

TriAD: A Distributed Shared-Nothing RDF Engine based on Asynchronous Message Passing

Sairam Gurajada (MPII), Stephan Seufert (MPII), Iris Miliaraki (Yahoo), Martin Theobald (UAntwerp)

Motivation

RDF is ubiquitous...

- Many organizations now support and publish RDF data
Eg. DBpedia (400M), YAGO2 (130M), Freebase,...

How to scale and yet be efficient..?

.. in managing and querying RDF data

Our Approach: TriAD RDF store

- Scalability – main-memory backed distributed setting
- Efficiency – 1. Asynchronous query processing
2. Join-ahead pruning

Problems with Current Approaches

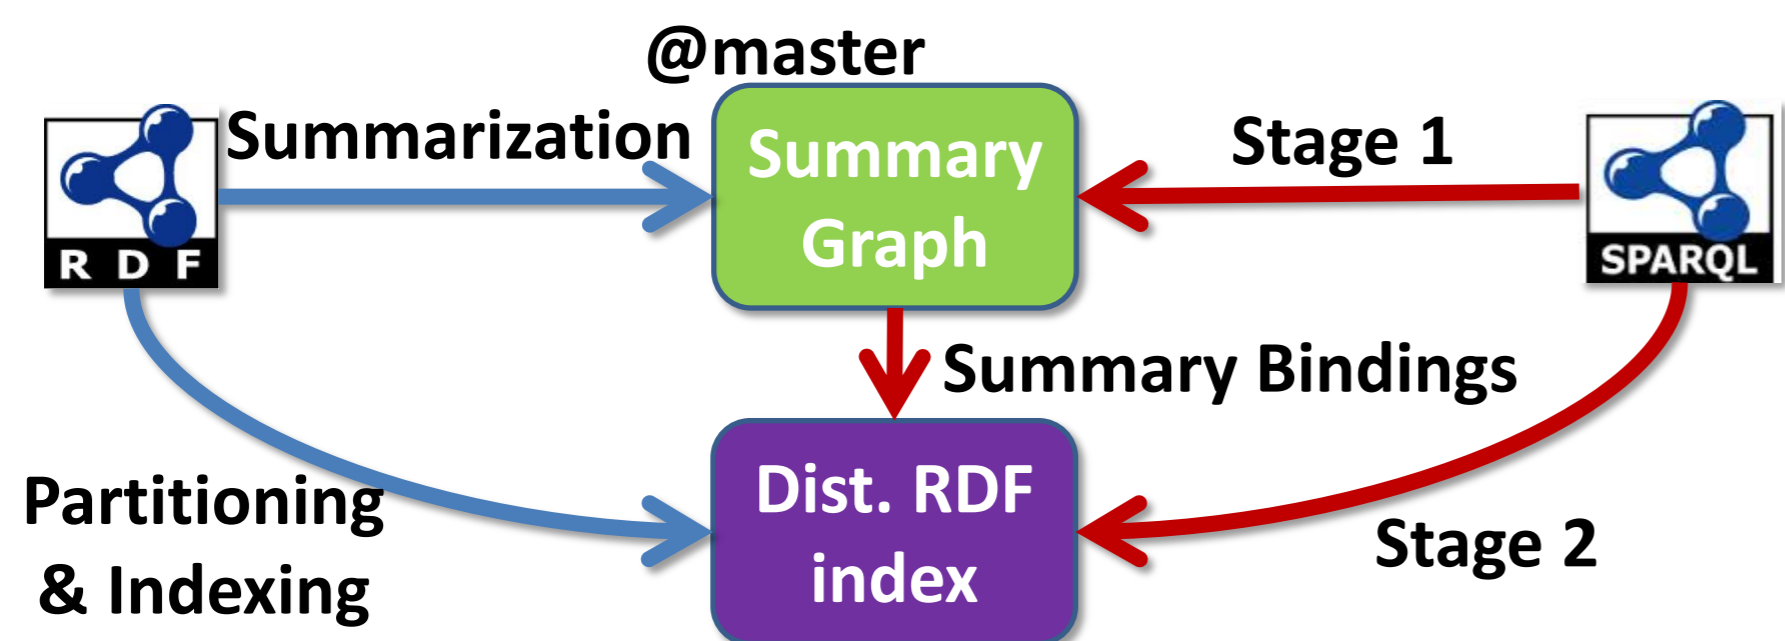
Problem 1: Synchronous processing of joins

