

The Odyssey Approach for Optimizing Federated SPARQL Queries

Gabriela Montoya¹, Hala Skaf-Molli², and Katja Hose¹

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

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

October 25th, 2017

Optimizing Federated SPARQL Queries

```
SELECT DISTINCT * WHERE {  
  ?film dbo:director ?director .           (tp1)  
  ?film rdf:type dbo:Film .                (tp2)  
  ?movie owl:sameAs ?film .              (tp3)  
  ?movie dcterms:title ?title .           (tp4)  
  ?movie movie:film_subject film_subject:444 (tp5)  
}
```

DBpedia, Drugbank, LMDB, ChEBI, Geonames, NYTimes, Jamendo, SWDF, KEGG

Optimizing Federated SPARQL Queries

```

SELECT DISTINCT * WHERE {
  ?film dbo:director ?director .           (tp1)
  ?film rdf:type dbo:Film .                (tp2)
  ?movie owl:sameAs ?film .              (tp3)
  ?movie dcterms:title ?title .           (tp4)
  ?movie movie:film_subject film_subject:444 (tp5)
}

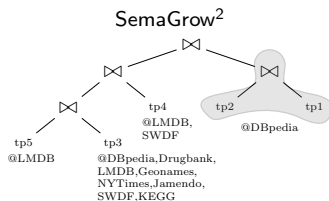
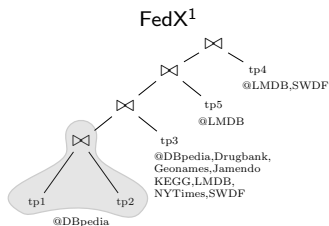
```

DBpedia, Drugbank, LMDB, ChEBI, Geonames, NYTimes, Jamendo, SWDF, KEGG

Subquery	Relevant Sources
?film dbo:director ?director .	DBpedia
?film rdf:type dbo:Film	DBpedia, Drugbank, LMDB, ChEBI Geonames, NYTimes, Jamendo, SWDF, KEGG
?movie owl:sameAs ?film	DBpedia, Drugbank, LMDB Geonames, NYTimes, Jamendo, SWDF, KEGG
?movie dcterms:title ?title	LMDB, SWDF
?movie movie:film_subject film_subject:444	LMDB

Existing approaches

Plan



Optimization Technique

Optimization Time (s)

Execution Time (s)

Heuristics

0.74

142

Dynamic Programming

4.75

6.93

¹A. Schwarte et al. "FedX: Optimization Techniques for Federated Query Processing on Linked Data". In: *ISWC'11*.

²A. Charalambidis et al. "SemaGrow: optimizing federated SPARQL queries". In: *SEMANTICS'15*.

Considering joins between triple patterns improves the optimization

Subquery	Relevant Sources
?movie owl:sameAs ?film . ?movie dcterms:title ?title . ?movie movie:film_subject film_subject:444	LMDB

Only entities that satisfy *owl:sameAs*, *dcterms:title*, **and** *movie:film_subject* are part of LMDB!
 Entities from all other sources that seem relevant for some triple patterns, actually will not contribute to the result!

Contributions

- ▶ Concise statistics describing links among triples while guaranteeing result completeness.
- ▶ A technique to compute such statistics in a federated setup.
- ▶ Affordable optimization based on dynamic programming.

Statistics computation at one location

```

film:1129 movie:film_subject film_subject:444 .
film:1129 dcterms:title "Kate & Leopold" .
film:1129 owl:sameAs dbr:Kate_&_Leopold .
film:16189 movie:film_subject film_subject:444 .
film:16189 dcterms:title "Journey to the Center of Time" .
film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .
dbr:Journey_to_the_Center_of_Time dbo:director dbr:David_L._Hewitt .
dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .
...

```

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

Statistics computation at one location

```

film:1129 movie:film_subject film_subject:444 .
film:1129 dcterms:title "Kate & Leopold" .
film:1129 owl:sameAs dbr:Kate_&_Leopold .
film:16189 movie:film_subject film_subject:444 .
film:16189 dcterms:title "Journey to the Center of Time" .
film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .
dbr:Journey_to_the_Center_of_Time dbo:director dbr:David_L._Hewitt .
dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .
...

