



HAL
open science

Federated SPARQL Queries Processing with Replicated Fragments

Gabriela Montoya, Hala Skaf-Molli, Pascal Molli, Maria-Esther Vidal

► **To cite this version:**

Gabriela Montoya, Hala Skaf-Molli, Pascal Molli, Maria-Esther Vidal. Federated SPARQL Queries Processing with Replicated Fragments. The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Oct 2015, Bethlehem, United States. hal-01169601

HAL Id: hal-01169601

<https://nantes-universite.hal.science/hal-01169601v1>

Submitted on 29 Jun 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Federated SPARQL Queries Processing with Replicated Fragments

Gabriela Montoya^{1,2}, Hala Skaf-Molli¹, Pascal Molli¹, and Maria-Esther Vidal³

¹ LINA – UFR de Sciences et Techniques, Nantes University
{first.last@univ-nantes.fr}

² Unit UMR6241 of the Centre National de la Recherche Scientifique (CNRS), France

³ Universidad Simón Bolívar, Venezuela mvidal@ldc.usb.ve

Abstract. Federated query engines allow to consume linked data from SPARQL endpoints. Replicating data fragments from different sources allows to re-organize data to better fit federated query processing of data consumers. However, existing federated query engines poorly support replication. In this paper, we propose a replication-aware federated query engine that extends state-of-art federated query engine ANAPSID and FedX with FEDRA, a source selection strategy that approximates the source selection problem with fragments replication (SSP-FR). For a given set of endpoints with replicated fragments and a SPARQL query, the problem is to find the endpoints to contact in order to minimize the number of tuples to transfer from endpoints to the federated query engines. We devise the FEDRA source selection algorithm that approximates SSP-FR. We implement FEDRA in the state-of-the-art federated query engines FedX and ANAPSID, and empirically evaluate their performance. Experimental results suggest that FEDRA efficiently solves SSP-FR, reducing the number of selected SPARQL endpoints as well as the size of query intermediate results.

Keywords: Linked Data, Federated Query Processing, Source Selection, Fragment Replication

1 Introduction

SPARQL endpoints enable to consume RDF data exploiting the *expressiveness* of the SPARQL query language. Nevertheless, recent studies reveal that existing public SPARQL endpoints main limitation is availability [4].

In distributed databases [16], a common practice to overcome availability problems is to replicate data near data consumers. Replication can be achieved by complete dataset replication, active caching, pre-fetching or fragmentation [12].

RDF data consumers can replicate *subsets* of RDF datasets or *replicated fragments*, and make them accessible through SPARQL endpoints. This will provide the support for an efficient RDF data re-organization according to the *needs* and *computational resource capacity* of data consumers, while these data can be still *accessed* using SPARQL endpoints. Unfortunately, although SPARQL endpoints can *transparently* access replicated fragments, as well as maintain their

consistency [12], federated query engines are not tailored to exploit the benefits of replicated fragments.

Federated SPARQL engines [1, 5, 8, 17, 20] allow data consumers to execute SPARQL queries against a federation of SPARQL endpoints. However, these engines are just designed to select the SPARQL endpoints that ensure both a complete answer and an efficient execution of the query. In presence of replication, existing federated query engines may retrieve data from every relevant endpoint, and produce a large number of intermediate results that trigger many requests to the endpoints. Thus, federated query engines may exhibit poor *performance* while *availability* of the selected SPARQL endpoints is negatively impacted.

Although the problem of managing RDF data *overlapping* during federated query processing has been addressed in [11, 18], the problem of managing *replication* in a federation of RDF datasets still remains open. DAW [18] is able to detect overlapping between datasets and optimize source selection based on that. However, because DAW is not designed to manage data replication, there is no support for explicitly define and use replicated fragments. In consequence, DAW may select redundant data sources and generate a high number of intermediate results as we will report in our experiments.

In this paper, we build a replication-aware SPARQL federated query engine by integrating into state-of-the art federated query engines FedX [20] and ANAPSID [1], a source selection strategy called FEDRA that solves the source selection problem with fragment replication (SSP-FR). For a given set of SPARQL endpoints with replicated fragments and a SPARQL query, the problem is to minimize the number of transferred data from endpoints to the federated query engines, while preserving answer completeness and reducing data redundancy.

We empirically study federated query engines FedX and ANAPSID extended with FEDRA and DAW on synthetic and real datasets. The results suggest that FEDRA efficiently reduces intermediate results and data redundancy.

The paper is organized as follows. Section 2 describes background and motivations. Section 3 defines replicated fragments and presents the source selection problem for fragment replication. Section 4 presents the FEDRA source selection algorithm. Section 5 reports our experimental results. Section 6 summarizes related works. Finally, conclusions and future works are outlined in Section 7.

