$ or $S^+[\#2, \#1]$. Let f be a bijection on $\{\#1, \#2\}$: if B' is of the form $S^+[\#1, \#2]$, f is the identity, otherwise f permutes $\#1$ and $\#2$. Then for all s_ℓ in the definition of S, we add $f(s_\ell)$ to P's definition, and we add the corresponding rules $f(s_\ell) \rightarrow p^+(\#1, \#2)$ to $\Pi_{\mathbb{P}_i}$. Consider one such rule rule $f(s_\ell) \rightarrow p^+(\#1, \#2)$.

Let $a_1, \ldots, a_k$ and $a'_1, \ldots, a'_k$ be defined as in Case 1. Since there is a rewriting of $\{a_j \mid 0 < j \leq k\}$ with $R \in \mathcal{R}^+_L$ into B', and since the rule R was obtained from a rule R' in $\mathcal{R}$ by replacing the transitive predicate s in the rule head by the repeatable pattern S^+, it follows that

$$\mathcal{R} \models f(s(\#1, \#2)) \rightarrow a_1 \wedge \ldots \wedge a_k$$

Arguing as in Case 1, we obtain

$$\Pi_{\mathbb{P}_0}, \mathcal{R} \models f(s(\#1, \#2)) \rightarrow p^+(\#1, \#2)$$

From the induction hypothesis, we know that that the rules $s_\ell \rightarrow s^+(\#1, \#2)$ are entailed from $\Pi_{\mathbb{P}_0}, \mathcal{R}$, and the same obviously holds for the rules $f(s_\ell) \rightarrow f(s^+(\#1, \#2))$. By combining the preceding entailments, we obtain $\Pi_{\mathbb{P}_0}, \mathcal{R} \models f(s_\ell) \rightarrow p^+(\#1, \#2)$.

*Induction step for property **P2**:* we assume **P2** holds for some rank $0 \leq i < k$ and show that it holds also for $i + 1$.

Suppose now that $\mathcal{F}', \Pi_{\mathbb{P}_{i+1}} \models (Q')^{++}$, for some fact base $\mathcal{F}'$ and CQ Q' (over the original predicates). This means that there is a finite derivation sequence $\mathcal{F}' = \mathcal{F}_0^{++}, \ldots, \mathcal{F}_m^{++}$ such that $\mathcal{F}_m^{++} \models (Q')^{++}$ and such that for all $0 \leq \ell < m$, $\mathcal{F}_{\ell+1}^{++}$ is obtained from $\mathcal{F}_\ell^{++}$ either (i) by a sequence of applications of rules from $\Pi_{\mathbb{P}_i}$ or (ii) by a sequence of applications of rules from $\Pi_{\mathbb{P}_{i+1}} \setminus \Pi_{\mathbb{P}_i}$.

In case (i), we have $\mathcal{F}_\ell^{++}, \Pi_{\mathbb{P}_i} \models \mathcal{F}_{\ell+1^{++}}$. Letting $\mathcal{F}_r$ be the fact base obtained by replacing every predicate p^+ in $\mathcal{F}_r^{++}$ by the corresponding predicate p, and recalling that $\Pi_{\mathbb{P}_i}$ contains the rule $p(x, y) \rightarrow p^+(x, y)$, we have $\mathcal{F}_\ell, \Pi_{\mathbb{P}_i} \models \mathcal{F}_{\ell+1}^{++}$. Applying the induction hypothesis (treating $\mathcal{F}_{\ell+1}^{++}$ as a CQ), we obtain $\mathcal{F}_\ell, \mathcal{R} \models \mathcal{F}_{\ell+1}$.

