



# Efficient Ontology-Based Data Integration with Canonical IRIs

Guohui Xiao<sup>1</sup>✉, Dag Hovland<sup>2</sup>, Dimitris Bilidas<sup>3</sup>, Martin Rezk<sup>4</sup>,  
Martin Giese<sup>2</sup>, and Diego Calvanese<sup>1</sup>

<sup>1</sup> Faculty of Computer Science, Free-University of Bozen-Bolzano, Bolzano, Italy

{xiao,calvanese}@inf.unibz.it

<sup>2</sup> Department of Informatics, University of Oslo, Oslo, Norway

{hovland,martingi}@ifi.uio.no

<sup>3</sup> National and Kapodistrian University of Athens, Athens, Greece

d.bilidas@di.uoa.gr

<sup>4</sup> Rakuten, Tokyo, Japan

martin.rezk@rakuten.com

**Abstract.** In this paper, we study how to efficiently integrate multiple relational databases using an ontology-based approach. In ontology-based data integration (ODI) an ontology provides a coherent view of multiple databases, and SPARQL queries over the ontology are rewritten into (federated) SQL queries over the underlying databases. Specifically, we address the scenario where records with different identifiers in different databases can represent the same entity. The standard approach in this case is to use `sameAs` to model the equivalence between entities. However, the standard semantics of `sameAs` may cause an exponential blow up of query results, since all possible combinations of equivalent identifiers have to be included in the answers. The large number of answers is not only detrimental to the performance of query evaluation, but also makes the answers difficult to understand due to the redundancy they introduce. This motivates us to propose an alternative approach, which is based on assigning canonical IRIs to entities in order to avoid redundancy. Formally, we present our approach as a new SPARQL entailment regime and compare it with the `sameAs` approach. We provide a prototype implementation and evaluate it in two experiments: in a real-world data integration scenario in Statoil and in an experiment extending the Wisconsin benchmark. The experimental results show that the canonical IRI approach is significantly more scalable.

## 1 Introduction

Large organizations, both public and private, typically need to manage, in multiple information systems, large amounts of data stored across multiple heterogeneous data sources. To support decision making in such settings, there is the need to access in an integrated way the data managed by the different systems, which in general are stored in different databases. A key challenge for data integration

in such settings is that the different systems have often been designed separately serving different purposes, and thus lack a coherent view of the underlying data.

*Ontology-based data access* (OBDA) [24] is a successful paradigm that addresses this challenge by relying on semantic technologies to provide a uniform conceptual view over heterogeneous data. Specifically, an ontology describing the domain of interest is connected to a data source through a declarative specification, given in terms of *mappings* [8] that relate symbols (i.e., classes and properties) in the ontology to (SQL) views over the data. The ontology layer can be queried using SPARQL, and queries are automatically rewritten by an OBDA engine into SQL queries expressed over the underlying database. Thus, users no longer need an understanding of the structure and organization of the data in the source. *Ontology-based data integration* (OBDI) is an extension of OBDA in which the data is not stored in a single database, but in a multitude of databases that need to be queried in an integrated way, while still maintaining the same conceptual architecture based on mappings [3]. OBDI has been successfully deployed in several domains, such as oil and gas [18], maritime security [1], and cultural heritage [2].

An important aspect in OBDI is the fact that the same conceptual entity may be stored in different databases but represented by different identifiers. Consider e.g., the same person, represented in one database through name, surname, and date of birth, in a second database through the social security number, and in a third database through the taxcode. In order to combine the information coming from different systems so as to produce coherent results to queries, we need a way to relate different representations so that they can be recognized as representing the same entity. This is challenging both from the theoretical and from the practical points of view.

The standard approach in the Semantic Web to modeling the fact that different IRIs actually represent the same entity is to use `owl:sameAs`<sup>1</sup>, which defines an equivalence relation between entities. In our previous work [5], we have shown how an approach based on the use of `sameAs` can be adopted also for OBDI. A serious drawback of this approach is that the semantics of `sameAs` may cause an exponential blowup of query results, since all possible combinations of equivalent identifiers have to be included in the answer to a query. The large number of answers is detrimental to the performance of query evaluation, and also introduces redundancy making answers difficult to understand.

To address this problem, we propose here an alternative approach, based on assigning to each entity a *canonical IRI*. Such canonical IRI is the one of choice for relating occurrences of the entity that are identified differently in different data sources, and also for returning the entity to the user. The idea of exploiting a canonical representation of IRIs internally, for optimization purposes, has already been used in some reasoning engines. E.g., RDFox [22] uses an optimized approach for rule evaluation by forward chaining based on propagating a canonical representation for entities. Also Stardog adopts a similar strategy in query

---

<sup>1</sup> For space reasons, from now on we will use `sameAs` as an abbreviation for `owl:sameAs`.

answering, by only returning a canonical representation of equivalent entities, to avoid the combinatorial explosion in query results<sup>2</sup>.

However, lifting this idea to the OBDI scenario is non-trivial. On the one hand, the choice on which source should provide the canonical IRI for an entity depends on the actual (type of) entity, and on the data sources in which it appears. A more critical issue is due to the fact that in OBDI, entities with the corresponding IRIs are not materialized but kept *virtual*, since they are generated dynamically at query answering time from the data in the sources through the mappings. Hence it is not really feasible to explicitly substitute the entities by canonical ones. Instead our proposal is based on embedding such information in the mapping specification, while delegating to the designer of the OBDI system the responsibility of choosing which canonical IRIs to adopt.

For such a setting, we provide the following contributions:

- We formally define the *canonical IRI semantics* as a SPARQL entailment regime. We show that such semantics is equivalent to the **sameAs** semantics (taking into account equivalence of IRIs), thus guaranteeing soundness and completeness of the associated query answering algorithm (see Sect. 3).
- We show that query answering and the canonical IRI semantics can be reduced to a simple entailment regime by means of query rewriting. However, this query rewriting step is only of theoretical interest since the rewritten query is not efficiently executable (see Sect. 4).
- We propose an alternative practical *mapping rewriting-based* technique for query answering, inspired by algorithms for mapping saturation [2] (see Sect. 5).
- We provide a prototype implementation of our technique and evaluate it in two experiments. The first experiment is based on a real-world data integration scenario in the oil company Statoil [18]. The second experiment is based on the Wisconsin benchmark [9]. The experimental results show that the canonical IRI approach is significantly more scalable (see Sect. 6).

We start our technical development by introducing in Sect. 2 the necessary background.

## 2 Preliminaries

We present now the preliminary notions on RDF graphs, SPARQL, and ontology-based data access on which we build in the rest of the paper.

