RESEARCH ARTICLE

A Survey of RDF Data Management Systems

M. Tamer ÖZSU(✉)

University of Waterloo, Cheriton School of Computer Science, Canada

arXiv:1601.00707v1 [cs.DB] 5 Jan 2016

Abstract   RDF is increasingly being used to encode data for the semantic web and for data exchange. There have been a large number of works that address RDF data management. In this paper we provide an overview of these works.

Keywords   RDF, SPARQL, Linked Object Data.

1   Introduction

The RDF (**R**esource **D**escription **F**ramework) is a W3C standard that was proposed for modeling Web objects as part of developing the semantic web. However, its use is now wider than the semantic web. For example, Yago and DBPedia extract facts from Wikipedia automatically and store them in RDF format to support structural queries over Wikipedia [1, 2]; biologists encode their experiments and results using RDF to communicate among themselves leading to RDF data collections, such as Bio2RDF (bio2rdf.org) and Uniprot RDF (dev.isb-sib.ch/projects/ uniprot-rdf). Related to semantic web, LOD (Linking Open Data) project builds a RDF data cloud by linking more than 3000 datasets, which currently have more than 84 billion triples[1]. A recent work [3] shows that the number of data sources in LOD has doubled within three years (2011-2014).

RDF data sets have all four accepted characteristics of "big data": volume, variety, velocity, and veracity. We hinted at increasing volumes above. RDF data, as captured in the LOD cloud, is highly varied with many different types of data from very many sources. Although early (and still much of) RDF data sets are stationary, there is increasing interest in streaming RDF to the extent that W3C has now set up a community interested in addressing problems of high velocity streaming data (www.w3.org/community/rsp/). RDF-encoded social network graphs that continually change is one example of streaming RDF. Even a benchmark has been proposed for this purpose [4]. Finally, it is commonly accepted that RDF data is "dirty", meaning that it would contain inconsistencies. This is primarily as a result of automatic extraction and conversion of data into RDF format; it is also a function of the fact that, in the semantic web context, the same data are contributed by different sources with different understandings. A line of research focuses on data quality and cleaning of LOD data and RDF data in general [5, 6].

Because of these characteristics, RDF data management has been receiving considerable attention. In particular, the existence of a standard query language, SPARQL, as defined by W3C, has given impetus to works that focus on efficiently executing SPARQL queries over large RDF data sets. In this paper, we provide an overview of research efforts in management of RDF data. We note that the amount of work in this area is considerable and it is not our aim to review all of it. We hope to highlight the main approaches and refer to some of the literature as examples. Furthermore, due to space considerations, we limit our focus on the management of stationary RDF data, ignoring works on

[1] The statistic is reported in http://stats.lod2.eu/.

other aspects.

The organization of the paper is as follows. In the next section (Section 2) we provide a high level overview of RDF and SPARQL to establish the framework. In Section 3, the centralized approaches to RDF data management are discussed, whereas Section 4 is devoted to a discussion of distributed RDF data management. Efforts in querying the LOD cloud is the subject of Section 5.

2   RDF Primer

In this section, we provide an overview of RDF and SPARQL. The objective of this presentation is not to cover RDF fully, but to establish a basic understanding that will assist in the remainder of the paper. For a fuller treatment, we refer the reader to original sources of RDF [7, 8]. This discussion is based on [9] and [10].

RDF models each "fact" as a set of triples (**s**ubject, **p**roperty (or **p**redicate), **o**bject), denoted as $\langle s, p, o \rangle$, where *subject* is an entity, class or blank node, a *property*[2] denotes one attribute associated with one entity, and *object* is an entity, a class, a blank node, or a literal value. According to the RDF standard, an entity is denoted by a URI (Uniform Resource Identifier) that refers to a named *resource* in the environment that is being modelled. Blank nodes, by contrast, refer to anonymous resources that do not have a name[3]. Thus, each triple represents a named relationship; those involving blank nodes simply indicate that "something with the given relationship exists, without naming it" [7].

It is possible to annotate RDF data with semantic metadata using RDFS (RDF Schema) or OWL, both of which are W3C standards. This annotation primarily enables reasoning over the RDF data (called entailment), that we do not consider in this paper. However, as we will see below, it also impacts data organization in some cases, and the metadata can be used for semantic query optimization. We illustrate the fundamental concepts by simple examples using RDFS, which allows the definition of *classes* and *class hierarchies*. RDFS has built-in class definitions – the more important ones being

---

[2] In literature, the terms "property" and "predicate" are used interchangeably; in this paper, we will use "property" consistently.

[3] In much of the research, blank nodes are ignored. Unless explicitly stated otherwise, we will ignore them in this paper as well.

Prefixes:
mdb=http://data.linkedmdb.org/resource/ geo=http://sws.geonames.org/
bm=http://wifo5-03.informatik.uni-mannheim.de/bookmashup/
exvo=http://lexvo.org/id/
wp=http://en.wikipedia.org/wiki/

Subject	Property	Object
mdb: film/2014	rdfs:label	"The Shining"
mdb:film/2014	movie:initial_release_date	"1980-05-23"
mdb:film/2014	movie:director	mdb:director/8476
mdb:film/2014	movie:actor	mdb:actor/29704
mdb:film/2014	movie:actor	mdb: actor/30013
mdb:film/2014	movie:music_contributor	mdb: music_contributor/4110
mdb:film/2014	foaf:based_near	geo:2635167
mdb:film/2014	movie:relatedBook	bm:0743424425
mdb:film/2014	movie:language	lexvo:iso639-3/eng
mdb:director/8476	movie:director_name	"Stanley Kubrick"
mdb:film/2685	movie:director	mdb:director/8476
mdb:film/2685	rdfs:label	"A Clockwork Orange"
mdb:film/424	movie:director	mdb:director/8476
mdb:film/424	rdfs:label	"Spartacus"
mdb:actor/29704	movie:actor_name	"Jack Nicholson"
mdb:film/1267	movie:actor	mdb:actor/29704
mdb:film/1267	rdfs:label	"The Last Tycoon"
mdb:film/3418	movie:actor	mdb:actor/29704
mdb:film/3418	rdfs:label	"The Passenger"
geo:2635167	gn:name	"United Kingdom"
geo:2635167	gn:population	62348447
geo:2635167	gn:wikipediaArticle	wp:United_Kingdom
bm:books/0743424425	dc:creator	bm:persons/Stephen+King
bm:books/0743424425	rev:rating	4.7
bm:books/0743424425	scom:hasOffer	bm:offers/0743424425amazonOffer
lexvo:iso639-3/eng	rdfs:label	"English"
lexvo:iso639-3/eng	lvont:usedIn	lexvo:iso3166/CA
lexvo:iso639-3/eng	lvont:usesScript	lexvo:script/Latn

Fig. 1   Example RDF dataset. Prefixes are used to identify the data sources.

rdfs:Class and rdfs:subClassOf that are used to define a class and a subclass, respectively (another one, rdfs:label is used in our query examples below). To specify that an individual resource is an element of the class, a special property, rdf:type is used. For example, if we wanted to define a class called Movies and two subclasses ActionMovies and Dramas, this would be accomplished in the following way:

 Movies rdf:type rdfs:Class .
ActionMovies rdfs:subClassOf Movies .
Dramas rdfs:subClassOf Movies .