2 Background and Motivations

Existing SPARQL federated query engines do not support replicated data. To illustrate, we replicated the DBpedia dataset ⁴ and defined two federations. The first is composed of one mirror of DBpedia, and the second of two identical mirrors of DBpedia. We used FedX [20] and ANAPSID [1] to execute the query in Figure 1a against both federations. In the first federation, these engines produced the same query answers. On the other hand, for the second federation, these query engines have no knowledge about the relationships among the mirrors of

⁴ DBpedia Live at August 15th, 2013.

<p>(a) DBpedia Query</p> <pre> select distinct ?p ?m ?n ?d where { ?p dbprop:name ?m . ?p dbprop:nationality ?n . ?p dbprop:doctoralAdvisor ?d } </pre>	<p>(b) Query Execution</p> <table border="1"> <thead> <tr> <th rowspan="2">#DBpedia Replicas</th> <th colspan="2">Execution Time (ms)</th> <th colspan="2"># Results</th> </tr> <tr> <th>FedX</th> <th>ANAPSID</th> <th>FedX</th> <th>ANAPSID</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>1,392</td> <td>22,972</td> <td>8,921</td> <td>8,921</td> </tr> <tr> <td>2</td> <td>215,907</td> <td>$\geq 1,800,000$</td> <td>418</td> <td>8,921</td> </tr> </tbody> </table>	#DBpedia Replicas	Execution Time (ms)		# Results		FedX	ANAPSID	FedX	ANAPSID	1	1,392	22,972	8,921	8,921	2	215,907	$\geq 1,800,000$	418	8,921
#DBpedia Replicas	Execution Time (ms)		# Results																	
	FedX	ANAPSID	FedX	ANAPSID																
1	1,392	22,972	8,921	8,921																
2	215,907	$\geq 1,800,000$	418	8,921																

Fig. 1: DBpedia query and its execution time and number of results over one and two replicas of DBpedia for FedX and ANAPSID

DBpedia, and they contact both data sources. In this way, performance in terms of execution time and number of results, is seriously degraded as depicted in Figure 1b⁵.

Furthermore, if the DAW approach were used, data providers would have to compute summaries of DBpedia; this computation may take up to six hours and the size of the obtained index is 89Mo. Then, after downloading each DBpedia summary, DAW would detect the overlap between the two data sources. DAW could select different DBpedia data sources per triple pattern, and execute thus the join between retrieved data at the federated engine level.

Of course, if federated query engines would know that one endpoint is the mirror of the other, the source selection pruning could be done more efficiently, i.e., only one source were selected to execute the query. This problem is even more challenging if we consider that one endpoint can partially replicate data from several RDF datasets, i.e., replicate a dataset fragment.

Suppose a Web application performs federated queries on $A1$ (DBpedia) and $A2$ (LinkedMDB). In order to speed up the queries, a data consumer endpoint $C1$ with replicated fragments has been installed as in Figure 2. Fragments are defined as simple CONSTRUCT SPARQL queries with one triple pattern. Fragments allow re-organizing the RDF data *on* $C1$ to better address data consumer’s needs.

Even in this simple setup, processing our running query on a federation including $A1$, $A2$, and $C1$ raises the problem of source selection with fragment replication (SSP-FR). There are five options to select sources for executing this query; these choices produce different number of transferred tuples as shown in Figure 2: (i) If no information about replicated fragments is available, all sources may be selected to retrieve data for all the triple patterns. The number of intermediate results is given in the solution $s1$. This will be the behavior of a federated query engine like FedX that ensures answer completeness⁶. (ii) Endpoints $A1$ and $A2$ could be chosen, in this case the number of intermediate results is given in $s2$. The number of intermediate results in $s2$ is less than $s1$ since some joins could be executed at $A1$ and $A2$. (iii) Another choice may be to use the $C1$ endpoint in combination with either $A1$ or $A2$ ($s3$, $s4$). This produces the same

⁵ FedX retrieves less results with two mirrors of DBpedia because it reaches the endpoints maximum number of result rows

⁶ In order to preserve joins between different endpoints, each triple pattern should be posed to each endpoint individually.




Fig. 2: Client defines a federation composed of DBpedia (A1), LinkedMDB (A2), and one consumer (C1) endpoints with four replicated fragments.

number of intermediate results as in $s2$, but they have the advantage of accessing less public endpoints. (iv) A last choice could be to use the $C1$ endpoint to retrieve data for all the triple patterns ($s5$). This solution profits from replicated fragments to execute opportunistic joins at $C1$; thus, it is able to achieve the *best* performance in terms of the number of intermediate results.