### 2.1 RDF and SPARQL

SPARQL [14] is the W3C standard language designed to query RDF graphs. We consider a vocabulary of three pairwise disjoint and countably infinite sets of symbols: **I** for *IRIs*, **L** for *RDF literals*, and **V** for *variables*. In line with

<sup>2</sup> [https://www.stardog.com/docs/#\\_same\\_as\\_reasoning](https://www.stardog.com/docs/#_same_as_reasoning).

previous work on ontology-based data access, we do not consider blank nodes. The elements of  $\mathbf{T} = \mathbf{I} \cup \mathbf{L}$  are called *RDF terms*, and those of  $\mathbf{T} \times \mathbf{I} \times \mathbf{T}$  are called (*RDF*) *triples*. An (*RDF*) *graph* is a set of triples. A *triple pattern* is an element of  $(\mathbf{T} \cup \mathbf{V}) \times (\mathbf{I} \cup \mathbf{V}) \times (\mathbf{T} \cup \mathbf{V})$ . A *basic graph pattern* (*BGP*) is a finite set of triple patterns. We consider the fragment of SPARQL queries defined by  $Q$  in the following EBNF grammar<sup>3</sup>:

$$\begin{aligned} P &::= B \mid Q \mid P \text{ FILTER } F \mid P \text{ UNION } P \mid (P, P) \mid P \text{ OPT } P \\ Q &::= \text{SELECT } \{ \mathbf{V} \text{ AS } \mathbf{V} \} \text{ WHERE } P \end{aligned}$$

where  $B$  is a BGP and  $F$  is a *filter expression* (we refer to [14] for details). An expression  $\{?y_1 \text{ AS } ?x_1, \dots, ?y_n \text{ AS } ?x_n\}$  is called a *projection* with *answer variables*  $\{?x_1, \dots, ?x_n\}$ . We abbreviate “ $?x \text{ AS } ?x$ ” with  $?x$ . For a SPARQL query  $Q = \text{SELECT } R \text{ WHERE } Q_1$ , we use  $\text{sig}(Q)$  to denote the answer variables of  $R$ .

The semantics of SPARQL queries is given in terms of *solution mappings*, which are *partial* maps  $s: \mathbf{V} \rightarrow \mathbf{T}$  with (possibly empty) domain  $\text{dom}(s)$ . Here, following [20, 23, 25], we use the set-based semantics for SPARQL (rather than the bag-based one, as in the W3C specification). More specifically, for a BGP  $B$ , the *answer*  $\llbracket B \rrbracket_G$  to  $B$  over a graph  $G$  is  $\llbracket B \rrbracket_G = \{s: \text{var}(B) \rightarrow \mathbf{T} \mid s(B) \subseteq G\}$ , where  $\text{var}(B)$  is the set of variables occurring in  $B$  and  $s(B)$  is the result of substituting each variable  $u$  in  $B$  by  $s(u)$ . Then, the *answer* to a SPARQL query  $Q$  over a graph  $G$  is the set  $\llbracket Q \rrbracket_G$  of solution mappings defined by induction using the SPARQL algebra operators (filter, join, union, optional, and projection) starting from BGPs; cf. [17]. This semantics is known as *simple entailment*.

## 2.2 SPARQL Entailment Regimes

SPARQL entailment regimes allow for querying RDF graphs with richer reasoning capabilities [12]. Specifically, an entailment regime  $E$  specifies how to obtain from an RDF graph  $G$  an entailed graph  $eg^E(G)$ . Then, the answer  $\llbracket B \rrbracket_G^E$  to a BGP  $B$  under the entailment regime  $E$  is defined as  $\llbracket B \rrbracket_{eg^E(G)}$ . Similarly, The answer  $\llbracket Q \rrbracket_G^E$  to a SPARQL query  $Q$  under the entailment regime  $E$  is defined as  $\llbracket Q \rrbracket_{eg^E(G)}$ . Note that entailment regimes only modify the evaluation of BGPs but not that of other SPARQL operators.

We present now the standard W3C semantics for SPARQL queries over OWL 2 ontologies under different entailment regimes. Under the *OWL 2 direct semantics entailment regime*, one can query an RDF graph  $G$  that consist of two parts: the *intensional* sub-graph (i.e., TBox or ontology)  $\mathcal{T}$  representing the background knowledge in terms of class and property axioms, and an *extensional* sub-graph (i.e., ABox)  $\mathcal{A}$  representing the data as class and property assertions. We write such a graph  $G$ , which represents a *knowledge base*, as  $(\mathcal{T}, \mathcal{A})$  to emphasize the partitioning when necessary. Moreover, for convenience, we use the triple notations  $(s, \text{rdf:type}, C)$  and  $(s, p, o)$  and the ABox assertion notations  $C(s)$  and  $p(s, o)$  interchangeably.

<sup>3</sup> Recall that in EBNF “ $\{A\}$ ” means any number of repetitions of  $A$ .

We work with ontologies expressed in the OWL 2 QL profile [21] of OWL 2. Such profile induces the *OWL 2 QL* (or simply, *QL*) *entailment regime* in which, for an OWL 2 QL knowledge base  $G$  we have  $eg^{QL}(G) = \{t \mid G \models_{DL} t\}$ , where  $\models_{DL}$  denotes the standard OWL 2 entailment, defined in terms of description logics semantics, cf. [27]. Under the QL entailment regime, a SPARQL query  $Q$  formulated over an ontology  $\mathcal{T}$  is *first-order rewritable*, i.e.,  $Q$  can be rewritten into a query  $Q_{\mathcal{T}}$  such that for every ABox  $\mathcal{A}$ , evaluating  $Q$  over  $(\mathcal{T}, \mathcal{A})$  (under the QL entailment regime) is equivalent to evaluating  $Q_{\mathcal{T}}$  over  $\mathcal{A}$  (under the simple entailment regime) [4, 20].

### 2.3 SPARQL Entailment Regimes for sameAs

In addition to the input graph  $G = (\mathcal{T}, \mathcal{A})$ , we consider now a set  $\mathcal{A}^s$  of **sameAs** triples of the form **sameAs**( $a, b$ ), specifying the equivalence between individuals  $a$  and  $b$ .

We now define an entailment regime that interprets **sameAs** as the equivalence closure (i.e., the reflexive, transitive, and symmetric closure) of  $\mathcal{A}^s$ . For a set  $\mathcal{A}^s$  of **sameAs** assertions, we denote by  $(\mathcal{A}^s)^*$  the equivalence closure of  $\mathcal{A}^s$ . Given an RDF graph  $\mathcal{A}$  and a set  $\mathcal{A}^s$  of **sameAs** assertions, we define the *sameAs entailed graph*  $sag(\mathcal{A}, \mathcal{A}^s)$  of  $\mathcal{A}$  with respect to  $\mathcal{A}^s$  as

$$\begin{aligned} & \mathcal{A} \cup \{C(o') \mid C(o) \in G, \text{sameAs}(o, o') \in (\mathcal{A}^s)^*\} \\ & \cup \{R(o'_1, o'_2) \mid R(o_1, o_2) \in G, \text{sameAs}(o_1, o'_1) \in (\mathcal{A}^s)^*, \text{sameAs}(o_2, o'_2) \in (\mathcal{A}^s)^*\}. \end{aligned}$$

Given an entailment regime  $E$ , we can derive an extended entailment regime  $E + \text{sameAs}$ , which takes into account equivalences between entities inferred through **sameAs** statements, as follows:

$$eg^{E + \text{sameAs}}(\mathcal{T}, \mathcal{A} \cup \mathcal{A}^s) = eg^E(\mathcal{T}, sag(\mathcal{A}, \mathcal{A}^s))$$

In this paper, we are mostly interested in the SPARQL entailment regimes *QL* and *QL + sameAs*, but our results also apply to other entailment regimes  $E$  in which **sameAs** is treated as a standard object property, and new **sameAs** triples *cannot* be inferred.