In case (ii), we have $\mathcal{F}_\ell^{++}, (\Pi_{\mathbb{P}_{i+1}} \setminus \Pi_{\mathbb{P}_i}) \models \mathcal{F}_{\ell+1}^{++}$. From property **P1**, we obtain $\mathcal{F}_\ell^{++}, \Pi_{\mathbb{P}_0}, \mathcal{R} \models \mathcal{F}_{\ell+1}^{++}$. Using the rule $p(x, y) \rightarrow p^+(x, y)$ (that is present in $\Pi_{\mathbb{P}_0}$), the latter yields $\mathcal{F}_\ell, \Pi_{\mathbb{P}_0}, \mathcal{R} \models \mathcal{F}_{\ell+1}^{++}$. Finally, we note that if we can derive $p^+(a, b)$ from $\mathcal{F}_\ell, \Pi_{\mathbb{P}_0}, \mathcal{R}$, then we can also infer $p(a, b)$ from $\mathcal{F}_\ell, \mathcal{R}$ by using the transitivity rule for p instead of using $p(x, y) \rightarrow p^+(x, y)$ and the transitivity rule for p^+. Thus, we have $\mathcal{F}_\ell, \mathcal{R} \models \mathcal{F}_{\ell+1}$.

We have thus shown that for every $0 \leq \ell < m$, $\mathcal{F}_\ell, \mathcal{R} \models \mathcal{F}_{\ell+1}$. Since $\mathcal{F}' = \mathcal{F}_0$, by chaining these implications together, we obtain $\mathcal{F}', \mathcal{R} \models \mathcal{F}_m$. Using the same reasoning as above, we can infer $\mathcal{F}_m \models Q'$ from $\mathcal{F}_m^{++} \models (Q')^{++}$. Then, by combining these statements, we obtain $\mathcal{F}', \mathcal{R} \models Q'$. *(end proof of claim)*

Now let $\mathbb{Q}$ be the set of queries computed in Step 5 by performing all possible external direct rewritings w.r.t. $\mathbb{P}$ and rules from $\mathcal{R}_L^+$, starting from Q^+, and let $Q_\mathbb{Q}$ be the set of CQs associated with $\mathbb{Q}$ (defined as in Step 6). We start by proving the following claim, which relates external direct rewriting steps to sequences of classical direct rewritings.

Claim 2 *Let $\mathcal{Q}_{i+1}$ be a direct rewriting of $\mathcal{Q}_i$ w.r.t. $\mathbb{P}$. Then every full instance of $(\mathcal{Q}_{i+1}, \mathbb{P})$ is obtained from a sequence of (classical) direct rewritings of some full instance of $(\mathcal{Q}_i, \mathbb{P})$.*

Proof of claim. Let $(\mathcal{Q}_{i+1}, \mathbb{P})$ be obtained from an external rewriting of $(\mathcal{Q}_i, \mathbb{P})$ with rule $R = B \rightarrow H$. This means that there is a minimally unifiable instance $\mathcal{Q}^e$ and a unifier $\mu = (X, H, P_u)$ of $\mathcal{Q}^e$ with H (with associated substitution σ) such that $\mathcal{Q}_{i+1} = \sigma(\mathcal{Q}^e \setminus X) \cup \sigma(B)$.

Let us consider a partial instance $\mathcal{Q}_{i+1}^P$ of $(\mathcal{Q}_{i+1}, \mathbb{P})$ that fully instantiates $\sigma(\mathcal{Q}^e \setminus X)$ but does not instantiate $\sigma(B)$ (we say that it is a $\sigma(B)$-excluding instance). Note that $\mathcal{Q}_{i+1}^P$ can be built equivalently by choosing a full instance Q^e of $(\mathcal{Q}^e, \mathbb{P})$, removing the atoms of X, then by applying the substitution σ and adding $\sigma(B)$. We can see that the classical direct rewriting of Q^e according to μ produces $\mathcal{Q}_{i+1}^P$. Moreover, since every full instance of $(\mathcal{Q}^e, \mathbb{P})$ is a full instance of $(\mathcal{Q}_i, \mathbb{P})$, we know that Q^e is an instance of $(\mathcal{Q}_i, \mathbb{P})$.

