

The *Odyssey* Approach for Optimizing Federated SPARQL Queries

Gabriela Montoya¹(✉), Hala Skaf-Molli², and Katja Hose¹

¹ Aalborg University, Aalborg, Denmark
{`gmontoya,khose`}@`cs.aau.dk`

² Nantes University, Nantes, France
`hala.skaf@univ-nantes.fr`

Abstract. Answering queries over a federation of SPARQL endpoints requires combining data from more than one data source. Optimizing queries in such scenarios is particularly challenging not only because of (i) the large variety of possible query execution plans that correctly answer the query but also because (ii) there is only limited access to statistics about schema and instance data of remote sources. To overcome these challenges, most federated query engines rely on heuristics to reduce the space of possible query execution plans or on dynamic programming strategies to produce optimal plans. Nevertheless, these plans may still exhibit a high number of intermediate results or high execution times because of heuristics and inaccurate cost estimations. In this paper, we present *Odyssey*, an approach that uses statistics that allow for a more accurate cost estimation for federated queries and therefore enables *Odyssey* to produce better query execution plans. Our experimental results show that *Odyssey* produces query execution plans that are better in terms of data transfer and execution time than state-of-the-art optimizers. Our experiments using the FedBench benchmark show execution time gains of at least 25 times on average.

1 Introduction

Federated SPARQL query engines [1, 4, 7, 14, 17] answer SPARQL queries over a federation of SPARQL endpoints. Query optimization is a particularly complex and challenging task in a federated setting. The query optimizer minimizes processing and communication costs by selecting only relevant sources for a query. It decomposes the query into subqueries, and produces a query execution plan with good join ordering and physical operators. With limited access to statistics, however, most federated query engines rely on heuristics [1, 17] to reduce the huge space of possible plans or on dynamic programming (DP) [5, 7] to produce optimal plans. However, these plans may still exhibit a high number of intermediate results or high execution times because of inadequate heuristics or inaccurate estimations of cost functions [8].

In this paper, we propose *Odyssey*, a cost-based query optimization approach for federations of SPARQL endpoints. *Odyssey* defines statistics for representing

entities inspired by [12] and statistics for representing links among datasets while guaranteeing result completeness. In a federated setting, computing statistics naturally requires access to more than one dataset. To reduce the overhead, *Odyssey* uses entity synopsis to identify links among datasets. This comes at the risk of losing some accuracy in the link identification but still guarantees that no links will be missed during query optimization, i.e., there is a small risk that more sources are queried than strictly necessary but the query result will be complete.

Odyssey uses the computed statistics to estimate the sizes of intermediate results and dynamic programming to produce an efficient query execution plan with a low number of intermediate results. In summary, this paper makes the following contributions:

- Concise statistics of adequate granularity representing entities and describing links among datasets while guaranteeing result completeness.
- A lightweight technique to compute federated statistics in a federated setup that relies on entity synopsis.
- A query optimization algorithm based on dynamic programming using our statistics to find the best plan.
- Extensive evaluation using a well-accepted standard benchmark for federated query processing [16], including comparison against a broad range of state-of-the-art related work [5, 7, 15, 17]. The results show *Odyssey*'s superiority with a speed-up of up to 126 times and a reduction of transferred data of up to 118 times on average.

This paper is organized as follows. Section 2 presents related work, Sect. 3 describes the *Odyssey* approach and its algorithms. Section 4 discusses our experimental results. Finally, conclusions and future work are outlined in Sect. 5.

2 Related Work

Query optimization in state-of-the-art federated query engines, such as FedX [17] and ANAPSID [1], relies on heuristics. For instance, FedX [17] integrates the variable counting heuristic, where relative selectivity of triple patterns is heuristically estimated according to the presence of constants and variables in the triple patterns. These heuristics are lightweight but might not lead to the best query execution plan [18]. To find an optimal plan, several approaches [5, 7, 14, 19] rely on dynamic programming. However, given the high number of alternative query plans for SPARQL queries with many triple patterns, dynamic programming is very expensive [8]. Another important factor of query optimization is source selection. Several approaches [1, 7, 15, 17, 19] try to determine the relevance of a source by sending ASK queries, which increases the costs for a single query but might amortize in large federations for an overlapping query load. Another technique is to estimate whether combining the data of multiple sources can lead to any join results, e.g., by computing the intersection of the sources' URI authorities [15] or detailed statistics [10, 13].

Federated query optimization can also rely on cardinality estimations based on statistics and used, for instance, to reduce sizes of intermediate results. Most available statistics [3] use the Vocabulary of Interlinked Datasets *void* [2], which describes statistics at dataset level (e.g., the number of triples), at the property level (e.g., for each property, its number of different subjects), and at the class level (e.g., the number of instances of each class). However, approaches based on *void* [5, 7, 9] and other statistics, such as *QTrees* [10] and *PARTrees* [13], share the drawback of missing the best query execution plans because of errors in estimating cardinalities caused by relying on assumptions that often do not hold for arbitrary RDF datasets [12], e.g., a uniform data distribution and that the results of triple patterns are independent.

Characteristic sets (CS) [6, 12] aim at solving this problem in centralized systems by capturing statistics about sets of entities having the same set of properties. This information can then be used to accurately estimate the cardinality and join ordering of star-shaped queries. Typically, any set of joined triple patterns in a query can be divided into connected star-shaped subqueries. Subqueries in combination with the predicate that links them, define a characteristic pair (CP) [8, 11]. Statistics about such CPs can then be used to estimate the selectivity of two star-shaped subqueries. Such cardinality estimations can be combined with dynamic programming on a reduced space of alternative query plans. Whereas existing work on CSs and CPs were developed for centralized environments, this paper proposes a solution generalizing these principles for federated environments.

