

Pay-as-you-go OWL Query Answering Using a Triple Store

Yujiao Zhou and Yavor Nenov and Bernardo Cuenca Grau and Ian Horrocks

Department of Computer Science, University of Oxford, UK

Abstract

We present an enhanced hybrid approach to OWL query answering that combines an RDF triple-store with an OWL reasoner in order to provide scalable pay-as-you-go performance. The enhancements presented here include an extension to deal with arbitrary OWL ontologies, and optimisations that significantly improve scalability. We have implemented these techniques in a prototype system, a preliminary evaluation of which has produced very encouraging results.

1 Introduction

The use of RDF (Manola, Miller, and McBride 2004) and SPARQL (SPARQL 1.1 Overview 2013) to store and access semi-structured data is increasingly widespread. In many cases, an OWL ontology is used to formally specify the meaning of data (Motik et al. 2009b), as well as to enhance query answers with tuples that are only entailed by the combination of the ontology and data.

Unfortunately, computing such answers is of high worst case complexity, and although heavily optimised, existing systems for query answering w.r.t. RDF data and an unrestricted OWL 2 ontology can process only small to medium size datasets (Möller et al. 2013; Wandelt, Möller, and Wessel 2010; Kollia and Glimm 2013). This has led to the development of query answering procedures that are more scalable, but that can (fully) process only fragments of OWL 2, and several prominent fragments have now been standardised as OWL 2 profiles (Motik et al. 2009a). Such systems have been shown to be (potentially) highly scalable (Urbani et al. 2012; Bishop et al. 2011; Urbani et al. 2011; Wu et al. 2008), but if the ontology falls outside the relevant profile, then the answers computed by such a system may be incomplete: if it returns an answer, then all tuples in the answer are (usually) valid, but some valid tuples may be missing from the answer. When used with out-of-profile ontologies, a query answer computed by such a system can thus be understood as providing a *lower-bound* on the correct answer; however, they cannot in general provide any upper bound or even any indication as to how complete the computed answer is (Cuenca Grau et al. 2012).

More recently, a hybrid approach has been proposed that addresses some of these issues. This approach uses a triple store with OWL 2 RL reasoning capabilities to compute not only the usual lower bound answer, but also an upper bound answer, in the latter case by rewriting the input ontology into a strictly stronger OWL 2 RL ontology (Zhou et al. 2013b). If lower and upper bound answers coincide, they obviously provide a sound and complete answer. Otherwise, *relevant fragments* of the ontology and data can be extracted that are guaranteed to be sufficient to test the validity of tuples in the “gap” between the two answers; these fragments are typically much smaller than the input ontology and data, and may thus allow for checking the gap tuples using an OWL 2 reasoner such as Hermit (Motik, Shearer, and Horrocks 2009) or Pellet (Sirin et al. 2007). This extraction and checking step was, however, only proven to be correct for Horn ontologies, and also suffered from scalability issues both w.r.t. the extraction itself and the subsequent checking.

In this paper, we present several important enhancements to this hybrid approach. First, we show how the lower bound can be tightened by integrating scalable reasoning techniques for other OWL 2 profiles. Second, we show how to extend the relevant fragment extraction procedure so that it is correct for arbitrary OWL 2 ontologies, and how to use the triple store itself to compute these fragments. Finally, we show how summarisation techniques inspired by the SHER system (Dolby et al. 2007; 2009) can be used to tighten the upper bound on query answers, thus further reducing the requirement for fully-fledged OWL 2 reasoning.

We have implemented our procedure in a prototypical system using RDFox as triple store (Motik et al. 2014) and we present a preliminary evaluation over both benchmark and realistic data which suggests that the system can provide scalable pay-as-you-go query answering for a wide range of OWL 2 ontologies, RDF data and queries. In almost all cases, the system is able to completely answer queries without resorting to fully-fledged OWL 2 reasoning, and even when this is not the case, relevant fragment extraction and summarisation are effective in reducing the size of the problem to manageable proportions. This paper is accompanied with an appendix containing all proofs.

2 Preliminaries

We adopt standard first order logic notions, such as variables, constants, atoms, formulas, clauses, substitutions, satisfiability, and entailment (Bachmair and Ganzinger 2001). We also assume basic familiarity with OWL 2 (Motik et al. 2009b) and its profiles (Motik et al. 2009a).

Datalog Languages Extended datalog languages are well-known KR formalisms based on rules, and they have many connections with OWL 2. A generalised rule (or just a *rule*) is a function-free first order sentence of the form

$$\forall \mathbf{x} \left(\bigwedge_{j=0}^n B_j(\mathbf{x}) \rightarrow \bigvee_{i=0}^m \exists \mathbf{y}_i \varphi_i(\mathbf{x}, \mathbf{y}_i) \right)$$

where $B_j(\mathbf{x})$ are *body* atoms and φ_i are conjunctions of *head* atoms. The universal quantifiers are left implicit from now on. A rule is *Horn* if $m \leq 1$, and it is *datalog* if it is Horn and does not contain existential quantifiers. A *fact* is a ground atom and a *dataset* is a finite set of facts. A *knowledge base* \mathcal{K} consists of a finite set of rules and a dataset. We treat equality (\approx) in knowledge bases as an ordinary predicate, but assume that every knowledge base in which equality occurs contains the axioms of equality for its signature. Each OWL 2 ontology can be normalised as one such knowledge base using the correspondence of OWL and first order logic and a variant of the structural transformation (see (Motik, Shearer, and Horrocks 2009) for details); furthermore, each OWL 2 RL ontology corresponds to a datalog knowledge base. From now on, we interchangeably use ontologies and the knowledge bases they correspond to; in particular, we use OWL 2 RL and datalog interchangeably.

Queries We focus on *conjunctive query answering* as the key reasoning problem. A *query* is a formula of the form $q(\mathbf{x}) = \exists \mathbf{y} \varphi(\mathbf{x}, \mathbf{y})$ with $\varphi(\mathbf{x}, \mathbf{y})$ a conjunction of atoms. Usually, we omit the *distinguished* variables \mathbf{x} of queries and write just q . The query is *atomic* if $\varphi(\mathbf{x}, \mathbf{y})$ is a single atom. A tuple of individuals \mathbf{a} is a (*certain*) *answer* to q w.r.t. a set of sentences \mathcal{F} iff $\mathcal{F} \models q(\mathbf{a})$. The set of all answers to $q(\mathbf{x})$ w.r.t. \mathcal{F} is denoted by $\text{cert}(q, \mathcal{F})$.

Datalog Reasoning There are two main techniques for query answering over a datalog knowledge base \mathcal{K} . *Forward chaining* computes the set $\text{Mat}(\mathcal{K})$ of facts entailed by \mathcal{K} , called the *materialisation* of \mathcal{K} . A query q over \mathcal{K} can be answered directly over the materialisation. *Backward chaining* treats a query as a conjunction of atoms (a *goal*). An *SLD* (*Selective Linear Definite*) *resolvent* of a goal $A \wedge \psi$ with a datalog rule $\varphi \rightarrow C_1 \wedge \dots \wedge C_n$ is a goal $\psi\theta \wedge \varphi\theta$, with θ the MGU (Most General Unifier) of A and C_j , for some $1 \leq j \leq n$. An *SLD proof* of a goal G_0 in \mathcal{K} is a sequence of goals (G_0, \dots, G_n) with G_n the empty goal (\square), and each G_{i+1} a resolvent of G_i and a rule in \mathcal{K} .

3 Overview

Background Our approach to query answering extends (Zhou et al. 2013a; 2013b). The main idea is to exploit a datalog-based triple store to compute both lower bound (sound but possibly incomplete) and upper bound (complete

$$\begin{aligned} \text{Foreman}(x) &\rightarrow \text{Manag}(x) & (T_1) \\ \text{Superv}(x) &\rightarrow \text{Manag}(x) & (T_2) \\ \text{Manag}(x) &\rightarrow \text{Superv}(x) \vee \exists y.(\text{boss}(x, y) \wedge \text{Manag}(y)) & (T_3) \\ \text{Superv}(x) \wedge \text{boss}(x, y) &\rightarrow \text{Workman}(y) & (T_4) \\ \text{TeamLead}(x) \wedge \text{boss}(x, y) \wedge \text{Manag}(y) &\rightarrow & (T_5) \\ \text{Manag}(x) &\rightarrow \exists y.(\text{boss}(x, y)) & (T_6) \\ \\ \text{Manag}(\text{Sue}) & (D_1) & \text{boss}(\text{Dan}, \text{Ben}) & (D_4) \\ \text{Superv}(\text{Dan}) & (D_2) & \text{Manag}(\text{Jo}) & (D_5) \\ \text{Superv}(\text{Rob}) & (D_3) & \text{TeamLead}(\text{Jo}) & (D_6) \\ & & \text{boss}(\text{Jane}, \text{Rob}) & (D_7) \end{aligned}$$