Now consider any full instance Q_{i+1} of $(\mathcal{Q}_{i+1}, \mathbb{P})$. Note that it is a full instance of some $\sigma(B)$-excluding instance $(\mathcal{Q}_{i+1}^P, \mathbb{P})$. There are two cases to consider:

- If $\sigma(B)$ is an atom, then $Q_{i+1} = \mathcal{Q}_{i+1}^P$ and thus Q_{i+1} is obtained from a classical direct rewriting of an instance of $(\mathcal{Q}_i, \mathbb{P})$.
- Otherwise, if $\sigma(B)$ is a repeatable pattern, then Q_{i+1} is obtained from $(\mathcal{Q}_{i+1}^P, \mathbb{P})$ by expanding $\sigma(B)$ into a sequence of k standard patterns, and expanding each of them into some atom a_ℓ. Let $B_k = \{a_\ell \mid 1 \leq \ell \leq k\}$. Then, $\sigma(B)$ is generated in forward chaining from B_k with a sequence of applications of rules: k applications of transitivity rules, and k applications of the rules encoded in $\mathbb{P}$, each one stemming from a finite sequence of applications of rules of $\mathcal{R}$ (see Claim 1). Thus from the completeness of classical rewriting, B_k can be obtained from a sequence of classical direct rewritings from $\sigma(B)$, and thus Q_{i+1} is obtained from a sequence of classical direct rewritings of an instance of $(\mathcal{Q}_i, \mathbb{P})$.

(end proof of claim)

The following claim shows the soundness of the external rewriting in Step 5 and completes the proof of the lemma.

Claim 3 *If $\mathcal{F}, \Pi_\mathbb{P} \models Q_\mathbb{Q}$, then $\mathcal{F}, \mathcal{R} \models Q$.*

Proof of claim. Suppose that $\mathcal{F}, \Pi_\mathbb{P} \models Q_\mathcal{Q}$ with $\mathcal{Q} \in \mathbb{Q}$. We know that the PCQ $\mathcal{Q}$ is obtained from a finite sequence $\mathcal{Q}_0 = Q^+, \mathcal{Q}_1, \ldots, \mathcal{Q}_k = \mathcal{Q}$ of PCQs such that for all $0 \leq j < k$, $(\mathcal{Q}_{j+1}, \mathbb{P})$ is a direct external rewriting of $(\mathcal{Q}_j, \mathbb{P})$. We will show by induction on j that $\mathcal{F}, \Pi_\mathbb{P} \models Q_{\mathcal{Q}_j}$ implies $\mathcal{F}, \mathcal{R} \models Q$ for every $0 \leq j \leq k$.

The base case ($j = 0$) is a direct consequence of Claim 1. For the induction step, we assume the property is true at rank i, and we show that it is true at rank $i + 1$.

Suppose that $\mathcal{F}, \Pi_\mathbb{P} \models Q_{\mathcal{Q}_{i+1}}$. From Proposition 2, it follows that there is a full instance Q_{i+1} of $(\mathcal{Q}_{i+1}, \mathbb{P})$ such that $\mathcal{F} \models Q_{i+1}$. By Claim 2, there is a full instance Q_i of $(\mathcal{Q}_i, \mathbb{P})$ such that Q_{i+1} is obtained from a sequence of classical rewritings from Q_i. Thus (from the correctness of the classical rewriting), there is a fact base $\mathcal{F}'$ such that $\mathcal{F}, \mathcal{R} \models \mathcal{F}'$ and $\mathcal{F}' \models Q_i$. Applying Proposition 2, we obtain $\mathcal{F}', \Pi_\mathbb{P} \models Q_{\mathcal{Q}_i}$. Now from our induction hypothesis, it follows that $\mathcal{F}', \mathcal{R} \models Q$, hence $\mathcal{F}, \mathcal{R} \models Q$. *(end proof of claim)* □

Theorem 2 *Let Q be a CQ, $(\mathcal{F}, \mathcal{R})$ be a linear+trans KB, and ($\Pi_\mathbb{P}$,$Q_\mathbb{Q}$) be the output of the algorithm. Then: $\mathcal{F}, \mathcal{R} \models Q$ iff $\mathcal{F}, \Pi_\mathbb{P} \models Q'$ for some $Q' \in Q_\mathbb{Q}$.*

Proof: Follows from Lemma 4 and Lemma 5. □

The following two lemmas show that the non-excluded minimally-unifiable instances are sufficient to ensure completeness when the input query is atomic or when the input rule set satisfies the safety condition.