3 The *Odyssey* Approach

Inspired by the latest advances in statistics for centralized triple stores [8, 11, 12], *Odyssey* uses statistics about individual datasets to derive detailed statistics for optimizing federated queries. In the following, we first describe the foundations of our statistics on individual datasets (Sect. 3.1) and then propose a novel method for computing such statistics in a federated environment based on entity descriptions (Sect. 3.2). As the detailed entity descriptions cause too much overhead in a federated setup, we propose a method for reducing the sizes of the descriptions (Sect. 3.3). Finally, we present the *Odyssey* approach for query optimization and its main steps (Sect. 3.4): source selection, join ordering, and query decomposition.

3.1 Dataset Statistics on Individual Datasets

Star-Shaped Subqueries. To estimate the cardinality and costs of BGPs sharing the same subject (or object), i.e., *star-shaped subqueries*, we exploit the principle that entities sharing the same set of properties are similar. In this context, we refer to the set of an entity’s properties as its characteristic set (CS) and use $cs_s(e)$ to denote the CS of entity e in

dataset s or $cs(e)$ if s is clear from the context. For instance, in DBpedia 3.5.1 $cs(dbr:Gary_Goetzman) = C_1 = \{dbo:birthDate, foaf:name, rdf:type, dbo:activeYearsStartYear, rdfs:label, skos:subject\}$. In total, 260 entities share this set of properties and therefore CS C_1 .

Listing 1.1. Statistics for CS C_1

```
{ count: 260,
  elems:
  [{ pred: dbo:birthDate, occurrences : 260 },
   { pred: foaf:name, occurrences: 326 },
   { pred: rdf:type, occurrences: 1023 },
   { pred: dbo:activeYearsStartYear, occurrences: 260 },
   { pred: rdfs:label, occurrences: 260 },
   { pred: skos:subject, occurrences: 1336 }]}
```

CSs can be computed by scanning once a dataset's triples sorted by subject; after all the entity properties have been scanned, the entity's CS is identified. For each CS C , we compute statistics, i.e., the number of entities sharing C ($count(C)$) and the number of triples with predicate p occurring with these entities ($occurrences(p, C)$). Listing 1.1 shows the statistics for the above mentioned example CS C_1 . Entities of C_1 occur on average in 1 triple with property $dbo:birthDate$ and in 3.94 triples with property $rdf:type$.

For a star-shaped query, only CSs including all of the query's properties are relevant as entities that only satisfy a subset of these properties cannot contribute to the answer.

Listing 1.2. Find persons that have been active

```
SELECT DISTINCT ?person WHERE {
  ?person dbo:birthDate ?date .           (tp1)
  ?person dbo:activeYearsStartYear ?sy .  (tp2)
  ?person foaf:name ?name                 (tp3)
}
```

For star-shaped queries asking for the set of unique entities described by some properties (query with DISTINCT modifier), the exact number of answers can be determined precisely (no estimation). For example, the cardinality of the query given in Listing 1.2 can be obtained by adding up the $count(C)$ of all CSs containing the properties $dbo:birthDate$, $dbo:activeYearsStartYear$, and $foaf:name$. In DBpedia 3.5.1, there are 7,059 CSs that include these three properties, and the total number of entities with these CSs is 83,438. Formally, the number of entities for a given set of properties P , $cardinality(P)$, is computed based on the CSs C_j that include all the properties in P as:

$$cardinality(P) = \sum_{P \subseteq C_j} count(C_j) \quad (1)$$

For queries without the DISTINCT modifier, we need to account for duplicates by considering the number of triples with predicate $p_i \in P$ that an entity is associated with on average:

$$estimatedCardinality(P) = \sum_{P \subseteq C_j} \left(count(C_j) * \prod_{p_i \in P} \frac{occurrences(p_i, C_j)}{count(C_j)} \right) \quad (2)$$

In DBpedia 3.5.1, as mentioned above, there are 7,059 CSs relevant for the query in Listing 1.2 with 83,438 entities as answer. These 83,438 entities are described by 109,830 triples with predicate *foaf:name*, 83,448 with predicate *dbo:birthDate*, and 110,460 with predicate *dbo:activeYearsStartYear*. If the query is considered without the DISTINCT modifier, i.e., considering duplicated results, we estimate: 148,486 matching entities in the result, which is very close to the real number (149,440).

Once the relevant CSs for a query have been identified, they can be used to find the join order minimizing the sizes of intermediate results. For the query in Listing 1.2, we start by estimating the cardinalities for each subquery with two out of the three triple patterns using Formula 1: {tp1, tp2}: 98,281, {tp1, tp3}: 209,731, and {tp2, tp3}: 127,712. The triple pattern not included in the cheapest subquery ({tp1, tp2}) is executed last (tp3). We proceed recursively with the cheapest subquery and determine the cardinalities for its subsets: {tp1}: 232,608 and {tp2}: 143,004. Again, the triple pattern not included in the cheapest subquery (tp1) will be executed last of the currently considered set of triple patterns. As a result, we will execute the join between tp2 and tp1 first and afterwards compute the join with tp3. We also get the order in which the triple patterns should be evaluated for the first join: first tp2 and then tp1.