Figure 1: Example knowledge base \mathcal{K}^{ex} .

but possibly unsound) query answers. Given a knowledge base \mathcal{K} and a query q , this is achieved by computing datalog knowledge bases $\mathcal{L}(\mathcal{K})$ and $\mathcal{U}(\mathcal{K})$ s.t. $\text{cert}(q, \mathcal{L}(\mathcal{K})) \subseteq \text{cert}(q, \mathcal{U}(\mathcal{K}))$. In (Zhou et al. 2013a), $\mathcal{L}(\mathcal{K})$ was simply the subset of datalog rules in \mathcal{K} and hence $\mathcal{K} \models \mathcal{L}(\mathcal{K})$. In turn, $\mathcal{U}(\mathcal{K})$ is the result of consecutively applying the transformations Σ , Ξ and Ψ defined next. Σ renders a knowledge base into a set of clauses via standard Skolemisation. Ξ transforms a set of clauses into a satisfiable set of Horn clauses, by first adding a nullary atom \perp to each negative clause, and then splitting each non-Horn clause into several clauses (one for each positive literal). Finally, Ψ transforms the output of Ξ by replacing each functional term by a fresh constant and replacing clauses by their equivalent datalog rules. Both, $\mathcal{L}(\mathcal{K})$ and $\mathcal{U}(\mathcal{K})$ are independent from query and data (they depend on the rules in \mathcal{K}).

We demonstrate these basic ideas using the knowledge base \mathcal{K}^{ex} in Figure 1, which we use as a running example. The query q^{ex} asks for individuals that manage a workman.

$$q^{ex}(x) = \exists y(\text{boss}(x, y) \wedge \text{Workman}(y))$$

We have that $\text{cert}(q^{ex}, \mathcal{K}^{ex}) = \{\text{Dan}, \text{Rob}, \text{Jo}\}$. The lower bound knowledge base $\mathcal{L}(\mathcal{K}^{ex})$ consists of facts D_1 - D_7 and the following datalog rules.

$$\begin{aligned} \text{Foreman}(x) &\rightarrow \text{Manag}(x) & (L_1) \\ \text{Superv}(x) &\rightarrow \text{Manag}(x) & (L_2) \\ \text{Superv}(x) \wedge \text{boss}(x, y) &\rightarrow \text{Workman}(y) & (L_4) \\ \text{TeamLead}(x) \wedge \text{boss}(x, y) \wedge \text{Manag}(y) &\rightarrow & (L_5) \end{aligned}$$

The upper bound $\mathcal{U}(\mathcal{K}^{ex})$ extends $\mathcal{L}(\mathcal{K}^{ex})$ with the rules given next, where c_1 and c_2 are fresh constants:

$$\begin{aligned} \text{Manag}(x) &\rightarrow \text{Superv}(x) & (U_1^3) \\ \text{Manag}(x) &\rightarrow \text{boss}(x, c_1) & (U_2^3) \\ \text{Manag}(x) &\rightarrow \text{Manag}(c_1) & (U_3^3) \\ \text{Manag}(x) &\rightarrow \text{boss}(x, c_2) & (U_6) \end{aligned}$$

It is easy to verify that $\text{cert}(q^{ex}, \mathcal{L}(\mathcal{K}^{ex})) = \{\text{Dan}\}$ and that $\text{cert}(q^{ex}, \mathcal{U}(\mathcal{K}^{ex})) = \{\text{Sue}, \text{Dan}, \text{Rob}, \text{Jo}\}$.

For a knowledge base \mathcal{K} like the one above, for which $\text{cert}(q, \mathcal{L}(\mathcal{K})) \subsetneq \text{cert}(q, \mathcal{U}(\mathcal{K}))$, we check for each ‘‘gap’’

tuple $\mathbf{a} \in \text{cert}(q, \mathcal{U}(\mathcal{K})) \setminus \text{cert}(q, \mathcal{L}(\mathcal{K}))$ whether $\mathbf{a} \in \text{cert}(q, \mathcal{K})$. This could be achieved simply by using a fully-fledged OWL 2 reasoner to check if $\mathcal{K} \models q(\mathbf{a})$, but this is typically not feasible for large datasets (Zhou et al. 2013a). In order to address this issue, (Zhou et al. 2013b) use backwards chaining reasoning over $\mathcal{U}(\mathcal{K})$ to extract a (typically small) relevant subset \mathcal{K}_f of the original knowledge base \mathcal{K} such that $\mathbf{a} \in \text{cert}(q, \mathcal{K})$ iff $\mathbf{a} \in \text{cert}(q, \mathcal{K}_f)$; a fully-fledged OWL 2 reasoner is then used to compute $\text{cert}(q, \mathcal{K}_f)$. Unfortunately, the fragment extraction technique is only answer preserving for Horn ontologies, thus the technique as a whole can only guarantee to compute sound and complete answers when \mathcal{K} is Horn. Moreover, the extraction technique used in (Zhou et al. 2013a) can lead to memory exhaustion in practice, and even when successful, computing $\text{cert}(q, \mathcal{K}_f)$ can still be challenging for fully-fledged OWL 2 reasoners.

Our Contribution In the following sections we show how this technique can be extended so as to deal with arbitrary ontologies, and to improve the quality and scalability of lower and upper bound computations. In Section 4, we show how the lower bound answer can be tightened by additionally applying other scalable query answering procedures that also rely on datalog reasoning. In Section 5, we show how the fragment extraction technique in (Zhou et al. 2013b) can be extended to deal with non-Horn ontologies, and how the datalog engine itself can be used to perform the extraction. Finally, in Section 6, we show how a summarisation technique inspired by the SHER system can be used to tighten the upper bound by ruling out “obvious” non-answers, and how we can further exploit this idea to reduce the number of calls to the fully-fledged OWL 2 reasoner.

Our approach is *pay-as-you-go* in the sense that the bulk of the computation is delegated to a highly scalable datalog engine. Although our main goal is to answer queries over OWL 2 ontologies efficiently, our technical results are very general and hence our approach is not restricted to OWL. More precisely, given a first-order KR language L that can be captured by generalised rules and over which we want to answer conjunctive queries, our only assumption is the availability of a fully-fledged reasoner for L and a datalog reasoner, which are both used as a “black box”.

Related techniques. The SCREECH system (Tserendorj et al. 2008) first exploits the KAON2 reasoner (Hustadt, Motik, and Sattler 2007) to rewrite a *SHIQ* ontology into disjunctive datalog while preserving atomic queries, and then transforms \vee into \wedge ; the resulting over-approximation can be used to compute upper bound query answers. However, this technique is restricted to *SHIQ* ontologies and atomic queries; furthermore, the set of rules obtained from KAON2 can be expensive to compute, as well as of exponential size. Both the QUILL system (Pan, Thomas, and Zhao 2009) and the work of (Wandelt, Möller, and Wessel 2010) under-approximate the ontology into OWL 2 QL; however, neither approximation is independent of both query and data, and using OWL 2 QL increases the chances that the approximated ontology will be unsatisfiable.

The SHER system uses summarisation (see Section 6 for details) to efficiently compute an upper bound answer, with

exact answers then being computed via successive relaxations (Dolby et al. 2007; 2009). The technique has been shown to be scalable in practice, but it is only known to be applicable to *SHLN* and atomic queries, and is less modular than our approach. In contrast, our approach can profitably exploit the summarisation technique, and could even improve scalability for the hardest queries by replacing Hermit with SHER when the extracted fragment is *SHLN*.

4 Tightening Lower Bounds

A direct way to compute lower bound answers given \mathcal{K} and q is to select the datalog fragment $\mathcal{L}(\mathcal{K})$ of \mathcal{K} and compute $\text{cert}(q, \mathcal{L}(\mathcal{K}))$ using the datalog engine. If \mathcal{K} is an OWL 2 knowledge base, however, it is possible to further exploit the datalog engine to compute a larger set of lower bound answers. To this end, of particular interest is the *combined approach* introduced to handle query answering in \mathcal{ELHO}_\perp^r (Stefanoni, Motik, and Horrocks 2013)—a logic that captures most of OWL 2 EL. Given an \mathcal{ELHO}_\perp^r knowledge base \mathcal{K}' and a query q , the combined approach first computes the upper bound datalog program $\mathcal{U}(\mathcal{K}')$ and the corresponding answers $\text{cert}(q, \mathcal{U}(\mathcal{K}'))$. A subsequent filtering step Φ , which is efficiently implementable, guarantees to eliminate all spurious tuples; the resulting answer $\Phi(\text{cert}(q, \mathcal{U}(\mathcal{K}')))$ is thus sound and complete w.r.t. q and \mathcal{K}' .