### 2.4 Ontology-Based Data Access and Integration

In ontology-based data access (OBDA), we start from an *OBDA specification*  $\mathcal{P} = (\mathcal{T}, \mathcal{M}, \mathcal{S})$ , consisting of a set  $\mathcal{T}$  of OWL 2 axioms, a set  $\mathcal{M}$  of *mapping assertions*, and a relational database schema  $\mathcal{S}$ . An *OBDA instance*  $(\mathcal{P}, D)$  is given by an OBDA specification  $\mathcal{P}$  and a relational database instance  $D$  compliant with  $\mathcal{S}$ .

Mapping assertions allow one to define how an ABox should be populated with values retrieved by means of SQL queries. Each mapping assertion has one of the forms

$$C(f(x)) \leftarrow sql(y) \qquad P(f(x), f'(x')) \leftarrow sql(y)$$

such that  $x \subseteq y$  and  $x' \subseteq y$ ,  $sql$  is an arbitrary SQL query over  $\mathcal{S}$  projecting columns  $y$ , and  $f$  and  $f'$  are functions constructing RDF terms out of values retrieved from the database. In a concrete mapping language like R2RML [8], these functions are specified as templates for IRIs and literals. For example,  $\langle \text{http://statoil.com/wellbore/\{id\}} \rangle$  is an IRI template where “{id}” is a placeholder, and it generates, e.g., the IRI  $\langle \text{http://statoil.com/wellbore/25} \rangle$  when {id} is instantiated with “25”. In practice, it is desired that each IRI can be constructed by at most one IRI template, and we make such assumption here; formally, the union of all the IRI templates occurring in an OBDA specification is an injective function. By applying all mapping assertions in  $\mathcal{M}$  to  $D$ , one can derive a (virtual) RDF graph  $\mathcal{A}_{\mathcal{M},D}$  [24]. Then, SPARQL query answering over an OBDA instance  $(\mathcal{P}, D)$  is defined as query answering over  $(\mathcal{T}, \mathcal{A}_{\mathcal{M},D})$ .

We recall that, ontology-based data integration (OBDI) considers a set of data sources instead of a single one, but otherwise the formal treatment is identical to that of OBDA.

The inspiration for working with equality in OBDI came from problems we encountered in the context of the EU-funded project Optique [11], when querying data in the Norwegian oil company Statoil. This is also the background of the *real world* experiments in Sect. 6.1. In the following, we present our running example, which provides a simplified version of the experimental setting of Optique. The content of the tables below is real data from the Norwegian Petroleum Directorate FactPages<sup>4</sup>.

*Example 1.* Assume two databases `national` and `corporate` with one table each:

national.wellbore			corporate.drillingops		
name	opPurp	wlbFld	name	driStDt	reason
1/3-1	WILDCAT		NO_1/3-1	06-07-1968	WILDCAT
2/4-2	WILDCAT	EKOFISK	NO_2/4-2	18-09-1969	
1/3-10	APPRAISAL	OSELVAR	NO_1/2-1	20-03-1989	WILDCAT
1/2-1	WILDCAT	BLANE	NO_1/3-A-1.H	22-07-2011	PRODUCTION

In both tables, the column `name` is the only key (i.e., names are unique within each table). These columns have almost the same values in the two databases, except that in the `corporate` database they are prefixed with “NO\_” (country code for Norway). An example set of mapping assertions is:

---

```

:NationalWellbore/{name} :inField {wlbFld} ; :purpose {opPurp} .
← SELECT name, wlbFld, opPurp FROM national.wellbore

:CorporateWellbore/{name} :drillingStarted {driStDt} ; :purpose {reason} .
← SELECT name, driStDt, reason FROM corporate.drillingops

```

---

<sup>4</sup> <http://factpages.npd.no>.

Examples of triples in the ABox defined by this OBDA scenario are<sup>5</sup>

```
(:NationalWellbore/1/2-1,      :inField,      BLANE),
(:CorporateWellbore/NO_1/3-A-1_H, :drillingStarted, 22-07-2011),
(:NationalWellbore/1/3-10,      :purpose,      APPRAISAL),
(:CorporateWellbore/NO_1/2-1,    :purpose,      WILDCAT).
```

Compare the property `:purpose`, which is mapped to several databases, with the properties `:inField` and `:drillingStarted`, which are mapped to a single database each. Properties like `:purpose` may or may not have values for the same entity in different datasets. So the SPARQL query  $(?w, :purpose, ?p)$  returns both  $\{?w \mapsto :NationalWellbore/1/3-1, ?p \mapsto WILDCAT\}$  and  $\{?w \mapsto :CorporateWellbore/NO_1/3-1, ?p \mapsto WILDCAT\}$  as answers, while the query  $((?w, :inField, ?f), (?w, :drillingStarted, ?d))$  has no answers.

There are different ways of improving this situation. Perhaps the most obvious and direct one is to just change the mapping set such that the wellbores are mapped into equal IRIs. This is not always possible, as a data source may not be under the control of the mapping author; consider, e.g., Linked Open Data. Moreover, when integrating two OBDA systems where there already are IRIs defined, instead of modifying existing mapping sets, it is in general more convenient to add mappings stating the equivalences between IRIs, so as to virtually generate the `sameAs` relation, e.g.,  $(:NationalWellbore/1/3-1, sameAs, :CorporateWellbore/NO_1/3-1)$ . Such mappings also serve as a documentation for the way the data is linked. ■

### 3 Canonical IRI Semantics

We observe that the standard `sameAs` semantics is not scalable for query answering in general. Suppose that  $\mu = \{?x_1 \mapsto a_1, \dots, ?x_n \mapsto a_n\}$  is a solution mapping of a SPARQL query  $q$  over an ABox that includes assertions  $\{sameAs(a_i, b_i) \mid 1 \leq i \leq n\}$ . Then one can replace any  $a_i$  by  $b_i$  in  $\mu$ , resulting in a semantically equivalent solution mapping, which is still a valid answer to  $q$ . In this case, there are  $2^n$  such possibilities. This example shows that in general, `sameAs` could cause an exponential blowup of query results. Such large numbers of answers are not only detrimental to the performance of query evaluation, but also make the query answers difficult to understand due to the redundancy they introduce. Indeed, our motivation is to integrate multiple datasets, and what we care most are the entities, not the IRIs that represent them. This motivates us to propose an alternative approach, which is based on assigning *canonical IRIs* to entities.

We assume that, in addition to  $\mathcal{A}$ , we have a set  $\mathcal{A}^c$  of *canonical IRI assertions* using the property `canIriOf`. We make the following assumption on  $\mathcal{A}^c$ , which states that one IRI cannot have more than two canonical representations:

<sup>5</sup> The IRI encoding of special symbols like “/” is omitted for readability.