Definition 1 [RDF data set] Let $\mathcal{U}, \mathcal{B}, \mathcal{L}$, and $\mathcal{V}$ denote the sets of all URIs, blank nodes, literals, and variables, respectively. A tuple $(s, p, o) \in (\mathcal{U} \cup \mathcal{B}) \times \mathcal{U} \times (\mathcal{U} \cup \mathcal{B} \cup \mathcal{L})$ is an *RDF triple*. A set of RDF triples form a *RDF data set*.

An example RDF data set is shown in Figure 1 where the data comes from a number of sources as defined by the URI prefixes.

RDF data can be modeled as an RDF graph, which is formally defined as follows.

Definition 2 [RDF graph] A *RDF graph* is a six-tuple

$G = \langle V, L_V, f_V, E, L_E, f_E \rangle$, where

1. $V = V_c \cup V_e \cup V_l$ is a collection of vertices that correspond to all subjects and objects in RDF data, where V_c, V_e, and V_l are collections of class vertices, entity vertices, and literal vertices, respectively.

2. L_V is a collection of vertex labels.

3. A *vertex labeling function* $f_V : V \rightarrow L_V$ is an bijective function that assigns to each vertex a label. The label of a vertex $u \in V_l$ is its literal value, and the label of a vertex $u \in V_c \cup V_e$ is its corresponding URI.

4. $E = \{\overrightarrow{u_1, u_2}\}$ is a collection of directed edges that connect the corresponding subjects and objects.

5. L_E is a collection of edge labels.

6. An *edge labeling function* $f_E : E \rightarrow L_E$ is an bijective function that assigns to each edge a label. The label of an edge $e \in E$ is its corresponding property.

An edge $\overrightarrow{u_1, u_2}$ is an *attribute property* edge if $u_2 \in V_l$; otherwise, it is a *link* edge.

Figure 2 shows an example of an RDF graph. The vertices that are denoted by boxes are entity or class vertices, and the others are literal vertices.
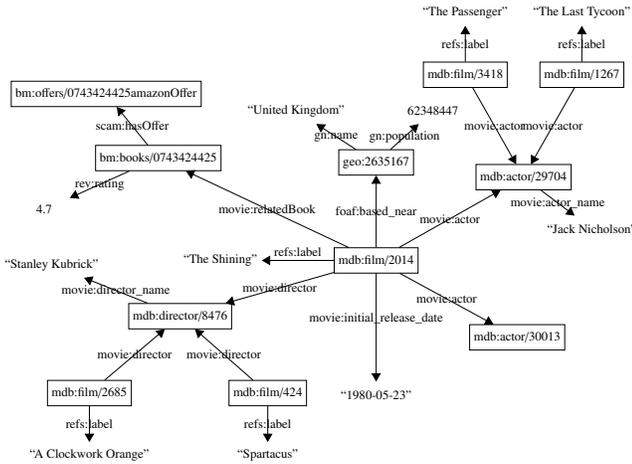

Fig. 2 RDF graph corresponding to the dataset in Figure 1

The W3C standard language for RDF is SPARQL, which can be defined as follows [11] (for a more formal definition, we refer to the W3C specification [12]).

Definition 3 [SPARQL query] Let $\mathcal{U}, \mathcal{B}, \mathcal{L}$, and $\mathcal{V}$ denote the sets of all URIs, blank nodes, literals, and variables, respectively. A SPARQL expression is expressed recursively

- A *triple pattern* $(\mathcal{U} \cup \mathcal{B} \cup \mathcal{V}) \times (\mathcal{U} \cup \mathcal{V}) \times (\mathcal{U} \cup \mathcal{B} \cup \mathcal{L} \cup \mathcal{V})$ is a SPARQL expression,
- (optionally) If P is a SPARQL expression, then $P \ FILTER \ R$ is also a SPARQL expression where R is a built-in SPARQL filter condition,
- (optionally) If P_1 and P_2 are SPARQL expressions, then $P_1 \ AND|OPT|OR \ P_2$ are also SPARQL expressions.

A set of triple patterns is called *basic graph pattern* (BGP) and SPARQL expressions that only contain these are called *BGP queries*. These are the subject of most of the research in SPARQL query evaluation.

An example SPARQL query that finds the names of the movies directed by "Stanley Kubrick" and have a related book that has a rating greater than 4.0 is specified as follows:

```
SELECT ?name
WHERE {
    ?m rdfs:label ?name. ?m movie:director ?d.
    ?d movie:director_name "Stanley Kubrick".
    ?m movie:relatedBook ?b. ?b rev:rating ?r.
    FILTER(?r > 4.0)
}
```

In this query, the first three lines in the WHERE clause form a BGP consisting of five triple patterns. All triple patterns in this example have *variables*, such as "?m", "?name" and "?r", and "?r" has a filter: FILTER(?r > 4.0).

A SPARQL query can also be represented as a *query graph*:

Definition 4 [SPARQL query graph] A *query graph* is a seven-tuple $Q = \langle V^Q, L_V^Q, E^Q, L_E^Q, f_V^Q, f_E^Q, FL \rangle$, where

1. $V^Q = V_c^Q \cup V_e^Q \cup V_l^Q \cup V_p^Q$ is a collection of vertices that correspond to all subjects and objects in a SPARQL query, where V_p^Q is a collection of variable vertices (corresponding to variables in the query expression), and V_c^Q and V_e^Q and V_l^Q are collections of class vertices, entity vertices, and literal vertices in the query graph Q, respectively.

2. E^Q is a collection of edges that correspond to properties in a SPARQL query.

3. L_V^Q is a collection of vertex labels in Q and L_E^Q is the edge labels in E^Q.
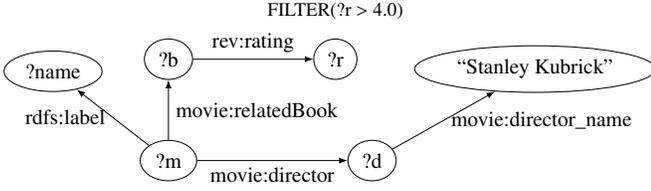
Fig. 3 SPARQL query graph corresponding to query Q_1

4. $f_V^Q : V^Q \to L_V^Q$ is a bijective vertex labeling function that assigns to each vertex in Q a label from L_V^Q. The label of a vertex $v \in V_p^Q$ is the variable; that of a vertex $v \in V_l^Q$ is its literal value; and that of a vertex $v \in V_c^Q \cup V_e^Q$ is its corresponding URI.