The combined approach is clearly compatible with ours. Given an OWL 2 knowledge base \mathcal{K} and query q , the procedure we use consists of the following steps. First, similarly to (Zhou et al. 2013a), we select the datalog fragment $\mathcal{K}_1 = \mathcal{L}(\mathcal{K})$, and compute the materialisation $\text{Mat}(\mathcal{K}_1)$ using the datalog engine. Second, we select the subset \mathcal{K}_2 of \mathcal{K} corresponding to \mathcal{ELHO}_\perp^r axioms and Skolemise existential quantifiers to constants to obtain $\mathcal{U}(\mathcal{K}_2)$. Then, we further compute the answers $\text{cert}(q, \mathcal{U}(\mathcal{K}_2) \cup \text{Mat}(\mathcal{K}_1))$. Finally, we apply the aforementioned filtering step Φ to obtain the final set of lower bound answers. Note that \mathcal{K}_1 and \mathcal{K}_2 are neither disjoint nor contained within each other.

The \mathcal{ELHO}_\perp^r fragment for our running example \mathcal{K}^{ex} consists of axioms T_1, T_2, T_4 and T_6 , and the resulting new lower bound answer of q^{ex} is the set $\{Dan, Rob\}$.

5 Computing Relevant Fragments

If lower and upper bound answers coincide, and $\mathcal{U}(\mathcal{K})$ does not entail the nullary atom \perp , then we have fully answered the query q . Otherwise, we consider the remaining candidate answers S and compute a (hopefully small) fragment of \mathcal{K} that is sufficient to determine whether \mathcal{K} is satisfiable, as well as whether each tuple in S is indeed an answer to q .

The Horn Case (Zhou et al. 2013b) propose an algorithm for computing such a fragment for the specific case where \mathcal{K} is Horn. The algorithm proceeds in two steps. First, if $\mathcal{U}(\mathcal{K})$ is found to entail \perp , the algorithm computes a fragment \mathcal{K}_\perp of \mathcal{K} . If \mathcal{K} is found satisfiable, then the algorithm computes $\mathcal{K}_{[q, S]}$ for the relevant candidate answers S and checks each of them using the fully-fledged reasoner. The relevant fragment extraction is done by an inspection of all SLD proofs in $\mathcal{U}(\mathcal{K})$ of \perp (for the first step) and each answer in S (for the second step). The two fragments are defined as follows.

Table 1: An SLD proof of $q^{ex}(Rob)$ in $\mathcal{U}(\mathcal{K}^{ex})$

$b(R, y) \wedge W(y)$	
$M(R) \wedge W(y)$	by U_6
$S(R) \wedge W(c_1)$	by U_3^3
$W(c_1)$	by D_2
$S(x) \wedge m(x, c_1)$	by L_4
$b(R, c_2)$	by D_2
$M(R)$	by U_6
$S(R)$	by L_2
\square	by D_2

Definition 1. Let \mathcal{K} be a knowledge base, $q(\mathbf{x})$ be a query, and S be a set of tuples. Then \mathcal{K}_\perp (resp. $\mathcal{K}_{[q,S]}$) is the set of all $\alpha \in \mathcal{K}$ for which there exists $\beta \in \mathcal{U}(\alpha)$ involved in an SLD proof of \perp (resp. $Q(\mathbf{a})$, for some $\mathbf{a} \in S$) in $\mathcal{U}(\mathcal{K})$.

Consider the SLD proof of $q^{ex}(Rob)$ in $\mathcal{U}(\mathcal{K}^{ex})$ presented in Table 1, where predicates and constants are abbreviated to their first letters. By Definition 1, T_2, T_4, T_6 , and D_2 are included in $\mathcal{K}^{ex}_{[q^{ex},\{Rob\}]}$, which will entail $q^{ex}(Rob)$. Note that, in general we need to consider *all* proofs of $q^{ex}(Rob)$ in $\mathcal{U}(\mathcal{K}^{ex})$, since $\mathcal{U}(\mathcal{K}^{ex})$ overapproximates \mathcal{K}^{ex} .

This technique is, however, only complete for Horn knowledge bases. Indeed, recall that $Jo \in \text{cert}(q^{ex}, \mathcal{K}^{ex})$, and note that every fragment of \mathcal{K}^{ex} that entails $q^{ex}(Jo)$ must include rule T_5 . According to Definition 1, $\mathcal{K}_{[q,\{Jo\}]}$ will include T_5 if and only if L_5 is used in an SLD proof of $q^{ex}(Jo)$ in $\mathcal{U}(\mathcal{K}^{ex})$. However, no such proof will involve L_5 , since the goal $q^{ex}(Jo)$ does not involve \perp , and there is no way of eliminating \perp from a goal using the rules in $\mathcal{U}(\mathcal{K}^{ex})$ since they do not contain \perp in their bodies.

The General Case We show that this technique can be extended to the general case by taking into account the interaction between non-Horn rules and rules with \perp in the head. In particular, we show that extending $\mathcal{K}_{[q,S]}$ with \mathcal{K}_\perp when checking candidate answers suffices to ensure completeness.

Theorem 1. Let \mathcal{K} be a knowledge base, $q(\mathbf{x})$ a conjunctive query, and S a set of tuples. Then, \mathcal{K} is satisfiable iff \mathcal{K}_\perp is satisfiable; furthermore, if \mathcal{K} is satisfiable, then,

$$\mathcal{K} \models q(\mathbf{a}) \text{ iff } \mathcal{K}_{[q,S]} \cup \mathcal{K}_\perp \models q(\mathbf{a}) \text{ for each } \mathbf{a} \in S.$$

Consider the proofs of \perp and $q^{ex}(Jo)$ presented in Table 2. According to Definition 1, $\{T_3, \dots, T_6, D_5, D_6\}$ is a subset of $\mathcal{K}_\perp \cup \mathcal{K}_{[q^{ex},\{Jo\}]}$, and, hence, one can readily check that $q^{ex}(Jo)$ is entailed by $\mathcal{K}_\perp \cup \mathcal{K}_{[q^{ex},\{Jo\}]}$.

The proof of Theorem 1 is involved, and details are deferred to our online appendix. Nonetheless, we sketch the arguments behind the proof. A first observation is that, w.l.o.g. we can restrict ourselves to the case where q is atomic.

Lemma 1. Let \mathcal{K} be a knowledge base, $q(\mathbf{x}) = \exists \mathbf{y} \varphi(\mathbf{x}, \mathbf{y})$ be a conjunctive query, S be a set of tuples, Q be a fresh predicate, and let $\mathcal{K}' = \mathcal{K}_{[q,S]} \cup \mathcal{K}_\perp$, then

$$\mathcal{K}' \models q(\mathbf{a}) \text{ iff } \mathcal{K}' \cup \{\varphi(\mathbf{x}, \mathbf{y}) \rightarrow Q(\mathbf{x})\} \models Q(\mathbf{a})$$

Table 2: SLD proofs in $\mathcal{U}(\mathcal{K}^{ex})$ of \perp and $q^{ex}(Jo)$

\perp		$b(J, y) \wedge W(y)$	
$T(x) \wedge b(x, y) \wedge M(y)$	by L_5	$M(J) \wedge W(c_2)$	by U_6
$b(J, y) \wedge M(y)$	by D_6	$W(c_2)$	by D_5
$M(J) \wedge M(c_2)$	by U_6	$S(x) \wedge m(x, c_2)$	by L_4
$M(c_2)$	by D_5	$M(x) \wedge m(x, c_2)$	by U_3^1
$M(x)$	by U_3^3	$b(J, c_2)$	by D_5
\square	by D_5	$M(J)$	by U_6
		\square	by D_5

The crux of the proof relies on the following properties of the transformation Ξ (the step in the definition of \mathcal{U} which splits each non-Horn clause C into different Horn clauses).

Lemma 2. Let \mathcal{N} be a set of first-order clauses. Then:

- if $C \in \mathcal{N}$ participates in a refutation in \mathcal{N} , then every $C' \in \Xi(C)$ is part of an SLD proof of \perp in $\Xi(\mathcal{N})$;
- if $C \in \mathcal{N}$ participates in a resolution proof in \mathcal{N} of an atomic query $Q(\mathbf{a})$, then each $C' \in \Xi(C)$ participates in an SLD proof of \perp or $Q(\mathbf{a})$ in $\Xi(\mathcal{N})$.

Thus, by Lemma 2, each resolution proof in a set of clauses \mathcal{N} can be mapped to SLD proofs in $\Xi(\mathcal{N})$ that “preserve” the participating clauses. The following lemma, which recapitulates results shown in (Zhou et al. 2013a), allows us to restate Lemma 2 for $\Psi \circ \Xi$ instead of Ξ .