**Assumption 1.** *The property  $\text{canIriOf}$  is inverse functional in  $\mathcal{A}^c$ , i.e.,:*

$$\{\text{canIriOf}(c_1, i), \text{canIriOf}(c_2, i)\} \subseteq \mathcal{A}^c \quad \text{implies} \quad c_1 = c_2.$$

**Definition 1 (Canonical IRI and Canonical Graph).** *For an IRI  $o$ , the canonical IRI  $\text{can}_{\mathcal{A}^c}(o)$  of  $o$  is  $c_o$ , if  $\text{canIriOf}(c_o, o) \in \mathcal{A}^c$  for some  $c_o$ , and  $o$  otherwise. Given  $\mathcal{A}$  and  $\mathcal{A}^c$ , the canonical graph  $\text{cg}(\mathcal{A}, \mathcal{A}^c)$  is*

$$\{C(\text{can}_{\mathcal{A}^c}(o)) \mid C(o) \in \mathcal{A}\} \cup \{P(\text{can}_{\mathcal{A}^c}(o_1), \text{can}_{\mathcal{A}^c}(o_2)) \mid P(o_1, o_2) \in \mathcal{A}\}.$$

Given an entailment regime  $E$  without special treatments of canonical IRI assertions, we can now define the *canonical IRI entailment regime  $E+\text{can}$*  as:

$$\text{eg}^{E+\text{can}}(\mathcal{T}, \mathcal{A} \cup \mathcal{A}^c) = \text{eg}^E(\mathcal{T}, \text{cg}(\mathcal{A}, \mathcal{A}^c))$$

Now we study the relationship between the canonical IRI semantics and the  $\text{sameAs}$  semantics. To do so, we first define the correspondence between  $\mathcal{A}^c$  and  $\mathcal{A}^s$ .

**Definition 2.** *We say that  $\mathcal{A}^c$  is compliant with  $\mathcal{A}^s$  if the following two conditions hold:*

1. *if  $\text{canIriOf}(o_1, o_2) \in \mathcal{A}^c$ , then  $\text{sameAs}(o_1, o_2) \in (\mathcal{A}^s)^*$ ;*
2. *if  $\text{sameAs}(o_1, o_2) \in (\mathcal{A}^s)^*$ , then  $\text{can}_{\mathcal{A}^c}(o_1) = \text{can}_{\mathcal{A}^c}(o_2)$ .*

Next, we extend canonical graphs to SPARQL queries. More precisely, given a SPARQL query  $Q$ , we obtain  $\text{cg}(Q, \mathcal{A}^c)$  by replacing each IRI  $o$  in  $Q$  by  $\text{can}_{\mathcal{A}^c}(o)$ .

In the following, we assume that SPARQL queries contain neither  $\text{sameAs}$  nor  $\text{canIriOf}$ . The next proposition shows that query answering under the canonical IRI and the  $\text{sameAs}$  entailment regimes are equivalent in the following sense:

**Theorem 1.** *Let  $(\mathcal{T}, \mathcal{A})$  be a KB,  $\mathcal{A}^c$  compliant with  $\mathcal{A}^s$ ,  $Q$  a SPARQL query,  $\text{sig}(Q) = \{?x_1, \dots, ?x_n\}$ , and  $E$  an entailment regime that does not imply new  $\text{sameAs}$  or  $\text{canIriOf}$  assertions. It holds that*

- *if  $\{?x_1 \mapsto o_1, \dots, ?x_n \mapsto o_n\} \in \llbracket Q \rrbracket_{\mathcal{T}, \mathcal{A} \cup \mathcal{A}^s}^{E+\text{sameAs}}$ , then*

$$\{?x_1 \mapsto \text{can}_{\mathcal{A}^c}(o_1), \dots, ?x_n \mapsto \text{can}_{\mathcal{A}^c}(o_n)\} \in \llbracket \text{cg}(Q, \mathcal{A}^c) \rrbracket_{\mathcal{T}, \mathcal{A} \cup \mathcal{A}^c}^{E+\text{can}};$$

- *if  $\{?x_1 \mapsto o_1, \dots, ?x_n \mapsto o_n\} \in \llbracket Q \rrbracket_{\mathcal{T}, \mathcal{A} \cup \mathcal{A}^c}^{E+\text{can}}$ , and  $\text{sameAs}(o_j, o'_j) \in (\mathcal{A}^s)^*$ , for  $j \in J \subseteq \{1, \dots, n\}$ , then*

$$\{?x_j \mapsto o_j \mid j \notin J\} \cup \{?x_j \mapsto o'_j \mid j \in J\} \in \llbracket Q \rrbracket_{\mathcal{T}, \mathcal{A} \cup \mathcal{A}^s}^{E+\text{sameAs}}.$$

**Example 2.** Referring to Example 1, a possible set  $\mathcal{A}_c$  of canonical IRI assertions (satisfying Assumption 1) is:

$$\begin{aligned} & \{(:\text{Wellbore}/1, \text{canIriOf}, :\text{NationalWellbore}/1/3-1), \\ & (: \text{Wellbore}/1, \text{canIriOf}, :\text{CorporateWellbore}/\text{NO.1}/3-1), \\ & (: \text{Wellbore}/2, \text{canIriOf}, :\text{NationalWellbore}/2/4-2), \\ & (: \text{Wellbore}/2, \text{canIriOf}, :\text{CorporateWellbore}/\text{NO.2}/4-2) \} \end{aligned}$$

An example of a set  $\mathcal{A}_s$  such that  $\mathcal{A}_c$  is compliant with  $\mathcal{A}_s$  is:

$$\{(\text{:NationalWellbore}/1, \text{sameAs}, \text{:CorporateWellbore}/1), \\ (\text{:NationalWellbore}/2, \text{sameAs}, \text{:CorporateWellbore}/2)\}$$

Here we have chosen the canonical representative to always be of the (new) form  $\text{:Wellbore}/\{id\}$ . This is a choice made by the mapping author, and the canonical representative could also have been one of the existing IRIs. If we let  $D$  be the database and  $\mathcal{M}$  the mappings from Example 1, the query

$$Q = \text{SELECT } ?w, ?f, ?d \text{ WHERE } ((?w, \text{:inField}, ?f), (?w, \text{:drillingStarted}, ?d))$$

has a non-empty answer set over the canonical graph  $cg(\mathcal{A}_{\mathcal{M}, D}, \mathcal{A}_c)$ , including the answer  $\{?w \mapsto \text{:Wellbore}/1, ?f \mapsto \text{EKOFISK}, ?d \mapsto 06-07-1968\}$ . In Sect. 5 we will see mapping assertions that populate this ABox. ■

## 4 Handling Canonical IRI Semantics by Query Rewriting

In this section, we develop a query rewriting algorithm for canonical IRI semantics so that the canonical graph does not need to be materialized. This rewriting algorithm is mostly of theoretical interest but not meant to be implemented, since the structure of rewritten queries is too complex.

