

ELITE: An *Entailment*-based Federated Query Engine for Complete and Transparent Semantic Data Integration

Andreas Nolle
Albstadt-Sigmaringen University
Jakobstraße 6
72458 Albstadt, Germany
Email: nolle@hs-albsig.de

German Nemirovski
Albstadt-Sigmaringen University
Jakobstraße 6
72458 Albstadt, Germany
Email: nemirovskij@hs-albsig.de

Abstract—In recent years the core of the semantic web has evolved to a conceptual layer built by a set of ontologies mapped onto data distributed in numerous data sources, interlinked, interpreted and processed in terms of its semantics. One of the central issues in this context became the federated querying of such linked data. This paper presents the federated query engine ELITE that facilitates a complete and transparent integration and querying of distributed autonomous data sources. To achieve this aim a combination of existing approaches for Ontology-based Data Access (OBDA) and federated query processing on Linked Open Data (LOD) are applied. Consolidating technologies like entailment regimes, the DL-Lite formalism, query rewriting, mapping relational data to RDF and an improved implementation of R-Tree based indexing contributes to the unique features of this federation engine. ELITE thereby enables the integration of various kinds of data sources, for example as relational databases or triple stores, simplicity of query design, guaranteed completeness of query results and highly efficient query processing. The federation engine has been developed and evaluated in the domain of carbon reduction in urban planning.

Keywords—*Federated Querying, Entailment Regime, Query Rewriting, Ontology-based Data Access (OBDA), DL-Lite, OWL QL, SPARQL, Linked Open Data (LOD), Indexing, R-Tree Index*

I. INTRODUCTION

Due to initiatives like Linking Open Data (LOD) [1] the total amount of semantic data online publicly available grows continuously. In recent time querying of distributed data has become an issue rising significant attention in the semantic (web) research. Approaches emerging in the context of semantic data integration range from systems for processing federated queries on linked RDF data to ontology-based data access focusing on linking relational data to ontologies.

Systems and methodologies for federation of linked data belong to the first category, e.g. LOD, that are basically stored in RDF triple stores exposing a SPARQL service. Federation is one of three main paradigms for the design of a linked data infrastructure identified by Görlitz & Staab [2]. Its aim is to analyze and process the query initially formulated by the client in order to i) identify query parts related to single sources, ii) rewrite the original query for each source, iii) forward the rewritten queries to the corresponding sources and iv) merge the query results carried out by each source. The main obstacle of this approach is the requirement to operate with

a set of vocabularies referred by the distributed sources like DBpedia [3], FOAF [4] or YAGO [5] in LOD. Consequently, most systems implementing the federation approach expect that clients formulate queries by using all vocabularies of each respective integrated data source to access distributed data. This fact restricts the flexibility of clients and forces users to think in terms that are alien to their usual domain of discourse.

The approaches of the second category aim at building a service tending to present a conceptual view of the domain of discourse formulated in terms of an ontology that is univocally mapped onto the data layer. Clients accessing the service abstracted from how the data layer is organized, formatted, and structured. They formulate semantic queries referring to the ontology and hence, to the native vocabulary of the domain of discourse. As the conceptual view typically consists of the ontology TBox only, there is no need for storing ABox individuals since semantic queries are translated in runtime into the query language (usually SQL) comprehensive for the data sources building the data layer. Furthermore, application of formalisms designed especially for OBDA purposes facilitate effective query rewriting e.g. from SPARQL into SQL, and therefore high efficiency of query evaluation ensuring sound and complete query results. However, since existing OBDA approaches focus on accessing non-RDF data in relation to its semantics, they don't take into account sources that already provide RDF data like triple stores. [6]–[8]

The approach proposed in this paper consolidates the advantageous features of the two approaches described above and as far as we know frees them of their most significant disadvantages. It uses the conceptual view for the domain of discourse formally specified as ontology, in a way similar to the OBDA approach. The central ontology in turn is linked and mapped onto vocabularies of each single data source integrated into the federated infrastructure. However in contrast to OBDA systems our approach exploits query execution strategies described in federation approaches for linked data combined with reasoning applied centrally to each query launched by the client. Due to the described architecture the system guarantees completeness of query results and enables clients to operate on the vocabulary strongly related to the domain of discourse. Besides, the system supports highly efficient query processing by means of an improved implementation of R-Tree-based indexing. The proposed approach has been implemented in

ELITE and evaluated for the domain of carbon reduction in urban planning.