Lemma 3. Let \mathcal{H} be a set of first-order Horn clauses, $Q(\mathbf{x})$ be an atomic query, and \mathbf{a} be a tuple of constants. If a clause C participates in an SLD proof of $Q(\mathbf{a})$ in \mathcal{H} , then $\Psi(C)$ participates in an SLD proof of $Q(\mathbf{a})$ in $\Psi(\mathcal{H})$.

With these Lemmas in hand, we can exploit refutational completeness of resolution and the entailment preservation properties of Skolemisation to show Theorem 1.

Fragment Computation The computation of the relevant fragments requires a scalable algorithm for “tracking” all rules and facts involved in SLD proofs for datalog programs. We present a novel technique that delegates this task to the datalog engine itself. The main idea is to extend the datalog program with additional rules that are responsible for the tracking; in this way, relevant rules and facts can be obtained directly from the materialisation of the modified program.

Definition 2. Let \mathcal{K} be a datalog knowledge base and let F be a set of facts in $\text{Mat}(\mathcal{K})$. Then, $\Delta(\mathcal{K}, F)$ is the datalog program containing the rules and facts given next:

- each rule and fact in \mathcal{K} ;
- a fact $\bar{P}(\mathbf{a})$ for each fact $P(\bar{\mathbf{a}})$ in F ;
- the following rules for each rule $r \in \mathcal{K}$ of the form $B_1(\mathbf{x}_1), \dots, B_m(\mathbf{x}_m) \rightarrow H(\mathbf{x})$, and every $1 \leq i \leq m$:

$$\bar{H}(\mathbf{x}) \wedge B_1(\mathbf{x}_1) \wedge \dots \wedge B_m(\mathbf{x}_m) \rightarrow \mathbf{S}(\mathbf{c}_r) \quad (1)$$

$$\bar{H}(\mathbf{x}) \wedge B_1(\mathbf{x}_1), \dots \wedge B_m(\mathbf{x}_m) \rightarrow \bar{B}_i(\mathbf{x}_i) \quad (2)$$

where \mathbf{c}_r is a fresh constant for each rule r , and \mathbf{S} is a globally fresh predicate.

The auxiliary predicates \bar{P} are used to record facts involved in proofs; in particular, if $\bar{P}(\bar{\mathbf{c}})$ is contained in

$Mat(\Delta(\mathcal{K}, F))$, we can conclude that $P(\vec{c})$ participates in an SLD proof in \mathcal{K} of a fact in F . Furthermore, each rule $r \in \mathcal{K}$ is represented by a fresh constant \mathbf{c}_r , and \mathbf{S} is a fresh predicate that is used to record rules of \mathcal{K} involved in proofs. In particular, if $\mathbf{S}(\mathbf{c}_r)$ is contained in $Mat(\Delta(\mathcal{K}, F))$, we can conclude that rule r participates in an SLD proof in \mathcal{K} of a fact in F . The additional rules (1) and (2) are responsible for the tracking and make sure that the materialisation of $\Delta(\mathcal{K}, F)$ contains the required information. Indeed, if there is an instantiation $B_1(\mathbf{a}_1) \wedge \dots \wedge B_m(\mathbf{a}_m) \rightarrow H(\mathbf{a})$ of a rule $r \in \Delta$, then, by virtue of (1), \mathbf{c}_r will be added to \mathbf{S} , and, by virtue of (2), each $\bar{B}_i(\mathbf{a}_i)$, for $1 \leq i \leq m$, will be derived. Correctness is established as follows.

Theorem 2. *Let \mathcal{K} be a datalog knowledge base and let F be a set of facts in $Mat(\mathcal{K})$. Then, a fact $P(\mathbf{a})$ (resp. a rule r) in \mathcal{K} participates in an SLD proof of some fact in F iff $\bar{P}(\mathbf{a})$ (resp. $\mathbf{S}(\mathbf{c}_r)$) is in $Mat(\Delta(\mathcal{K}, F))$.*

6 Summarisation and Answer Dependencies

Once the relevant fragment has been computed, we still need to check, using the fully-fledged reasoner, whether it entails each candidate answer. This can be computationally expensive if either the fragment is large and complex, or there are many candidate answers to verify.

Data Summarisation To address these issues, we first exploit summarisation techniques (Dolby et al. 2007) to efficiently prune candidate answers. The main idea behind summarisation is to “shrink” the data in a knowledge base by merging all constants that instantiate the same unary predicates. Since summarisation is equivalent to extending the knowledge base with equality assertions between constants, the summarised knowledge base entails the original one by the monotonicity of first-order logic.

Definition 3. *Let \mathcal{K} be a knowledge base. A type T is a set of unary predicates; for a constant a in \mathcal{K} , we say that $T = \{A \mid A(a) \in \mathcal{K}\}$ is the type for a . Furthermore, for each type T , let c_T be a globally fresh constant uniquely associated with T . The summary function over \mathcal{K} is the substitution σ mapping each constant a in \mathcal{K} to c_T , where T is the type for a . Finally, the knowledge base $\sigma(\mathcal{K})$ obtained by replacing each constant a in \mathcal{K} with $\sigma(a)$ is called the summary of \mathcal{K} .*

By summarising a knowledge base in this way, we thus overestimate the answers to queries (Dolby et al. 2007).

Proposition 3. *Let \mathcal{K} be a knowledge base, and let σ be the summary function over \mathcal{K} . Then, for every conjunctive query q we have $\sigma(\text{cert}(q, \mathcal{K})) \subseteq \text{cert}(\sigma(q), \sigma(\mathcal{K}))$.*

Summarisation can be exploited to detect spurious answers in S : if a candidate answer is not in $\text{cert}(\sigma(q), \sigma(\mathcal{K}))$, then it is not in $\text{cert}(q, \mathcal{K})$. Since summarisation can significantly reduce the size of a knowledge base, we can efficiently detect non-answers even if checking each of them over the summary requires calling the OWL reasoner.

Corollary 1. *Let \mathcal{K} be a knowledge base, let q be a query, let S be a set of tuples, and let $\mathcal{K}' = \mathcal{K}_{[q, S]} \cup \mathcal{K}_\perp$. Furthermore, let σ be the summary function over \mathcal{K}' . Then, $\sigma(\mathcal{K}') \not\models \sigma(q(\mathbf{a}))$ implies $\mathcal{K} \not\models q(\mathbf{a})$ for each $\mathbf{a} \in S$.*

Data	DL	Axioms	Facts
LUBM(n)	<i>SHI</i>	93	$10^5 n$
UOBM ⁻ (n)	<i>SHLN</i>	314	$2 \times 10^5 n$
FLY	<i>SRI</i>	144,407	6,308
DBPedia ⁺	<i>SHON</i>	1,757	12,119,662
NPD	<i>SHIF</i>	819	3,817,079

Table 3: Statistics for test data

Strategy	Solved	Univ	t_{avg}
RL Bounds	14	1000	10.7
+ EL Lower Bound	22	1000	7.0
+ Sum, Dep	24	100	41.9

Table 4: Result for LUBM

Exploiting dependencies between answers In this last step, we try to further reduce the calls to the fully-fledged reasoner by exploiting dependencies between the candidate answers in S . Consider tuples \mathbf{a} and \mathbf{b} in S and the dataset \mathcal{D} in the relevant fragment $\mathcal{K}_{[q, S]} \cup \mathcal{K}_\perp$; furthermore, suppose we can find an endomorphism h of \mathcal{D} in which $h(\mathbf{a}) = \mathbf{b}$. If we can determine (by calling the fully-fledged reasoner) that \mathbf{b} is a spurious answer, then so must be \mathbf{a} ; as a result, we no longer need to call the reasoner to check \mathbf{a} . We exploit this idea to compute a dependency graph having candidate answer tuples as nodes and an edge (\mathbf{a}, \mathbf{b}) whenever an endomorphism in \mathcal{D} exists mapping \mathbf{a} to \mathbf{b} . Computing endomorphisms is, however, a computationally hard problem, so we have implemented a sound (but incomplete) greedy algorithm that allows us to approximate the dependency graph.

7 Evaluation

We have implemented a prototype based on RDFox and Hermit (v. 1.3.8). In our experiments we used the LUBM and UOBM benchmarks, as well as the Fly Anatomy ontology, DBPedia and NPD FactPages (their features are summarised in Table 3). Data, ontologies, and queries are available online.¹ We compared our system with Pellet (v. 2.3.1) and TrOWL (Thomas, Pan, and Ren 2010) on all datasets. While Pellet is sound and complete for OWL 2, TrOWL relies on approximate reasoning and does not provide correctness guarantees. Tests were performed on a 16 core 3.30GHz Intel Xeon E5-2643 with 125GB of RAM, and running Linux 2.6.32. For each test, we measured materialisation times for upper and lower bound, the time to answer each query, and the number of queries that can be answered using different techniques. All times are in seconds.