```

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

Statistics computation at one location

Characteristic Sets (CS)³

$$C_{D,1} = \{ \text{movie:film_subject}, \text{dcterm:s:title}, \text{owl:sameAs} \}$$

$$\text{count}(C_{D,1})=1$$

```

film:1129 movie:film_subject film_subject:444 .
film:1129 dcterm:s:title "Kate & Leopold" .
film:1129 owl:sameAs dbr:Kate_&_Leopold .
film:16189 movie:film_subject film_subject:444 .
film:16189 dcterm:s:title "Journey to the Center of Time" .
film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .
dbr:Journey_to_the_Center_of_Time dbo:director dbr:David_L._Hewitt .
dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .
...

```

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

Statistics computation at one location

Characteristic Sets (CS)³

$$C_{D,1} = \{ \text{movie:film_subject}, \text{dcterm:s:title}, \text{owl:sameAs} \}$$

$$\text{count}(C_{D,1})=2$$

```
film:1129 movie:film_subject film_subject:444 .
film:1129 dcterm:s:title "Kate & Leopold" .
film:1129 owl:sameAs dbr:Kate_&_Leopold .
film:16189 movie:film_subject film_subject:444 .
film:16189 dcterm:s:title "Journey to the Center of Time" .
film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .
dbr:Journey_to_the_Center_of_Time dbo:director dbr:David_L._Hewitt .
dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .
...
```

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

Statistics computation at one location

Characteristic Sets (CS)³

$C_{D,1} = \{ \text{movie:film_subject}, \text{dcterms:title}, \text{owl:sameAs} \}$

$\text{count}(C_{D,1})=2$

$C_{D,2} = \{ \text{dbo:director}, \text{rdf:type} \}$ $\text{count}(C_{D,2})=1$

film:1129 movie:film_subject film_subject:444 .

film:1129 dcterms:title "Kate & Leopold" .

film:1129 owl:sameAs dbr:Kate.&_Leopold .

film:16189 movie:film_subject film_subject:444 .

film:16189 dcterms:title "Journey to the Center of Time" .

film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .

dbr:Journey_to_the_Center_of_Time dbo:director dbr:David_L._Hewitt .

dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .

...

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

Other basic statistics are also stored

Characteristic Sets (CS)³

$C_{D,1} = \{ \text{movie:film_subject}, \text{dcterm:s:title}, \text{owl:sameAs} \}$

$\text{count}(C_{D,1})=2$; $\text{occurrences}(\text{dcterm:s:title}, C_{D,1})=2$

$C_{D,2} = \{ \text{dbo:director}, \text{rdf:type} \}$ $\text{count}(C_{D,2})=1$

`film:1129 movie:film_subject film_subject:444 .`

`film:1129 dcterm:s:title "Kate & Leopold" .`

`film:1129 owl:sameAs dbr:Kate.&_Leopold .`

`film:16189 movie:film_subject film_subject:444 .`

`film:16189 dcterm:s:title "Journey to the Center of Time" .`

`film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .`

`dbr:Journey_to_the_Center_of_Time dbo:director dbr:David_L._Hewitt .`

`dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .`

...

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

CSs are connected to other CSs

Characteristic Sets (CS)³

$C_{D,1} = \{ \text{movie:film_subject}, \text{dcterm:s:title}, \text{owl:sameAs} \}$

$\text{count}(C_{D,1})=2; \text{ocurrences}(\text{dcterm:s:title}, C_{D,1})=2$

$C_{D,2} = \{ \text{dbo:director}, \text{rdf:type} \}$ $\text{count}(C_{D,2})=1$

film:1129 movie:film_subject film_subject:444 .

film:1129 dcterm:s:title "Kate & Leopold" .

film:1129 owl:sameAs dbr:Kate.&_Leopold .

film:16189 movie:film_subject film_subject:444 .

film:16189 dcterm:s:title "Journey to the Center of Time" .

film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .

dbr:Journey_to_the_Center_of_Time dbo:director dbr:David.L._Hewitt .

dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .

