A Rule System for Querying Persistent RDFS Data^{*}

Giovambattista Ianni¹, Thomas Krennwallner², Alessandra Martello¹,
and Axel Polleres³

¹ Dipartimento di Matematica, Università della Calabria, I-87036 Rende (CS), Italy
{ianni, a.martello}@mat.unical.it

² Institut für Informationssysteme, Technische Universität Wien, A-1040 Vienna, Austria
tkren@kr.tuwien.ac.at

³ Digital Enterprise Research Institute, National University of Ireland, Galway
axel.polleres@deri.org

Abstract. We present GiaBATA, a system for storing, aggregating, and querying Semantic Web data, based on declarative logic programming technology, namely on the *dlvhex* system, which allows us to implement a fully SPARQL compliant semantics, and on *DLV^{DB}*, which extends the DLV system with persistent storage capabilities. Compared with off-the-shelf RDF stores and SPARQL engines, we offer more flexible support for rule-based RDFS and other higher entailment regimes by enabling custom reasoning via rules, and the possibility to choose the reference ontology on a per query basis. Due to the declarative approach, GiaBATA gains the possibility to apply well-known logic-level optimization features of logic programming (LP) and deductive database systems. Moreover, our architecture allows for extensions of SPARQL by non-standard features such as aggregates, custom built-ins, or arbitrary rulesets. With the resulting system we provide a flexible toolbox that embeds Semantic Web data and ontologies in a fully declarative LP environment.

1 Introduction

Thanks to W3C standards such as RDF(S), OWL, and recently RIF, – see www.w3.org/2001/sw/ for an overview –, an increasing bulk of machine-readable data and knowledge encoded in rules and ontologies are becoming available on the Web. This vast amount of information uses rather lightweight ontologies in contrast to the full power of expressive description logic reasoning, which imposes new demands on systems to deal with Semantic Web data. A Semantic Web system should be able to *handle and evaluate arbitrarily complex queries on large amounts of RDF instance data* that are not manageable in memory. Thus, current RDF stores (see Sec. 2) provide persistent storage facilities and typically support SPARQL [1,2]. On top of SPARQL, data processing requires *ontological inferences as well as custom, possibly nonmonotonic rules*, which are crucial, e.g., for modelling mappings between vocabularies. Moreover, it should be possible to *query combinations of different data graphs, ontologies*,

^{*} This work has been partially supported by the Italian Research Ministry (MIUR) project Interlink I104CG8AGG, the Austrian Science Fund (FWF) project P20841, by Science Foundation Ireland (SFI) project Lón (SFI/02/CE1/I131), and by the EU FP6 project inContext (IST-034718).

<pre>@prefix bob: <http://bob.org/foaf.rdf#> . bob:me foaf:name "Bob"; foaf:knows _:a . _:a foaf:name "Alice" ; rdfs:seeAlso <http://alice.org/foaf.rdf> .</pre>	<pre>@prefix alice: <http://alice.org/foaf.rdf#> . alice:me a foaf:Person; foaf:name "Alice" ; foaf:knows [foaf:name "Bob"] ; foaf:knows [foaf:name "Charles"] .</pre>
(a) Graph bob.org/foaf.rdf	(b) Graph alice.org/foaf.rdf
<pre>@prefix foaf: <http://xmlns.com/foaf/0.1/> foaf:knows rdfs:range foaf:Person. myOnt:isFriendOf rdfs:subPropertyOf foaf:knows. ...</pre>	
(c) Graph example.org/myOnt.rdfs, an extension of the FOAF ontology	

Fig. 1. Two RDF graphs, and an extract of the FOAF ontology

and rules dynamically harvested from the Web. In this respect, most available RDF databases only offer limited support for dynamic inference, custom reasoning via rules, or querying upon datasets of choice.

We propose an architecture aiming at closing this gap by providing the GiaBATA system which caters for knowledge-intensive applications on top of Semantic Web data by combining (a) SPARQL querying on varying datasets and varying schemes, (b) rule-based RDFS and (limited) OWL inference, and (c) a persistent storage system based on deductive database technology. Based on systems described in Sec. 2, our proposed RDF store fulfils the above goals and is easily extensible by features such as using arbitrary aggregates and built-in predicates as well as adding custom rules to SPARQL [3].