As the number of transferred tuples increases, the availability of the contacted SPARQL endpoints can be affected. A replication aware federated query engine could select the best sources to reduce the size of intermediate results while preserving answer completeness. In this paper, we formally address the following problem: Given a SPARQL query and a set of relevant SPARQL endpoints with replicated fragments, choose the SPARQL endpoints to contact in order to produce a complete query answer and transfer the minimum amount of data? We aim to develop an algorithm that produces solution $s5$ whenever possible, providing as output the sources to be used by a federated query engine.

3 Definitions and Problem Description

This section introduces definitions and the source selection problem with fragment replication (SSP-FR).

3.1 Definitions

Fragments are used to replicate RDF data. The data of a fragment is defined by means of the dataset public endpoint, or *authoritative endpoint*, and a CONSTRUCT query with *one* triple pattern.




Fig. 3: Client defines a federation composed of $C1, C2$, and $C3$ that replicates fragments $f2 - f7$

Definition 1 (Fragment). A fragment is a tuple $f = \langle u, s \rangle$

- u is the non-null URI of the authoritative endpoint where f is available;
- s is a *CONSTRUCT* query with one triple pattern.

Without loss of generality, s is limited to one triple pattern as in [12, 22]; this reduces the complexity of *fragment containment* problem as described in Definition 2. Additionally, we assume replicated fragments comprise RDF data accessible from public endpoints, i.e., the authoritative endpoints of the replicated fragments are disjoint with data consumer endpoints. This will allow data consumers to re-organize RDF data replicated from different public endpoints to fit in this way, their needs and requirements.

In this work, we make the following hypotheses: (i) Fragments are replicated from public endpoints, and there is just one level of replication. (ii) Fragments are read-only and perfectly synchronized; the fragment synchronization problem is studied in [12], while querying fragments with divergence is addressed in [15]. (iii) For the sake of simplicity, we suppose that RDF data accessible through the endpoints are described as fragments.

To illustrate, consider the federation given in Figure 3. This federation extends the setup in Figure 2. Suppose three Web applications perform queries against DBpedia and LinkedMDB. To speed up query processing, data consumer endpoints: $C1$, $C2$, and $C3$ with replicated fragments have been configured.

At startup, the federated query engine loads the fragments description for each of the federation endpoints, and computes both the *fragment* and *containment* mappings. The **fragment mappings** is a function that maps fragments to a set of endpoints; the **containment mapping** is based on *containment* relation ($f_l \sqsubseteq f_k$) described in the Definition 2.

Two fragments loaded from two different endpoints C_i, C_j that have the same authoritative endpoint and equivalent construct queries are concatenated in the fragment mapping. For example, the federated engine loads fragments $\langle ?film\ db:director\ ?director, http://dbpedia.org/sparql \rangle$ from C_1, C_2, C_3 , computes equivalence, and adds in its fragment mappings $\langle ?film\ db:director\ ?director, http://dbpedia.org/sparql \rangle \rightarrow \{C_1, C_2, C_3\}$.

Query containment and equivalence have been studied extensively. We adapt the definition given in [10] for the case of a triple pattern query.

Definition 2 (Triple Pattern Containment). *Let $TP(D)$ denote the result of execution of the triple pattern TP over an RDF dataset D . Let TP_1 and TP_2 be two triple patterns. We say that TP_1 is contained in TP_2 , denoted by $TP_1 \sqsubseteq TP_2$, if for any RDF dataset D , $TP_1(D) \subseteq TP_2(D)$. We say that TP_1 is equivalent to TP_2 denotes $TP_1 \equiv TP_2$ if $TP_1 \sqsubseteq TP_2$ and $TP_2 \sqsubseteq TP_1$.*

In the case of triple patterns, testing containment⁷ amounts to finding a substitution of the variables in the triple patterns⁸. $TP_1 \sqsubseteq TP_2$, iff there is a substitution θ such that applying θ to TP_2 returns the triple pattern TP_1 . Testing triple pattern containment has a complexity of $O(1)$. Solving the decision problem of triple pattern containment between TP_1 and TP_2 , $TP_1 \sqsubseteq TP_2$, requires to check if TP_1 imposes at least the same restrictions as TP_2 on the subject, predicate, and object positions, i.e., TP_1 should have at most the same number of variables as TP_2 .

For the federation in Figure 3, $f_5 \sqsubseteq f_4$ because f_4 and f_5 share the same authoritative endpoint and there is a substitution θ defined as $\{(?genre, film_genre : 14), (?movie, ?movie)\}$ and applying θ to f_4 returns f_5 . After identifying a substitution θ for all pair-wise fragments, it is straightforward to compute a **containment mapping** for a federation of SPARQL endpoints.

We can rely on fragment descriptions and the containment property to determine relevant fragments to a query. Relevant fragments contain relevant RDF data to each of the triple patterns of the query. A fragment is *relevant* to a query Q , if it is relevant to *at least one* triple pattern of the query.