The document is structured as follows: After Section II containing a survey of related work we introduce the system architecture of our approach in Section III. The implementation of entailment regimes through rewriting of SPARQL queries is described in Section IV. Section V presents an improved R-Tree-based indexing. Before concluding in Section VII we discuss evaluation results in Section VI.

II. RELATED WORK

This paper is written for readers familiar with the paradigm of semantic web [9], basics of knowledge representation by means of first order logics and description logics [10] and, having experience in application of SPARQL [11]. We recommend the mentioned references for more information about these topics.

In order to achieve complete results on SPARQL queries not only the explicitly specified data and its relations have to be taken into account but also the knowledge that can be inferred by reasoning on the RDF graph. The definition of such extended interpretation on query evaluation is part of the proposed recommendation of SPARQL 1.1 and is called entailment regimes [12]. [13]

As stated in the introduction ODBA, is one of the approaches fundamental to the work described in this paper. It enables a mapping of relational data to ontologies and therefore access to this data in terms of its semantics. The OBDA platforms –ontop– [14] and MASTRO-I [15], [16] facilitate high efficiency of query evaluation through a syntactically restricted ontology language. Through this they enable reasoning for complete query results with reduced complexity as well as highly optimized techniques for query rewriting. In contrast, other RDB2RDF tools merely focusing on the real-time conversion of relational data to RDF, commonly have restricted support of entailment regimes and weak performance in query answering especially for large data sets. Apart from the most popular D2RQ platform [17], [18] tools like Virtuoso RDF views [19], Triplify [20] or Revelytics Spyder [21] also belong to these systems. Extended but not complete lists of other systems can be found online [22], [23]. [14], [24]

Another approach strongly related to the present work focuses on facilitation of the infrastructure for linked data, and in particular on optimization of federated SPARQL queries. One of the most complete surveys of this issue is given by Görlitz & Staab [2]. The authors distinguish between different architecture variations for implementing a query-based search on LOD. As the other architectures like peer-to-peer systems have crucial disadvantages in connection with LOD, we will only consider the federation architecture. Due to the lack of space in this paper we don't focus on other architectures mentioned by the authors. In federative systems only metadata, indices and/or data statistics of each data source are kept centrally and represent data stored in each source. These items are processed to guarantee completeness of query results on the one hand, and to facilitate high efficiency of querying distributed linked (open) data sources on the other hand. Currently available are several implementations of this approach, such as FedX [25], [26], SemWIQ [27], DARQ [28],

SQUIN [29], UniStore [30], SPLENDID [31], ANAPSID [32], Avalanche [33] or AliBaba [34]. Apart from these systems, federated queries can be processed by endpoints implementing the SPARQL version 1.1. SPARQL 1.1 has been extended by a construct for specification of federated queries [35]. However most of the systems mentioned above as well as SPARQL 1.1 engines accept only federated queries, referring to the same vocabularies which are referred in the data sources to be queried. This fact puts strong limitations on usability for clients who formulate these queries and therefore have to handle all these vocabularies.

To our best knowledge at least six approaches exist that adopt various techniques to overcome this limitation. For instance, the approach described by Vidal et al. [36] uses explicit rule definitions to map elements of the domain specific ontology to the central one. However due to the formalism used for the rule specifications application of entailment regimes is not considered in its entirety. Another approach by Correndo et al. [37] uses RDF to express rewriting rules and enables query rewriting for each single source. This approach does not require a central ontology. Instead it enables translations among ontologies referred in the sources to be queried. The third system called LOQUS [38] provides a conceptual view (central ontology) to the distributed sources. LOQUS rewrites the original query into a set of source specific ones by using links and mappings specified in the central ontology, but doesn't take into account entailment regimes. The advancement of this system called ALOQUS introduced by Joshi et al. [39] facilitates querying of linked data based on ontology alignments without the need to have detailed information about the data sources. ALOQUS additionally supports an automatic mapping between ontologies to get the alignments but also omits reasoning steps. A further approach elaborated by Li & Heflin [40], [41] uses a reasoner to produce results related to a query and therefore ensure completeness. However since the reasoner is located centrally and is invoked at the end of the query execution process all selected data have to be entirely loaded in the reasoner. The sixth approach contributed by Makris et al. [42], [43] utilizes ontology mappings to rewrite original queries in terms of the target ontologies in order to access federated sources and is therefore the closest one to our approach.