Motivating Example. Take the two data graphs about our well-known friends Alice, Bob, and Charles in Fig. 1(a)+(b). Both use vocabulary defined in an extended FOAF ontology (<http://xmlns.com/foaf/spec/>); the relevant part is in Fig. 1(c).

For instance, a SPARQL query could be used to extract names of persons mentioned in graphs that belong to friends of Bob. We assume that Bob provides links to the graphs associated to the persons he knows using `rdfs:seeAlso` statements:

```
select ?N from <http://example.org/myOnt.rdfs> from <http://bob.org/foaf.rdf>
      from named <http://alice.org/foaf.rdf>
where { bob:me myOnt:isFriendOf ?X . ?X foaf:seeAlso ?G .
      graph ?G { ?P a foaf:Person; foaf:name ?N } }
```

(1)

The dataset $DS = (G, N)$ of a SPARQL query is defined by the *default graph* G obtained from the RDF merge [4] (denoted $\uplus$) of all graphs mentioned in **from** clauses plus a set of named graphs $N = \{(g_1, G_1), \dots, (g_k, G_k)\}$, i.e., pairs of IRIs given in **from named** clauses and their corresponding graphs. In our example, the dataset is being specified explicitly by a set of **from** and **from named** clauses, i.e., the dataset of query (1) would be $DS_1 = (G_{\text{myOnt}} \uplus G_{\text{bob.org}}, \{(\langle \text{http://alice.org/} \rangle, G_{\text{alice.org}})\})$.

Since neither of the three source graphs contain a triple matching the pattern `bob:me myOnt:isFriendOf ?X` in the **where** clause, most existing RDF stores with SPARQL support would return an empty result on (1). This is because typically SPARQL engines only take simple RDF entailment into account when answering queries, which boils down to basic graph pattern matching. When taking ontological inference by the statements of the `myOnt` ontology into account, however, one would expect to obtain three matchings for the variable `?N`: $\{?N/\text{"Alice"}, ?N/\text{"Bob"}, ?N/\text{"Charles"}\}$.

In order to obtain this expected answer, SPARQL's basic graph pattern matching needs to be extended, see [1, Section 12.6]. In principle, this means that the graph pattern in the **where** clause needs to be matched against a graph enlarged by applying inferences

prescribed by the given entailment regime.¹ SPARQL extensions providing entailment regimes other than simple RDF entailment are still an open research problem.²

Approximations of entailment regimes such as RDFS and OWL [5,6,7], however, can be modeled by finite Datalog rulesets that cover sound but not necessarily complete RDFS and OWL inferences. For instance, the RDF Semantics document contains an informative set of entailment rules, a subset of which is implemented by most available RDF stores. However, these rules could be dynamically evaluated upon query time, or inferences could be materialized at graph loading time, producing a persistent extended graph which queries are issued on. Materialization upon loading time has drawbacks though: when the dataset includes a number of different graphs, it is not clear upfront (i) which ontology should be taken into account for which data graph, and (ii) to which graph the inferred triples “belong,” which in turn complicates querying named graphs.

As for (i), assume that a user agent wants to issue another query on graph `bob.org` with the original FOAF ontology instead of our modified version within the default graph. Since the property `myOnt:isFriendOf` is stated to be a subproperty of `foaf:knows`, the triple `bob:me myOnt:isFriendOf _:a`, which could be inferred from $G_{\text{myOnt}} \sqcup G_{\text{bob.org}}$, would most probably contribute in an undesired way to such a different query. Current RDF stores however prevent to parametrize inference with an ontology of choice, since typically inferences are computed upon loading time *once and for all*.

As for (ii), queries upon datasets including named graphs are even more problematic. Even if inference was supported properly, the answer to (1) over dataset DS_1 is just $\{?N/“Alice”\}$. The `myOnt` ontology is only merged into the default graph, but not into the named graph. Thus, there is no way to infer that “Bob” and “Charles” are actually names of `foaf:Persons` within the named graph $G_{\text{alice.org}}$. Indeed, SPARQL does not allow to explicitly specify a dataset to be merged with named graphs.