...

dbr:Journey_to_the_Center_of_Time \rightarrow owl:sameAs \rightarrow $C_{D,1}$

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

CSs are connected to other CSs

Characteristic Sets (CS)³

$C_{D,1} = \{ \text{movie:film_subject}, \text{dcterm:s:title}, \text{owl:sameAs} \}$

$\text{count}(C_{D,1})=2; \text{ocurrences}(\text{dcterm:s:title}, C_{D,1})=2$

$C_{D,2} = \{ \text{dbo:director}, \text{rdf:type} \}$ $\text{count}(C_{D,2})=1$

film:1129 movie:film_subject film_subject:444 .

film:1129 dcterm:s:title "Kate & Leopold" .

film:1129 owl:sameAs dbr:Kate.&_Leopold .

film:16189 movie:film_subject film_subject:444 .

film:16189 dcterm:s:title "Journey to the Center of Time" .

film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .

dbr:Journey_to_the_Center_of_Time dbo:director dbr:David_L._Hewitt .

dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .

...

dbr:Journey_to_the_Center_of_Time \rightarrow owl:sameAs \rightarrow $C_{D,1}$

Characteristic Pairs (CP)⁴

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

CSs are connected to other CSs

Characteristic Sets (CS)³

$C_{D,1} = \{ \text{movie:film_subject}, \text{dcterm:s:title}, \text{owl:sameAs} \}$

$\text{count}(C_{D,1})=2; \text{ocurrences}(\text{dcterm:s:title}, C_{D,1})=2$

$C_{D,2} = \{ \text{dbo:director}, \text{rdf:type} \}$ $\text{count}(C_{D,2})=1$

film:1129 movie:film_subject film_subject:444 .

film:1129 dcterm:s:title "Kate & Leopold" .

film:1129 owl:sameAs dbr:Kate.&_Leopold .

film:16189 movie:film_subject film_subject:444 .

film:16189 dcterm:s:title "Journey to the Center of Time" .

film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .

dbr:Journey_to_the_Center_of_Time dbo:director dbr:David_L._Hewitt .

dbr:Journey_to_the_Center_of_Time rdf:type dbo:Film .

...

dbr:Journey_to_the_Center_of_Time \rightarrow owl:sameAs \rightarrow $C_{D,1}$

Characteristic Pairs (CP)⁴

$(C_{D,1}, C_{D,2}, \text{owl:sameAs}) \text{count}((C_{D,1}, C_{D,2}, \text{owl:sameAs})) = 1$

³T. Neumann and G. Moerkotte. "Characteristic Sets: Accurate Cardinality Estimation for RDF Queries with Multiple Joins". In: *ICDE'11*.

⁴A. Gubichev and T. Neumann. "Exploiting the query structure for efficient join ordering in SPARQL queries". In: *EDBT'14*.

Cardinality computation

```

SELECT DISTINCT ?film ?movie WHERE {
  ?film dbo:director ?director .           (tp1)
  ?film rdf:type ?type .                   (tp2)
  ?movie owl:sameAs ?film .              (tp3)
  ?movie dcterms:title ?title .           (tp4)
  ?movie movie:film_subject ?subject      (tp5)
}

```

$$cardinality((P_k, P_l, p)) = \sum_{P_k \subseteq C_i \wedge P_l \subseteq C_j} count((C_i, C_j, p)) \quad (1)$$

$P_k = \{ owl : sameAs, dcterms : title, movie : film_subject \}$

$P_l = \{ dbo : director, rdf : type \}$

$p = owl : sameAs$

Cardinality estimation

```

SELECT DISTINCT * WHERE {
  ?film dbo:director ?director .      (tp1)
  ?film rdf:type ?type .              (tp2)
  ?movie owl:sameAs ?film .         (tp3)
  ?movie dcterms:title ?title .       (tp4)
  ?movie movie:film_subject ?subject (tp5)
}