Despite our extensive studies and to the best of our knowledge none of the existing approaches consider the complexity of reasoning and query answering. Concerning this Calvanese et al. [7], [8] propose the description logics family DL-Lite with the aim of low reasoning complexity. They have shown that reasoning tasks for DL-Lite are in PTIME and query answering is in LOGSPACE (more precisely even in AC^0), each in size of the TBox and ABox correspondingly. DL-Lite offers maximum expressiveness due to the fact that these efficient reasoning and query answering over large data sets are facilitated. In contrast, reasoning on more expressive description logics is in worst-case EXPTIME and query answering in NLOGSPACE. For these reasons application of DL-Lite in the central ontology would facilitate efficient reasoning and query answering. Since the OWL 2 QL profile [44] bases upon the DL-Lite family, we use OWL 2 QL as specification language. Through that we achieve not just an efficient rewriting of the original query into data source specific queries but also an efficient accessing of relational data sources using OBDA tools




Fig. 1. System Architecture of ELITE

like `-ontop-` and its integrated reasoner Quest [14], [45], [46]. Furthermore we don't require any additional definition of mapping or transformation rules like other mentioned approaches.

III. SYSTEM ARCHITECTURE

The proposed architecture has been developed with the goals mentioned in the introduction (Section I). These are i) efficient query processing in particular complete evaluation of conjunctive federated queries, ii) ability for clients to formulate queries abstracted from the technological and structural characteristics of the data sources and iii) ability to federate data sources using different schemas, e.g. RDF triple stores and relational data base systems. One key feature of the proposed federative approach is the application of a single central ontology that on the one hand represents vocabulary and the knowledge structure of the domain of discourse and, on the other hand serves as a mediator for heterogeneous data sources. All queries formulated by a client and potentially targeting data distributed in single sources refer to this ontology. However the client is not constrained to use only the central vocabulary but is also able to use elements (e.g. concepts) of data source specific vocabularies. The system architecture of ELITE is outlined in Fig. 1.

The central ontology (A in Fig. 1) is basically a TBox that comprises apart of own items, i.e. concepts, roles and axioms defined in the own namespaces, items semantically related to the domain of discourse and specified in openly accessible vocabularies (B), e.g. those used in the LOD sources. The integration of such items into the central ontology is implemented by their linking through using RDFs or OWL properties. For instance by equate the concept `ResBuilding` of the central ontology with the concept `DwellingHouse` of an external ontology ($\text{ResBuilding} \equiv \text{DwellingHouse}$), or by subsumption of an external concept `DwellingHouse` ($\text{DwellingHouse} \sqsubseteq \text{Building}$).

Such mapping and linking facilitate federated querying of related data sources using their own vocabularies. Issues related to the design of the central ontology are out of scope of this paper, we just would like to mention that the efficient design of the central ontology is basically carried out manually, e.g. using a dedicated process model, as it is described by Nemirovski et al. [47], or using automated ontology matching systems as shown by Shvaiko & Euzenat [48].

The proposed approach works for triple stores as well as for relational databases under the condition that all integrated sources (C) provide a SPARQL-endpoint to query data. For conversion of relational data to the RDF format tools like D2RQ and `-ontop-` are used. These tools are applied locally for each source to map relational data to items of the central ontology as it is shown for example by Poggi et al. [6]. Apart from the central ontology (A) another centrally located component of ELITE architecture is the index that is initialized (0) after the centralized ontology and the data sources are defined. The most important principles of index structure and its usage for identification of data sources related to a query are described in Section V.

The querying process starts when a query formulated by an external application or by a user is launched at the central SPARQL endpoint of ELITE (1). Elements of this query, e.g. basic graph patterns (BGP), refer to items of the central ontology and elements of other vocabularies linked to the central ontology as described above. In the first step the initial query is rewritten (2) by means of entailment regimes, whereby the semantics of the central ontology as well as of the linked external ontologies are taken into account. The application of entailment regimes (described in Section IV) results in multiple queries reflecting all possibilities to express the same statement using the knowledge implicated by all considered vocabularies. The disjunction of evaluation results of these queries by each data source delivers the initial query result complete in terms of all considered vocabularies. Yet not all BGPs contained in these queries are relevant to all integrated sources. Therefore not all queries, more precisely its patterns generated through application of entailment regimes, can be answered by any sources. To avoid sending queries causing empty results, such queries are excluded by the index look-up procedure from further processing. In the next step (3) this procedure is applied for every single BGP in all queries to determine the sources which contain data related to each BGP. The index also delivers an estimated size of results that would be returned by each data source for each BGP (see Section V for a detailed description of indexing). Using estimated sizes and dependencies among patterns of each query, an optimized query execution plan is generated (4). After results to all pre-selected queries executed by pre-selected sources according to the previously generated plan are retrieved, the federated result is returned to the client that has launched the initial query (5).