Hadoop-based systems process joins using iterative & synchronous MapReduce jobs

Problem 2: Pruning dangling triples

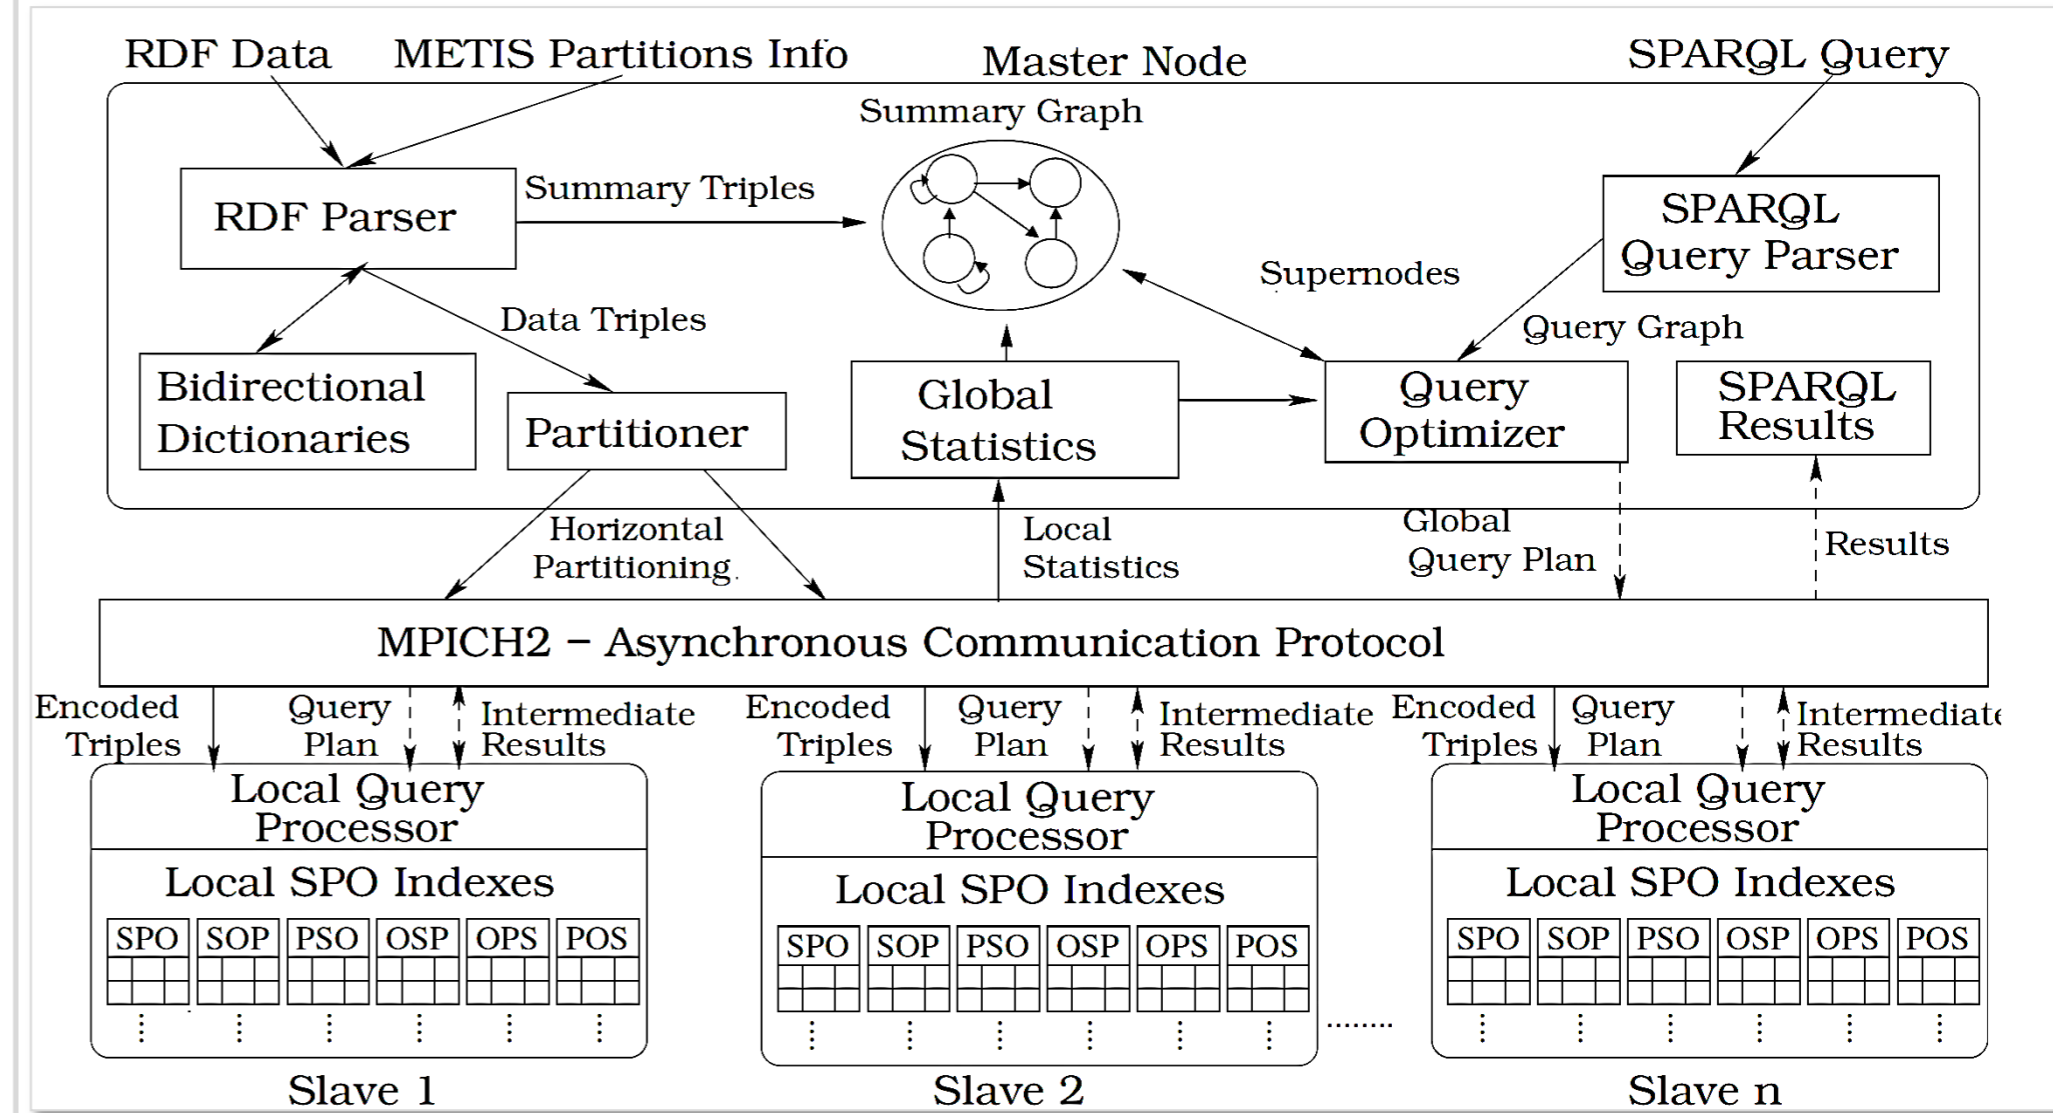
- Dangling triples are the triples that appear in intermediate relations but not part of final joins
- Approaches: Graph exploration, Sideways Information Passing (SIP)
- Graph exploration – works well for selective queries
SIP – needs synchronization among join operators

Our Approach (in a nut shell)