Definition 3 (Fragment relevance). *Let f be a fragment defined by a triple pattern TP_1 . Let TP_2 be a triple pattern of a query Q . f is relevant to Q if $TP_2 \sqsubseteq TP_1$ or $TP_1 \sqsubseteq TP_2$.*

Table 1a shows the relevant fragments to the triple patterns in query Q , and the endpoints that provide these fragments. For example, the triple pattern $tp1$

⁷ Containment testing is adapted from [9].

⁸ The substitution operator preserves URIs and literals, i.e., only variables are substituted.

Q triple pattern	RF	Endpoints
tp ₁ ?director dbo:nationality ?nat	f6	C1
	f7	C2
tp ₂ ?film dbo:director ?director	f2	C1, C2, C3
tp ₃ ?movie owl:sameAs ?film	f3	C2, C3
tp ₄ ?movie linkedmdb:genre ?genre	f4	C1, C3
	f5	C2

(a) Relevant Fragments to Q

TP	$D_0(tp)$	$D_1(tp)$	$D_2(tp)$
tp ₁	{ C1,C2 }	{ C1,C2 }	{ C1,C2 }
tp ₂	{ C1,C2,C3 }	{ C1 }	{ C3 }
tp ₃	{ C2,C3 }	{ C2 }	{ C3 }
tp ₄	{ C1,C2,C3 }	{ C3 }	{ C3 }
Triples to transfer	421,675	170,078	8,953

(b) Answer completeness preservation

Table 1: SSP-FR for query Q over a federation of $C1$, $C2$, and $C3$ of Figure 3.

has two relevant fragments: $f6$ and $f7$, and triple pattern $tp4$ has two relevant fragments: $f4$ and $f5$. Fragment $f4$ can produce the complete answer of $tp4$ because $f5 \sqsubseteq f4$, while both $f6$ and $f7$ are required to answer $tp1$.

3.2 Source Selection Problem with Fragment Replication (SSP-FR)

Given a SPARQL query Q , a set of SPARQL endpoints E , the set of fragments F that have been replicated by at least one endpoint in E , a fragment mapping $endpoints()$, a containment mapping \sqsubseteq .

The Source Selection Problem with Fragment Replication (SSP-FR) is to assign to each triple pattern in Q , the set of endpoints from E that need to be contacted to answer Q . A solution of SSP-FR corresponds to a mapping D that satisfies the following properties:

1. **Answer completeness preservation:** sources selected in D do not reduce the query engine answer completeness.
2. **Data redundancy minimization:** $cardinality(D(tp))$ is minimized for all triple pattern tp in Q , i.e., redundant data is minimized.
3. **Data transfer minimization:** executing the query using the sources selected in D minimizes the number of transferred data.

We illustrate SSP-FR on running query Q of Figure 3. Table 1a presents relevant fragments for each triple pattern. Table 1b shows three $D(tp)$ that ensure the completeness preservation property. Even if D_1 and D_2 minimize the number of selected endpoints per triple pattern, only D_2 minimizes the transferred data. Indeed, executing $tp1, tp2, tp3$ against replicated fragments that are located in the same data consumer endpoint will greatly reduce the size of intermediate results.

The approach proposed by Saleem et al. [18] is not designed for solving SSP-FR. Indeed, it does not take into account replicated data, and may produce a solution as D_1 . The FEDRA algorithm exploits properties of the replicated fragments and is able to find solution D_2 .

Algorithm 1 FEDRA Source Selection algorithm

Require: Q : SPARQL Query; F : set of Fragments; $\text{endpoints} : \text{Fragment} \rightarrow \text{set of Endpoint}$; $\sqsubseteq : \text{triplePattern} \times \text{triplePattern}$
Ensure: selectedEndpoints : map from Triple Pattern to set of Endpoint.
1: **function** SOURCESELECTION($Q, F, \text{endpoints}, \sqsubseteq$)
2: $\text{triplePatterns} \leftarrow$ get triple patterns in Q
3: **for each** $tp \in \text{triplePatterns}$ **do**
4: $R(tp) \leftarrow \text{RELEVANTFRAGMENTS}(tp, F)$ \triangleright Relevant fragments as in Definition 3
5: $R(tp) \leftarrow \{ \{f : f \in R(tp) : tp \sqsubseteq f\} \cup \{ \{f\} : f \in R(tp) : f \sqsubseteq tp \wedge \neg(\exists g : g \in R(tp) : f \sqsubset g \sqsubseteq tp) \}$
6: $E(tp) \leftarrow \{ (\bigcup \text{endpoints}(f) : f \in fs) : fs \in R(tp) \}$
7: $\text{basicGP} \leftarrow$ get basic graph patterns in Q
8: **for each** $\text{bgp} \in \text{basicGP}$ **do**
9: $\text{UNIONREDUCTION}(\text{bgp}, E)$ \triangleright endpoints reduction for multiple fragments triples
10: $\text{BGPREDUCTION}(\text{bgp}, E)$ \triangleright endpoints reduction for the bgp triples
11: **for each** $tp \in \text{domain}(E)$ **do**
12: $\text{selectedEndpoints}(tp) \leftarrow$ for each set in $E(tp)$ include one element
13: **return** selectedEndpoints