Arbitrary Queries. To estimate the cardinality for queries with more complex shapes, we need to consider the connections (links) between entities with different CSs. Entity *dbr:Evan_Almighty*, for example, is linked to *dbr:Tom_Shadyac* via property *dbo:director* by triple (*dbr:Evan_Almighty*, *dbo:director*, *dbr:Tom_Shadyac*).

The links between CSs via properties can formally be described by characteristic pairs (CPs), they are defined as $(cs_s(e1), cs_s(e2), p)$ for entities $e1$ and $e2$ if $(e1, p, e2) \in s$. The statistics – $count((C_i, C_j, p))$ – capture the number of links between a pair of CSs (C_i and C_j) using a particular property p . For example, given the CSs of *dbr:Evan_Almighty* and *dbr:Tom_Shadyac* as C_1 and C_2 the number of links via property *dbo:director* is given by: $count((C_1, C_2, dbo:director))$.

Listing 1.3. Find movies and their directors

```
SELECT DISTINCT ?film ?director WHERE {
  ?film dbo:runtime ?runtime .           (tp1)
  ?film dbo:director ?director .         (tp2)
  ?film dbo:budget ?budget .             (tp3)
  ?director dbo:birthDate ?date .        (tp4)
  ?director dbo:activeYearsStartYear ?sy . (tp5)
  ?director foaf:name ?name               (tp6)
}
```

The number of unique results (pairs of entities with set of properties P_k and P_l , query with DISTINCT modifier) can be exactly computed (not estimated) using the formula:

$$cardinality((P_k, P_l, p)) = \sum_{P_k \subseteq C_i \wedge P_l \subseteq C_j} count((C_i, C_j, p)) \quad (3)$$

For the query in Listing 1.3 property *dbo:director* links several pairs of CSs representing movies and actors. Hence, we need to compute $\sum_{f_1 \wedge f_2} count((C_i, C_j, dbo:director))$, where f_1 is $(\{dbo:runtime, dbo:director, dbo:budget\} \subseteq C_i)$ and f_2 is $(\{dbo:birthDate, dbo:activeYearsStartYear, foaf:name\} \subseteq C_j)$; one of the operands of this sum is $count((C_1, C_2, dbo:director))$ mentioned in the example above. For this query, DBpedia 3.5.1 contains 1,509 CPs linking entities from two CSs via property *dbo:director*.

If a query does not involve the DISTINCT modifier, result cardinality estimation considers the property occurrences in the CSs:

$$\begin{aligned}
 estimatedCardinality((P_k, P_l, p)) = & \sum_{P_k \subseteq C_i \wedge P_l \subseteq C_j} (count((C_i, C_j, p)) \\
 & * \prod_{p_k \in P_k - \{p\}} \left(\frac{ocurrences(p_k, C_i)}{count(C_i)}\right) * \prod_{p_l \in P_l} \left(\frac{ocurrences(p_l, C_j)}{count(C_j)}\right))
 \end{aligned}
 \tag{4}$$

Assuming that the order of joins within star-shaped subqueries has already been optimized based on the CSs as described above, we treat each star-shaped subquery as a single meta-node to reduce complexity. We estimate the cardinalities of joins between the meta-nodes using the statistics on CPs and use dynamic programming (DP) to determine the optimal join order that minimizes the sizes of intermediate results. Although the presentation in this section focuses on subject-subject joins, the same principle can be applied to other types of joins, e.g., object-object.

3.2 Federated Statistics

In general, entities might occur in multiple datasets in a federation S . Hence, we define a *federated characteristic set* (FCS) as follows: $fc_{sS}(e) = \bigcup_{s \in S} cs_s(e)$, S might be omitted if clear from the context. However, triples describing the same entity are typically part of a single dataset so that most CSs can be computed




Fig. 1. Federated computation of statistics

over each dataset independently from the others¹. The *federated characteristic pair* (FCP) of entities $e1$ and $e2$ via property p in federation S is defined as $(fcs_S(e1), fcs_S(e2), p)$. For FCSs FC_i and FC_j and property p , we compute statistics $count(FC_i)$, $occurrences(p, FC_i)$, and $count((FC_i, FC_j, p))$ as before for CSs and CPs. For simplicity, the following sections focus on FCPs connecting CSs instead of FCSs. The generalization using FCSs is straightforward.

Whereas single dataset statistics can be computed once and provided by the sources in the same way they currently provide void statistics [2], FCSs and FCPs require more effort and centralized knowledge about all entities in the considered datasets. A naive way to compute FCSs and FCPs is evaluating expensive SPARQL queries with FILTER expressions involving NOT EXISTS, but this can take weeks for a dataset with thousands of CSs. It is much more efficient if the sources directly share information about local subjects and objects with the federated query engine: $local_subjects_s(C)$ contains the IRIs of entities with CS C for source s , while $local_objects_s(p, C)$ contains the IRIs of entities linked via predicate p to subjects with CS C . Such information can, for instance, be obtained efficiently while computing CSs and CPs locally and then shared with the federated query engine.