We recall that BGP queries under the simple entailment regime are monotonic (in the input graph  $G$ ): if  $s \in \llbracket Q \rrbracket_G$  then  $s \in \llbracket Q \rrbracket_{G \cup \Delta G}$ . However, observe that BGP queries under canonical IRI semantics are non-monotonic: for example let  $G = \{(a, \text{rdf:type}, C)\}$ ,  $\Delta G = \{(c_a, \text{canIriOf}, a)\}$ , and  $Q = (?x, \text{rdf:type}, C)$ ; then  $\{?x \mapsto a\} \in \llbracket Q \rrbracket_G^{\text{can}}$  but  $\{?x \mapsto a\} \notin \llbracket Q \rrbracket_{G \cup \Delta G}^{\text{can}}$ . This means that for some (and actually for every) BGPs  $Q$ , there is no union of BGPs  $Q'$  such that for all RDF graphs  $G$ ,  $\llbracket Q \rrbracket_G^{\text{can}} = \llbracket Q' \rrbracket_G$ . The reason is that  $\llbracket Q \rrbracket_G^{\text{can}}$  is always non-monotonic in  $G$ , while  $\llbracket Q' \rrbracket_G$  is always monotonic in  $G$ . Thus we cannot rewrite a BGP under the canonical IRI semantics to a union of BGPs under the standard SPARQL semantics and expect to obtain the same answers. In the following, we show that query answering under the  $E+can$  entailment regime can be reduced to query answering under the  $E$  entailment regime, by using non-monotonic SPARQL construct like `NOT EXISTS`.

To define such a translation, we first introduce a subquery, denoted  $qc[?xc, ?x]$ , that returns  $can_{\mathcal{A}^c}(?x)$  for  $?xc$ . This is achieved by encoding the two cases of Definition 1:

$$qc[?xc, ?x] = \text{SELECT } ?xc, ?x \text{ WHERE } ((?xc, \text{canIriOf}, ?x) \text{ UNION} \\ (\text{SELECT } ?x \text{ AS } ?xc, ?x \text{ WHERE } ((?x, ?p_1, ?o) \text{ UNION } (?s, ?p_2, ?x)) \\ \text{FILTER NOT EXISTS } (?y, \text{canIriOf}, ?x)))$$

**Lemma 1.** *Let  $\mathcal{A}$  and  $\mathcal{A}^c$  be as above and  $a$  an individual occurring in  $\mathcal{A}$ . Then*

$$\{?xc \mapsto c_a, ?x \mapsto a\} \in \llbracket qc[?xc, ?x] \rrbracket_{\mathcal{A} \cup \mathcal{A}^c} \text{ iff } c_a = can_{\mathcal{A}^c}(a).$$

**Definition 3.** The canonical-iri rewriting  $\psi(Q)$  of a SPARQL query  $Q$  is obtained from  $Q$  by replacing each triple pattern  $t$  with the sub-query  $\psi(t)$  obtained from  $t$  as follows:

- (1) for each variable  $?x$ , introduce a fresh variable  $?x_c$ ,
- (2) for each occurrence of a variable  $?x$ , introduce a fresh variable  $?x_o$ , change  $?x$  to  $?x_o$ , and join with  $qc[?x_c, ?x_o]$ , and
- (3) for each variable  $?x$ , add a projection  $?x_c$  AS  $?x$ .

*Example 3.* Consider Example 1 and the query  $Q = \text{SELECT } ?w, ?f, ?d \text{ WHERE } ((?w, :inField, ?f), (?w, :drillingStarted, ?d))$ . Then

$$\begin{aligned} \psi(Q) = & \text{SELECT } ?w, ?f, ?d \text{ WHERE } ( \\ & (\text{SELECT } ?wc1 \text{ AS } ?w, ?fc \text{ AS } ?f \text{ WHERE} \\ & \quad ((?wo1, :inField, ?fo), qc[?wc1, ?wo1], qc[?fc, ?fo])), \\ & (\text{SELECT } ?wc2 \text{ AS } ?w, ?dc \text{ AS } ?d \text{ WHERE} \\ & \quad ((?wo2, :drillingStarted, ?d), qc[?wc2, ?wo2], qc[?dc, ?do])) \end{aligned}$$

We note that Step 2 in Definition 3 is necessary to deal with each *occurrence* of a variable separately, as shown in the following example:

*Example 4.* Let  $\mathcal{A} = \{(a, P, b)\}$ ,  $\mathcal{A}^c = \{(c, \text{canIriOf}, a), (c, \text{canIriOf}, b)\}$ , and  $B = (?x, P, ?x)$  a BGP. It is easy to see that the following query returns the expected answers:

$$\psi(B) = \text{SELECT } ?xc \text{ AS } ?x \text{ WHERE } ((?x', P, ?x''), qc[?xc, ?x'], qc[?xc, ?x'']).$$

**Theorem 2.** For any  $E, \mathcal{T}, \mathcal{A}, \mathcal{A}^c$ , and  $Q$ ,  $\llbracket Q \rrbracket_{\mathcal{T}, \mathcal{A} \cup \mathcal{A}^c}^{E+can} = \llbracket \psi(Q) \rrbracket_{\mathcal{T}, \mathcal{A} \cup \mathcal{A}^c}^E$ .

We observe that although the size of the rewritten query is linear in the size of the original one, the approach is impractical. Intuitively, the number of variables in the NOT EXISTS clauses and in the inner UNION clauses is linear in the number of variables in the original query, and a (naive) query execution engine would need to enumerate over all individuals in the ABox for each such variable. Hence, the query execution time would essentially grow exponentially in the number of variables in the query.

## 5 Handling Canonical IRI Statements in OBDA

We now propose a practical approach to supporting canonical IRI semantics for ontology-based integration of cross-linked datasets. We assume that the mapping set  $\mathcal{M}$  may include a subset  $\mathcal{M}^c$  consisting of all mapping assertions populating `canIriOf` triples.

central.aliasTable		
id	nationalName	corporateName
1	NO_1/3-1	1/3-1
2	NO_2/4-2	2/4-2
3	NO_1/2-1	1/2-1
4	NO_1/3-A-1_H	
5		1/3-10

:Wellbore/{id} canIriOf :NationalWellbore/{nationalName} . ← SELECT id, nationalName FROM central.aliasTable
:Wellbore/{id} canIriOf :CorporateWellbore/{corporateName} . ← SELECT id, corporateName FROM central.aliasTable

**Fig. 1.** Table and mappings for Example 5

*Example 5.* We extend the OBDA scenario from Example 1 with mappings for `canIriOf` into tables in a database `central`, both shown in Fig. 1. The canonical IRI mapping assertions relate wellbores in databases `national` and `corporate` by employing an existing aliasing table in database `central`. By materializing these mappings and the database, we obtain the assertions in Example 2. *Master tables* like `central` are available in the Statoil use case, and are typical in corporate scenarios.

The canonical IRI mappings do not depend on the existence of such master tables, as arbitrary SQL queries in the source part of mappings can be used. However, regarding both maintainability and performance, our experience is that materializing the resolution of equality into a master table is a better choice. ■

In an OBDA setting, we propose a practical method based on compiling the consequences of canonical IRI semantics into mappings. This method is inspired by the mapping saturation algorithm used for classical OBDA [20, 26].