Algorithm 2 Union reduction algorithm

Require: tps : set of TriplePattern; E : mapping from TriplePattern to set of set of Endpoint
14: **procedure** UNIONREDUCTION(tps, E)
15: $\text{triplesWithMultipleFragments} \leftarrow \{ tp : tp \in \text{tps} \wedge \text{cardinality}(E(tp)) > 1 \}$
16: **for each** $tp \in \text{triplesWithMultipleFragments}$ **do**
17: $\text{commonSources} \leftarrow (\bigcap f : f \in E(tp))$ \triangleright get sources in all subsets in $E(tp)$
18: **if** $\text{commonSources} \neq \emptyset$ **then**
19: $E(tp) \leftarrow \{ \text{commonSources} \}$

4 Fedra: an Algorithm for SSP-FR

The goal of FEDRA is to reduce data transfer by taking advantage of the replication of relevant fragments for several triple patterns on the same endpoint. Algorithm 1 proceeds in four main steps: I. Identify relevant fragments for triple patterns, a triple pattern of a Basic Graph Pattern (BGP) can be contained in one fragment or a union of fragments (cf lines 4-5). II. Localize relevant replicated fragments on the endpoints, e.g., for Q : $tp1_{C1 \cup C2} \bowtie tp2_{C1, C2, C3} \bowtie tp3_{C2, C3} \bowtie tp4_{C1, C3}$ (line 6). III. Prune endpoints for the unions (line 9). IV. Prune endpoints for the BGPs using a set covering heuristic (line 10).

Next, we illustrate how Algorithm 1 works on our running query Q and data consumer endpoints $C1, C2, C3$ from Figure 3⁹.

First, for each triple pattern, FEDRA computes relevant fragments in $R(tp)$, and groups them if they provide the same relevant data. For $tp1$, $R(tp1) \rightarrow \{ \{f6\}, \{f7\} \}$. For $tp4$, as $f5 \sqsubseteq f4$, $f5$ is safely removed at line 5, and $R(tp4) \rightarrow \{ \{f4\} \}$. Second, FEDRA localizes fragments on endpoints in $E(tp)$. Figure 4 shows the execution plans encoded in $R(tp)$ and $E(tp)$. Triple patterns like $tp1$, with more than one relevant fragment, represent unions in the execution plan.

Procedure UNIONREDUCTION (cf. Algorithm 2) reduces subsets of endpoints with relevant fragments into one set with all the common endpoints if at least one

⁹ Notice that DBpedia is not included in the federation for processing Q , so using only fragments $f6$ and $f7$ to retrieve data for $tp1$ will not produce all the answers that would be produced using DBpedia.

Algorithm 3 Basic graph pattern reduction algorithm

Require: tps : set of TriplePattern; E : mapping from TriplePattern to set of set of Endpoint
 20: **procedure** BGPREDUCTION(tps, E)
 21: triplesWithOneFragment $\leftarrow \{ tp : tp \in tps \wedge \text{cardinality}(E(tp)) = 1 \}$
 22: (S, C) \leftarrow minimal set covering instance using triplesWithOneFragment \triangleleft E
 23: C' \leftarrow MINIMALSETCOVERING(S, C)
 24: selected \leftarrow get endpoints encoded by C'
 25: **for each** tp \in triplesWithOneFragment **do**
 26: E(tp) \leftarrow E(tp) \cap selected




Fig. 4: Execution plan encoded in data structures R and E; multiple subsets represent union of different fragment (ex. $\{f_6\}$, $\{f_7\}$); elements of the subset represent alternative location of fragments (ex. $\{C1, C3\}$); bold sources are the selected sources after set covering is used to reduce the sources

common endpoint exists. In our running example there is no common endpoint that replicates both f_6 and f_7 . However, if, for example, f_7 were also replicated at C_1 , then only C_1 would be selected to execute tp_1 .

Procedure BGPREDUCTION (cf. Algorithm 3) transforms the join part of $E(tp)$ (cf. Figure 4) into a set covering problem (cf. line 22). Each triple pattern is an element of the set to cover, e.g., tp_2 , tp_3 , tp_4 correspond to s_2 , s_3 , s_4 (cf. Figure 5a). And for each endpoint in $E(tp)$, we include the subset of triple patterns associated with that endpoint, e.g., for endpoint C_1 we include the subset $\{s_2, s_4\}$ as relevant fragments tp_2 and tp_4 are replicated by C_1 (cf. Figure 5b). Line 23 relies on an existing heuristic [13] to find the minimum set covering. In our example, it computes $C' = \{\{s_2, s_3, s_4\}\}$. Line 24 computes the selected endpoints, in our example, $\text{selected} = \{C_3\}$.