5. $f_E^Q : V^Q \to L_E^Q$ is a bijective vertex labeling function that assigns to each edge in Q a label from L_E^Q. An edge label can be a property or an edge variable.

6. FL are constraint filters.

The query graph for Q_1 is given in Figure 3.

The semantics of SPARQL query evaluation can, therefore, be defined as subgraph matching using graph homomorphism whereby all subgraphs of an RDF graph G are found that are homomorphic to the SPARQL query graph Q. In this context, OPT represents the optional triple patterns that may be matched.

Definition 5 [SPARQL graph match] Consider an RDF graph G and a query graph Q that has n vertices $\{v_1, ..., v_n\}$. A set of n distinct vertices $\{u_1, ..., u_n\}$ in G is said to be a *match* of Q, if and only if there exists a bijective function F, where $u_i = F(v_i)$ $(1 \leqslant i \leqslant n)$, such that :

1. If v_i is a literal vertex, v_i and u_i have the same literal value;

2. If v_i is an entity or class vertex, v_i and u_i have the same URI;

3. If v_i is a variable vertex, u_i should satisfy the filter constraint over parameter vertex v_i if any; otherwise, there is no constraint over u_i;

4. If there is an edge from v_i to v_j in Q, there is also an edge from u_i to u_j in G. If the edge label in Q is p (i.e., property), the edge from u_i to u_j in G has the same label. If the edge label in Q is a parameter, the edge label should satisfy the corresponding filter constraint; otherwise, there is no constraint over the edge label from u_i to u_j in G.
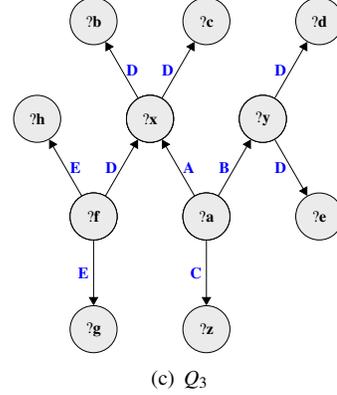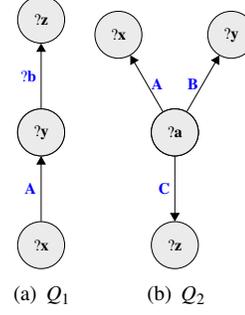

Fig. 4 Sample SPARQL queries

It is usual to talk about SPARQL query types based on the shape of the query graph. Typically, three query types are observed: (i) linear (Figure 4a), where the variable in the object field of one triple pattern appears in the subject of another triple pattern (e.g., ?y in Q_1) (ii) star-shaped (Figure 4b), where the variable in the object field of one triple pattern appears in the subject of multiple other triple patterns (e.g., ?a in Q_2), and (iii) snowflake-shaped (Figure 4c), which is a combination of multiple star queries.

---

3 Data Warehousing Approaches

In this section we consider the approaches that take a centralized approach where the entire data is maintained in one RDF database. These fall into five categories: those that map the RDF data directly into a relational system, those that use a relational schema with extensive indexing (and a native storage system), those that denormalize the triples table into clustered properties, those that use column-store organization, and those that exploit the native graph pattern matching semantics of SPARQL.

Many of the approaches compress the long character strings to integer values using some variation of dictionary encoding in order to avoid expensive string operations. Each string is mapped to an integer in a mapping table, and that integer is then used in the RDF triple table(s). This facilitates fast indexing and access to values, but involves a level of indirection through the mapping table to get at the original strings. Therefore, some of these systems (e.g., Jena [13]) employ encoding only for strings that are longer than a threshold. We ignore encoding in this paper, for clarity of presentation, and represent the data in its original string form.

3.1   Direct Relational Mappings

RDF triples have a natural tabular structure. A direct approach to handle RDF data using relational databases is to create single table with three columns (Subject, Property, Object) that holds the triples (there usually are additional auxiliary tables, but we ignore them here). The SPARQL query can then be translated into SQL and executed on this table. It has been shown that SPARQL 1.0 can be full translated to SQL [14, 15]; whether the same is true for SPARQL 1.1 with its added features is still open. This approach aims to exploit the well-developed relational storage, query processing and optimization techniques in executing SPARQL queries. Systems such as Sesame SQL92SAIL[4] [16] and Oracle [17] follow this approach.

Assuming that the table given in Figure 1 is a relational table, the example SPARQL query given earlier can be translated to the following SQL query (where s,p,o correspond to column names: Subject, Property, Object):

```
SELECT T1.object
FROM   T as T1, T as T2, T as T3,
       T as T4, T as T5
WHERE T1.p="rdfs:label"
AND T2.p="movie:relatedBook"
AND T3.p="movie:director"
AND T4.p="rev:rating"
AND T5.p="movie:director_name"
AND T1.s=T2.s
AND T1.s=T3.s
AND T2.o=T4.s
AND T3.o=T5.s
AND T4.o > 4.0
AND T5.o="Stanley Kubrick"
```

---

[4] Sesame is built to interact with any storage system since it implements a Storage and Inference Layer (SAIL) to interface with the particular storage system on which it sits. SQL92SAIL is the specific instantiation to work on relational systems.

An immediate problem that can be observed with this approach is the high number of self-joins – these are not easy to optimize. Furthermore, in large data sets, this single triples table becomes very large, further complicating query processing.

3.2   Single Table Extensive Indexing

One alternative to the problems created by direct relational mapping is to develop native storage systems that allow extensive indexing of the triple table. Hexastore [18] and RDF-3X [19, 20] are examples of this approach. The single table is maintained, but extensively indexed. For example, RDF-3X creates indexes for all six possible permutations of the subject, property, and object: (spo, sop,ops,ops,sop,pos). Each of these indexes are sorted lexicographically by the first column, followed by the second column, followed by the third column. These are then stored in the leaf pages of a clustered B^+-tree.

The advantage of this type of organization is that SPARQL queries can be efficiently processed regardless of where the variables occur (subject, property, object) since one of the indexes will be applicable. Furthermore, it allows for index-based query processing that eliminates some of the self-joins – they are turned into range queries over the particular index. Even when joins are required, fast merge-join can be used since each index is sorted on the first column. The obvious disadvantages are, of course, the space usage, and the overhead of updating the multiple indexes if data is dynamic.