Lemma 6 *Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ, $R \in R_L^+$, Q be an instance of interest of $(\mathcal{Q}, \mathbb{P})$ and $\mu = (Q', H, P_u)$ be an external unifier of Q with R such that two external terms w.r.t. μ from a given pattern $P^+[t_1, t_2]$ are unified together and with no existential variable.*

Every minimally-unifiable instance $(\mathcal{Q}, \mathbb{P})$ w.r.t. μ that replaces $P^+[t_1, t_2]$ as in the External Rewriting cases (i), (ii), or (iii) will lead to a direct rewriting $(\mathcal{Q}'_i, \mathbb{P})$ that is more specific than $(\mathcal{Q}, \mathbb{P})$. Furthermore, for any classical direct rewriting $\mathcal{Q}''_i$ of $\mathcal{Q}'_i$ with R, either $(\mathcal{Q}, \mathbb{P}) \geq (\mathcal{Q}''_i, \mathbb{P})$ or there is a classical direct rewriting $\mathcal{Q}'$ of the minimally-unifiable instance of $\mathcal{Q}$ that replaces $P^+[t_1, t_2]$ as in case (iv) with R and $(\mathcal{Q}', \mathbb{P}) \geq (\mathcal{Q}''_i, \mathbb{P})$.

Proof: Without loss of generality, let us write $\mathcal{Q} = q[t_1, t_2] \wedge P^+[t_1, t_2]$ where $q[t_1, t_2]$ denotes a set of atoms where t_1 and t_2 may occur. We denote by x_s and x_e ($s < e$) the external terms of $P^+[t_1, t_2]$ w.r.t. μ, and by $A[x_s, x_e]$ the sequence of atoms expanded from $P^+[t_1, t_2]$ involved in the unifier. Since we assume that no existential variable is unified with variables x_s and x_e, no atom from q can be part of the unifier. Consider the following minimally-unifiable instances:

1. $\mathcal{Q}_1 = q[t_1, t_2] \wedge P^+[t_1, x_s] \wedge A[x_s, x_e] \wedge P^+[x_e, t_2]$
2. $\mathcal{Q}_2 = q[t_1, t_2] \wedge [t_1 = x_s, x_e] \wedge P^+[x_e, t_2]$
3. $\mathcal{Q}_3 = q[t_1, t_2] \wedge P^+[t_1, x_s] \wedge X[x_s, x_e = t_2]$

By unifying x_s and x_e together, we obtain the following instances:

1. $q[t_1, t_2] \wedge P^+[t_1, x_s] \wedge A[x_s, x_s] \wedge P^+[x_s, t_2]$
2. $q[t_1, t_2] \wedge A[t_1, t_1] \wedge P^+[t_1, t_2]$
3. $q[t_1, t_2] \wedge P^+[t_1, t_2] \wedge A[t_2, t_2]$

Let $\mathcal{Q}'_i$ be the direct rewriting of $\mathcal{Q}_i$ w.r.t. μ with R. It is easy to see that $\mathcal{Q} \subseteq \mathcal{Q}'_2$ and $\mathcal{Q} \subseteq \mathcal{Q}'_3$, thus, $(\mathcal{Q}'_2, \mathbb{P})$ and $(\mathcal{Q}'_3, \mathbb{P})$ are more specific than $(\mathcal{Q}, \mathbb{P})$.