Finally, (Algorithm 1, line 12) chooses among endpoints that provide the same fragment and reduces data redundancy. For query Q, the whole algorithm returns D_2 of Table 1b.

Proposition 1. Algorithm 1 has a time complexity of $O(n \cdot m^2)$, with n the number of triple patterns in the query, m the number of fragments, k the number of endpoints, l the number of basic graph patterns in the query, and $m \gg k \wedge k \gg l$.

For each dataset, we setup a ten consumer SPARQL endpoint federation. A consumer SPARQL endpoint is implemented using Jena Fuseki 1.1.1¹⁰ Each consumer endpoint selects 100 random queries. Each triple pattern of the query is executed as a SPARQL construct query with the LDF client¹¹. The results are stored locally if not present in at least three consumer endpoints and a fragment definition is created.

In order to measure the number of transferred data, the federated query engine access data consumer endpoints through a proxy.

Implementations: FedX 3.0¹² and ANAPSID¹³ have been modified to call FEDRA and DAW [18] source selection strategies during query processing. Thus, each engine can use the selected sources to perform its own optimization strategies. Because FedX is implemented in Java and ANAPSID is implemented in Python, FEDRA and DAW¹⁴ are implemented in both Java 1.7 and Python 2.7.3.. Thus, FEDRA and DAW are integrated in FedX and ANAPSID, reducing the performance impact of including these new source selection strategies. Proxies are implemented in Java 1.7. using the Apache HttpComponents Client library 4.3.5.¹⁵ We used R project¹⁶ to compute the Wilcoxon signed rank test [24].

Evaluation Metrics: *i) Number of Selected Sources (NSS):* is the sum of the number of sources that has been selected per triple pattern. *ii) Number of Transferred Tuples (NTT):* is the number of tuples transferred from all the endpoints to the query engine during a query execution.

5.1 Data Redundancy Minimization

To measure the reduction of the number of selected sources, 100 queries were randomly chosen, and the source selection was performed for these queries for each federation using ANAPSID and FedX with and without FEDRA or DAW. For each query, the sum of the number of selected sources per triple pattern was computed. Boxplots are used to present the results (cf. Figure 6). Both FEDRA and DAW significantly reduce the number of selected sources, however, the reduction achieved by FEDRA is greater than the achieved by DAW.

To confirm it, we formulated the null hypothesis: “FEDRA selects the same number of sources as DAW does”, and performed a Wilcoxon signed rank test, p-values were inferior or equal to 1.4e-05 for all federations and engines¹⁷. Obtained p-values allow to discard the null hypothesis DAW and FEDRA reduction are similar and accept the alternative hypothesis that FEDRA reduction is greater than the one achieved by DAW. FEDRA source selection strategy identifies the

¹⁰ <http://jena.apache.org/>, January 2015.

¹¹ <https://github.com/LinkedDataFragments>, March 2015

¹² <http://www.fluidops.com/fedx/>, September 2014.

¹³ <https://github.com/anapsid/anapsid>, September 2014.

¹⁴ We had to implement DAW as its code is not available.

¹⁵ <https://hc.apache.org/>, October 2014.

¹⁶ <http://www.r-project.org/>

¹⁷ All the p-values are available in the web page: <https://sites.google.com/site/fedrasourceselection>




Fig. 6: Number of Selected Sources for execution of ANAPSID (A) and FedX (F) using FEDRA (F+), DAW (D+), and the engine source selection

relevant fragments and endpoints that provide the same data. Only one of them is actually selected; in consequence, a huge reduction on the number of selected sources of up to 400% per query is achieved.

5.2 Data transfer minimization

To measure the reduction in the number of transferred tuples, queries were executed using proxies that measure the number of transmitted tuples from endpoints to the engines. Because queries that timed out have no significance on number of transferred tuples, we removed all these queries from the study¹⁸. Results (cf. Figure 7) show that FEDRA source selection strategy leads to executions with considerably less intermediate results in all the federations except SWDF federation. In some queries of the SWDF federation, FEDRA +FedX sends exclusive groups composed of triple patterns with no join variable, and

¹⁸ Up to six queries out of 100 queries did not successfully finish in 1,800 seconds, details available at the web page




Fig. 7: Number of Transferred Tuples during execution with ANAPSID (A) and FedX (F) using FEDRA (F+), DAW (D+), and the engine source selection

these exclusive groups leads to a significant number of intermediate results that counters FEDRA positive impact over other queries.