3.3   Property Tables

Property tables approach exploits the regularity exhibited in RDF datasets where there are repeated occurrence of patterns of statements. Consequently, it stores "related" properties in the same table. The first system that proposed this approach is Jena [13]; IBM's DB2RDF [21] also follows the same strategy. In both of these cases, the resulting tables are mapped to a relational system and the queries are converted to SQL for execution.

Jena defines two types of property tables. The first type, which can be called *clustered property table*, group together the properties that tend to occur in the same (or
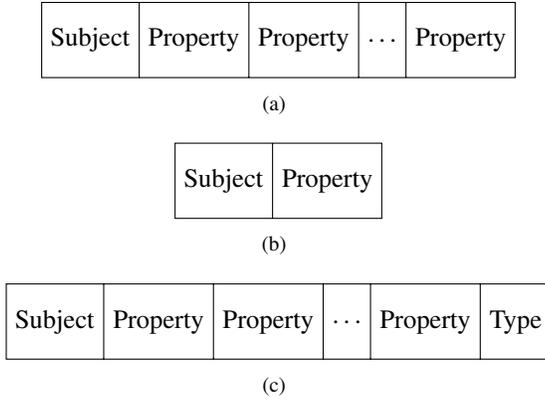
Subject	Property	Property	$\cdots$	Property

(a)

Subject	Property

(b)

Subject	Property	Property	$\cdots$	Property	Type

(c)

Fig. 5 Clustered property table design

Subject	Spill	$Prop_1$	val_1	$Prop_2$	val_2	$\cdots$	$Prop_k$	val_k

(a) DPH

l_id	value

(b) DS

Fig. 7 DB2RDF table design

similar) subjects. It defines different table structures for single-valued properties versus multi-valued properties. For single-valued properties, the table contains the subject column and a number of property columns (Figure 5(a)). The value for a given property may be null if there is no RDF triple that uses the subject and that property. Each row of the table represents a number of RDF triples – the same number as the non-null property values. For these tables, the subject is the primary key. For multi-valued properties, the table structure includes the subject and the multi-valued property (Figure 5(b)). Each row of this table represents a single RDF triple; the key of the table is the compound key (subject,property). The mapping of the single triple table to property tables is a database design problem that is done by a database administrator.

Jena also defines a *property class table* that cluster the subjects with the same *type* of property into one property table (Figure 5(c)). In this case, all members of a class (recall our discussion of class structure within the context of RDFS) together in one table. The "Type" column is the value of rdf:type for each property in that row.

The example dataset in Figure 1 may be organized to create one table that includes the properties of subjects that are films, one table for properties of directors, one table for properties of actors, one table for properties of books and so on. Figure 6 shows one of these tables corresponding to the film subject. Note that the "actor" property is multi-valued (since there are two of them in film/2014), so a separate table is created for it.

IBM DB2RDF [21] also follows the same strategy, but with a more dynamic table organization (Figure 7). The table, called *direct primary hash* (DPH) is organized
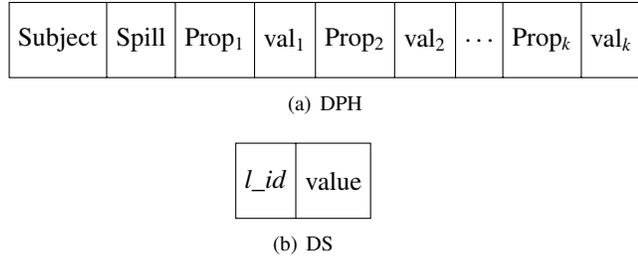
by each subject, but instead of manually identifying "similar" properties, the table accommodates k property columns, each of which can be assigned a different property in different rows. Each property column is, in fact, two columns: one that holds the property label, and the other that holds the value. If the number of properties for a given subject is greater than k, then the extra properties are spilled onto a second row and this is marked on the "spill" column. For multivalued properties, a *direct secondary hash* (DSH) table is maintained – the original property value stores a unique identifier l_id, which appears in the DS table along with the values.

DB2RDF accomplishes the mapping from the single triples table into the DPH and DS tables automatically; the objective is to minimize the number of columns that are used in DPH while minimizing spills (since these cause expensive self-joins) that result from multiple properties being mapped to the same column. Note that, across all subjects, a property is always mapped to the same column; however, a given column can contain more than one property in different rows. The objective of the mapping is to ensure that the columns can be overloaded with properties that do not occur together, but that properties that occur together are assigned to different columns.

The advantage of property table approach is that joins in star queries (i.e., subject-subject joins) become single table scans. Therefore, the translated query has fewer joins. The disadvantages are that in either of the two forms discussed above, there could be a significant number of null values in the tables (see the number of NULLs in Figure 6), and dealing with multivalued properties requires special care. Furthermore, although star queries can be handled efficiently, this approach may not help much with other query types. Finally, when

Subject	label	initial_release_date	director	music_contributor	based_near	relatedBook	language
film/2014	"The Shining"	"1980-05-23"	director/8476	music_contributor/4110	2635167	0743424425	iso639-3/eng
film/2685	"A Clockwork Orange"	NULL	director/8476	NULL	NULL	NULL	NULL
film/424	"Spartakus"	NULL	director/8476	NULL	NULL	NULL	NULL
film/1267	"The Last Tycoon"	NULL	NULL	NULL	NULL	NULL	NULL
film/3418	"The Passenger"	NULL	NULL	NULL	NULL	NULL	NULL

Subject	actor
film/2014	actor/29704
film/2014	actor/30013
film/1267	actor/29704
film/3418	actor/29704

Fig. 6 Property table organization of subject "mdb:film" from the example dataset (prefixes are removed)

manual assignment is used, clustering "similar" properties is non-trivial and bad design decisions exacerbate the null value problem.

3.4 Binary Tables

Binary tables approach [22, 23] follows column-oriented database schema organization and defines a two-column table for each property containing the subject and object. This results in a set of tables each of which are ordered by the subject. This is a typical column-oriented database organization and benefits from the usual advantages of such systems such as reduced I/O due to reading only the needed properties and reduced tuple length, compression due to redundancy in the column values, etc. In addition, it avoids the null values that is experienced in property tables as well as the need for manual or automatic clustering algorithms for "similar" properties, and naturally supports multivalued properties – each become a separate row as in the case of Jena's DS table. Furthermore, since tables are ordered on subjects, subject-subject joins can be implemented using efficient merge-joins. The shortcomings are that the queries require more join operations some of which may be subject-object joins that are not helped by the merge-join operation. Furthermore, insertions into the tables has higher overhead since multiple tables need to be updated. It has been argued that the insertion problem can be mitigated by batch insertions, but in dynamic RDF repositories the difficulty of insertions is likely to remain a significant problem. The proliferation of the number of tables may have a negative impact on the scalability (with respect to the number of properties) of binary tables approach [24].