LUBM Materialisation is fast on LUBM (Guo, Pan, and Heflin 2005): it takes 274s (286s) to materialise the basic lower (upper) bound entailments for LUBM(1000). These bounds match for all 14 standard LUBM queries, and we have used 10 additional queries for which this is not the case; we tested our system on all 24 queries (see Table 4 for a summary of the results). The refined lower bound was

¹<http://www.cs.ox.ac.uk/isg/people/yujiao.zhou/#resources>

Strategy	Solved	Univ	$t_{avg.}$
RL Bounds	9	500	5.8
+ EL Lower Bound	12	500	9.7
+ Summarisation	14	60	40.2
+ Dependencies	15	1	2.7

Table 5: Result for $UOBM^-$

materialised in 307s, and it matches the upper bound for 8 of the 10 additional queries; thus, our system could answer 22 of the 24 queries over LUBM(1000) efficiently in 11s on average.² For the remaining 2 queries, we could scale to LUBM(100) in reasonable time. On LUBM(100) the gaps contain 29 and 14 tuples respectively, none of which were eliminated by summarisation; however, exploiting dependencies between gap tuples reduced the calls to Hermit to only 3 and 1 respectively, with the majority of time taken in extraction (avg. 86s) and Hermit calls (avg. 384.7s). On LUBM(1000), Pellet ran out of memory. For LUBM(100), it took on average 8.2s to answer the standard queries with an initialisation overhead of 388s. TrOWL timed out after 1h on LUBM(100).

UOBM UOBM is an extension of LUBM (Ma et al. 2006). Query answering over UOBM requires equality reasoning (e.g., to deal with cardinality constraints), which is not natively supported by RDFox,³ so we have used a slightly weakened ontology $UOBM^-$ for which equality is not required. Materialisation is still fast on $UOBM^-$ (500): it takes 305s (356s) to materialise the basic lower (upper) bound entailments. We have tested the 15 standard queries (see Table 5). The basic lower and upper bounds match for 9 queries, and this increases to 12 when using our refined lower bound (the latter took 312s to materialise); our system is efficient for these queries, with an avg. query answering time of less than 10s over $UOBM^-$ (500). For 2 of the remaining queries, summarisation prunes all candidate answers. Avg. times for these queries were under 40s for $UOBM^-$ (60). For the one remaining query, summarisation rules out 6245 among 6509 answers in the gap, and the dependency analysis groups all the remaining individuals. Hermit, however, takes 20s to check the representative answer for $UOBM^-$ (1), and 2000s for $UOBM^-$ (10). Pellet times out even on $UOBM^-$ (1). TrOWL took 237s on average to answer 14 out of the 15 queries over $UOBM^-$ (60).⁴ A comparison with our system reveals that TrOWL answers may be neither sound nor complete for most test queries.

Fly Anatomy Ontology Fly Anatomy is a complex ontology, rich in existential axioms, and including a dataset with over 6,000 facts. We tested it with five queries provided by the developers the ontology. It took 1.7s (5.9s) to materialise lower (upper) bound entailments. The basic lower bounds for all queries are empty, whereas the refined lower bounds (which take 5.7s to materialise) match with the upper bound

in all cases; as a result, we can answer the queries in 0.1s on average. Pellet fails to answer queries given a 1h timeout, and TrOWL returns only empty answers.

DBPedia⁺ In contrast to Fly, the DBPedia dataset is relatively large, but the ontology is simple. To provide a more challenging test, we have used the LogMap ontology matching system (Jiménez-Ruiz et al. 2012) to extend DBPedia with the tourism ontology which contains both disjunctive and existential axioms. Since the tested systems report errors on datatypes, we have removed all axioms and facts involving datatypes. It takes 37.2s (40.7s) to materialise the basic lower (upper) bound entailments. The upper bound was unsatisfiable and it took 55.2s to check satisfiability of the \mathcal{K}_\perp fragment. We queried for instances of all 441 atomic concepts. Bounds matched in 439 cases (using the refined lower bound), and these queries were answered in 0.3s on average. Summarisation filtered out all gap tuples for the other 2 queries; the answer time for these was 58s. Pellet takes 280.9s to initialise and answers each query in an average time of 16.2s. TrOWL times out after 1h.

NPD FactPages The NPD FactPages ontology describes petroleum activities on the Norwegian continental shelf. The ontology is not Horn, and it includes existential axioms. As in the case of DBPedia, we removed axioms involving datatypes. Its dataset has about 4 million triples; it takes 8s (10s) to materialise the lower (upper) bound entailments. The upper bound is unsatisfiable, and it took 60s to check satisfiability of \mathcal{K}_\perp . We queried for the instances of the 329 atomic concepts, and could answer all queries using a combination of lower and upper bounds and summarisation in 5s on average. Queries with matching bounds (294 out of 329) could be answered within 0.035s. Pellet took 127s to initialise, and average query answering time was 3s. TrOWL took 1.3s to answer queries on average; answers were complete for 320 out of the 329 queries.

8 Discussion

We have proposed an enhanced hybrid approach for query answering over arbitrary OWL 2 ontologies. The approach integrates scalable and complete reasoners to provide pay-as-you-go performance: 772 of the 814 test queries could be answered using highly scalable lower and upper bound computations, 39 of the remaining 42 queries yielded to scalable extraction and summarisation techniques, and even for the remaining 3 queries our fragment extraction and dependency techniques greatly improved scalability. Our approach is complementary to other optimisation efforts, and could immediately benefit from alternative techniques for efficiently computing lower bounds and/or a more efficient OWL reasoner. Furthermore, our technical results are very general, and hold for any language L captured by generalised rules and over which we want to answer queries; our only assumption is the availability of a fully-fledged query engine for L and one for datalog, both used as a “black box”.

There are still many possibilities for future work. For the immediate future, our main focus will be improving the fragment extraction and checking techniques so as to improve scalability for the hardest queries.

²Avg. query answering times measured after materialisation.

³RDFox axiomatises equality as a congruence relation.

⁴An exception is reported for the remaining query.

Acknowledgements

This work was supported by the Royal Society, the EPSRC projects Score!, Exoda, and MaSI³, and the FP7 project OPTIQUE.