IV. ENTAILMENT BY QUERY REWRITING

In order to take knowledge inferred from persistently stored RDF triples in the query execution process into account the implementation of such entailment regimes can be done in three different ways. The first technique is to materialize all triples that are inferred according to the used entailment regimes e.g. RDFS or OWL QL entailment regimes, that is also called forward-chaining. The alternative technology is called backward-chaining. It operates by inferring of new triples through query rewriting in runtime. The third option is a hybrid combination of both approaches. These techniques (query rewriting and materialization of inferred triples) are further described by Glimm [13]. Though the query evaluation can be more efficient by using forward-chaining, it isn't really compatible with the federation approach. Since most of linked (open) data sources grant read only access to the clients, an external location for storing the inferred triples has to be found. Furthermore all changes of the data sources have to be reflected by the materialized data to keep it up to date. Contrary to this backward-chaining is more flexible with reference to handling dynamic data and without the need of an additional location to store the materialized data.

Since in our approach the entailment is done centralized, all inferable knowledge as well as all conceivable external concepts and properties referred in the central ontology are taken into account in the patterns of the rewritten SPARQL queries. Therefore the federative system doesn't have to rely on the entailment by the SPARQL services of single data sources to guarantee the completeness of query results. Even if the local SPARQL services don't implement entailment regimes at all, the query results will be complete in terms of the central ontology. Yet the central query rewriting does not only help in cases when local SPARQL services are not (fully) support entailment. Furthermore, even though a SPARQL service of a single source implements the entailment regimes, it only can infer knowledge based on the vocabularies (ontologies) available at the corresponding source. If the central ontology that specifies the domain of discourse is not among these vocabularies, completeness of query results in terms of the central ontology can't be guaranteed based on local entailments only.

At this point we want to thank the developers of Quest who provided us with the source code of the system at our disposal. We extend it by implementing a translator to convert datalog programs, Quests internal rule representation of the rewritten queries into SPARQL queries by analyzing and inspecting the generated rules. As Quest contains various rewriting implementations we have chosen the Tree Witness Rewriter implemented by Kontchakov et al. [49]–[51] because the authors have shown that this rewriter produces better evaluation results than to other ones. Assuming for example that the TBox of the central ontology contains the following statements (namespace specifications are omitted):

```
ResBuilding ⊆ Building (1)
CommBuilding ⊆ Building
∃hasFloor_Area ⊆ Building
∃hasFloor_Area- ⊆ B_Floor_Area
∃hasValue ⊆ B_Floor_Area
Range(hasValue) ≡ rdf:decimal
```

```
∃hasRFloorArea ⊆ ResBuilding
∃hasRFloorArea- ⊆ ResFloorArea
ResFloorArea ⊆ B_Floor_Area
hasRFloorArea ⊆ hasFloor_Area
```

The distributed sources $DS_A(2)$ and $DS_B(3)$ contain the following data (ABox):

```
ResBuilding(Hundertwasserhaus) (2)
ResFloorArea(Hwh3356)
:Hundertwasserhaus :hasRFloorArea :Hwh3356
:Hwh3356 :hasValue 3356.0^^xsd:decimal
:Casa_Mila :hasRFloorArea :CM1000
:CM1000 :hasValue 1000^^xsd:decimal
```

```
CommBuilding(Tanzende_Tuerme) (3)
B_Floor_Area(TT33357)
:Tanzende_Tuerme :hasFloor_Area :TT33357
:TT33357 :hasValue 33357.0^^xsd:decimal
```

To get all buildings the user only has to define the following query:

```
SELECT ?building { (4)
  ?building a :Building . }
```

After entailment by query rewriting we get the following datalog programs:

```
q(building) :- hasRFloorArea(building,_) (5)
q(building) :- Building(building)
q(building) :- CommBuilding(building)
q(building) :- hasFloor_Area(building,_)
q(building) :- ResBuilding(building)
```

These programs can be rewritten accordingly as SPARQL queries that have to be federated. To get all floor area values of all buildings the user only has to define the following query:

```
SELECT ?building ?buildingFloorArea { (6)
  ?building :hasFloor_Area _:bfa .
  _:bfa :hasValue ?buildingFloorArea . }
```