For example, the binary table representation of the

Subject	Object
film/2014	"The Shining"
film/2685	"A Clockwork Orange"
film/424	"Spartacus"
film/1267	"The Last Tycoon"
film/3418	"The Passenger"
iso639-3/eng	"English"

(a) rdfs:label

Subject	Object
film/2014	actor/29704
film/2014	actor/30013
film/1267	actor/29704
film/3418	actor/29704

(b) movie:actor

Fig. 8 Binary table organization of properties "refs:label" and "movie:actor" from the example dataset (prefixes are removed)

example dataset given in Figure 1 would create one table for each unique property – there are 18 of them. Two of these tables are shown in Figure 8.

3.5 Graph-based Processing

Graph-based RDF processing approaches fundamentally implement the semantics of RDF queries as defined in Section 2. In other words, they maintain the graph structure of the RDF data (using some representation such as adjacency lists), convert the SPARQL query to a query graph, and do subgraph matching using homomorphism to evaluate the query against the RDF graph. Systems such as that proposed by Bönström et al [25], gStore [9, 26], and chameleon-db [27] follow this approach.

The advantage of this approach is that it maintains the original representation of the RDF data and enforces the intended semantics of SPARQL. The disadvantage is the cost of subgraph matching – graph homomorphism is NP-complete. This raises issues with respect to the scalability

of this approach to large RDF graphs; typical database techniques including indexing can be used to address this issue. In the remainder, we present the approach within the context of one system, gStore.

gStore is a graph-based triple store system that can answer different kinds of SPARQL queries – exact queries, queries with wildcards (i.e., where partial information is known about a query object such as knowing the year of birth but not the full birthdate), and aggregate queries that are included in SPARQL 1.1 – over dynamic RDF data repositories. It uses adjacency list representation of graphs. An important feature of gStore is to encode each each entity and class vertex into a fixed length bit string. One motivation for this encoding is to deal with fixed length bit string rather than variable length character strings – this is similar to the dictionary encoding mentioned above. The second, and more important, motivation is to capture the "neighborhood" information for each vertex in the encoding that can be exploited during graph matching. This results in the generation of a *data signature graph* G^*, in which each vertex corresponds to a class or an entity vertex in the RDF graph. Specifically, G^* is induced by all entity and class vertices in the original RDF graph G together with the edges whose endpoints are either entity or class vertices. Figure 9(b) shows the data signature graph G^* that corresponds to RDF graph G in Figure 2. An incoming SPARQL query is also represented as a *query graph Q* that is similarly encoded into a *query signature graph Q^**. The encoding of query graph depicted in Figure 3 into a query signature graph Q_2^* is shown in Figure 9(a).

The problem now turns into finding matches of Q^* over G^*. Although both the RDF graph and the query graph are smaller as a result of encoding, the NP-completeness of the problem remains. Therefore, gStore uses a filter-and-evaluate strategy to reduce the search space over which matching is applied. The objective is to first use a false-positive pruning strategy to find a set of candidate subgraphs (denoted as CL), and then validate these using the adjacency list to find answers (denoted as RS). Accordingly, two issues need to be addressed. First, the encoding technique should guarantee that $RS \subseteq CL$ – the encoding described above provably achieves this. Second, an efficient subgraph matching algorithm is required to find matches of Q^*
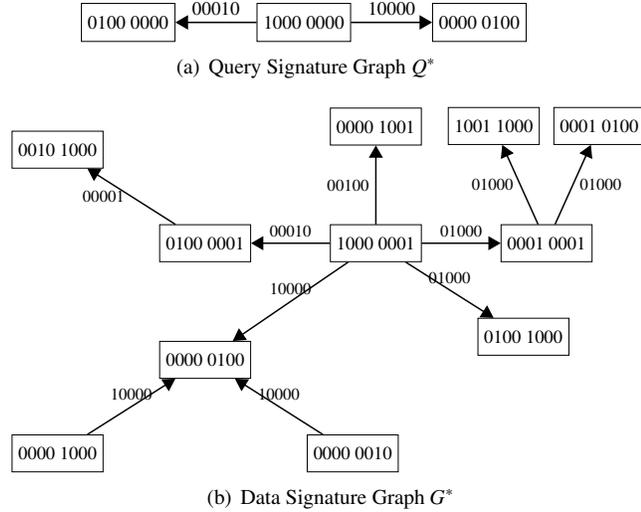

(a) Query Signature Graph Q^*

(b) Data Signature Graph G^*

Fig. 9 Signature graphs

over G^*. For this, gStore uses an index structure called VS*-tree that is a summary graph of G^*. VS*-tree is used to efficiently process queries using a pruning strategy to reduce the search space for finding matches of Q^* over G^*.

---

4 Distributed RDF Processing

In the previous section we focused on centralized, single machine approaches to RDF data management and SPARQL query processing. In this section, we focus on distributed approaches. This section is taken from [28]. We identify and discus four classes of approaches: cloud-based solutions, partitioning-based approaches, federated SPARQL evaluation systems, and partial evaluation-based approach.

4.1 Cloud-based Approaches

There have been a number of works (e.g., [29–36]) focusing on managing large RDF datasets using existing cloud platforms; a very good survey of these is is provided by Saoudi and Manolescu [37]. Many of these approaches follow the MapReduce paradigm; in particular they use HDFS, and store RDF triples in flat files in HDFS. When a SPARQL query is issued, the HDFS files are scanned to find the matches of each triple pattern, which are then joined using one of the MapReduce join implementations (see [38] for more

detailed description of these). The most important difference among these approaches is how the RDF triples are stored in HDFS files; this determines how the triples are accessed and the number of MapReduce jobs. In particular, SHARD [30] directly stores the data in a single file and each line of the file represents all triples associated with a distinct subject. HadoopRDF [31] and PredicateJoin [32] further partition RDF triples based on the property and store each partition within one HDFS file. EAGRE [33] first groups all subjects with similar properties into an entity class, and then constructs a compressed RDF graph containing only entity classes and the connections between them. It partitions the compressed RDF graph using the METIS algorithm [39]. Entities are placed into HDFS according to the partition set that they belong to.

Besides the HDFS-based approaches, there are also some works that use other NoSQL distributed data stores to manage RDF datasets. JenaHBase [29] and H_2RDF [35, 36] use some permutations of subject, property, object to build indices that are then stored in HBase (http://hbase.apache.org). Trinity.RDF [34] uses the distributed memory-cloud graph system Trinity [40] to index and store the RDF graph. It uses hashing on the vertex values to obtain a disjoint partitioning of the RDF graph that is placed on nodes in a cluster.

These approaches benefit from the high scalability and fault-tolerance offered by cloud platforms, but may suffer lower performance due to the difficulties of adapting MapReduce to graph computation.