The federated query engine can then use this information to compute FCSs and FCPs. Consider, for instance, the two datasets Lmdb and DBpedia in Fig. 1; based on the CSs (Fig. 1(a)), the sources compute *entity descriptions* ($local_subjects_i$ and $local_objects_i$ in Fig. 1(b)). Entity *film:28350* has properties $\{\text{movie:language}, \dots, \text{owl:sameAs}\} = C_{Lmdb,1}$. Hence, $film:28350 \in local_subjects_{Lmdb}(C_{Lmdb,1})$. There is a triple with *dbr:Evan_Almighty* as value of property *owl:sameAs* for an entity with CS $C_{Lmdb,1}$ (*film:28350*) so $dbr:Evan_Almighty \in local_objects_{Lmdb}(\text{owl:sameAs}, C_{Lmdb,1})$ (Fig. 1(b)). The overlap between the set of entities $local_subjects_{DBpedia}(C_{DBpedia,1})$ and $local_objects_{Lmdb}(\text{owl:sameAs}, C_{Lmdb,1})$ represent linked entities between Lmdb and DBpedia via property *owl:sameAs*. Hence, we obtain FCP $(C_{Lmdb,1}, C_{DBpedia,1}, \text{owl:sameAs})$ (Fig. 1(c)). $count((C_{Lmdb,1}, C_{DBpedia,1}, \text{owl:sameAs}))$ corresponds to the cardinality of the intersection between all the $local_objects_{DBpedia}$ and $local_subjects_{Lmdb}$ linked by property *owl:sameAs*.

Algorithm 1 describes in more detail how to compute FCPs only based on the pre-computed statistics $local_objects_{d1}$ and $local_subjects_{d2}$ (`newFunction(0)` returns a new function with default value 0). First, all common entities in $local_objects_{d1}$ and $local_subjects_{d2}$ are identified in line 7. These common entities represent links between CSs $C_{d1,i}$ and $C_{d2,j}$ via property p and are captured by a FCP (lines 9–10).

Listing 1.4. Find Lmdb movies that are also DBpedia movies

```
SELECT ?film ?movie WHERE {
  ?film dbo:budget ?budget .
  ?film dbo:director ?director .
  ?movie owl:sameAs ?film .
  ?movie imdb:sequel ?seq
}
```

¹ FCSs describing entities across multiple datasets are very rare. In FedBench, for instance, they affect less than 0.5% of all CSs.

Algorithm 1. Compute FCPs Algorithm

Input: *local_objects_{d1}* and *local_subjects_{d2}* for datasets *d1* and *d2*
Output: A set of FCPs (FCPs) with links from *d1* to *d2*; *count(fcp)* for each fcp in FCPs

```

1: function COMPUTEFCPs(local_subjectsd2, local_objectsd1)
2:   FCPs ← { }
3:   count ← newFunction(0)
4:   for (p, Cd1,i) ∈ domain(local_objectsd1) do
5:     entities ← local_objectsd1(p, Cd1,i)
6:     for Cd2,j ∈ domain(local_subjectsd2(Cd2,j) do
7:       entities ← entities ∩ local_subjectsd2(Cd2,j)
8:       if entities ≠ ∅ then
9:         FCPs ← FCPs ∪ { (Cd1,i, Cd2,j, p) }
10:        count((Cd1,i, Cd2,j, p)) ← count((Cd1,i, Cd2,j, p)) + cardinality(entities)
11:      end if
12:    end for
13:  end for
14:  return FCPs, count
15: end function

```

FCPs can be used for cardinality estimation and join ordering using the same principles as described in Sect. 3.1. Consider a federation consisting of DBpedia (160,061 CSs) and LMDb (8,466 CSs) with 22,592 FCPs and query in Listing 1.4. We can use Formula 4 with the FCPs connecting LMDb to DBpedia via the owl:sameAs property to estimate the result cardinality: 171. This is close to the real cardinality (293).

3.3 Reducing the Sizes of Entity Descriptions

As the entity descriptions (*local_subjects_d* and *local_objects_d*) introduced above are often very expensive to compute, maintain, and exchange, we propose a technique to reduce their sizes. We organize the entity descriptions in a tree structure that summarizes the entities used as subject or object in any of the dataset’s triples. Inspired by [10, 13, 15], we factorize common prefixes, transform suffixes into integers, and summarize sets of integers in buckets, i.e., a set synopsis consisting of minimum value (mn), maximum value (mx), $[mn, mx]$, number of elements, *num*, and their set of two least significant bytes (*lsb*). *lsb*(*i*) is computed as $i \bmod 2^{16}$ and is included to improve the synopsis’ accuracy.

The tree structure is organized in three levels. The top level summarizes the prefixes of entity IRIs occurring as subjects and objects in the dataset. Suffixes are mapped to integers using a hash function, these integers are summarized in the middle and bottom levels. The middle level includes buckets where parent nodes subsume the synopsis of their children (containment relationship between parent and child ranges and summation between parent and child *num*) and aids in efficiently accessing the bottom level. The bottom level (leaves) stores (in *local_subjects* and *local_objects*) only the integer’s *lsb* to reduce the storage space while improving the synopsis’ accuracy.

In Fig. 2 we present a fragment of the reduced descriptions for LMDb. The reduced descriptions include all the entities that are subject or object in the dataset’s triples. In particular, it includes the entity with IRI <http://data.linkedmdb.org/resource/film/28350> (Fig. 2(c)). This IRI prefix identifies

the subtree that summarizes the entity (light gray ellipses in Fig. 2(a)), while the hash code of its suffix (resource/film/28350), 1093595742, is used to identify the leaf that includes its lsb (-3490), i.e., with 1093595742 between its minimum and maximum values (gray rectangle in Fig. 2(b)). Its lsb is in $local_subjects(C_{LMDB,1})$ and $local_objects(mol:link_source, C_{LMDB,2})$ in the identified leaf (trapezium in Fig. 2(b)). This tree structure exhibits size reduction and eases the computation of FCPs by allowing to discard large portions of the descriptions contrary to descriptions in Fig. 1(b), where all the $local_subjects$ and $local_objects$ need to be pair-wise tested for overlap.