In order to overcome both point (i) and (ii) above, the current nonintuitive semantics of SPARQL with RDFS can be mitigated by slightly extending SPARQL’s syntax. We give the possibility of specifying ontologies which are used for the entire dataset, including named graphs, by means of the new directive **using ontology**. For instance, to specify the dataset $DS_2 = (G_{\text{myOnt}} \sqcup G_{\text{bob.org}}, \{(<\text{http://alice.org}/>, G_{\text{myOnt}} \sqcup G_{\text{alice.org}})\})$, we allow the following syntax replacing the first line of query (1):

```
select ?N using ontology <http://example.org/myOnt.rdfs> ...
```

Hence, the **using ontology** construct allows for coupling the given dataset with the proper class/property hierarchy of the `myOnt` data schema. The possibility of query dependent and dynamic ontology inference shows the effectiveness of this approach.

2 System Description and Evaluation

GiaBATA is based on a combination of the DLV^{DB} [8] and `dlvhex` [9] systems, and caters for persistent storage of both data and ontology graphs. The former system is a variant of the DLV system with built-in database support. The latter is a solver for the so called HEX-programs [9]. `dlvhex` features an extensible plugin system which

¹ The entailment rules for getting the expected results for (1) are RDFS3 and RDFS7 from [4].

² cf. <http://www.polleres.net/sparqltutorial/>, Unit 5b, for details

Fig. 2. The GiaBATA system

we used for developing a rewriter-plugin able to translate SPARQL queries to HEX-programs. Our prototype supports the standard SPARQL specifications extended with the syntax outlined above, arbitrary entailment regimes (currently, we support RDFS), and the addition of custom rules. We provide a SPARQL-protocol compliant Web service interface. Fig. 2 shows a high-level view of our system architecture. It consists of four main components: (i) a SPARQL to datalog rewriter as part of a plugin for *dlvhex* (see [10,11]), (ii) a schema rewriter which manipulates the rewritten datalog rules, adds auxiliary definitions in order to match the underlying database schema, and implements the chosen semantics by adding rules to the input of the module in (iii) the DLV^{DB} datalog engine, which rewrites the input program to native SQL queries, and thereby accesses triples persistently stored in (iv) a DBMS of choice storing the RDF data according to a particular storage schema. This architecture provides a fully declarative implementation of SPARQL and brings support for further (deductive) database and LP optimization and extensibility.

By and large, the chosen database schema for the triples exploits a main table storing quadruples composed of subject, predicate, object, and the source graph of each triple. This representation has been improved [12] by additional relations mapping URIs/literal string to integer values, which avoids string matching in favour of integer comparison.

We have investigated the persistent storage facilities as well as SPARQL features of some of the state-of-the-art RDF stores, namely, Sesame RDF database 2.0, Owlim v3.0b8, AllegroGraph 3.0.1, and Jena/Joseki 3.2.³ All systems support persistent storage, either based on RDBMS or on other dedicated backends, and deal with RDFS inferencing and partial or non-standard OWL fragments. While a persistence strategy is provided, both reasoning and query evaluation are usually performed in main memory, adopting a materialization strategy. In our tests we focused on querying under ρ DF inference regime [5], requiring the capability to match patterns against named graphs. The investigation revealed that combining RDFS and named graph queries lead to unexpected behavior. While inference is properly handled as long as the query ranges over the whole dataset, RDFS reasoning fails in case of queries using explicit default or named graphs. That leaves dynamic named graph queries with inference infeasible in current systems.

3 Related Work and Outlook of the Demo

There are earlier approaches to integrate ontological inference in a logic programming environment. As opposed to SWI Prolog's Semantic Web library [13] which also offers

³ <http://www.openrdf.org/>, <http://www.ontotext.com/owlim/>, <http://jena.hpl.hp.com/>, and <http://agraph.franz.com/>

SPARQL support, we support Answer Set Programming as underlying LP paradigm, instead of Prolog. In contrast to OntoDLV [14], which supports proprietary ontology and query languages, the focus of GiaBATA is on compliance with Web standards such as SPARQL, RDF(S), OWL, and RIF. We currently work on the efficiency of FILTER expressions by rewriting them to SQL queries over the underlying database, see [15,16]. Moreover, one of the next steps includes to support arbitrary database schemes. Applying optimization techniques from both the DB and LP area should boost evaluation performance. Extensions of SPARQL by aggregates and custom built-ins presented in earlier works [11] carry over to our persistent storage version with minor modifications. We are going to add support for extended graphs, that is SPARQL “views” as defined also in [11].