4.2 Partitioning-based Approaches

The partition-based approaches [41–45] divide an RDF graph G into several fragments and place each at a different site in a parallel/distributed system. Each site hosts a centralized RDF store of some kind. At run time, a SPARQL query Q is decomposed into several subqueries such that each subquery can be answered locally at one site, and the results are then agregated. Each of these papers proposes its own data partitioning strategy, and different partitioning strategies result in different query processing methods.

In GraphPartition [41], an RDF graph G is partitioned into n fragments, and each fragment is extended by including N-hop neighbors of boundary vertices. According to the partitioning strategy, the diameter of the graph corresponding to each decomposed subquery should not be larger than N to enable subquery processing at each local site. WARP [42] uses some frequent structures in workload to further extend the results of GraphPartition. Partout [43] extends the concepts of minterm predicates in relational database systems, and uses the results of minterm predicates as the fragmentation units. Lee et. al. [44] define the partition unit as a vertex and its neighbors, which they call a "vertex block". The vertex blocks are distributed based on a set of heuristic rules. A query is partitioned into blocks that can be executed among all sites in parallel and without any communication. TriAD uses METIS [39] to divide the RDF graph into many partitions and the number of result partitions is much more than the number of sites. Each result partition is considered as a unit and distributed among different sites. At each site, TriAD maintains six large, in-memory vectors of triples, which correspond to all SPO permutations of triples. Meanwhile, TriAD constructs a summary graph to maintain the partitioning information.

All of the above methods implement particular partitioning and distribution strategies that align with their specific requirements. When there is freedom to partition and distribute the data, this works fine, but there are circumstances when partitioning and distribution may be influenced by other requirements. For example, in some applications, the RDF knowledge bases are partitioned according to topics (i.e., different domains), or are partitioned according to different data contributors; in other cases, there may be administrative constraints on the placement of data. In these cases, these approaches may not have the freedom to partition the data as they require. Therefore, partition-tolerant SPARQL processing may be desirable.

4.3 Federated Systems

Federated queries run SPARQL queries over multiple SPARQL endpoints. A typical example is linked data, where different RDF repositories are interconnected, providing a *virtually integrated distributed database*. Federated SPARQL query processing is a very different environment than what we target in this paper, but we discuss these systems for completeness.

A common technique is to precompute metadata for

each individual SPARQL endpoint. Based on the metadata, the original SPARQL query is decomposed into several subqueries, where each subquery is sent to its relevant SPARQL endpoints. The results of subqueries are then joined together to answer the original SPARQL query. In DARQ [46], the metadata is called *service description* that describes which triple patterns (i.e., property) can be answered. In [47], the metadata is called Q-Tree, which is a variant of RTree. Each leaf node in Q-Tree stores a set of source identifers, including one for each source of a triple approximated by the node. SPLENDID [48] uses Vocabulary of Interlinked Datasets (VOID) as the metadata. HiBISCuS [49] relies on "capabilities" to compute the metadata. For each source, HiBISCuS defines a set of capabilities which map the properties to their subject and object authorities. TopFed [50] is a biological federated SPARQL query engine whose metadata comprises of an N3 specification file and a Tissue Source Site to Tumour (TSS-to-Tumour) hash table, which is devised based on the data distribution.

In contrast to these, FedX [51] does not require preprocessing, but sends "SPARQL ASK" to collect the metadata on the fly. Based on the results of "SPARQL ASK" queries, it decomposes the query into subqueries and assign subqueries with relevant SPARQL endpoints.

Global query optimization in this context has also been studied. Most federated query engines employ existing optimizers, such as dynamic programming [52], for optimizing the join order of local queries. Furthermore, DARQ [46] and FedX [51] discuss the use of semijoins to compute a join between intermediate results at the control site and SPARQL endpoints.

4.4 Partial Query Evaluation Approaches

Partial function evaluation is a well-known programming language strategy whose basic idea is the following: given a function $f(s, d)$, where s is the known input and d is the yet unavailable input, the part of f's computation that depends only on s generates a partial answer. This has been used for distributed SPARQL processing in the Distributed gStore system [28].

An RDF graph is partitioned using some graph partitioning algorithm such as METIS [39] (the particular graph partitioning algorithm does not matter as the approach is oblivious to it) into vertex-disjoint fragments (edges that cross fragments are replicated in source and target fragments). Each site receives the full SPARQL query Q and executes it on the local RDF graph fragment providing data parallel computation. In this particular setting, the partial evaluation strategy is applied as follows: each site S_i treats fragment F_i as the known input in the partial evaluation stage; the unavailable input is the rest of the graph ($\overline{G} = G \setminus F_i$).

There are two important issues to be addressed in this framework. The first is to compute the partial evaluation results at each site S_i given a query graph Q – in other words, addressing the graph homomorphism of Q of F_i; this is called the *local partial match* since it finds the matches internal to fragment F_i. Since ensuring edge disjointness is not possible in vertex-disjoint partitioning, , there will be *crossing edges* between graph fragments. The second task is the assembly of these local partial matches to compute crossing matches. Two different assembly strategies are proposed: *centralized assembly*, where all local partial matches are sent to a single site, and *distributed assembly*, where the local partial matches are assembled at a number of sites in parallel.

5 Querying Linked Data

As noted earlier, a major reason for the development of RDF is to facilitate the semantic web. An important aspect of this enterprise is the Web of Linked Data (WLD) that connects multiple web data sets encoded in RDF. In one sense, this is the web data integration, and querying the WLD is an important challenge.

WLD uses RDF to build the semantic web by following four principles:

1. All web resources are locally identified by their URIs.
2. Information about web resources/entities are encoded as RDF triples. In other words, RDF is the semantic web data model.
3. Connections among data sets are established by data links.
4. Sites that host RDF data need to be able to service HTTP requests to serve up linked resources.

Our earlier reference was to Linked Open Data (LOD), which enforces a fifth requirement on WLD:

5. The linked data content should be open.

Examples of LOD datasets include DBpedia and Freebase. The following formalizes Web of Linked Data, ignoring the openness requirement.

The starting point is a Linked Document (LD) that is a web document with embedded RDF triples that encode web resources. These web documents are possibly interconnected to get the graph structure.

Definition 6 [Web of Linked Data] [11] Given an set $\mathcal{D}$ of Linked Documents (LD), a Web of Linked data is a tuple $W = (D, adoc, data)$ where:

- $D \subseteq \mathcal{D}$,
- *adoc* is a partial mapping from URIs to D, and
- *data* is a total mapping from D to finite sets of RDF triples.

Each of these documents may contain RDF triples that form *data links* to other documents, which is formalized as follows.