Let Q_1 be a full instance of $(\mathcal{Q}'_1, \mathbb{P})$. We construct a full instance Q of $(\mathcal{Q}, \mathbb{P})$ as follows. First note that $q[t_1, t_2]$ is common to both $\mathcal{Q}'_1$ and $\mathcal{Q}$, so we will expand all patterns in $q[t_1, t_2]$ exactly as in Q_1. Now let k_1 (resp. k_2) be the number of children of $P^+[t_1, x_s]$ (resp. $P^+[x_s, t_2]$) in the instantiation of Q_1, and expand $P^+[t_1, t_2]$ in Q into $k = k_1 + k_2$ children: $P[t_1 = x_0, x_1], \ldots, P[x_{k-1}, x_k = t_2]$. Expand each $P[x_i, x_{i+1}]$ with $i < k_1$ as is expanded the i^{th} child of $P^+[t_1, x_s]$ in Q_1; and each $P[x_i, x_{i+1}]$ with $k_1 \leq i < k$ as is expanded the $(i - k_1 + 1)^{th}$ child of $P^+[x_s, t_2]$ in Q_1. By construction, there is an homomorphism from Q to Q_1. We have thus shown that $(\mathcal{Q}, \mathbb{P}) \geq (\mathcal{Q}'_1, \mathbb{P})$.

Furthermore, let $\mathcal{Q}''_i$ be a classical direct rewriting of $\mathcal{Q}'_i$ with a rule $R' = B' \to H'$ w.r.t. unifier $\mu' = (Q', H', P'_u)$, where $1 \leq i \leq 3$. If at least one atom involved in μ' occurs in $\mathcal{Q}'_i \setminus \sigma(B)$ (where σ is the substitution associated with μ), then, let $\mu'' = \{Q'', H', P''_u\}$ where $Q'' = Q' \setminus \sigma(B)$ and P''_u is the restriction of P'_u to terms occurring in $Q'' \cup H'$. Since $Q'' \neq \emptyset$ and all terms from $\sigma(B)$ cannot connect two different terms from $q[t_1, t_2]$ (indeed, the only term shared between $\sigma(B)$ and $q[t_1, t_2]$ is either t_1 or t_2), $\sigma(B)$ can be seen as a loop on t_1 (or t_2), therefore we can remove $\sigma(B)$ while preserving the unifier, i.e., μ'' is a unifier of $\mathcal{Q}$ with R'. Moreover, since P''_u and Q'' are only restrictions of P'_u and Q' respectively, it holds that $\mu'' \geq \mu'$. Then, we denote by $\mathcal{Q}''$ the direct rewriting of $\mathcal{Q}$ with R' w.r.t. μ'' and obtain $\mathcal{Q}'' \geq \mathcal{Q}'_i$. The other possibility is that all atoms involved in μ' occur in $\sigma(B)$, then, $\mathcal{Q}''_2$ (resp. $\mathcal{Q}''_3$) is more specific than $\mathcal{Q}$ since $\mathcal{Q} \subseteq \mathcal{Q}''_2$ (resp. $\mathcal{Q} \subseteq \mathcal{Q}''_3$). Moreover, for any instance Q''_1 of $\mathcal{Q}''_1$, one can easily build an instance Q' of $\mathcal{Q}$ in the same way as above, and see that $Q' \geq Q''_1$. Thus, we have $(\mathcal{Q}, \mathbb{P}) \geq (\mathcal{Q}''_i, \mathbb{P})$. □

Lemma 7 *Let $(\mathcal{Q}, \mathbb{P})$ be a PCQ, $R \in \mathcal{R}_L^+$, Q be an instance of interest of $(\mathcal{Q}, \mathbb{P})$ and $\mu = (Q', H, P_u)$ be an external unifier of Q with R such that one external term w.r.t. μ from a given pattern $P^+[t_1, t_2]$ is unified with an existential variable, and where all atoms in Q' are obtained from the expansion of a repeatable pattern.*

If $\mathcal{Q}$ is atomic, or if $\mathcal{R}_L$ is a set of safe linear rules, then every minimally-unifiable instance of $(\mathcal{Q}, \mathbb{P})$ w.r.t. μ that replaces all $P_i^+[t_1, t_2]$ as in the External Rewriting cases (ii) or (iii) will lead to a direct rewriting $(\mathcal{Q}'_i, \mathbb{P})$ that is more specific than $(\mathcal{Q}, \mathbb{P})$. Furthermore, for any direct rewriting $\mathcal{Q}''_i$ of $\mathcal{Q}'_i$ with R, either $\mathcal{Q} \geq \mathcal{Q}''_i$ or there is a direct rewriting $\mathcal{Q}'$ of a minimally-unifiable instance of $\mathcal{Q}$ w.r.t. μ that replaces at least one repeatable pattern as in case (iv) and is such that $\mathcal{Q}' \geq \mathcal{Q}''_i$.