The live discussion will provide a better insight into GiaBATA showing use cases that cover the key features of the system in accessing, storing and querying Semantic Web data. For features explanation we will refer to largely adopted benchmark suites,⁴ while ad hoc showcases will help to understand the advantages of our proposed **using ontology** extension. Finally, the users will get an overview of GiaBATA’s capabilities and will learn how to interact with it.

References

1. Prud’hommeaux, E., Seaborne A. (ed.): SPARQL Query Language for RDF, W3C Recommendation (January 2008)
2. Pérez, J., Arenas, M., Gutierrez, C.: Semantics and complexity of SPARQL. In: Cruz, I., Decker, S., Allemang, D., Preist, C., Schwabe, D., Mika, P., Uschold, M., Aroyo, L.M. (eds.) ISWC 2006. LNCS, vol. 4273, pp. 30–43. Springer, Heidelberg (2006)
3. Polleres, A., Scharffe, F., Schindlauer, R.: SPARQL++ for mapping between RDF vocabularies. In: Meersman, R., Tari, Z. (eds.) OTM 2007, Part I. LNCS, vol. 4803, pp. 878–896. Springer, Heidelberg (2007)
4. Hayes, P.: RDF semantics, W3C Recommendation (February 2004)
5. Muñoz, S., Pérez, J., Gutiérrez, C.: Minimal deductive systems for RDF. In: Franconi, E., Kifer, M., May, W. (eds.) ESWC 2007. LNCS, vol. 4519, pp. 53–67. Springer, Heidelberg (2007)
6. ter Horst, H.J.: Combining RDF and part of OWL with rules: Semantics, decidability, complexity. In: Gil, Y., Motta, E., Benjamins, V.R., Musen, M.A. (eds.) ISWC 2005. LNCS, vol. 3729, pp. 668–684. Springer, Heidelberg (2005)
7. Ianni, G., Martello, A., Panetta, C., Terracina, G.: Efficiently querying RDF(S) ontologies with Answer Set Programming. *J. Logic Computation* (August 2008)
8. Terracina, G., Leone, N., Lio, V., Panetta, C.: Experimenting with recursive queries in database and logic programming systems. *Theor. Pract. Log. Prog.* 8(2), 129–165 (2008)
9. Eiter, T., Ianni, G., Schindlauer, R., Tompits, H.: Effective Integration of Declarative Rules with External Evaluations for Semantic Web Reasoning. In: Sure, Y., Domingue, J. (eds.) ESWC 2006. LNCS, vol. 4011, pp. 273–287. Springer, Heidelberg (2006)
10. Polleres, A.: From SPARQL to rules (and back). In: WWW 2007, pp. 787–796. ACM Press, New York (2007)
11. Polleres, A., Schindlauer, R.: dlhex-sparql: A SPARQL-compliant query engine based on dlhex. In: ALPSWS 2007, CEUR WS Proceedings, vol. 287, pp. 3–12 (September 2007)

⁴ LUBM Benchmark available at <http://swat.cse.lehigh.edu/projects/lubm/>

12. Abadi, D.J., Marcus, A., Madden, S., Hollenbach, K.J.: Scalable Semantic Web Data Management Using Vertical Partitioning. In: VLDB 2007, pp. 411–422. ACM Press, New York (2007)
13. Wielemaker, J., Hildebrand, M., van Ossenbruggen, J.: Prolog as the Fundament for Applications on the Semantic Web. In: ALPSWS 2007 (2007)
14. Ricca, F., Gallucci, L., Schindlauer, R., Dellarmi, T., Grasso, G., Leone, N.: OntoDLV: An ASP-based System for Enterprise Ontologies. *J. Logic Computation* (August 2008)
15. Lu, J., Cao, F., Ma, L., Yu, Y., Pan, Y.: An Effective SPARQL Support over Relational Databases. In: SWDB-ODBS 2007, co-located with VLDB 2007, pp. 57–76 (2007)
16. Theoharis, Y., Christophides, V., Karvounarakis, G.: Benchmarking Database Representations of RDF/S Stores. In: Gil, Y., Motta, E., Benjamins, V.R., Musen, M.A. (eds.) ISWC 2005. LNCS, vol. 3729, pp. 685–701. Springer, Heidelberg (2005)