Definition 7 [Data Link] [11] A Web of Linked Data $W = (D, adoc, data)$ (as defined in Definition 6) contains a data link from document $d \in D$ to document $d' \in D$ if there exists a URI u such that

- u is mentioned in an RDF triple $t \in data(d)$, and
- $d' = adoc(u)$

The semantics of SPARQL queries over the WLD becomes tricky. One possibility is to adopt *full web semantics* that specifies the scope of evaluating a SPARQL query expression to be all linked data. There is no known (terminating) query execution algorithm that can guarantee result completeness under this semantics. The alternative is a family of *reachability-based semantics* that define the scope of evaluating a SPARQL query in terms of the documents that can be reached: given a set of seed URIs and a reachability condition, the scope is all data along the paths of the data links from the seeds and that satisfy the reachability condition. The family is defined by different reachability conditions. In this case, there are computationally feasible algorithms.

There are three approaches to SPARQL query execution over WLD [53]: traversal-based, index-based, and hybrid. *Traversal approaches* [54, 55] basically implement a reachability-based semantics: starting from seed URIs, they recursively discover relevant URIs by traversing specific data links at query execution runtime. For these algorithms, the selection of the seed URIs is critical for performance. The advantage of traversal approaches is their simplicity (to implement) since they do not need to maintain any data structures (such as indexes). The disadvantages are the latency of query execution since these algorithms "browse" web documents, and repeated data retrieval from each document introduces significant latency. They also have limited possibility for parallelization – they can be parallelized to the same extent that crawling algorithms can.

The *index-based approaches* use an index to determine relevant URIs, thereby reducing the number of linked documents that need to be accessed. A reasonable index key is triple patterns [56] in which case the "relevant" URIs for a given query are determined by accessing the index, and the query is evaluated over the data retrieved by accessing those URIs. In these approaches, data retrieval can be fully parallelized, which reduces the negative impact of data retrieval on query execution time. The disadvantages of the approach are the dependence on the index – both in terms of the latency that index construction introduces and in terms of the restriction the index imposes on what can be selected – and the freshness issues that result from the dynamicity of the web and the difficulty of keeping the index up-to-date.

Hybrid approaches [57] perform a traversal-based execution using prioritized listing of URIs for look-up. The initial seeds come from a pre-populated index; new discovered URIs that are not in the index are ranked according to number of referring documents.

6 Conclusions

In this paper, we gave a high level overview of RDF data management, focusing on the various approaches that have been adopted. The discussion focused on centralized RDF data management (what is called "data warehousing approach" in this paper), distributed RDF systems, and querying over the LOD data. There are many additional works on RDF that are omitted in this paper. Most notably, works on formal semantics of SPARQL, reasoning over RDF data, data integration using RDF, streaming RDF processing, and SPARQL

processing under dynamic workloads [27] are topics that are not covered.

Acknowledgements

References

1. Suchanek F M, Kasneci G, Weikum G. Yago: a core of semantic knowledge. In: Proc. 16th Int. World Wide Web Conf. 2007, 697–706

2. Bizer C, Lehmann J, Kobilarov G, Auer S, Becker C, Cyganiak R, Hellmann S. Dbpedia - a crystallization point for the web of data. J. Web Sem., 2009, 7(3): 154–165

3. Schmachtenberg M, Bizer C, Paulheim H. Adoption of best data practices in different topical domains. In: Proc. 13th Int. Semantic Web Conf. 2014, 245–260

4. Zhang Y, Duc P M, Corcho O, Calbimonte J P. SRBench: A streaming RDF/SPARQL benchmark. In: Proc. 11th Int. Semantic Web Conf. 2012, 641–657

5. Zaveri A, Rula A, Maurino A, Pietrobon R, Lehmann J, Auer S. Quality assessment for linked data: A survey. Web Semantics J., 2012, 1–5

6. Tang N. Big RDF data cleaning. In: Proc. Workshops of 31st Int. Conf. on Data Engineering. 2015, 77–79

7. Klyne G, Carroll J J, McBride B. RDF 1.1 concepts and abstract syntax. Available at http://www.w3.org/TR/rdf11-concepts/; last accessed 17 November 2015

8. Harris S, Seaborne A. SPARQL 1.1 query language. Available at http://www.w3.org/TR/sparql11-query/; last accessed 17 November 2015

9. Zou L, Özsu M T, Chen L, Shen X, Huang R, Zhao D. gStore: A graph-based SPARQL query engine. VLDB J., 2014, 23(4): 565–590

10. Hartig O, Özsu M T. Reachable subwebs for traversal-based query execution. 2014, 541–546 (Companion Volume)

11. Hartig O. SPARQL for a web of linked data: Semantics and computability. In: Proc. 9th Extended Semantic Web Conf. 2012, 8–23

12. W3C . SPARQL query language for RDF – Formal definitions. Accessible at http://www.w3.org/2001/sw/DataAccess/rq23/sparql-defns.html#defn_GroupGraphPattern, 2006. last accessed 21 December 2015

13. Wilkinson K. Jena property table implementation. Technical Report HPL-2006-140, HP Laboratories Palo Alto, October 2006

14. Angles R, Gutierrez C. The expressive power of SPARQL. In: Proc. 7th Int. Semantic Web Conf. 2008, 114–129

15. Sequeda J F, Arenas M, Miranker D P. OBDA: query rewriting or materialization? in practice, both! In: Proc. 13th Int. Semantic Web Conf. 2014, 535—551

16. Broekstra J, Kampman A, Harmelen v F. Sesame: A generic architecture for storing and querying RDF and RDF schema. In: Proc. 1st Int. Semantic Web Conf. 2002, 54–68

17. Chong E, Das S, Eadon G, Srinivasan J. An efficient SQL-based RDF querying scheme. In: Proc. 31st Int. Conf. on Very Large Data Bases. 2005, 1216–1227

18. Weiss C, Karras P, Bernstein A. Hexastore: sextuple indexing for semantic web data management. Proc. VLDB Endowment, 2008, 1(1): 1008–1019

19. Neumann T, Weikum G. RDF-3X: a RISC-style engine for RDF. Proc. VLDB Endowment, 2008, 1(1): 647–659

20. Neumann T, Weikum G. The RDF-3X engine for scalable management of RDF data. VLDB J., 2009, 19(1): 91–113

21. Bornea M A, Dolby J, Kementsietsidis A, Srinivas K, Dantressangle P, Udrea O, Bhattacharjee B. Building an efficient RDF store over a relational database. In: Proc. ACM SIGMOD Int. Conf. on Management of Data. 2013, 121–132

22. Abadi D J, Marcus A, Madden S R, Hollenbach K. Scalable semantic web data management using vertical partitioning. In: Proc. 33rd Int. Conf. on Very Large Data Bases. 2007, 411–422

23. Abadi D J, Marcus A, Madden S, Hollenbach K. SW-Store: a vertically partitioned DBMS for semantic web data management. VLDB J., 2009, 18(2): 385–406

24. Sidirourgos L, Goncalves R, Kersten M, Nes N, Manegold S. Column-store support for RDF data management: not all swans are white. Proc. VLDB Endowment, 2008, 1(2): 1553–1563

25. Bönström V, Hinze A, Schweppe H. Storing RDF as a graph. In: Proc. 1st Latin American Web Congress. 2003, 27 – 36

26. Zou L, Mo J, Chen L, Özsu M T, Zhao D. gStore: answering SPARQL queries via subgraph matching. Proc. VLDB Endowment, 2011, 4(8): 482–493

27. Aluç G. Workload Matters: A Robust Approach to Physical RDF Database Design. PhD thesis, University of Waterloo, 2015

28. Peng P, Zou L, Özsu M T, Chen L, Zhao D. Processing SPARQL queries over distributed RDF graphs. VLDB J., 2015. Forthcoming.