It is important to notice that in a case when entailment regimes are not implemented at SPARQL services, neither centrally nor locally, it is the task of the client to formulate the query in a way that makes sure that all related RDF graph patterns are taken into account. For this purpose the client has to be aware of all formal vocabulary structures her/his query refers to, and semantics that can be applied to these vocabularies. Consequentially the complexity of queries and of client applications would increase significantly. For instance, the query for selecting all floor area values of all buildings (6) would look like this:

```
SELECT ?building ?buildingFloorArea { (7)
  { ?building :hasFloor_Area _:bfa . }
  UNION
  { ?building :hasRFloorArea _:bfa . }
  _:bfa :hasValue ?buildingFloorArea . }
```




Fig. 2. Systematical Representation of our Index Structure & Spatial Join of two BGPs

V. INDEXING

Since a query may address complex relations across several data sources it is not unusual that a single source is only able to answer a part of the query. However the complete query can't be forwarded to a source that is not able to deliver results for each query pattern and hence would return an empty result set. Instead, the federation engine has to identify relations between data sources and query patterns, to generate, based on these relations, dedicated sub-queries and forward them to corresponding sources. To select all sources relevant for each single part of a query an index catalog implementing a look-up service and able to estimate the size of the sub-query results of each source is required.

The comparative analysis of currently available indexing methods for distributed RDF data can't be provided in this paper due to space limitations. For a survey of the state of the art we point out the work of Görlitz & Staab [2]. Here we only describe the indexing approach developed for ELITE. Our index is an extension of the approach called QTree developed by Harth et al. [52]. QTree provides an optimized index to summarize RDF data. The index is based on R-Trees [53] and summarizes the data in so called Minimum Bounding Boxes (MBBs). MBBs are approximation cubes of three-dimensional points built with the hash values for each triple element (subject, predicate and object). The query result estimation for single SPARQL operations like joins is done by performing the same operation on the MBBs. Afterwards, the results are considered for selecting data sources related to particular BGPs of a query and for determining the sequence of sub-query processing (query execution plans). Since the approach of QTree has some limitations, e.g. inaccurate granularity of MBBs, and therefore is only able to process simple queries for small sets of data this approach has been adopted by Prasser et al. [54] who have implemented PARTree, a more efficient indexing suited for complex querying of large RDF data sources.

Similar to PARTree our approach follows the strategy to index each single data source independently but in contrast to PARTree it uses more fine-grained partitioning of RDF triples. A systematical representation of the described index architecture is shown below in part (A) of Fig. 2. The index partitions are identified by the hash value of the predicate's URI (p_{hash}), the hash value of the prefix - more precisely the namespace - of the subject (s_prefix_{hash}) and object (o_prefix_{hash}) and additionally by the type of the subject (s_type_{hash}) and object (o_type_{hash}). In this context the type of a particular ABox individual is defined as the URI of the TBox concept connected to the individual by the `rdf:type` relation. For the last triple of the assertion below the type of the subject Hundertwasserhaus would be `ResBuilding`:

```
ResBuilding(Hundertwasserhaus) (8)
FloorArea(Hwh3356)
:Hundertwasserhaus :hasRFloorArea :Hwh3356
```

If no type of an individual can be derived the assigned hash value for s_type_{hash} or o_type_{hash} is 0. Since the `rdf:type` for TBox elements does not exist, this would be also the case for the objects in type assertions like the first two statements of the example above. Using this fine-grained segmentation for triple partitions leads to more accurate results compared to QTree as well as PARTree and therefore to a more exact data source selection as well as more precise result estimations. For type assertions like e.g. the first one shown in (8) the described segmentation results in an one-dimensional index structure, because the type URIs of the predicate and the object as well as the namespaces for the subject are unique in each partition.

Calling the index look-up service for a BGP results in selections of MBBs for each source containing triples fitting to this pattern. For this purpose the corresponding hash values are calculated for all pattern elements defined by an URI and all found MBBs of each source are returned. For instance

two BGBs that are connected with the same subject variable, estimated results can be calculated by joining the spatial results (MBBs). A graphical representation of this example is given in part (B) of Fig. 2.

Besides the partitioning another main difference to PARTree and QTree is that our index isn't hold in main memory but instead stored in a spatial database. Due to this solution we benefit from the optimized features of spatial database systems to manage spatial objects, to perform spatial operations like joins and to achieve high scalability of the index. To fill the index with summarized data of each data source we crawl them through of SPARQL queries specially constructed for this purpose.