```

$$\begin{aligned}
 \text{estimatedCardinality}((P_k, P_l, p)) = & \sum_{P_k \subseteq C_i \wedge P_l \subseteq C_j} (\text{count}((C_i, C_j, p)) \\
 & * \prod_{p_k \in P_k - \{p\}} \left(\frac{\text{ocurrences}(p_k, C_i)}{\text{count}(C_i)} \right) * \prod_{p_l \in P_l} \left(\frac{\text{ocurrences}(p_l, C_j)}{\text{count}(C_j)} \right)) \quad (2)
 \end{aligned}$$

$P_k = \{owl : sameAs, dcterms : title, movie : film_subject\}$

$P_l = \{dbo : director, rdf : type\}$

$p = owl : sameAs$

Cardinality estimation

```

SELECT DISTINCT * WHERE {
  ?film dbo:director ?director .           (tp1)
  ?film rdf:type dbo:Film .               (tp2)
  ?movie owl:sameAs ?film .             (tp3)
  ?movie dcterms:title ?title .           (tp4)
  ?movie movie:film_subject film_subject:444 (tp5)
}

```

$$\begin{aligned}
 \text{estimatedCardinality}((P_k, P_l, p)) = & \max_{P_k \subseteq C_i \wedge P_l \subseteq C_j} \left(\text{count}((C_i, C_j, p)) \right. \\
 & * \prod_{p_k \in P_k - \{p\}} \left(\frac{\text{ocurrences}(p_k, C_i)}{\text{count}(C_i)} \right) * \prod_{p_l \in P_l} \left(\frac{\text{ocurrences}(p_l, C_j)}{\text{count}(C_j)} \right) \left. \right) \quad (3)
 \end{aligned}$$

$P_k = \{owl : sameAs, dcterms : title, movie : film_subject\}$

$P_l = \{dbo : director, rdf : type\}$

$p = owl : sameAs$

Federated computation of statistics

1. At each source the CSs are computed

film:1129 movie:film_subject film_subject:444 .

film:1129 dcterms:title "Kate & Leopold" .

film:1129 owl:sameAs dbr:Kate_&_Leopold .

film:16189 movie:film_subject film_subject:444 .

film:16189 dcterms:title "Journey to the Center of Time" .

film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .

...

$$C_{\text{LMDB},i} = \{ \text{movie:film_subject}, \text{owl:sameAs}, \text{dcterms:title}, \dots \}$$

...

Federated computation of statistics

1. At each source the CSs are computed
2. At each source local statistics are computed

film:1129 movie:film_subject film_subject:444 .

film:1129 dcterms:title "Kate & Leopold" .

film:1129 owl:sameAs dbr:Kate_&_Leopold .

film:16189 movie:film_subject film_subject:444 .

film:16189 dcterms:title "Journey to the Center of Time" .

film:16189 owl:sameAs dbr:Journey_to_the_Center_of_Time .

...

$$C_{\text{LMDB},i} = \{ \text{movie:film_subject}, \text{owl:sameAs}, \text{dcterms:title}, \dots \}$$

...

$$\text{local_subjects}_{\text{LMDB}}(C_{\text{LMDB},i}) = \{ \text{film:1129}, \text{film:16189}, \dots \}$$

$$\text{local_objects}_{\text{LMDB}}(\text{owl:sameAs}, C_{\text{LMDB},i}) = \{ \text{dbr:Kate_}\&_Leopold, \text{dbr:Journey_to_the_Center_of_Time}, \dots \}$$

Federated computation of statistics

3. Local statistics and CSs are transferred to the federated query engine

```

local_subjectsLMDB(CLMDB,i) = { film:1129, film:16189, ... }
local_objectsLMDB(owl:sameAs, CLMDB,i) = { ..., dbr:Journey_to_the_Center_of_Time, ... }
local_subjectsDBpedia(CDBpedia,j) = { ..., dbr:Journey_to_the_Center_of_Time, ... }
local_objectsDBpedia(dbo:director, CDBpedia,j) = { dbr:David_L..Hewitt, ... }
...