29. Khadilkar V, Kantarcioglu M, Thuraisingham B M, Castagna P. Jena-HBase: A distributed, scalable and efficient RDF triple store. In: Proc. International Semantic Web Conference Posters & Demos Track. 2012

30. Rohloff K, Schantz R E. High-performance, massively scalable distributed systems using the mapreduce software framework: the shard triple-store. In: Proc. Int. Workshop on Programming Support Innovations for Emerging Distributed Applications. 2010. Article No. 4

31. Husain M F, McGlothlin J, Masud M M, Khan L R, Thuraisingham B. Heuristics-based query processing for large RDF graphs using cloud computing. IEEE Trans. Knowl. and Data Eng., 2011, 23(9): 1312–1327

32. Zhang X, Chen L, Wang M. Towards efficient join processing over large RDF graph using mapreduce. In: Proc. 24th Int. Conf. on Scientific and Statistical Database Management. 2012, 250–259

33. Zhang X, Chen L, Tong Y, Wang M. EAGRE: Towards scalable I/O efficient SPARQL query evaluation on the cloud. In: Proc. 29th Int. Conf. on Data Engineering. 2013, 565–576

34. Zeng K, Yang J, Wang H, Shao B, Wang Z. A distributed graph engine for web scale RDF data. Proc. VLDB Endowment, 2013, 6(4): 265–276

35. Papailiou N, Konstantinou I, Tsoumakos D, Koziris N. H_2RDF: adaptive query processing on RDF data in the cloud. 2012, 397–400 (Companion Volume)

36. Papailiou N, Tsoumakos D, Konstantinou I, Karras P, Koziris N. H_2RDF+: an efficient data management system for big RDF graphs. In: Proc. ACM SIGMOD Int. Conf. on Management of Data. 2014, 909–912

37. Kaoudi Z, Manolescu I. RDF in the clouds: A survey. VLDB J., 2015, 24: 67–91

38. Li F, Ooi B C, Özsu M T, Wu S. Distributed data management using MapReduce. ACM Comput. Surv., 2014, 46(3): Article No. 31

39. Karypis G, Kumar V. Analysis of multilevel graph partitioning. 1995. Article No. 29

40. Shao B, Wang H, Li Y. Trinity: a distributed graph engine on a memory cloud. In: Proc. ACM SIGMOD Int. Conf. on Management of Data. 2013, 505–516

41. Huang J, Abadi D J, Ren K. Scalable SPARQL querying of large RDF graphs. Proc. VLDB Endowment, 2011, 4(11): 1123–1134

42. Hose K, Schenkel R. WARP: Workload-aware replication and partitioning for RDF. In: Proc. Workshops of 29th Int. Conf. on Data Engineering. 2013, 1–6

43. Galarraga L, Hose K, Schenkel R. Partout: a distributed engine for efficient RDF processing. 2014, 267–268 (Companion Volume)

44. Lee K, Liu L. Scaling queries over big rdf graphs with semantic hash partitioning. Proc. VLDB Endowment, 2013, 6(14): 1894–1905

45. Gurajada S, Seufert S, Miliaraki I, Theobald M. TriAD: A distributed shared-nothing RDF engine based on asynchronous message passing. In: Proc. ACM SIGMOD Int. Conf. on Management of Data. 2014, 289–300

46. Quilitz B. Querying distributed RDF data sources with SPARQL. In: Proc. 5th European Semantic Web Conf. 2008, 524–538

47. Harth A, Hose K, Karnstedt M, Polleres A, Sattler K, Umbrich J. Data summaries for on-demand queries over linked data. In: Proc. 19th Int. World Wide Web Conf. 2010, 411–420

48. Görlitz O, Staab S. SPLENDID: SPARQL endpoint federation exploiting VOID descriptions. In: Proc. ISWC 2011 Workshop on Consuming Linked Data. 2011

49. Saleem M, Ngomo A N. HiBISCuS: Hypergraph-based source selection for SPARQL endpoint federation. In: Proc. 11th Extended Semantic Web Conf. 2014, 176–191

50. Saleem M, Padmanabhuni S S, Ngomo A N, Iqbal A, Almeida J S, Decker S, Deus H F. TopFed: TCGA tailored federated query processing and linking to LOD. J. Biomedical Semantics, 2014, 5: 47

51. Schwarte A, Haase P, Hose K, Schenkel R, Schmidt M. FedX: Optimization techniques for federated query processing on linked data. In: Proc. 10th Int. Semantic Web Conf. 2011, 601–616

52. Astrahan M, Blasgen M, Chamberlin D, Eswaran K, Gray J, Griffiths P, King W, Lorie R, McJones P, Mehl J, Putzolu G, Traiger I, Wade B, Watson V. System r: Relational approach to database management. ACM Trans. Database Syst., 1976, 1(2): 97–137

53. Hartig O. An overview on execution strategies for linked data queries. Datenbankspektrum, 2013, 13(2): 89–99

54. Hartig O. SQUIN: a traversal based query execution system for the web of linked data. In: Proc. ACM SIGMOD Int. Conf. on Management of Data. 2013, 1081–1084

55. Ladwig G, Tran T. SIHJoin: Querying remote and local linked data. In: Proc. 8th Extended Semantic Web Conf. 2011, 139–153

56. Umbrich J, Hose K, Karnstedt M, Harth A, Polleres A. Comparing data summaries for processing live queries over linked data. World Wide Web J., 2011, 14(5-6): 495–544

57. Ladwig G, Tran T. Linked data query processing strategies. In: Proc. 9th Int. Semantic Web Conf. 2010, 453–469

M. Tamer Özsu is Professor of Computer Science at the University of Waterloo. Dr. Özsu's current research focuses on large scale data distribution, and management of unconventional data (e.g., graphs, RDF, XML, streams).

He is a Fellow of ACM and IEEE, an elected member of the Science Academy of Turkey, and a member of Sigma Xi and AAAS.