Fig. 2. Reduced entity descriptions for LMDB in Fig. 1. The tree factorizes common prefixes in the top level (in the ellipses) and summarizes the suffixes in the middle (in the rectangles) and bottom (in the trapezium) levels

Computation costs are greatly reduced by pruning large portions of the tree and comparing only a few pairs of leaves, the ones that have common prefixes and overlapping representation of the suffixes. An important feature of these summaries is that entities present in more than one dataset are always detected.

These trees are considerably lighter than the entity descriptions discussed in Sect. 3.2, but they might reduce accuracy. For FedBench’s DBpedia 3.5.1 subset, a dataset with 43,126,772 triples that occupies 6.1 GB, the $local_subjects$ and $local_objects$ occupy 1.37 GB and the tree occupies only 68 MB². They have compression ratios of 4.45 and 91.86, respectively. Regarding the quality, the tree summary allows for computing all the FCSs and FCPs.

To reduce the resources used by the tree, we have reduced the number of CSs as suggested in [8, 12] to 10,000. Only the CSs that are shared by the largest number of entities are kept, and the others are removed and merged into the remaining CSs if possible. For instance, by selecting from the remaining CSs

² Implementation based on Java’s HashSet and HashMap was used to measure their sizes.

that include all the properties of the removed CS, the one with the smallest number of properties and combining their *count* and *ocurrences*, or by splitting the removed CS into two disjoint property sets that can be merged with other CSs. This may reduce the accuracy of the query cardinality estimation, but it allows to bound the resources used to store and access these statistics.

Entity summaries can be kept up-to-date in two ways. For datasets that are rarely updated, the subtree representing the entities with the prefix affected by the updates, e.g., Fig. 2(b) in our example, can be re-computed. For datasets that are often updated, leaves should support removal of entities, this can easily be done by storing the multiplicity of each least significant byte so they are removed only if all the entities with that least significant byte have been removed from the dataset.

3.4 Optimizing Federated Queries

Query optimization in *Odyssey* can logically be divided into the following steps: (i) preprocessing and source selection, (ii) join ordering, and (iii) query decomposition. Arbitrary queries can be handled incrementally by optimizing its subqueries. In the following, we address the optimization of queries with bound predicates, *Odyssey* relies on existing optimizers to handle other queries.

Preprocessing and Source Selection. We first parse the query and identify its star-shaped subqueries. Then, properties in each star-shaped subquery are used to identify relevant CSs and sources. For example, the subquery composed by tp3 and tp4 in Fig. 3(a) has relevant CSs that include both owl:sameAs and movie:sequel. In the FedBench federation described in Table 1, these CSs are only part of LMDB. Therefore, LMDB is the only relevant source for this subquery. Afterwards, we use CPs/FCPs to identify relevant sources for the links between the star-shaped subqueries.

Join Ordering. Once we have identified the set of relevant sources, we can estimate cardinalities of subqueries and find the best join ordering. We first optimize the order of joins and triple patterns within each star-shaped subquery




Fig. 3. Query Q_F and its optimized plan

using CS statistics ($count(C_i)$ and $occurrences(p, C_i)$) as explained in Sect. 3.1. Afterwards, as described in Sect. 3.1, each subquery is treated as a meta-node and we estimate cardinalities of the joins between these meta-nodes using the formulas presented in Sect. 3.1 to estimate subquery costs and apply DP. For Q_F (Fig. 3(a)), three star-shaped subqueries are identified and treated as meta-nodes to estimate the cardinalities of their joins (Fig. 4, left). Figure 4 (right) shows the estimated cardinality and cost of the subqueries, solid arrows indicate which smaller subqueries were combined by the DP algorithm to form larger subqueries. As the number of subqueries is usually considerably lower than the number of triple patterns, applying DP becomes affordable.

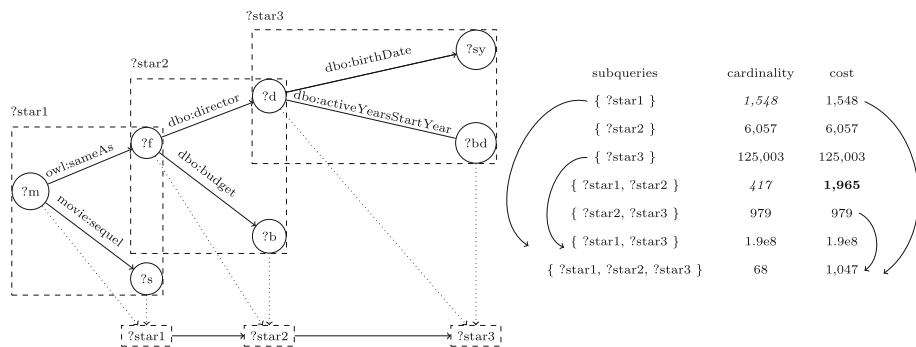


Fig. 4. Example query optimization

In our current implementation, the cost function is solely defined on the cardinalities of intermediate results and how many results need to be transferred from endpoints during execution. This favors query plans with selective subqueries. For instance, the cost of the join between meta-nodes ?star1 and ?star2 (**1,965**) includes the result size (417) and the sum of all transferred intermediate results (1,548). This cost function assumes that all endpoints have the same characteristics. We can easily extend this cost function by additional parameters that can be fine-tuned to represent the characteristics of each endpoint individually, e.g., communication delays, response times, etc.