References

- Bachmair, L., and Ganzinger, H. 2001. Resolution theorem proving. In *Handbook of Automated Reasoning*. Elsevier. 19–99.
- Bishop, B.; Kiryakov, A.; Ognyanoff, D.; Peikov, I.; Tashev, Z.; and Velkov, R. 2011. OWLim: A family of scalable semantic repositories. *J. of Web Semantics* 2(1):33–42.
- Cuenca Grau, B.; Motik, B.; Stoilos, G.; and Horrocks, I. 2012. Completeness guarantees for incomplete ontology reasoners: Theory and practice. *J. of Artificial Intelligence Research* 43(1):419–476.
- Dolby, J.; Fokoue, A.; Kalyanpur, A.; Kershenbaum, A.; Schonberg, E.; Srinivas, K.; and Ma, L. 2007. Scalable semantic retrieval through summarization and refinement. In *Proc. of the 22nd International Joint Conference on Artificial Intelligence*.
- Dolby, J.; Fokoue, A.; Kalyanpur, A.; Schonberg, E.; and Srinivas, K. 2009. Scalable highly expressive reasoner (SHER). *J. of Web Semantics* 7(4):357 – 361.
- Gallier, J. H. 1985. *Logic for computer science: foundations of automatic theorem proving*. Harper & Row Publishers, Inc.
- Guo, Y.; Pan, Z.; and Heflin, J. 2005. LUBM: A benchmark for OWL knowledge base systems. *J. of Web Semantics* 3(2-3):158–182.
- Hustadt, U.; Motik, B.; and Sattler, U. 2007. Reasoning in description logics by a reduction to disjunctive Datalog. *J. of Automated Reasoning* 39(3):351–384.
- Jiménez-Ruiz, E.; Cuenca Grau, B.; Zhou, Y.; and Horrocks, I. 2012. Large-scale interactive ontology matching: Algorithms and implementation. In *Proc. of 20th European Conference on Artificial Intelligence*.
- Kollia, I., and Glimm, B. 2013. Optimizing SPARQL query answering over ontologies. *J. of Artificial Intelligence Research* 48:253–303.
- Ma, L.; Yang, Y.; Qiu, Z.; Xie, G.; Pan, Y.; and Liu, S. 2006. Towards a complete OWL ontology benchmark. *The Semantic Web: Research and Applications* 125–139.
- Manola, F.; Miller, E.; and McBride, B. 2004. RDF Primer. *W3C recommendation* 10:1–107.
- Möller, R.; Neuenstadt, C.; Özgür L. Özçep; and Wandelt, S. 2013. Advances in accessing big data with expressive ontologies. In *Proc. of the Joint German/Austrian Conference on Artificial Intelligence*, 118–129.
- Motik, B.; Cuenca Grau, B.; Horrocks, I.; Wu, Z.; Fokoue, A.; and Lutz, C. 2009a. OWL 2 Web Ontology Language Profiles. *W3C Recommendation*.
- Motik, B.; Patel-Schneider, P. F.; Parsia, B.; Bock, C.; Fokoue, A.; Haase, P.; Hoekstra, R.; Horrocks, I.; Ruttenberg, A.; Sattler, U.; et al. 2009b. OWL 2 web ontology language: Structural specification and functional-style syntax. *W3C recommendation* 27:17.
- Motik, B.; Nenov, Y.; Piro, R.; Horrocks, I.; and Olteanu, D. 2014. Parallel materialisation of datalog programs in centralised, main-memory RDF systems. In *Proc. of the 28th AAAI Conference on Artificial Intelligence*.
- Motik, B.; Shearer, R.; and Horrocks, I. 2009. Hypertableau reasoning for description logics. *J. of Artificial Intelligence Research* 36(1):165–228.
- Pan, J.; Thomas, E.; and Zhao, Y. 2009. Completeness guaranteed approximations for OWL-DL query answering. *J. of Description Logic* 477.
- Sirin, E.; Parsia, B.; Cuenca Grau, B.; Kalyanpur, A.; and Katz, Y. 2007. Pellet: A practical OWL-DL reasoner. *J. of Web Semantics* 5(2):51–53.
2013. SPARQL 1.1 Overview. W3C Recommendation. Available at <http://www.w3.org/TR/sparql11-overview/>.
- Stefanoni, G.; Motik, B.; and Horrocks, I. 2013. Introducing nominals to the combined query answering approaches for EL. In *Proc. of the 27th AAAI Conference on Artificial Intelligence*.
- Thomas, E.; Pan, J. Z.; and Ren, Y. 2010. TrOWL: Tractable owl 2 reasoning infrastructure. In *Proc. of the 7th Extended Semantic Web Conference*.
- Tserendorj, T.; Rudolph, S.; Krötzsch, M.; and Hitzler, P. 2008. Approximate OWL-reasoning with Screech. In *Proceedings of the 2nd International Conference on Web Reasoning and Rule Systems*.
- Urbani, J.; Van Harmelen, F.; Schlobach, S.; and Bal, H. 2011. QueryPIE: Backward reasoning for OWL Horst over very large knowledge bases. In *Proc. of the 10th International Semantic Web Conference*.
- Urbani, J.; Kotoulas, S.; Maassen, J.; van Harmelen, F.; and Bal, H. 2012. WebPIE: A web-scale parallel inference engine using MapReduce. *J. of Web Semantics* 10(1).
- Wandelt, S.; Möller, R.; and Wessel, M. 2010. Towards scalable instance retrieval over ontologies. *Int. J. Software and Informatics* 4(3):201–218.
- Wu, Z.; Eadon, G.; Das, S.; Chong, E. I.; Kolovski, V.; Anamalai, M.; and Srinivasan, J. 2008. Implementing an inference engine for RDFS/OWL constructs and user-defined rules in Oracle. In *Proc. of the 24th International Conference on Data Engineering*.
- Zhou, Y.; Cuenca Grau, B.; Horrocks, I.; Wu, Z.; and Banerjee, J. 2013a. Making the most of your triple store: Query answering in OWL 2 using an RL reasoner. In *Proc. of the 22nd International World Wide Web Conference*.
- Zhou, Y.; Nenov, Y.; Cuenca Grau, B.; and Horrocks, I. 2013b. Complete query answering over horn ontologies using a triple store. In *Proc. of the 12th International Semantic Web Conference*.

A Correctness of Fragment Extraction

We will now show that the relevant subsets extraction is in fact sound and complete for any OWL 2 ontologies by proving that Theorem 1 is correct. We first look at several definitions that we will use later on.

Definition 4. With \mathcal{B} we denote a resolution calculus, consisting of the following rules:

$$\frac{C \vee A \quad D \vee \neg B}{C\sigma \vee D\sigma} \quad (\text{Resolution})$$

where $\sigma = \text{MGU}(A, B)$, the clauses $C \vee A$ and $D \vee \neg B$ are called the premises, and $C\sigma \vee D\sigma$ is called the conclusion of the resolution step.

$$\frac{C \vee A \vee B}{C\sigma \vee A\sigma} \quad (\text{Factoring})$$

where $\sigma = \text{MGU}(A, B)$ and the clause $C\sigma \vee A\sigma$ is called the conclusion of the factoring step.

Definition 5. Let Γ be a set of first-order clauses. A resolution derivation of a goal clause G in Γ is a sequence of clauses H_1, \dots, H_l that satisfies the following conditions.

1. $H_l = G$.
2. For each $1 \leq i \leq l$, H_i is the conclusion of an inference by \mathcal{B} from premises in $\Gamma \cup \{H_1, \dots, H_{i-1}\}$. If a premise is $H_p \in \{H_1, \dots, H_{i-1}\}$, H_i is called a successor of H_p .
3. Every H_i has at least one successor in the derivation for each $i < l$.

Particularly, a refutation is a derivation with the empty clause in the end. The support of a derivation $\text{sup}(G_1, \dots, G_l)$ is defined as the set of clauses in Γ that have been used to derive a clause G_i , for some $1 \leq i \leq l$.

Definition 6. Let Γ be a set of first-order clauses, an SLD proof of a goal G in Γ is a sequence of clauses

$$G_0 = \neg G \xrightarrow{C_1} G_1 \rightsquigarrow \dots \rightsquigarrow G_{l-1} \xrightarrow{C_l} G_l$$

where G_i is the empty clause and $C_i \in \Gamma$ and each G_{i+1} is the conclusion of a resolution step from G_i and C_{i+1} .

Obviously, there is a one to one correspondence between the SLD proofs and the SLD proofs defined in the preliminaries, where C_i is from the clauses that corresponds to the datalog rules and each goal clause here is a negation of the goal there. We will use the new definition in the appendix.

From Lemma 1 it follows that we only need to show theorem in case of atomic queries. So we restrict ourselves to atomic queries from now on.

Theorem 1. Let \mathcal{K} be a knowledge base, $q(\mathbf{x})$ a conjunctive query, and S a set of tuples. Then, \mathcal{K} is satisfiable iff \mathcal{K}_\perp is satisfiable; furthermore, if \mathcal{K} is satisfiable, then,

$$\mathcal{K} \models q(\mathbf{a}) \text{ iff } \mathcal{K}_{[q,S]} \cup \mathcal{K}_\perp \models q(\mathbf{a}) \text{ for each } \mathbf{a} \in S.$$

Proof. The “if” direction for the above two claims is trivial due to monotonicity of first-order logic. Hence, we only need to show the “only if” direction. We then prove the two claims separately.

- If \mathcal{K} is unsatisfiable, then $\mathcal{K}^\Sigma = \Sigma(\mathcal{K})$ is also unsatisfiable since Skolemisation is satisfiability-preserving. So there is a refutation in \mathcal{K}^Σ . Let G_1, \dots, G_l be a refutation in \mathcal{K}^Σ . By applying Lemma 2 and Lemma 3, for each $C \in \text{sup}(G_1, \dots, G_l)$, we have that every $C' \in \Psi \circ \Xi(C)$ is in the support of an SLD proof of \perp in $\Psi \circ \Xi(\mathcal{K}_\Sigma)$. So $\text{sup}(G_1, \dots, G_l) \subseteq \Sigma(\mathcal{K}_\perp)$ and hence $\Sigma(\mathcal{K}_\perp)$ is unsatisfiable, which is equisatisfiable with \mathcal{K}_\perp .
- If \mathcal{K} is satisfiable, then $\mathcal{K}^\Sigma = \Sigma(\mathcal{K})$ is satisfiable. For each $\mathbf{a} \in S$ s.t. $\mathcal{K} \models q(\mathbf{a})$, we also have $\mathcal{K}^\Sigma \models q(\mathbf{a})$. Then $q(\mathbf{a})$ is derivable by \mathcal{B} in \mathcal{K}^Σ . Assume G_1, \dots, G_l is a resolution derivation of $q(\mathbf{a})$ in \mathcal{K}_Σ . Similarly, by applying Lemma 2 and Lemma 3, for every $C \in \text{sup}(G_1, \dots, G_l)$, we have that $C' \in \Psi \circ \Xi(C)$ is in the support of an SLD proof of \perp or $q(\mathbf{a})$ in $\Psi \circ \Xi(C)$. So $\text{sup}(G_1, \dots, G_l) \subseteq \Sigma(\mathcal{K}_{[q,S]} \cup \mathcal{K}_\perp)$. Then $\Sigma(\mathcal{K}_{[q,S]} \cup \mathcal{K}_\perp) \models q(\mathbf{a})$, i.e. $\Sigma(\mathcal{K}_{[q,S]} \cup \mathcal{K}_\perp \cup \{\neg q(\mathbf{a})\})$ is unsatisfiable. So $\mathcal{K}_{[q,S]} \cup \mathcal{K}_\perp \cup \{\neg q(\mathbf{a})\}$ is unsatisfiable, and hence $\mathcal{K}_{[q,S]} \cup \mathcal{K}_\perp \models q(\mathbf{a})$.