VI. EVALUATION

The goal of our evaluation was to prove the completeness of query results generated by ELITE, and furthermore to demonstrate that conventional systems, which neither use query rewriting techniques nor apply materialization of inferred data aren't able to deliver comparable results. To evaluate ELITE we have used part of an ontology developed for the SEMANCO project [55] operating in the domain of carbon reduction in urban planning. This ontology is conforming to the rules defined by the DL-Lite_A formalism. DL-Lite_A introduced by Poggi et al. [6] is a member of the DL-Lite family and comprises the computational feature of the DL-lite core variant with higher expressiveness.

Since none of the six strongly related approaches mentioned in Section II are publicly accessible we have compared ELITE to available state of the art federation systems with respect to completeness of query results. For evaluation purposes several data sets containing data properties of individuals specified within the selected ontology have been distributed over two data sources. Utilizing D2RQ we mapped the relational data to the appropriated ontology elements.

Using this infrastructure a comparative evaluation with the FedX system for federation of distributed data has been carried out. As stated in Section IV, query results generated by D2RQ using FedX only take into account data defined explicitly by the D2RQ mappings. This hypothesis has been confirmed by the evaluation results. Since inferring has been carried out neither by local installations of D2RQ nor by FedX, no other matches to the query patterns originally submitted have been identified. In contrast to this, caused by the implemented entailment by query rewriting, ELITE has delivered complete results taking into account all implicit knowledge. As expected, evaluation results (Table I) demonstrate that ELITE provides an effective method for entailment-based federated query processing and delivers complete results when other systems fail.

TABLE I. SELECTED EVALUATION RESULTS

	FedX	ELITE	subsistent
query 1 (4)	20	40	40
query 2 (6)	10	20	20
query 3 (7)	20	20	20

VII. CONCLUSION AND FUTURE WORK

This paper introduced the federated query engine ELITE that enables querying of linked data on distributed autonomous data sources transparently. Application of entailment regimes by query rewriting facilitates completeness of query results in terms of a central ontology that represents a native vocabulary of the domain of discourse. By the ability to generate complete query results regardless of the supported SPARQL endpoint features by integrated data sources, ELITE has an outstanding position compared to other federation approaches. Beside that this paper elucidated an improved implementation of R-Tree based indexing. This technology and application of the DL-Lite formalism for ontology design are crucial in relation to high efficiency.

Our next steps towards further improvement of ELITE characteristics will be the implementation of a query execution plan, optimization and identification of alternative techniques for federated querying. Since in this paper we only have described the evaluation results concerning completeness of query results, our future target will be evaluation of efficiency. We will evaluate selected benchmarks and try to quantify contributions of particular technologies to the efficiency of query processing.

ACKNOWLEDGMENT

We thank Dr. Mariano Rodríguez-Muro for providing us the source code of Quest as well as for supporting. The main contribution of this work has been developed within the SEMANCO project, which is being carried out with the support of the Seventh Framework Programme "ICT for Energy Systems" 20112014, under the grant agreement no. 287534.

REFERENCES

- [1] LinkingOpenData - W3C Wiki, <http://www.w3.org/wiki/SweoIG/TaskForces/CommunityProjects/LinkingOpenData>.
- [2] O. Görlitz and S. Staab, "Federated data management and query optimization for linked open data," in *New Directions in Web Data Management 1*. Springer, 2011, pp. 109–137.
- [3] DBpedia, <http://dbpedia.org>.
- [4] The Friend of a Friend (FOAF) project, <http://www.foaf-project.org>.
- [5] YAGO2s: A High-Quality Knowledge Base, <http://mpii.de/yago>.
- [6] A. Poggi, D. Lembo, D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati, "Linking data to ontologies," in *Journal on data semantics X*. Springer, 2008, pp. 133–173.
- [7] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati, "Tractable reasoning and efficient query answering in description logics: The DL-Lite family," *Journal of Automated reasoning*, vol. 39, no. 3, pp. 385–429, 2007.
- [8] A. Artale, D. Calvanese, R. Kontchakov, and M. Zakharyashev, "The DL-Lite family and relations," *Journal of Artificial Intelligence Research*, vol. 36, no. 1, pp. 1–69, 2009.
- [9] P. Hitzler, M. Krotzsch, and S. Rudolph, *Foundations of semantic web technologies*. Chapman and Hall/CRC, 2011.
- [10] F. Baader, *The description logic handbook: theory, implementation, and applications*. Cambridge: Cambridge University Press, 2003.
- [11] SPARQL 1.1 Overview, <http://www.w3.org/TR/sparql11-overview>.
- [12] SPARQL 1.1 Entailment Regimes, <http://www.w3.org/TR/sparql11-entailment>.
- [13] B. Glimm, "Using SPARQL with RDFS and OWL entailment," in *Reasoning Web. Semantic Technologies for the Web of Data*. Springer, 2011, pp. 137–201.