Query Decomposition. Finally, we optimize the SPARQL queries that are actually sent to the endpoints and try to minimize their number. For instance, we combine all triple patterns and logical subqueries to a particular endpoint into a single SPARQL query to a particular endpoint whenever possible. For instance, meta-nodes ?star2 and ?star3 in Fig. 4 are combined into one subquery (Fig. 3(b)) and evaluated by the DBpedia endpoint.

Table 1. FedBench [16] dataset statistics: number of distinct triples (#DT), predicates (#P), CSs (#CS), and CPs (#CP); computation time in seconds of *Odyssey*, HiBISCuS, and void statistics

Dataset	#DT	#P	# CS	#CP	#FCP	<i>Odyssey</i>	HiBISCuS	void
ChEBI	4,772,706	28	978	9,958	19,360	82.91	96.02	73.89
KEGG	1,090,830	21	67	239	13,822	30.15	95.23	12.84
Drugbank	517,023	119	3,419	12,589	103,070	1,299.9	76.4	6.98
DBpedia subset	42,855,253	1,063	10,000	1,069,431	6,583	2,739	770.48	1,465.36
Geonames	107,949,927	26	673	7,707	322,672	1,885.97	609.52	39,694.07
Jamendo	1,049,647	26	42	190	1,259	31.25	99.17	14.66
SWDF	103,595	118	547	6,713	17,557	7.27	69.21	2.03
LMDB	6,147,916	222	8,466	94,188	359,340	947.16	317.21	355.45
NYTimes	335,119	36	47	158	3.96	10.01	72.56	4.22
Federated						620.35		
Total						7,654.27	2,205.8	41,629.5

4 Evaluation

In this section, we present the results of our experimental study that compares our approach, *Odyssey*, with state-of-the-art federated query engines: HiBISCuS (FedX-HiBISCuS, cold and warm cache) [15], SemaGrow [5], FedX (cold and warm cache) [17], and SPLENDID [7]. Full implementations, statistics, and results are available at <https://github.com/gmontoya/federatedOptimizer>.

Datasets and Queries: We use the real datasets and queries proposed in the FedBench benchmark [16]. Queries are divided into three groups Linked Data (LD1-LD11), Cross Domain (CD1-CD7), and Life Science (LS1-LS7). They have 2–7 triple patterns and star and hybrid shapes. They have between 1 and 9,054 answers. Basic statistics about the datasets are listed in Table 1. We ran each query ten times and report the averages over the last nine runs. Standard deviation is included as error bars on the plots.

Implementation: *Odyssey* is implemented in Java using the Jena library to parse and transform queries into queries with SPARQL 1.1 service clauses. Our implementation uses the FedX 3.1 framework with deactivated native optimization to execute *Odyssey*'s query plans.

Hardware Configuration: For our experiments we used virtual machines (VMs). A VM using up to 4 GB of RAM to run the federated query engine and nine VMs with 2 processors, 8 GB of RAM and CPU 2294.250 MHz to host Virtuoso endpoints with the datasets described in Table 1 (one dataset and endpoint per VM).

Statistics Computation: As DBpedia has a very high number of CSs (160,061), we reduced them to 10,000 by merging (as suggested in [8, 12] and explained in Sect. 3.3) without significant losses in the quality of estimations. Details on creation times of statistics are listed in Table 1. *Odyssey*'s statistics




Fig. 5. Optimization time in ms (OT, log scale). CD1 and LS2 have variable predicates and *Odyssey* relies on FedX to find plan.

can be more expensive to compute for datasets with more than 3,419 CSs and cheaper than HiBISCuS’s for datasets with less than 67 CSs. In total, *Odyssey*’s statistics are computed five times faster than voidD’s.

Evaluation Metrics: (i) *Optimization time (OT)*: is the elapsed time since the query is issued until the optimized query plan is produced, (ii) *number of selected sources (NSS)*: is the number of sources that have been selected to answer a query, (iii) *number of subqueries (NSQ)*: is the number of subqueries that are included in the query plan, (iv) *execution time (ET)*: is the time elapsed since the evaluation of the query plan starts until the complete answer is produced (with a timeout of 1,800s), (v) *number of transferred tuples (NTT)*: is the number of tuples transferred from all the endpoints to the query engine during query evaluation.

Result Completeness: All approaches produce the complete result set for non-timed out queries, except SPLENDID for query LS7.

4.1 Experimental Results

Optimization Time. Figure 5 shows the optimization time (OT) for the studied approaches. Because of the detailed statistics and dynamic programming, one might expect *Odyssey* to suffer from a considerable overhead in OT. As our experimental results show, however, *Odyssey*’s query planner is competitive to most other approaches with a slight advantage for FedX-Warm as this system has cached information about the query relevant sources. For instance, *Odyssey* is up to 69 times faster (SemaGrow) than other approaches on average.

Number of Selected Sources. As Fig. 6 shows, *Odyssey* selects only a small number of relevant sources; for instance, at least 1.81 times less (FedX-Cold/Warm and SemaGrow) and up to 1.93 times less (HiBISCuS-Cold/Warm)




Fig. 6. Number of selected sources (NSS)

on average. For some queries, e.g., LS4, existing approaches already select the optimal number of sources. For LD7, *Odyssey* selects a larger number of sources than the optimum because our approach does not perform ASK queries during execution to prune irrelevant sources. Sometimes *Odyssey* overestimates the set of relevant sources – but on the other hand it never misses any relevant sources. For LS1, most approaches select just one (10^0) source because there is only one dataset that has triples with the predicate in the query.