Despite that, globally FEDRA shows an effective reduction of the number of transferred tuples. To confirm it, we formulated the null hypothesis: “using sources selected by FEDRA leads to transfer the same number of tuples as using sources selected by DAW”; and performed a Wilcoxon signed rank test, p-values were inferior or equal to 0.002 for all federations and engines except SWDF federation + FedX engine. In consequence, for all combinations of federation and engines except SWDF+FedX, we can discard the null hypothesis DAW and FEDRA number of transferred tuples are similar and accept the alternative hypothesis that FEDRA achieves a greater reduction of the number of transferred tuples than DAW. The reduction of the number of transferred tuples is mainly due to FEDRA source selection strategy aims to find opportunities to execute joins in the endpoints, and mostly, it leads to significant reduction in the intermediate results size of up to four orders of magnitude.

6 Related Work

In distributed databases, data fragmentation and replication improve data availability and query performance [16]. Linked Data [6] is intrinsically a federation of autonomous participants where federated queries are unknown to a single participant, and a tight coordination of data providers is difficult to achieve. This makes data fragmentation [16] and distributed query processing [14] of distributed databases not a feasible solution [21] for Linked Data.

Recently, the Linked Data fragments approach (LDF) [23, 22] proposes to improve Linked Data availability by moving query execution load from servers to clients. A client is able to execute locally a restricted SPARQL query by downloading fragments required to execute the query from an LDF server through a simple HTTP request. This strategy allows clients to cache fragments locally and decreases the load on the LDF server. LDF chooses a clear tradeoff by shifting query processing to clients, at the cost of slower query execution [22]. In experiments, we presented how to federate several SPARQL consumer endpoints that replicates fragments from LDF servers. Re-organizing fragments on data consumers opens the opportunity to process federated queries even with LDF servers.

Col-graph [12] enables data consumers to materialize triple pattern fragments and to expose them through SPARQL endpoints to improve data quality. A data consumer can update her local fragments and share updates with data providers and consumers. Col-graph proposes a coordination free protocol to maintain the consistency of replicated fragments. Currently, FEDRA can process federated queries over Col-graph collaboration networks if the topology of Col-graph is restricted to two layers without cycles. Moreover, FEDRA does not yet consider divergence between fragments produced by concurrent editing. Querying fragments with divergence is addressed in [15].

Recently, HiBISCuS [19] a source selection approach has been proposed to reduce the number of selected sources. The reduction is achieved by annotating sources with their authority URIs, and pruning sources that cannot have triples that match any of the query triple patterns. HiBISCuS differs from our aim of both selecting sources that are required to the answer, and avoiding the selection of sources that only provide redundant replicated fragments. While not directly related to replication, HiBISCuS index could be used in conjunction with FEDRA to perform join-aware source selection in presence of replicated fragments.

Recently, QBB[11] and DAW[18] propose duplicate-aware strategies for selecting sources for federated query engines. Both approaches use sketches to estimate the overlapping among sources.

DAW uses a combination of Min-Wise Independent Permutations (MIPs) [7], and triple selectivity information to estimate the overlap between the results of different sources. Based on how many new query results are expected to be found, sources that are below predefined benefits, are discarded and not selected. Compared to DAW, FEDRA does not require to compute data summaries because FEDRA relies on fragment definitions and fragment containment to manage replication. Computing containments based on fragment descriptions is less expensive

than computing data summaries; moreover, data updates are more frequent than fragment description updates. FEDRA minimizes the number of endpoints and data transfer and produces complete query answers. Consequently, if DAW and FEDRA could find the same number of sources to execute a query, FEDRA source selection considers the query basic graph patterns to delegate join execution to the endpoints and reduce intermediate results size. This key feature cannot be achieved by DAW as it performs source selection only at the triple pattern level.

7 Conclusions

In this paper, we illustrated how replicating fragments allow to re-organize data from different data sources to better fit queries of data consumers. Then, we proposed a replication-aware federated query engine by extending state-of-art federated query engine ANAPSID and FedX with FEDRA, a source selection strategy that approximates SSP-FR.

FEDRA exploits fragment localities to reduce intermediate results. Experimental results demonstrate that FEDRA achieves significant reduction of intermediate results while preserving query answer completeness.

This work opens several perspectives. First, we made the hypothesis that replicated fragments are perfectly synchronized and cannot be updated. We can leverage this hypothesis and manage the problem of federated query processing with divergence [15].

Several variants of SSP-FR can also be developed. SSP-FR do not make differences between endpoints, the cost of accessing endpoints is considered equal. SSP-FR and FEDRA can be modified to compute the source selection that minimize public endpoint access as addressed in [15].