```

Federated computation of statistics

- Local statistics and CSs are transferred to the federated query engine
- Set overlap is computed to estimate the statistics of federated CSs and CPs

$\text{local_subjects}_{\text{LMDB}}(C_{\text{LMDB},i}) = \{ \text{film:1129, film:16189, ... } \}$

$\text{local_objects}_{\text{LMDB}}(\text{owl:sameAs}, C_{\text{LMDB},i}) = \{ \dots, \text{dbr:Journey_to_the_Center_of_Time}, \dots \}$

$\text{local_subjects}_{\text{DBpedia}}(C_{\text{DBpedia},j}) = \{ \dots, \text{dbr:Journey_to_the_Center_of_Time}, \dots \}$

$\text{local_objects}_{\text{DBpedia}}(\text{dbo:director}, C_{\text{DBpedia},j}) = \{ \text{dbr:David_L_Hewitt}, \dots \}$

...

$(C_{\text{LMDB},i}, C_{\text{DBpedia},j}, \text{owl:sameAs})$

Odyssey's Query Optimization

```
SELECT DISTINCT * WHERE {  
  ?film dbo:director ?director .           (tp1)  
  ?film rdf:type dbo:Film .                 (tp2)  
  ?movie owl:sameAs ?film .               (tp3)  
  ?movie dcterms:title ?title .            (tp4)  
  ?movie movie:film_subject film_subject:444 (tp5)  
}
```

Odyssey's Query Optimization

1. Identify the star-shaped subqueries

```
SELECT DISTINCT * WHERE {  
  ?film dbo:director ?director .           (tp1)  
  ?film rdf:type dbo:Film .                 (tp2)  
  ?movie owl:sameAs ?film .               (tp3)  
  ?movie dcterms:title ?title .            (tp4)  
  ?movie movie:film_subject film_subject:444 (tp5)  
}
```


Odyssey's Query Optimization

1. Identify the star-shaped subqueries
2. Identify the relevant CSs and estimate cardinality

```

SELECT DISTINCT * WHERE {
  ?film dbo:director ?director .           ( tp1 )
  ?film rdf:type dbo:Film .               ( tp2 )
  ?movie owl:sameAs ?film .             ( tp3 )
  ?movie dcterms:title ?title .           ( tp4 )
  ?movie movie:film_subject film_subject:444 ( tp5 )
}

```

$C_{DBpedia,j} = \{ \text{dbo:director, rdf:type, ...} \}$
 estimatedCardinality($\{ \text{dbo:director, rdf:type} \}$)=162
 $C_{LMDB,i} = \{ \text{movie:film_subject, owl:sameAs, dcterms:title, ...} \}$
 estimatedCardinality($\{ \text{movie:film_subject, owl:sameAs, dcterms:title} \}$)=4

Odyssey's Query Optimization

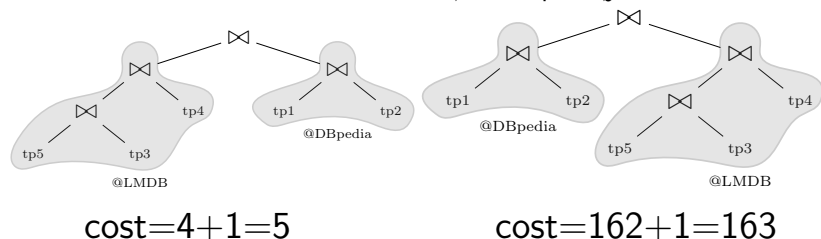
3. Identify the relevant CPs and estimate cardinality.

$$\text{estimatedCardinality}((C_{\text{LMDB},i}, C_{\text{DBpedia},j}, \text{owl:sameAs})) = 1$$

Odyssey's Query Optimization

3. Identify the relevant CPs and estimate cardinality.
4. A cost function is used to select the plan that leads to transfer the lower number of tuples.