□

In the following proof we will use Lemma 4, which was shown in (Gallier 1985).

Lemma 4. Let \mathcal{K}^H be a set of first-order Horn clauses. If there is a resolution derivation N_1, \dots, N_l of a fact G in \mathcal{K}^H , then there is an SLD proof of G in \mathcal{K}^H of the form $G_0 = \neg G \xrightarrow{C_1} G_1 \rightsquigarrow \dots \rightsquigarrow G_{l-1} \xrightarrow{C_l} G_l$ s.t. $\text{sup}(N_1, \dots, N_l) = \{C_1, \dots, C_l\}$.

Lemma 2. Let \mathcal{N} be a set of first-order clauses. Then:

- if $C \in \mathcal{N}$ participates in a refutation in \mathcal{N} , then every $C' \in \Xi(C)$ is part of an SLD proof of \perp in $\Xi(\mathcal{N})$;

- if $C \in \mathcal{N}$ participates in a resolution proof in \mathcal{N} of an atomic query $Q(\mathbf{a})$, then each $C' \in \Xi(C)$ participates in an SLD proof of \perp or $Q(\mathbf{a})$ in $\Xi(\mathcal{N})$.

Proof. According to Lemma 4, it suffices to prove the following two claims.

- (C1) if $C \in \mathcal{N}$ participates in a refutation in \mathcal{N} , then every $C' \in \Xi(C)$ is part of a resolution derivation of \perp in $\Xi(\mathcal{N})$;
- (C2) if $C \in \mathcal{N}$ participates in a resolution proof in \mathcal{N} of an atomic query $Q(\mathbf{a})$, then each $C' \in \Xi(C)$ participates in a resolution derivation of \perp or $Q(\mathbf{a})$ in $\Xi(\mathcal{N})$.

We next prove (C1). Let G_1, \dots, G_l be a refutation in \mathcal{N} . We construct H_1, \dots, H_l a sequence of sets of clauses as follows.

- If G_i is the conclusion of a resolution step from two clauses C, D with the substitution σ on the literal L in C and the literal $\neg L$ in D , then we define $S_i(X)$ as follows for $X = C$ or $X = D$.

$$S_i(X) \triangleq \begin{cases} \Xi(X) & X \in \mathcal{N} \\ H_j & X = G_j \text{ for some } 1 \leq j < i \end{cases}$$

Assume that $S_i(C) = \{C_1, \dots, C_s, C_{s+1}, \dots, C_{s'}\}$ and $S_i(D) = \{D_1, \dots, D_t, D_{t+1}, \dots, D_{t'}\}$ with $C_1\sigma, \dots, C_{s'}\sigma$ containing L and $D_1\sigma, \dots, D_{t'}\sigma$ containing $\neg L$.

$$H_i \triangleq \{\text{conclusion}(C_i\sigma, D_j\sigma) \mid 1 \leq i \leq s, 1 \leq j \leq t\} \cup \{C_{s+1}\sigma, \dots, C_{s'}\sigma, D_{t+1}\sigma, \dots, D_{t'}\sigma\}$$

- If G_i is the conclusion of a factoring step from the clause C with a substitution σ , $H_i \triangleq S_i(C)\sigma$.

We further define S^- to be the result of removing all \perp s in the set of clauses S .

Property (Aggregation Property). *For each $1 \leq i \leq l$ that:*

- (A1) $G_i = \bigvee H_i^-$;
- (A2) $X = \bigvee (S_i(X))^-$ for a premise X .

Proof. We proof the case by induction on the index i .

- Base case: $i = 1$.
In this case, all premises come from \mathcal{N} , and hence $S_1(C) = \Xi(C)$ for any premise C . So (A2) holds due to the definition of $\Xi(\cdot)$. We consider the following cases for (A1).
 - If G_1 is the conclusion of a resolution step from two clauses C, D from \mathcal{N} , $S_1(C) = \Xi(C)$ and $S_1(D) = \Xi(D)$. Let L' be a literal from C or D .
 - * If $L' = L$ or $L' = \neg L$, L' doesn't appear in G_1 because G_1 is the conclusion by a resolution step on L' . According to the definition of H_1 , it doesn't appear in H_1 or H_1^- either.
 - * Otherwise, L' remains in G_1 and H_1^- .
In addition, \perp doesn't appear in G_1 or H_1^- . So G_1 and H_1^- contains exactly the same set of literals.
 - If G_1 is the conclusion of a factoring step from a clause C in \mathcal{N} , $\bigvee H_1^- = \bigvee (S_1(C))^- \sigma = C\sigma = G_1$.
- Inductive case: for each $j < i$, $G_j = \bigvee H_j^-$ and $X = \bigvee (S_j(X))^-$.
Property (A2) holds because of the definition of $\Xi(\cdot)$ and the induction hypothesis on (A1). For (A1) we consider the following cases.
 - If G_i is the conclusion of a resolution step from two clauses C, D with the substitution σ on the literal L in C and the literal $\neg L$ in D . Let L' be a literal from $C\sigma$ or $D\sigma$. According to (A2), L' appears in some clauses in $S_i(C)\sigma$ or $S_i(D)\sigma$.
 - * $L' = L$ or $L' = \neg L$, L' doesn't appear in G_i because G_i is the conclusion by a resolution step on L' . According to the definition of H_i and H_i^- , it doesn't appear in H_i or H_i^- either.
 - * Otherwise, L' remains in G_i and H_i^- .
In addition, \perp doesn't appear in G_i or H_i^- . So G_i and H_i^- contains exactly the same set of literals, and thus $G_i = \bigvee H_i^-$.
 - If G_i is the conclusion of a factoring step from C with the substitution σ , we have that $G_i = C\sigma$ and $H_i = S_i(C)\sigma$. Then we have $G_i = \bigvee H_i^-$ according to (A2).

□

Based on the aggregation property (A1) and the fact that G_l is the empty clause, we have H_l contains only \perp .

Property (Successor Property). *For each $1 \leq i \leq l$ that every clause in H_i has at least one successor or appears in H_l .*

Proof. We proceed by induction on the index i .

- Base case: $i = l$. The property holds trivially.
- Inductive case: for each $j > i$, every clause in H_j has at least one successor or appears in H_l .
If $i < l$, G_i has a successor G_k whose existence is guaranteed by Definition 5 for some $k > i$.
 - G_k is the conclusion of a resolution step from G_i and C with the substitution σ on the literal L in G_i and the literal $\neg L$ in C . Let H_i be $\{D_1, \dots, D_t, D_{t+1}, \dots, D_{t'}\}$ with each clause in $\{D_1\sigma, \dots, D_t\sigma\}$ contains the literal L and each in $\{D_{t+1}, \dots, D_{t'}\}$ doesn't. According to the side aggregation property, we have $C\sigma = \bigvee (S_i(C))^- \sigma$. Since $C\sigma$ contains $\neg L$, so does $(S_i(C))^-$. Therefore, every clause in $\{D_1, \dots, D_t\}$ is able to be resolved with a clause in $S_i(C)$. So each of them is a premise of a clause in H_k and hence they have a successor. Since $D_{t+1}, \dots, D_{t'}$ are copied to H_k , then the property holds for them due to the induction hypothesis.
 - If G_k is the conclusion of a factoring step from G_i with the substitution σ , $H_k = H_i\sigma$. So every clause in H_i has a success in H_k . So the property also holds in this case.

□

Let $C \in \text{sup}(G_1, \dots, G_l)$, then it is a premise of G_k for some $1 \leq k \leq l$. One can show as in the proof of the successor property that every clause $C' \in \Xi(C)$ is either copied into H_k or has a conclusion in H_k . The successor property guarantees that we can extract a *weakened resolution derivation* N_1, \dots, N_l s.t. $N_i \in H_i$ for each $1 \leq i \leq l$ and C' is a premise of N_k . A weakened derivation means each N_i can be either a conclusion as in Definition 5 or be copied from N_j for $j < i$. Every weakened derivation can be trivially rewritten into a resolution derivation by eliminating repeating clauses. Since $N_l = \perp$, N_1, \dots, N_l is in fact a weakened proof of \perp . So C' is in the support of a proof of \perp in $\Xi(\mathcal{N})$. We have hence proved (C1).