References

1. M. Acosta, M.-E. Vidal, T. Lampo, J. Castillo, and E. Ruckhaus. Anapsid: An adaptive query processing engine for sparql endpoints. In *International Semantic Web Conference*, pages 18–34, 2011.
2. G. Aluç, O. Hartig, M. T. Özsu, and K. Daudjee. Diversified stress testing of RDF data management systems. In *The Semantic Web - ISWC 2014 - 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part I*, pages 197–212, 2014.
3. G. Aluç, M. T. Ozsu, K. Daudjee, and O. Hartig. chameleon-db: a workload-aware robust rdf data management system. *University of Waterloo, Tech. Rep. CS-2013-10*, 2013.
4. C. B. Aranda, A. Hogan, J. Umbrich, and P.-Y. Vandenbussche. Sparql web-querying infrastructure: Ready for action? In *International Semantic Web Conference (2)*, pages 277–293, 2013.
5. C. Basca and A. Bernstein. Avalanche: Putting the spirit of the web back into semantic web querying. In A. Polleres and H. Chen, editors, *ISWC Posters&Demos*, volume 658 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2010.

6. C. Bizer, T. Heath, and T. Berners-Lee. Linked data - the story so far. *Int. J. Semantic Web Information Systems*, 5(3):1–22, 2009.
7. A. Z. Broder, M. Charikar, A. M. Frieze, and M. Mitzenmacher. Min-wise independent permutations. *Journal Computer and System Sciences*, 60(3):630–659, 2000.
8. O. Görlitz and S. Staab. Splendid: Sparql endpoint federation exploiting void descriptions. In O. Hartig, A. Harth, and J. Sequeda, editors, *COLD*, 2011.
9. C. Gutierrez, C. A. Hurtado, A. O. Mendelzon, and J. Pérez. Foundations of semantic web databases. *Journal of Computer and System Sciences*, 77(3):520–541, 2011.
10. A. Y. Halevy. Answering queries using views: A survey. *The VLDB Journal*, 10(4):270–294, Dec. 2001.
11. K. Hose and R. Schenkel. Towards benefit-based RDF source selection for SPARQL queries. In *Proceedings of the 4th International Workshop on Semantic Web Information Management, SWIM*, page 2, 2012.
12. L.-D. Ibáñez, H. Skaf-Molli, P. Molli, and O. Corby. Col-graph: Towards writable and scalable linked open data. In *International Semantic Web Conference*, 2014.
13. D. S. Johnson. Approximation algorithms for combinatorial problems. In *Proceedings of the fifth annual ACM symposium on Theory of computing*, pages 38–49. ACM, 1973.
14. D. Kossmann. The state of the art in distributed query processing. *ACM Computer Survey*, 32(4):422–469, 2000.
15. G. Montoya, H. Skaf-Molli, P. Molli, and M.-E. Vidal. Fedra: Query Processing for SPARQL Federations with Divergence. Technical report, Université de Nantes, May 2014.
16. M. T. Özsu and P. Valduriez. *Principles of distributed database systems*. Springer, 2011.
17. B. Quilitz and U. Leser. Querying distributed RDF data sources with SPARQL. In *The Semantic Web: Research and Applications, 5th European Semantic Web Conference, ESWC*, pages 524–538, 2008.
18. M. Saleem, A.-C. N. Ngomo, J. X. Parreira, H. F. Deus, and M. Hauswirth. Daw: Duplicate-aware federated query processing over the web of data. In *International Semantic Web Conference*, pages 574–590, 2013.
19. M. Saleem and A. N. Ngomo. Hibiscus: Hypergraph-based source selection for SPARQL endpoint federation. In *11th International Conference ESWC 2014*, pages 176–191, 2014.
20. A. Schwarte, P. Haase, K. Hose, R. Schenkel, and M. Schmidt. Fedx: Optimization techniques for federated query processing on linked data. In *International Semantic Web Conference*, pages 601–616, 2011.
21. J. Umbrich, K. Hose, M. Karnstedt, A. Harth, and A. Polleres. Comparing data summaries for processing live queries over linked data. *World Wide Web*, 14(5-6):495–544, 2011.
22. R. Verborgh, O. Hartig, B. De Meester, G. Haesendonck, L. De Vocht, M. Vander Sande, R. Cyganiak, P. Colpaert, E. Mannens, and R. Van de Walle. Querying datasets on the Web with high availability. In *13th International Semantic Web Conference*, 2014.
23. R. Verborgh, M. Vander Sande, P. Colpaert, S. Coppens, E. Mannens, and R. Van de Walle. Web-scale querying through Linked Data Fragments. In *7th Workshop on Linked Data on the Web*, 2014.
24. F. Wilcoxon. Individual comparisons by ranking methods. In *Breakthroughs in Statistics*, pages 196–202. Springer, 1992.