Number of Subqueries. As Fig. 7 shows, *Odyssey* uses considerably fewer subqueries than other approaches, at least 2.62 times less (HiBISCuS-Cold/Warm) and up to 3.41 times less (SPLENDID) on average. The fact that *Odyssey* always produces the correct and complete answers confirms that *Odyssey* correctly identifies and exploits cases for which it is advantageous to combine subqueries. *Odyssey*'s reduction of the number of relevant sources has a positive




Fig. 7. Number of subqueries (NSQ)




Fig. 8. Execution time in ms (ET, log scale)

impact on the number of subqueries (NSQ), *Odyssey*'s pruning of non-relevant sources allows for combining triple patterns into subqueries without affecting the result completeness. Some queries, like LD2, LD4, and LD9, include triple patterns that can be evaluated by a unique endpoint of the federation and existing approaches already decompose the query into the optimal NSQ. Only for LD7, FedX-Cold/Warm, SPLENDID, and SemaGrow decompose the query into fewer subqueries than *Odyssey*, this is because they use ASK queries to assess a source's relevance. *Odyssey* could be enhanced with this strategy.

Execution Time. Some approaches failed to answer all queries before the timeout (1,800s): SPLENDID (2 queries) and SemaGrow (4 queries). Even when considering only those queries that completed before the timeout, *Odyssey* is on average 126.26 times faster than SPLENDID and 28.30 times faster than SemaGrow. Figure 8 shows the execution times (ET) for the studied approaches. *Odyssey* is on average at least 25.46 times faster (FedX-Warm). Only for a few queries *Odyssey* is (slightly) slower than other approaches, e.g., LS3. As for the other metrics, *Odyssey*'s ET can be improved if ASK queries were used during query execution to further reduce the relevant sources similarly as it is done by other approaches. For five of the queries, *Odyssey* is one of the fastest approaches and for 11 queries, *Odyssey* is the fastest approach. *Odyssey*'s achieved reductions on the NSS and NSQ have a positive impact on the ET; as fewer endpoints are queried fewer times, *Odyssey* produces results faster than most approaches in most cases.

Number of Transferred Tuples. Figure 9 shows the number of transferred tuples (NTT) for the studied approaches. *Odyssey* transfers fewer tuples than other approaches. Even when considering only those queries that completed before the timeout, *Odyssey* transfers on average 1.15 times fewer tuples faster than SemaGrow and 108.4 times fewer tuples than SPLENDID. For the




Fig. 9. Number of transferred tuples (NTT, log scale, $10^0 = 1$)

approaches that completed all the queries, *Odyssey* transfers at least 117.55 fewer tuples (HiBISCuS-Cold/Warm) on average. Most approaches are competitive in terms of NTT. The largest difference is observed for LS6, where *Odyssey* clearly outperforms the other approaches transferring 500 times fewer tuples. In contrast to other approaches, *Odyssey* not only reduces the number of requests sent to the endpoints but also avoids non-selective subqueries, which significantly reduces network traffic and the local query load at the endpoints.

4.2 Combining *Odyssey* with Existing Optimizers

We have also integrated *Odyssey* techniques directly into the FedX optimizer and obtained:

- *Odyssey*-FedX-Cold, which relies on CSs and CPs to select sources and decomposes the query but uses FedX join ordering.
- FedX-Cold-*Odyssey*, which relies on the FedX optimizer for source selection but uses *Odyssey* for query decomposition and join ordering.

Figure 10 compares the execution times (ET) of these two implementations with *Odyssey*, FedX-Cold, and FedX-Warm. In most cases the combined approaches are considerably faster than native FedX. In a few cases, however, their ET can increase considerably. In these cases, queries include a highly selective subquery with one triple pattern and using FedX’s heuristic to execute subqueries with more than one triple pattern first leads to plans that are more expensive than others. On average, the combined approaches are 26.86 and 3.99 times faster than FedX-Cold.

For query LD7, *Odyssey* and FedX-Cold/Warm exhibit similar ETs whereas FedX-Cold-*Odyssey* is considerably faster. For this query it happens that the advantages of both *Odyssey* and FedX coincide, i.e., we can take advantage of the good join ordering by *Odyssey* but also of the additional pruning based on ASK queries by FedX.




Fig. 10. Execution time in ms (ET, log scale) of query plans optimized using *Odyssey* and FedX

Even if *Odyssey*'s OT can be higher in comparison to existing approaches, *Odyssey* produces better plans composed of fewer subqueries and fewer selected sources per triple pattern without compromising result completeness. Benefits of these features have been evidenced with significantly faster ETs and less transferred data from endpoints to the federated query engine.

5 Conclusion

In this paper, we have presented *Odyssey*, an approach for optimizing federated SPARQL queries based on statistics. These statistics detail information about the data provided by remote endpoints as well as the links between them. This enables more accurate cost estimations, query optimization, and selection of relevant sources. Our extensive experimental evaluation shows that *Odyssey* produces query execution plans that are better in terms of data transfer and execution time than state-of-the-art optimizers. In our future work, we plan to further improve *Odyssey* by considering in which situations exactly it is worthwhile to use additional aspects of other optimizers, such as ASK queries and associated statistics. Another interesting direction of future work is to further reduce the computation time and sizes of the entity descriptions and provide efficient strategies to update the descriptions and statistics.