$(C_{\text{LMDB},i}, C_{\text{DBpedia},j}, \text{owl:sameAs})$
 estimatedCardinality $((C_{\text{LMDB},i}, C_{\text{DBpedia},j}, \text{owl:sameAs}))=1$

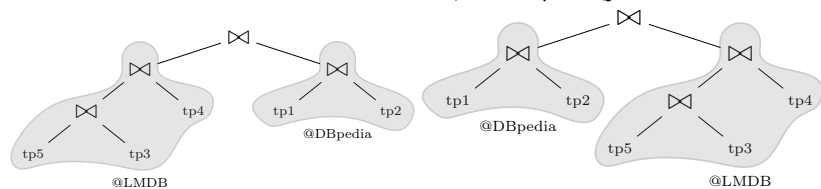


Odyssey's Query Optimization

3. Identify the relevant CPs and estimate cardinality.
4. A cost function is used to select the plan that leads to transfer the lower number of tuples.
5. If necessary, compute join ordering using Dynamic Programming.

$(C_{\text{LMDB},i}, C_{\text{DBpedia},j}, \text{owl:sameAs})$

$\text{estimatedCardinality}((C_{\text{LMDB},i}, C_{\text{DBpedia},j}, \text{owl:sameAs}))=1$



$$\text{cost}=4+1=5$$

$$\text{cost}=162+1=163$$

Odyssey provides a better plan

```

SELECT DISTINCT * WHERE {
  ?film dbo:director ?director .           (tp1)
  ?film rdf:type dbo:Film .                (tp2)
  ?movie owl:sameAs ?film .              (tp3)
  ?movie dcterms:title ?title .           (tp4)
  ?movie movie:film_subject film_subject:444 (tp5)
}

```

⁵A. Schwarte et al. "FedX: Optimization Techniques for Federated Query Processing on Linked Data". In: *ISWC'11*.

⁶A. Charalambidis et al. "SemaGrow: optimizing federated SPARQL queries". In: *SEMANTICS'15*.

⁷G. Montoya et al. "The Odyssey Approach for Optimizing Federated SPARQL Queries". In: *ISWC'17*.

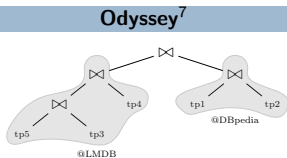
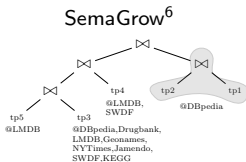
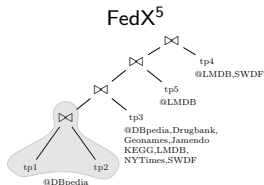
Odyssey provides a better plan

```

SELECT DISTINCT * WHERE {
  ?film dbo:director ?director .           (tp1)
  ?film rdf:type dbo:Film .               (tp2)
  ?movie owl:sameAs ?film .             (tp3)
  ?movie dcterms:title ?title .          (tp4)
  ?movie movie:film_subject film_subject:444 (tp5)
}

```

Plan

OT
ET0.74s
142s4.75s
6.93s0.22s
1.30s

⁵A. Schwarte et al. "FedX: Optimization Techniques for Federated Query Processing on Linked Data". In: *ISWC'11*.

⁶A. Charalambidis et al. "SemaGrow: optimizing federated SPARQL queries". In: *SEMANTICS'15*.

⁷G. Montoya et al. "The Odyssey Approach for Optimizing Federated SPARQL Queries". In: *ISWC'17*.

Experimental setup

- ▶ Queries and datasets from FedBench⁸.
- ▶ Comparison with existing approaches FedX⁹, SemaGrow¹⁰, SPLENDID¹¹, HiBISCuS¹².
- ▶ A Virtuoso 7.2.4.2 endpoint was deployed for each dataset.
- ▶ Plots present the average over nine executions with a timeout of 30m.

⁸M. Schmidt et al. "FedBench: A Benchmark Suite for Federated Semantic Data Query Processing". In: *ISWC'11*.