Proof: Let $(\mathcal{Q}, \mathbb{P})$, R, Q and μ be as in the lemma statement, and let $P_1^+[t_1^1, t_2^1], \ldots, P_k^+[t_1^k, t_2^k]$ be the repeatable patterns that are relevant for μ. For each $1 \leq i \leq k$, we denote by $P_i[t_1^i = x_0^i, x_1^i], \ldots, P_i[x_{k-1}^i, x_{k_i}^i = t_2^i]$ the sequence of standard patterns expanded from $P_i^+[t_1^i, t_2^i]$, and we let $x_{s_i}^i$ and $x_{e_i}^i$ ($s_i < e_i$) be the external terms of $P_i^+[t_1^i, t_2^i]$ w.r.t. μ. We assume without loss of generality that it is $x_{e_i}^i$ that is unified with an existential variable, and let $A_i[x_{s_i}^i, x_{e_i}^i = t_2^i]$ denote the atoms expanded from $P_i[x_j^i, x_{j+1}^i]$ with $s_i \leq j < e_i$.

Since the unifier μ is single-piece, all repeatable patterns relevant for μ have to share some variable. For simplicity, we assume that they all share their second term, i.e. $t_2^i = t_2^j$ for all $1 \leq i, j \leq k$. (The argument is entirely similar, just more notationally involved, if this assumption is not made.) Let us use t_2 for this shared term. Then we can write $\mathcal{Q}$ as follows:

$$\mathcal{Q} = q[t_1^1, \ldots, t_1^k] \wedge \bigwedge_{1 \leq i \leq k} P_i^+[t_1^i, t_2]$$

Note that t_2 cannot occur in q.

Because we have chosen the second term to be shared in all repeatable patterns, we only need to consider the minimally-unifiable instance $\mathcal{Q}_M$ of $(\mathcal{Q}, \mathbb{P})$ w.r.t. μ that replaces each $P_i^+[t_1^i, t_2]$ by $P_i^+[t_1^i, x_s^i], A_i[x_s^i, x_e^i = t_2]$, i.e. External Rewriting case (ii). Thus, we have

$$\mathcal{Q}_M = q[t_1^1, \ldots, t_1^k] \wedge \bigwedge_{1 \leq i \leq k} (P_i^+[t_1^i, x_s^i] \wedge A_i[x_s^i, t_2]).$$

Let σ be the substitution associated with μ. From the safety condition (see Section 6), we know that there is a pair of positions $\{p_1, p_2\}$ for the predicate p of H, such that for all atoms $p(\vec{t})$ occurring in a pattern definition the terms #1 and #2 occurs in positions $\{p_1, p_2\}$. We further note that the external terms in the concerned patterns are t_2 (which unifies with an

existential variable in H) and the terms x_s^i (which unify with a non-existential variable), and each of these external terms must be obtained by instantiating term $\#1$ or $\#2$. Since the $A_i[x_s^i, t_2]$ are unified together, and share the same predicate p, it follows that all of the x_s^i must occur in the same position (either p_1 or p_2) of p; t_2 will occur in the other position among p_1 and p_2. We therefore obtain;

$$\sigma(x_s^1) = \sigma(x_s^2) = \ldots = \sigma(x_s^k) = x',$$

where x' is the term in B that unifies with all of the x_s^i. (Note that if $\mathcal{Q}$ is an atomic query, there is a single A_i, so the previous statement obviously holds, even without the safety condition.) Thus, $\mathcal{Q}_M$ becomes:

$$q[t_1^1, \ldots, t_1^k] \wedge \bigwedge_{1 \leq i \leq k} (P_i^+[t_1^i, x'] \wedge A_i[x', t_2]).$$