- [14] M. Rodríguez-Muro and D. Calvanese, "High performance query answering over DL-Lite ontologies," in *Proc. of the 13th Int. Conf. KR*, 2012.
- [15] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, A. Poggi, and R. Rosati, "MASTRO-I: Efficient integration of relational data through DL ontologies," in *Proceedings of the 20th International Workshop on Description Logics*. Citeseer, 2007.
- [16] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, A. Poggi, M. Rodríguez-Muro, R. Rosati, M. Ruzzi, and D. F. Savo, "The Mastro system for ontology-based data access," *Semantic Web*, vol. 2, no. 1, pp. 43–53, 2011.
- [17] C. Bizer and A. Seaborne, "D2RQ-treating non-RDF databases as virtual RDF graphs," in *Proceedings of the 3rd international semantic web conference (ISWC2004)*, 2004, p. 26.
- [18] C. Bizer and R. Cyganiak, "D2r server-publishing relational databases on the semantic web," in *5th international Semantic Web conference*, 2006, p. 26.
- [19] Mapping Relational Data to RDF with Virtuoso's RDF Views, <http://virtuoso.openlinksw.com/whitepapers/relational%20rdf%20views%20mapping.html>.
- [20] S. Auer, S. Dietzold, J. Lehmann, S. Hellmann, and D. Aumueller, "Triplify: light-weight linked data publication from relational databases," in *Proceedings of the 18th international conference on World wide web*. ACM, 2009, pp. 621–630.
- [21] Spyder, <http://www.revelytix.com/content/spyder>.
- [22] Implementations - RDB2RDF, <http://www.w3.org/2001/sw/rdb2rdf/wiki/Implementations>.
- [23] Links and Resources, <http://d2rq.org/resources>.
- [24] M. Rodríguez-Muro and D. Calvanese, "Quest, a system for ontology based data access," in *OWLED-2012 - OWL: Experiences and Directions Workshop 2012*, vol. 849. CEUR Electronic Workshop Proceedings, 2012.
- [25] A. Schwarte, P. Haase, K. Hose, R. Schenkel, and M. Schmidt, "FedX: a federation layer for distributed query processing on linked open data," in *The Semantic Web: Research and Applications*. Springer, 2011, pp. 481–486.
- [26] —, "FedX: Optimization techniques for federated query processing on linked data," in *The Semantic Web-ISWC 2011*. Springer, 2011, pp. 601–616.
- [27] A. Langegger, W. Wöß, and M. Blöchl, "A semantic web middleware for virtual data integration on the web," in *The Semantic Web: Research and Applications*. Springer, 2008, pp. 493–507.
- [28] B. Quilitz and U. Leser, "Querying distributed RDF data sources with SPARQL," in *The Semantic Web: Research and Applications*. Springer, 2008, pp. 524–538.
- [29] O. Hartig, C. Bizer, and J.-C. Freytag, "Executing SPARQL queries over the web of linked data," in *The Semantic Web-ISWC 2009*. Springer, 2009, pp. 293–309.
- [30] M. Karnstedt, K.-U. Sattler, and M. Hauswirth, "Scalable distributed indexing and query processing over Linked Data," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 10, pp. 3–32, 2012.
- [31] O. Görlitz and S. Staab, "SPLENDID: SPARQL endpoint federation exploiting VOID descriptions," in *Proceedings of the 2nd International Workshop on Consuming Linked Data, Bonn, Germany*, 2011.
- [32] M. Acosta, M.-E. Vidal, T. Lampo, J. Castillo, and E. Ruckhaus, "ANAPSID: AN Adaptive query ProcesSing engine for sparql endpoints," in *The Semantic Web-ISWC 2011*. Springer, 2011, pp. 18–34.
- [33] C. Basca and A. Bernstein, "Avalanche: putting the spirit of the web back into semantic web querying," in *The 6th International Workshop on SSWS at ISWC*, 2010.
- [34] AliBaba, <http://www.openrdf.org/alibaba.jsp>.
- [35] SPARQL 1.1 Federated Query, <http://www.w3.org/TR/sparql11-federated-query>.