In order to make sure that the RDF graph constructed from the mappings satisfies Assumption 1, we state a stronger assumption on the mappings:

**Assumption 2.** *For each IRI template  $iri$ , there is at most one mapping assertion of the form:  $iri_c(\mathbf{a}) \text{ canIriOf } iri(\mathbf{b}) \leftarrow sql(\mathbf{a}, \mathbf{b})$ .*

**Lemma 2.** *If  $\mathcal{M}^c$  is a set of mapping assertions satisfying Assumption 2, then  $\mathcal{A}_{\mathcal{M}^c, D}$  satisfies Assumption 1 for every database instance  $D$ .*

Now we are ready to present the mapping rewriting algorithm. Intuitively, it replaces all the individuals and IRI-templates in the mapping by their canonical representation.

**Definition 4.** *Let  $\mathcal{M} = \mathcal{M}' \cup \mathcal{M}^c$  be a set of mapping assertions. The canonical-iri rewriting  $cm(\mathcal{M}', \mathcal{M}^c)$  of  $\mathcal{M}$  is obtained by processing each  $m \in \mathcal{M}'$  as follows:*

- for each IRI template  $iri(\mathbf{a})$  occurring in  $m$ , if  $\mathcal{M}^c$  contains a mapping

$$iri_c(\mathbf{b}_0) \text{ canIriOf } iri(\mathbf{b}_1) \leftarrow sql(\mathbf{b}_0, \mathbf{b}_1)$$

*then we replace  $iri(\mathbf{a})$  in the target of  $m$  by  $iri_c(\mathbf{b}_0)$  and join the source query with  $sql(\mathbf{b}_0, \mathbf{b}_1)$ ,  $\mathbf{a} = \mathbf{b}_1$ .*

- for each occurrence of an IRI  $o$ , if  $o = iri(\mathbf{a})$  for some IRI template  $iri$ , then we process it as in the IRI template case.

*Example 6.* Let  $\mathcal{M}'$  be the mappings in Example 1 and  $\mathcal{M}_c$  the mappings in Example 5. Then  $cm(\mathcal{M}', \mathcal{M}_c)$  obtained according to Definition 4 is as follows:

---

```

:Wellbore/{id} :inField {wlbFld} ; :purpose {opPurp} .
← SELECT wlbFld, opPurp, id FROM national.wellbore, central.aliasTable
   WHERE name = nationalName

:Wellbore/{id} :drillingStarted {driStDt} ; :purpose {reason} .
← SELECT driStDt, reason, id FROM corporate.drillingops, central.aliasTable
   WHERE name = corporateName

```

---

It is easy to see that the algorithm described in Definition 4 has the same effect as applying  $cg$  to the RDF graph:

**Lemma 3.** *Let  $\mathcal{M} = \mathcal{M}' \cup \mathcal{M}^c$  be a set of mapping assertions, and  $D$  a database instance. Then  $cg(\mathcal{A}_{\mathcal{M}', D}, \mathcal{A}_{\mathcal{M}^c, D}) = \mathcal{A}_{cm(\mathcal{M}', \mathcal{M}^c), D}$ .*

It follows that the mapping rewriting algorithm is sound and complete.

**Theorem 3.** *Let  $E$  be an entailment regime,  $\mathcal{P} = (\mathcal{T}, \mathcal{M}, \mathcal{S})$  an OBDA specification where  $\mathcal{M} = \mathcal{M}' \cup \mathcal{M}^c$ ,  $(\mathcal{P}, D)$  an OBDA instance, and  $Q$  a SPARQL query. Then  $\llbracket Q \rrbracket_{\mathcal{T}, \mathcal{A}_{\mathcal{M}, D}}^{E+can} = \llbracket Q \rrbracket_{\mathcal{T}, \mathcal{A}_{cm(\mathcal{M}', \mathcal{M}^c), D}}^E$ .*

This theorem shows that the support for canonical IRI semantics in OBDA can be implemented by simply adding a canonical-iri rewriting step into the existing workflow during the startup of an OBDA system.

The canonical IRI approach shown above shifts effort from the rewriting to the mapping construction. That is, compared with the approach using `sameAs` triples, more information must be encoded in the mappings. Specifically, one must encode transitivity, and choose a representative for each equivalence class. An important question is therefore if this is applicable in practical scenarios. Our experience from the Statoil use case is that one can exploit existing *master tables* correlating the names used for objects in different databases. The mappings for canonical IRIs are simple mappings into these tables.

## 6 Implementation and Experiments

The canonical IRI approach described above has been implemented in *Ontop* [2], which now supports both `sameAs` and canonical IRI semantics. Observe that *Ontop* does not perform SQL federation, therefore it usually relies on systems such as Teiid<sup>6</sup> or Exareme [7] to integrate multiple databases. These systems act as *mediators* and expose to *Ontop* a set of tables coming from the different databases. In particular, *Ontop* has been tested with Exareme intensively. Exareme was initially developed as an engine for complex dataflow processing on elastic clouds [19], and was subsequently enriched with data federation capabilities. As a result, the Exareme SQL federation engine is able to decompose complex relational queries, use common sub-expression identification techniques

<sup>6</sup> <http://teiid.jboss.org>.

in order to save processing costs (e.g., to process only once identical query fragments coming from different subqueries of a union query), and decide which query fragment should be sent to each external database. Exareme then processes intermediate results coming from different databases in parallel in order to produce the final results.

We now present two sets of experiments evaluating the performance of queries over crossed-linked datasets. One experiment was conducted on a production environment of databases in Statoil, based on queries from geoscientists at the company. Since all the data sources are on production servers with confidential data, the load changes, and the OBDA setting is too complex to isolate different features of this approach, we also created a controlled OBDA environment in our own server to study our technique.

## 6.1 Real-World Experiments

We integrated 7 data sources (relational databases) used in Statoil, extending an existing ontology and the set of mappings, and creating the tables necessary for `sameAs` and `canIriOf`. The queries and ontology are published in [15] and a description of the corporate use case is given in [18]. One of the data sources is the *slegge* database, which is also described in [15] together with the mappings toward this database.

The experiments in Statoil were run with a catalogue of 76 SPARQL queries constructed from information needs written down by geologists and geoscientists in the company. The domain of the queries is that of subsurface exploration, with a focus on wellbore information. The most complex query had 23 triple patterns, using object and data properties coming from 5 data sources. The queries were executed with a 20 min timeout, both with `sameAs` approach and with the canonical IRI approach.

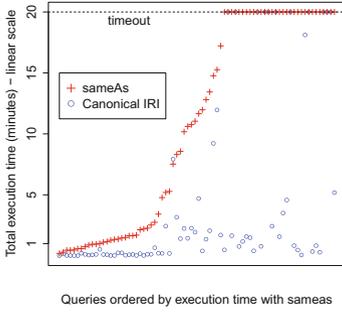
The *Ontop* rewriting engine and Exareme SQL federation engine all run on virtual machines deployed on the company intranet, as the data cannot be moved out. *Ontop* ran on a single machine, while the Exareme SQL federation ran on 8 other machines. The oracle databases are version 10g, and run on separate machines.<sup>7</sup>

We realize that this setup does not comply with the normal *clean* setup of a database experiment. However, the complexity (7 datasources) and realism (real questions and production databases) of the setup means the results have great value, although their precision is sub-optimal. Compare this with biology, where the *in vivo* experiments on live creatures, dealing with the full complexity of the organisms, may lead to results that cannot be seen in the *in vitro* experiments, and therefore are considered superior.