There is an isomorphism from $\mathcal{Q}$ to $\mathcal{Q}_M \setminus \{A_i \mid 1 \leq i \leq k\}$ that maps t_2 to x'. We then observe that $\{A_i \mid 1 \leq i \leq k\}$ is exactly the set of atoms that will be erased in the direct rewriting $\mathcal{Q}'_M = \mathcal{Q}_M \setminus \{A_i \mid 1 \leq i \leq k\} \cup \sigma(B)$, where σ is a substitution associated with μ. Therefore, $\mathcal{Q}$ is isomorphic to $\mathcal{Q}'_M \setminus \sigma(B)$, hence $(\mathcal{Q}, \mathbb{P}) \geq (\mathcal{Q}'_M, \mathbb{P})$. One can see that the same reasoning as in the previous proof can be applied here to show that any further direct rewriting $\mathcal{Q}''_M$ of $\mathcal{Q}'_M$ will lead to more specific queries. □

Theorem 3 *The modified query rewriting algorithm halts. Moreover, Theorem 2 (soudness and completeness) holds for the modified algorithm if either the input CQ is atomic, or the input rule set is safe.*

Proof: From Lemma 4, we know that if we do not exclude any rewriting the algorithm is sound and complete, and Lemma 6 and 7 show that for any rewriting $\mathcal{Q}$ that we exclude, there is another rewriting $\mathcal{Q}'$ obtainable using only non-excluded direct rewritings that is more general than $\mathcal{Q}$. Therefore, the modified algorithm (in case of an atomic CQ, or a safe rule set) is complete. Furthermore, excluding rewritings cannot comprise the soundness of the rewriting mechanism. □

Theorem 4 *Both (i) atomic CQ entailment over linear+trans KBs and (ii) CQ entailment over safe linear+trans KBs are* NL*-complete in data complexity.*

Proof: Consider a CQ Q, a linear+trans rule set $\mathcal{R}$, and a fact base $\mathcal{F}$. Suppose that either Q is atomic or $\mathcal{R}$ satisfies the safety condition. Using Theorem 3, we can compute a finite set $\Pi_\mathbb{P}$ of Datalog rules and a finite set $Q_\mathbb{Q}$ of CQs with the property that $\mathcal{F}, \mathcal{R} \models Q$ iff $\mathcal{F}, \Pi_\mathbb{P} \models Q'$ for some $Q' \in Q_\mathbb{Q}$. As $\Pi_\mathbb{P}$ and $Q_\mathbb{Q}$ do not depend on the fact base $\mathcal{F}$, they can be computed and stored using constant space w.r.t. $|\mathcal{F}|$.

To test whether $\mathcal{F}, \Pi_\mathbb{P} \models Q'$ for some $Q' \in Q_\mathbb{Q}$, we proceed as follows. For each rewriting $Q' \in Q_\mathbb{Q}$, we can consider every possible mapping π from the variables of Q' to the terms of $\mathcal{F}$. We then check whether the facts in $\pi(Q')$ are entailed from $\mathcal{F}, \Pi_\mathbb{P}$. For every atom $\alpha \in Q'$ over one of the original predicates, we can directly check if $\pi(\alpha) \in \mathcal{F}$, since the rules in $\Pi_\mathbb{P}$ can only be used to derive facts over the new predicates p^+. For every atom $p^+[t_1, t_2] \in Q'$ where p^+ is a new predicate, we need to check whether $\mathcal{F}, \Pi_\mathbb{P} \models p^+(\pi(t_1), \pi(t_2))$. Because of the shape of the rules in $\Pi_\mathbb{P}$, the latter holds just in the case that there is a path of constants $c_1, \ldots, c_n$ with $c_1 = \pi(t_1)$ and $c_n = \pi(t_2)$ such that for every $1 \leq i < n$, there is a rule $\rho_i = B_i \rightarrow p^+(\#1, \#2)$ and substitution σ_i of the variables in B_i by constants in $\mathcal{F}$ such that $\sigma_i(\#1) = c_i$, $\sigma_i(\#2) = c_{i+1}$, and $\sigma_i(B_i) \in \mathcal{F}$. To check for the existence of such a path, we guess the constants c_i in the path one at a time, together with the witnessing rule ρ_i and substitution σ_i, using a counter to ensure that the number of guessed constants does not exceed the number of constants in $\mathcal{F}$. Note that we need only logarithmically many bits for the counter, so the entire procedure runs in non-deterministic logarithmic space.