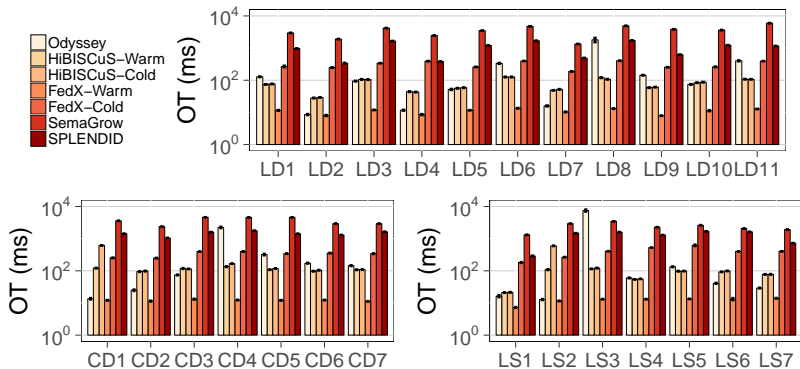
⁹A. Schwarte et al. "FedX: Optimization Techniques for Federated Query Processing on Linked Data". In: *ISWC'11*.

¹⁰A. Charalambidis et al. "SemaGrow: optimizing federated SPARQL queries". In: *SEMANTICS'15*.

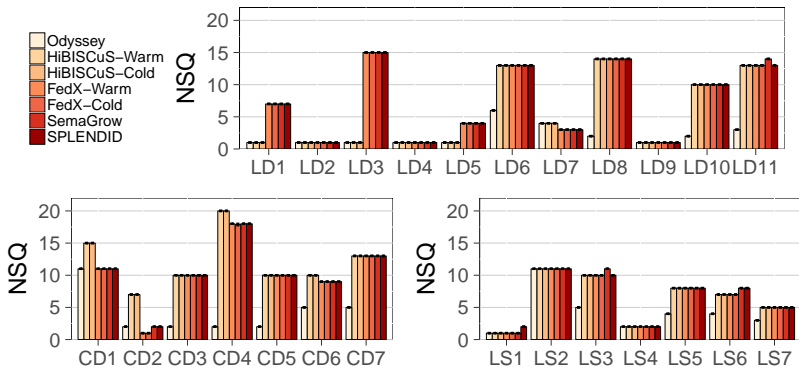
¹¹O. Görlitz and S. Staab. "SPLENDID: SPARQL Endpoint Federation Exploiting VOID Descriptions". In: *COLD'11*.

¹²M. Saleem and A. N. Ngomo. "HiBISCuS: Hypergraph-Based Source Selection for SPARQL Endpoint Federation". In: *ESWC'14*.

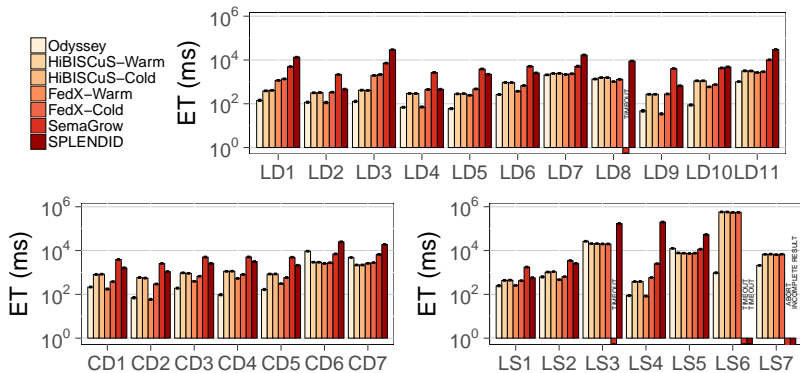
Odyssey's Optimization Time is comparable with other approaches'



Odyssey's plans have less subqueries than other approaches' plans



Odyssey's plans are overall executed faster than other approaches' plans



Conclusions

- ▶ Odyssey uses cardinality estimations that allow for better optimizations.
- ▶ Local statistics allow to discover connections between datasets in a federated setup.
- ▶ Odyssey's plans are in general better than existing approaches' plans.

Future works

- ▶ Odyssey's optimization can be improved using with runtime optimizations (ASK queries).
- ▶ Reduce the local statistics computation times and sizes.
- ▶ Provide efficient strategies to update the statistics.

Questions?