The total query execution times, for both the `sameAs` and canonical IRI approaches are shown in Fig. 2. The improvement from `sameAs` to canonical IRI

---

<sup>7</sup> Typical machine: HP ProLiant Server, 24 Intel Xeon CPUs (X560@2.67 GHz), 283 GB RAM.



	sameAs	Canonical IRI
Total queries	76	76
Timeouts	31	11
Min exec. time	12s	0.50s
Mean exec. time	11m	4.3m
Median exec. time	11m	0.77m

**Fig. 2.** Execution time and statistics for the queries in the federated setting at Statoil

is drastic. With the canonical IRI approach all queries, with three exceptions, are faster, there are fewer timeouts, and the majority of the queries execute within 3 min.

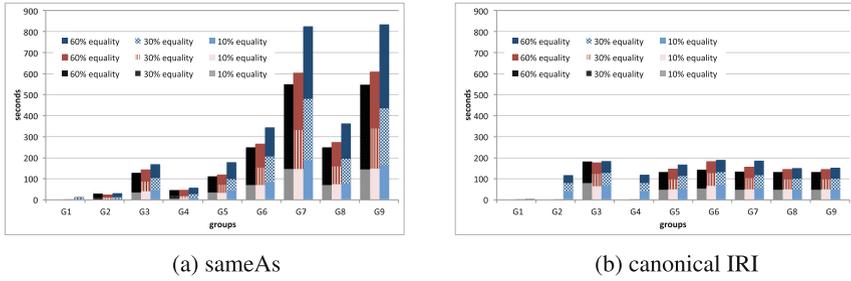
## 6.2 Controlled Experiments Using Wisconsin Benchmark

We also evaluated the canonical IRI approach in a more controlled setting. All files needed to reproduce this example are provided online<sup>8</sup>. This example is a simulated federated scenario with a single database consisting of 4 Wisconsin tables [9], representing different datasets, and 6 *linking tables*, see [5]. We reused the Wisconsin Benchmark tables, each of them containing 100M rows, and we created 3 new *canonical IRI tables* out of existing linking tables. We added the columns for provenance and canonical id, so that the mappings can generate *canIriOf* relations out of these 3 tables.

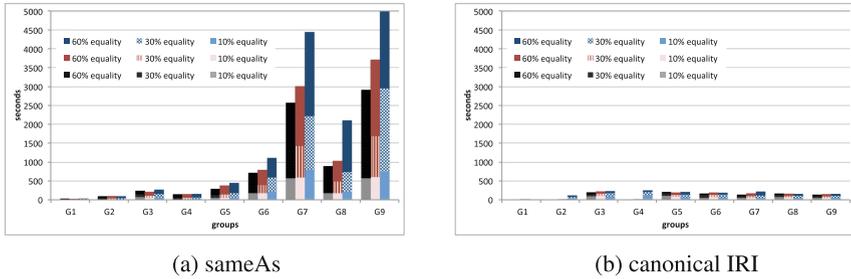
To evaluate the overhead of equality reasoning when answering SPARQL queries, we considered the following three parameters: (1) number of linked datasets (2–3); (2) selectivity of the query (returning 0.001%, 0.01%, 0.1% of the dataset); (3) number of equal objects between datasets (10%, 30%, 60%). In total, we ran 1332 queries, grouped in 9 groups: (G1) no properties, (G2) 1 data property and 0 object properties, (G3) 0 data properties and 1 object property, ..., (G9) 2 data properties and 2 object properties.

The results confirm an improvement, reducing the high cost of execution of the queries. In the previous setting based on *sameAs*, with the 6 linking tables and 2 linked-datasets scenario, with 120M equal objects (60%), in the worst case, most of the queries ran in around 3 min, while with the new canonical IRI approach, all queries can be executed in around 1 min. The query that performed worst in the previous setting (4 joins, 2 data properties, 2 object properties) returned 480 000 results and took around 6 min. In the new setting, for the same query, the number of results reduced to 60 000, avoiding duplicates. The resulting SQL query is simpler, and can be executed in only 50s. In the

<sup>8</sup> <https://github.com/ontop/ontop-examples/tree/master/eswc-2018-canonical-iri>.



**Fig. 3.** Execution times of most expensive queries with 2 datasets



**Fig. 4.** Execution times of most expensive queries with 3 datasets

3 linked-datasets scenario, the improvement is even more visible: the slowest executions took around 9 min in the previous setting, and less than 1.5 min in the new setting. The worst query in the previous setting took around 1.5 h, and returned 1 620 000 results. This query can now be executed in 53 s, and the number of results returned is significantly reduced to 60 000. The number of linked datasets is the variable that has most impact on query performance, but it is less influential in the canonical IRI setting. Another observation is that compared to the sameAs setting, the optimization for canonical IRI semantics is done only in the off-line phase and additional startup time is negligible. In fact, the startup time is only around 5 s.

In Figs. 3 and 4, we visualize the comparison of the execution times in these two settings. For each group, we show the execution times of the most expensive queries. The results confirm that under the canonical IRI semantics, we are able to run the queries significantly faster and the query answering times are more uniform.

## 7 Conclusions

This work is a natural continuation of [5], which was the first work to study the issue of equivalence of IRIs in OBDI and used the standard `sameAs` construct to model equivalence between entities. Our proposal based on canonical IRIs

improves on such an approach in terms of efficiency, and avoids the drawback of redundant answers.

Another important aspect in OBDI is how to efficiently evaluate the rewritten SQL queries over federated databases. Several federated database systems and prototypes have been presented in the literature. Early approaches include TSIMMIS [6], Garlic [13], and Tukwila [16]. Unpredictability regarding processing is the main issue that these systems have to cope with, using techniques such as adaptive query planning and query caching. BigDAWG [10] is a more recent approach that lays emphasis on the “one size does not fit all” principle, by trying to take advantage of specialized engines used as endpoints in order to efficiently process different types of data. How to natively deal with federated query evaluation in the setting of OBDI is subject of further work.

**Acknowledgement.** This research is supported by the project OBATS, funded by Free University of Bozen-Bolzano, by the Euregio IPN12 KAOS, funded by the “European Region Tyrol-South Tyrol-Trentino” (EGTC) under the first call for basic research projects, and by the Sirius Centre funded by the Norwegian Research Council.