Hardness for NL can be shown by an easy reduction from the NL-complete directed reachability problem. □

Theorem 5 *Both (i) atomic CQ entailment over linear+trans KBs and (ii) CQ entailment over safe linear+trans KBs are* PSPACE*-complete in combined complexity.*

Proof: Consider a CQ Q and linear+trans rule set $\mathcal{R} = \mathcal{R}_L \cup \mathcal{R}_T$, with $\mathcal{R}_L$ a set of linear rules and $\mathcal{R}_T$ a set of transitivity rules, and a fact base $\mathcal{F}$. Suppose that either condition (i) or (ii) of the theorem statement holds. It follows from Theorem 3 that the modified query rewriting algorithm halts and returns a set of a finite set $\Pi_\mathbb{P}$ of Datalog rules and a finite set $Q_\mathbb{Q}$ of CQs such that $\mathcal{F}, \mathcal{R} \models Q$ iff $\mathcal{F}, \Pi_\mathbb{P} \models Q'$ for some $Q' \in Q_\mathbb{Q}$.

To prove membership in PSPACE, we will not be able to explicitly construct $\Pi_\mathbb{P}$ and $Q_\mathbb{Q}$, since both may be of exponential size. Instead, we can proceed as follows:

1. Guess a CQ Q' of size at most $|Q|$.
2. Check that $Q' \in Q_\mathbb{Q}$ (modulo variable renaming).
3. Guess a mapping π from the variables of $\mathcal{Q}'$ to constants from $\mathcal{F}$.
4. Check that every atom in $\pi(\mathcal{Q}')$ is entailed from $\mathcal{F}, \Pi_\mathbb{P}$.
5. Return yes if the checks succeeded, otherwise no.

It should be clear that the non-deterministic procedure just described has a successful execution just in the case that $\mathcal{F}, \Pi_\mathbb{P} \models Q'$ for some $Q' \in Q_\mathbb{Q}$. Since NPSPACE=PSPACE, it suffices to show that the checks in Steps 2 and 4 can be implemented in polynomial space.

First consider the problem of deciding whether $Q' \in Q_\mathbb{Q}$. Starting from the initial query Q, we maintain a counter of how many direct external rewritings we have computed, and at each step we guess the next direct rewriting and increment the counter. Only the two most recent direct rewritings are kept in memory at any time (and by Proposition 6, both are of size at most $|Q|$). If at some point, we reach a query whose corresponding CQ is isomorphic to the guessed CQ Q', then we can stop. Otherwise, the counter will eventually reach the number of non-isomorphic queries built from the available atoms and patterns, and we can halt and return no. (Note that there are only exponentially many such queries, so the counter can be implemented using only polynomially many bits.) Importantly, to perform a direct external rewriting step, we must be able to decide whether some atom is present in a pattern definition P_i in $\mathbb{P}$. This can be done in a similar manner, by guessing a sequence of direct internal rewriting steps,

maintaining a counter of the number of steps, and stopping when either the desired atom is found or the counter exceeds the number of non-isomorphic atoms. We have thus shown how the check in Step 2 can be performed using only polynomial space.

For the entailment checks in Step 4, we can proceed in almost exactly the same manner as in the proof of Theorem 4. The only significant difference is that to check whether $\mathcal{F}, \Pi_{\mathbb{P}} \models p^+(\pi(t_1), \pi(t_2))$, we will need to identify the rules in $\Pi_{\mathbb{P}}$ with head predicate p^+. This can be done by checking whether the corresponding atom appears in the pattern definition of P in $\mathbb{P}$, using the same technique (based upon guessing a sequence of internal rewriting steps) as was used for Step 2.

Finally, we note that hardness for PSPACE is a direct consequence of the PSPACE-hardness of AQ entailment in the presence of linear rules. □

This figure "classes.jpg" is available in "jpg" format from:

http://arxiv.org/ps/1504.07443v1