Claim (C2) can be proved analogously. Let G_1, \dots, G_l be a resolution derivation of $Q(\mathbf{a})$ in \mathcal{N} . We can construct H_1, \dots, H_l in the same way as described above. But in this case, the aggregation property implies that H_l is either $Q(\mathbf{a})$ or \perp . Let $C \in \text{sup}(G_1, \dots, G_l)$, then it is a premise of G_k from \mathcal{N} for some $1 \leq k \leq l$. For each $C' \in \Xi(C)$, there is a weakened resolution derivation N_1, \dots, N_l s.t. $N_i \in H_i$ for each $1 \leq i \leq l$ and C' is a premise of N_k . So C' is in the support of a refutation or a derivation of $Q(\mathbf{a})$ in $\Xi(\mathcal{N})$. □

B Datalog-based Approach

Theorem 2. Let \mathcal{K} be a datalog knowledge base and let F be a set of facts in $\text{Mat}(\mathcal{K})$. Then, a fact $P(\mathbf{a})$ (resp. a rule r) in \mathcal{K} participates in an SLD proof of some fact in F iff $\bar{P}(\mathbf{a})$ (resp. $\mathbf{S}(\mathbf{c}_r)$) is in $\text{Mat}(\Delta(\mathcal{K}, F))$.

By slightly abuse of notations, for an atom $A = P(\mathbf{t})$, we define $\bar{A} = \bar{P}(\mathbf{t})$ in the following proof.

Proof of Theorem 2. We first prove the “only if” direction. Assume that a fact $P(\mathbf{a})$ (a rule r) participates in the following SLD proof of a fact $Q(\mathbf{c}) \in F$.

$$G_0 = \neg Q(\mathbf{c}) \xrightarrow{C_1, \theta_1} G_1 \rightsquigarrow \dots \xrightarrow{C_l, \theta_l} G_l$$

Let $\theta = \theta_1 \dots \theta_l$. W.o.l.g. we can assume that $\theta_i \theta = \theta$ for each $1 \leq i \leq l$. This can be obtained by renaming variables.

Property (Entailment Property). For each $1 \leq i \leq l$, $\mathcal{K} \models (A_1 \wedge \dots \wedge A_n)\theta$ where $G_i = \neg A_1 \vee \dots \vee \neg A_n$.

Proof. We prove the property by induction on the index i .

- Base case: If $i = l$, it holds trivially.
- Inductive case: Assume the following conditions.

$$G_{i-1} = \neg A_1 \vee \dots \vee \neg A_n \tag{3}$$

$$C_i = \neg B_1(\mathbf{x}_1) \vee \dots \vee \neg B_m(\mathbf{x}_m) \vee H(\mathbf{x}) \tag{4}$$

$$\theta_i = \text{MGU}(A_k, H(\mathbf{x})) \tag{5}$$

Then G_i is the following clause

$$(\neg A_1 \vee \dots \vee \neg A_{k-1} \vee \neg B_1(\mathbf{x}_1) \vee \dots \vee \neg B_m(\mathbf{x}_m) \vee \neg A_{k+1} \vee \dots \vee \neg A_n)\theta_i \tag{6}$$

The induction hypothesis of i ensures that $\mathcal{K} \models G_i\theta$, we then prove that the property holds for $i - 1$ as well.

Note that $\theta_i\theta = \theta$, then $\mathcal{K} \models (A_1 \wedge \dots \wedge A_{k-1} \wedge A_{k+1} \wedge \dots \wedge A_n)\theta$ and $\mathcal{K} \models (B_1(\mathbf{x}_1) \wedge \dots \wedge B_m(\mathbf{x}_m))\theta$. Since $\mathcal{K} \models C_i$, then $\mathcal{K} \models H(\mathbf{x})\theta$. Because $A_k\theta_i = H(\mathbf{x})\theta_i$, then $\mathcal{K} \models (A_1 \wedge \dots \wedge A_n)\theta$. So the property holds for $i - 1$ as well. □

Property (Derivable Property). For each $1 \leq i \leq l$, $\Delta(\mathcal{K}, F) \models (\bar{A}_1 \wedge \dots \wedge \bar{A}_n)\theta$ where $G_i = \neg A_1 \vee \dots \vee \neg A_n$.

Proof. We prove the property by induction on the index i .

- Base case: If $i = 0$, $\bar{Q}(\mathbf{c}) \in \Delta(\mathcal{K}, F)$ and then $\Delta(\mathcal{K}, F) \models \bar{Q}(\mathbf{c})$. So the property holds.
- Inductive case: Assume we have (3), (4) and (5) hold, then G_{i+1} is the clause shown in (6). The induction hypothesis of $i - 1$ ensures that $\Delta(\mathcal{K}, F) \models (\bar{A}_1 \wedge \dots \wedge \bar{A}_n)\theta$, we then prove the property also holds i . It is trivial that $(\neg A_1 \vee \dots \vee \neg A_{k-1} \vee \neg B_1(\mathbf{x}_1) \vee \dots \vee \neg B_m(\mathbf{y}_m) \vee \neg A_{k+1} \vee \dots \vee \neg A_n)\theta$ is an instance of G_i , because $\theta_i\theta = \theta$. According to the induction hypothesis, we have $\Delta(\mathcal{K}, F) \models (\bar{A}_1 \wedge \dots \wedge \bar{A}_n)\theta$. So it suffices to prove that $\Delta(\mathcal{K}, F) \models (\bar{B}_1(\mathbf{x}_1) \wedge \dots \wedge \bar{B}_m(\mathbf{x}_m))\theta$. The entailment property implies that $\mathcal{K} \models (B_1(\mathbf{x}_1) \wedge \dots \wedge B_m(\mathbf{x}_m))\theta$. So does $\Delta(\mathcal{K}, F)$. In addition, $\Delta(\mathcal{K}, F) \models \bar{A}_k\theta$ (i.e. $\Delta(\mathcal{K}, F) \models \bar{H}(\mathbf{x})\theta$). Together with the rules $\bar{H}(\mathbf{x}) \wedge B_1(\mathbf{x}_1) \wedge \dots \wedge B_m(\mathbf{x}_m) \rightarrow \bar{B}_j(\mathbf{x}_j)$ for each $1 \leq j \leq m$ in $\Delta(\mathcal{K}, F)$, we have $\Delta(\mathcal{K}, F) \models (\bar{B}_1(\mathbf{x}_1) \wedge \dots \wedge \bar{B}_m(\mathbf{x}_m))\theta$. □

If $P(\mathbf{a})$ participates in the above SLD proof, there is a $1 \leq i \leq n$, s.t. $C_i = P(\mathbf{a})$. Then G_{i-1} contains $\neg E$ with $E\theta_i = P(\mathbf{a})$. So the derivable property entails that $\Delta(\mathcal{K}, F) \models \bar{P}(\mathbf{a})$, and hence $\bar{P}(\mathbf{a}) \in Mat(\Delta(\mathcal{K}, F))$.

If a rule r participates in the above SLD proof. Assume that C_i is the corresponding clause of r of the form (4) and G_{i-1} is in the form of (3). Then G_i is the clause presented in (6). Because $\mathcal{K} \models (B_1(\mathbf{x}_1) \wedge \dots \wedge B_m(\mathbf{x}_m))\theta$, so does $\Delta(\mathcal{K}, F)$. In addition, the derivable property guarantees that $\Delta(\mathcal{K}, F) \models \bar{A}_k\theta$ (i.e. $\Delta(\mathcal{K}, F) \models \bar{H}(\mathbf{x})\theta$). Together with the rule $\bar{H}(\mathbf{x}) \wedge B_1(\mathbf{x}_1) \wedge \dots \wedge B_m(\mathbf{x}_m) \rightarrow \mathbf{S}(\mathbf{c}_r)$ in $\Delta(\mathcal{K}, F)$, we have $\Delta(\mathcal{K}, F) \models \mathbf{S}(\mathbf{c}_r)$, and hence $\mathbf{S}(\mathbf{c}_r) \in Mat(\Delta(\mathcal{K}, F))$.

The other direction can show that if a fact $\bar{P}(\mathbf{a})$ (a rule $\mathbf{S}(\mathbf{c}_r)$) is in $Mat(\Delta(\mathcal{K}, F))$, then there exists a hyper-resolution of a fact $Q(\mathbf{c})$ in F that involves $P(\mathbf{a})$ (or r). Every such proof corresponds to a resolution proof using the same rules and facts, and can hence be transformed using Lemma 4 into an SLD proof of $q(\mathbf{c})$.

Finally, it is clear that $Mat(\Delta(\mathcal{K}, F))$ can be computed in polynomial time because datalog evaluation is polynomial to the size of data, and the arity of predicates in the knowledge base \mathcal{K} is bounded. □