Acknowledgments. This research was partially funded by the Danish Council for Independent Research (DFR) under grant agreement no. DFF-4093-00301.

References

1. Acosta, M., Vidal, M.-E., Lampo, T., Castillo, J., Ruckhaus, E.: ANAPSID: an adaptive query processing engine for SPARQL endpoints. In: Aroyo, L., Welty, C., Alani, H., Taylor, J., Bernstein, A., Kagal, L., Noy, N., Blomqvist, E. (eds.) ISWC 2011. LNCS, vol. 7031, pp. 18–34. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25073-6_2](https://doi.org/10.1007/978-3-642-25073-6_2)
2. Alexander, K., Cyganiak, R., Hausenblas, M., Zhao, J.: Describing linked datasets. In: LDOW 2009 (2009)
3. Buil-Aranda, C., Hogan, A., Umbrich, J., Vandenbussche, P.-Y.: SPARQL web-querying infrastructure: ready for action? In: Alani, H., et al. (eds.) ISWC 2013. LNCS, vol. 8219, pp. 277–293. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-41338-4_18](https://doi.org/10.1007/978-3-642-41338-4_18)
4. Basca, C., Bernstein, A.: Querying a messy web of data with avalanche. *J. Web Semant.* **26**, 1–28 (2014)
5. Charalambidis, A., Troumpoukis, A., Konstantopoulos, S.: SemaGrow: optimizing federated SPARQL queries. In: SEMANTICS 2015, pp. 121–128 (2015)
6. Du, F., Chen, Y., Du, X.: Partitioned indexes for entity search over RDF knowledge bases. In: Lee, S., Peng, Z., Zhou, X., Moon, Y.-S., Unland, R., Yoo, J. (eds.) DASFAA 2012. LNCS, vol. 7238, pp. 141–155. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-29038-1_12](https://doi.org/10.1007/978-3-642-29038-1_12)
7. Görlitz, O., Staab, S.: SPLENDID: SPARQL endpoint federation exploiting VOID descriptions. In: COLD 2011 (2011)
8. Gubichev, A., Neumann, T.: Exploiting the query structure for efficient join ordering in SPARQL queries. In: EDBT 2014, pp. 439–450 (2014)
9. Hagedorn, S., Hose, K., Sattler, K., Umbrich, J.: Resource planning for SPARQL query execution on data sharing platforms. In: COLD, pp. 49–60 (2014)
10. Harth, A., Hose, K., Karnstedt, M., Polleres, A., Sattler, K., Umbrich, J.: Data summaries for on-demand queries over linked data. In: WWW 2010, pp. 411–420 (2010)
11. Meimaris, M., Papastefanatos, G., Mamoulis, N., Anagnostopoulos, I.: Extended characteristic sets: graph indexing for SPARQL query optimization. In: ICDE 2017 (2017)
12. Neumann, T., Moerkotte, G.: Characteristic sets: accurate cardinality estimation for RDF queries with multiple joins. In: ICDE 2011, pp. 984–994 (2011)
13. Prasser, F., Kemper, A., Kuhn, K.A.: Efficient distributed query processing for autonomous RDF databases. In: EDBT 2012, pp. 372–383 (2012)
14. Quilitz, B., Leser, U.: Querying distributed RDF data sources with SPARQL. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) ESWC 2008. LNCS, vol. 5021, pp. 524–538. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-68234-9_39](https://doi.org/10.1007/978-3-540-68234-9_39)
15. Saleem, M., Ngonga Ngomo, A.-C.: HiBISCuS: hypergraph-based source selection for SPARQL endpoint federation. In: Presutti, V., d’Amato, C., Gandon, F., d’Aquin, M., Staab, S., Tordai, A. (eds.) ESWC 2014. LNCS, vol. 8465, pp. 176–191. Springer, Cham (2014). doi:[10.1007/978-3-319-07443-6_13](https://doi.org/10.1007/978-3-319-07443-6_13)
16. Schmidt, M., Görlitz, O., Haase, P., Ladwig, G., Schwarte, A., Tran, T.: FedBench: a benchmark suite for federated semantic data query processing. In: Aroyo, L., Welty, C., Alani, H., Taylor, J., Bernstein, A., Kagal, L., Noy, N., Blomqvist, E. (eds.) ISWC 2011. LNCS, vol. 7031, pp. 585–600. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25073-6_37](https://doi.org/10.1007/978-3-642-25073-6_37)

17. Schwarte, A., Haase, P., Hose, K., Schenkel, R., Schmidt, M.: FedX: optimization techniques for federated query processing on linked data. In: Aroyo, L., Welty, C., Alani, H., Taylor, J., Bernstein, A., Kagal, L., Noy, N., Blomqvist, E. (eds.) ISWC 2011. LNCS, vol. 7031, pp. 601–616. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25073-6_38](https://doi.org/10.1007/978-3-642-25073-6_38)
18. Stocker, M., Seaborne, A., Bernstein, A., Kiefer, C., Reynolds, D.: SPARQL basic graph pattern optimization using selectivity estimation. In: WWW 2008, pp. 595–604 (2008)
19. Wang, X., Tiropanis, T., Davis, H.C.: LHD: optimising linked data query processing using parallelisation. In: LDOW (2013)