## References

1. Brüggemann, S., Bereta, K., Xiao, G., Koubarakis, M.: Ontology-based data access for maritime security. In: Sack, H., Blomqvist, E., d’Aquin, M., Ghidini, C., Ponzetto, S.P., Lange, C. (eds.) ESWC 2016. LNCS, vol. 9678, pp. 741–757. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-34129-3\\_45](https://doi.org/10.1007/978-3-319-34129-3_45)
2. Calvanese, D., Cogrel, B., Komla-Ebri, S., Kontchakov, R., Lanti, D., Rezk, M., Rodriguez-Muro, M., Xiao, G.: Ontop: answering SPARQL queries over relational databases. *Semant. Web J.* **8**(3), 471–487 (2017)
3. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Poggi, A., Rosati, R., Ruzzi, M.: Data integration through *DL-Lite<sub>A</sub>* ontologies. In: Schewe, K.-D., Thalheim, B. (eds.) SDKB 2008. LNCS, vol. 4925, pp. 26–47. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-88594-8\\_2](https://doi.org/10.1007/978-3-540-88594-8_2)
4. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in description logics: the *DL-Lite* family. *J. Autom. Reason.* **39**(3), 385–429 (2007)
5. Calvanese, D., Giese, M., Hovland, D., Rezk, M.: Ontology-based integration of cross-linked datasets. In: Arenas, M., Corcho, O., Simperl, E., Strohmaier, M., d’Aquin, M., Srinivas, K., Groth, P., Dumontier, M., Heflin, J., Thirunarayan, K., Staab, S. (eds.) ISWC 2015. LNCS, vol. 9366, pp. 199–216. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-25007-6\\_12](https://doi.org/10.1007/978-3-319-25007-6_12)
6. Chawathe, S.S., Garcia-Molina, H., Hammer, J., Ireland, K., Papakonstantinou, Y., Ullman, J.D., Widom, J.: The TSIMMIS project: integration of heterogeneous information sources. In: Proceedings of the 10th Meeting of the International Proceedings Society of Japan (IPSJ 1994), pp. 7–18 (1994)
7. Chronis, Y., Fofoulas, Y., Nikolopoulos, V., Papadopoulos, A., Stamatogiannakis, L., Svingos, C., Ioannidis, Y.E.: A relational approach to complex dataflows. In: Proceedings of the EDBT/ICDT Workshops. CEUR, vol. 1558. ceur-ws.org (2016)
8. Das, S., Sundara, S., Cyganiak, R.: R2RML: RDB to RDF mapping language. W3C Recommendation, W3C, September 2012. <http://www.w3.org/TR/r2rml/>

9. DeWitt, D.J.: The Wisconsin benchmark: past, present, and future. In: Gray, J. (ed.) *The Benchmark Handbook*. Morgan Kaufmann, Burlington (1992)
10. Duggan, J., Elmore, A.J., Stonebraker, M., Balazinska, M., Howe, B., Kepner, J., Madden, S., Maier, D., Mattson, T., Zdonik, S.: The BigDAWG polystore system. *SIGMOD Rec.* **44**(2), 11–16 (2015)
11. Giese, M., Soylu, A., Vega-Gorgojo, G., Waaler, A., Haase, P., Jiménez-Ruiz, E., Lanti, D., Rezk, M., Xiao, G., Özçep, Ö.L., Rosati, R.: Optique: zooming in on big data. *IEEE Comput.* **48**(3), 60–67 (2015)
12. Glimm, B., Ogbuji, C.: SPARQL 1.1 entailment regimes. W3C Recommendation, W3C, March 2013. <http://www.w3.org/TR/sparql11-entailment/>
13. Haas, L.M., Kossmann, D., Wimmers, E.L., Yang, J.: Optimizing queries across diverse data sources. In: *Proceedings of the VLDB 1997*, pp. 276–285 (1997)
14. Harris, S., Seaborne, A.: SPARQL 1.1 query language. W3C Recommendation, W3C, March 2013. <http://www.w3.org/TR/sparql11-query>
15. Hovland, D., Kontchakov, R., Skjæveland, M.G., Waaler, A., Zakharyashev, M.: Ontology-based data access to slegge. In: d’Amato, C., Fernandez, M., Tamma, V., Lecue, F., Cudré-Mauroux, P., Sequeda, J., Lange, C., Heflin, J. (eds.) *ISWC 2017*. LNCS, vol. 10588, pp. 120–129. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-68204-4\\_12](https://doi.org/10.1007/978-3-319-68204-4_12)
16. Ives, Z.G., Florescu, D., Friedman, M., Levy, A., Weld, D.S.: An adaptive query execution system for data integration. *SIGMOD Rec.* **28**(2), 299–310 (1999)
17. Kaminski, M., Kostylev, E.V., Cuenca Grau, B.: Query nesting, assignment, and aggregation in SPARQL 1.1. *ACM Trans. Database Syst.* **42**(3), 17:1–17:46 (2017)
18. Kharlamov, E., Hovland, D., Skjæveland, M.G., Bilidas, D., Jiménez-Ruiz, E., Xiao, G., Soylu, A., Lanti, D., Rezk, M., Zheleznyakov, D., Giese, M., Lie, H., Ioannidis, Y.E., Kotidis, Y., Koubarakis, M., Waaler, A.: Ontology based data access in Statoil. *J. Web Semant.* **44**, 3–36 (2017)
19. Killapi, H., Sitaridi, E., Tsangaris, M.M., Ioannidis, Y.: Schedule optimization for data processing flows on the cloud. In: *Proceeding ACM SIGMOD 2011*, pp. 289–300 (2011)
20. Kontchakov, R., Rezk, M., Rodríguez-Muro, M., Xiao, G., Zakharyashev, M.: Answering SPARQL queries over databases under OWL 2 QL entailment regime. In: Mika, P., et al. (eds.) *ISWC 2014*. LNCS, vol. 8796, pp. 552–567. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-11964-9\\_35](https://doi.org/10.1007/978-3-319-11964-9_35)
21. Motik, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokoue, A., Lutz, C.: *OWL 2 Web Ontology Language Profiles*, 2nd edn. W3C Recommendation, W3C, December 2012
22. Motik, B., Nenov, Y., Piro, R.E.F., Horrocks, I.: Handling owl:sameAs via rewriting. In: *Proceedings of AAAI 2015*, pp. 231–237. AAAI Press (2015)
23. Pérez, J., Arenas, M., Gutierrez, C.: Semantics and complexity of SPARQL. *ACM Trans. Database Syst.* **34**(3), 16:1–16:45 (2009)
24. Poggi, A., Lembo, D., Calvanese, D., De Giacomo, G., Lenzerini, M., Rosati, R.: Linking data to ontologies. *J. Data Semant.* **10**, 133–173 (2008)
25. Rodríguez-Muro, M., Rezk, M.: Efficient SPARQL-to-SQL with R2RML mappings. *J. Web Semant.* **33**, 141–169 (2015)
26. Sequeda, J.F., Arenas, M., Miranker, D.P.: OBDA: query rewriting or materialization? In practice, both! In: Mika, P., et al. (eds.) *ISWC 2014*. LNCS, vol. 8796, pp. 535–551. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-11964-9\\_34](https://doi.org/10.1007/978-3-319-11964-9_34)
27. W3C OWL Working Group. *OWL 2 Web Ontology Language document overview*, 2nd edn. W3C Recommendation, W3C, December 2012