- [36] V. M. Vidal, J. A. de Macêdo, J. C. Pinheiro, M. A. Casanova, and F. Porto, "Query processing in a mediator based framework for linked data integration," *International Journal of Business Data Communications and Networking (IJBDCN)*, vol. 7, no. 2, pp. 29–47, 2011.
- [37] G. Correndo, M. Salvadores, I. Millard, H. Glaser, and N. Shadbolt, "SPARQL query rewriting for implementing data integration over linked data," in *Proceedings of the 2010 EDBT/ICDT Workshops*. ACM, 2010, p. 4.
- [38] P. Jain, K. Verma, P. Z. Yeh, P. Hitzler, and A. P. Sheth, "LOQUUS: Linked Open Data SPARQL Querying System," Tech. rep., Kno. e. sis Center, Wright State University, Dayton, Ohio, 2010. Available from <http://www.pascal-hitzler.de/resources/publications/loqus-tr-2010.pdf>, Tech. Rep., 2010.
- [39] A. K. Joshi, P. Jain, P. Hitzler, P. Z. Yeh, K. Verma, A. P. Sheth, and M. Damova, "Alignment-based Querying of Linked Open Data," in *On the Move to Meaningful Internet Systems: OTM 2012*. Springer, 2012, pp. 807–824.
- [40] Y. Li and J. Heflin, "Using reformulation trees to optimize queries over distributed heterogeneous sources," in *The Semantic Web-ISWC 2010*. Springer, 2010, pp. 502–517.
- [41] Y. Li, "A Federated Query Answering System for Semantic Web Data," Ph.D. dissertation, Lehigh University, 2013.
- [42] K. Makris, N. Gioldasis, N. Bikakis, and S. Christodoulakis, "Sparql rewriting for query mediation over mapped ontologies," Tech. rep., Technical University of Crete, Tech. Rep., 2010.
- [43] —, "Ontology mapping and SPARQL rewriting for querying federated RDF data sources," in *On the Move to Meaningful Internet Systems, OTM 2010*. Springer, 2010, pp. 1108–1117.
- [44] OWL 2 Web Ontology Language Profiles (Second Edition), <http://www.w3.org/TR/owl2-profiles>.
- [45] Quest, http://ontop.inf.unibz.it/?page_id=7.
- [46] M. Rodríguez-Muro and D. Calvanese, "Quest, an OWL 2 QL reasoner for ontology-based data access," *OWLED 2012*, 2012.
- [47] G. Nemirovski, A. Nolle, Á. Sicilia, I. Ballarini, and V. Corado, "Data Integration Driven Ontology Design, Case Study Smart City," in *The Semantic Smart City Workshop (SemCity-13)*, 2013, accepted, to appear.
- [48] P. Shvaiko and J. Euzenat, "Ontology matching: state of the art and future challenges," vol. 25, pp. 158–176, 2013.
- [49] R. Kontchakov, C. Lutz, D. Toman, F. Wolter, and M. Zakharyashev, "The combined approach to ontology-based data access," in *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume Three*. AAAI Press, 2011, pp. 2656–2661.
- [50] S. Kikot, R. Kontchakov, V. Podolskii, and M. Zakharyashev, "Long Rewritings, Short Rewritings," in *The 2012 Int. Workshop on Description Logics*. CEUR Electronic Workshop Proceedings, 2012, pp. 235–245.
- [51] S. Kikot, R. Kontchakov, and M. Zakharyashev, "On (in) tractability of OBDA with OWL 2 QL," in *Proc. of the 24th Int. Workshop on Description Logic*, 2011, pp. 224–234.
- [52] A. Harth, K. Hose, M. Karnstedt, A. Polleres, K.-U. Sattler, and J. Umbrich, "Data summaries for on-demand queries over linked data," in *Proceedings of the 19th international conference on World wide web*. ACM, 2010, pp. 411–420.
- [53] A. Guttman, "R-trees: a dynamic index structure for spatial searching," in *Proceedings of the 1984 ACM SIGMOD international conference on Management of data*. ACM, 1984, pp. 47–57.
- [54] F. Prasser, A. Kemper, and K. A. Kuhn, "Efficient distributed query processing for autonomous RDF databases," in *Proceedings of the 15th International Conference on Extending Database Technology*. ACM, 2012, pp. 372–383.
- [55] SEMANCO, <http://semanco-project.eu>.