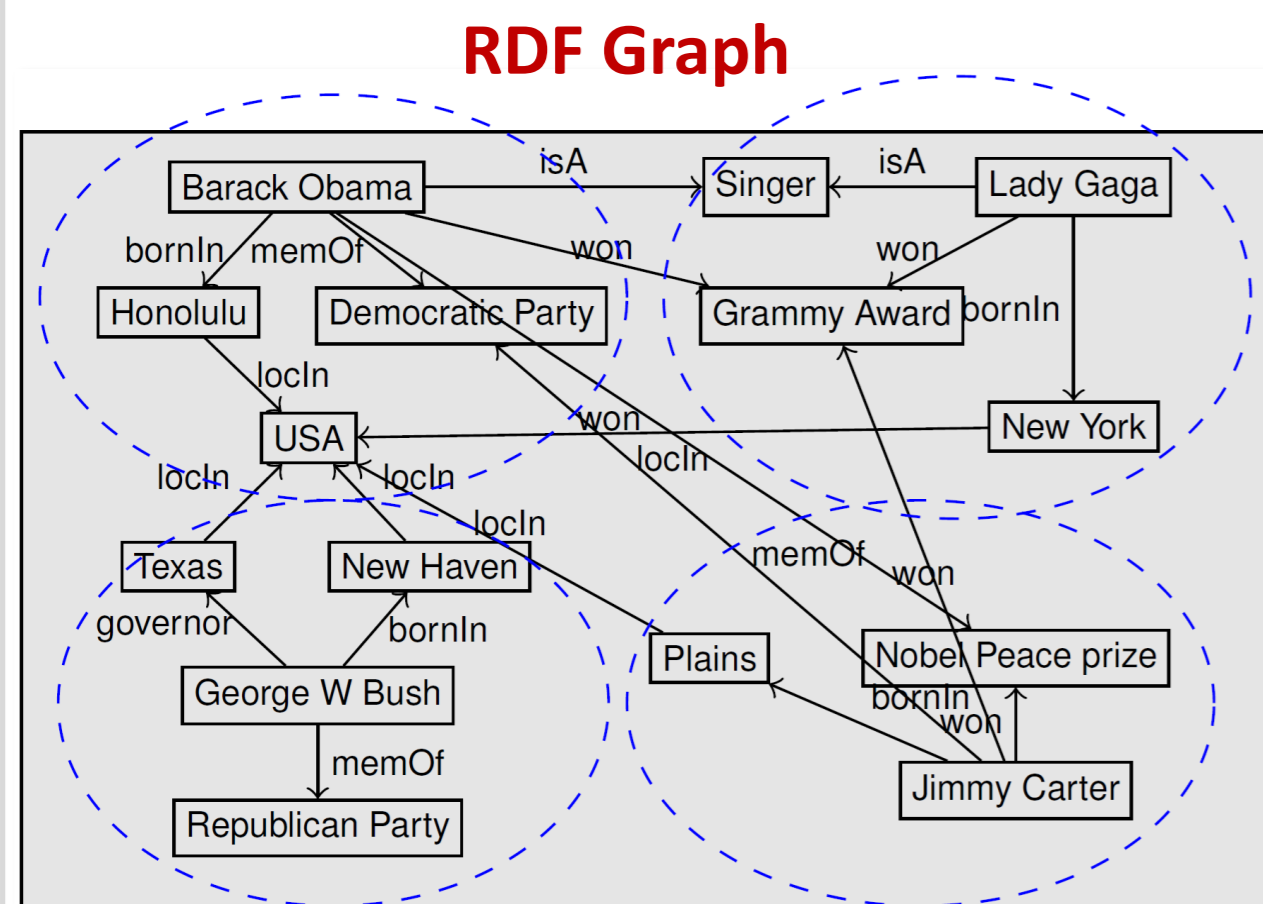
- Summary index (via summary bindings) is used for pruning dangling triples (**join-ahead pruning**)
- Asynchronous distributed query execution over RDF index

TriAD Architecture



Graph Summarization & Query Processing Workflow

Graph summarization



Cost function for optimal partitions

$$C_{Q,n} := C_S + C_{P,n}$$

$$= \frac{d|V_S|}{|E_D|} \cdot c_D + \frac{\lambda}{|V_S|} \cdot \frac{c_D}{n}$$

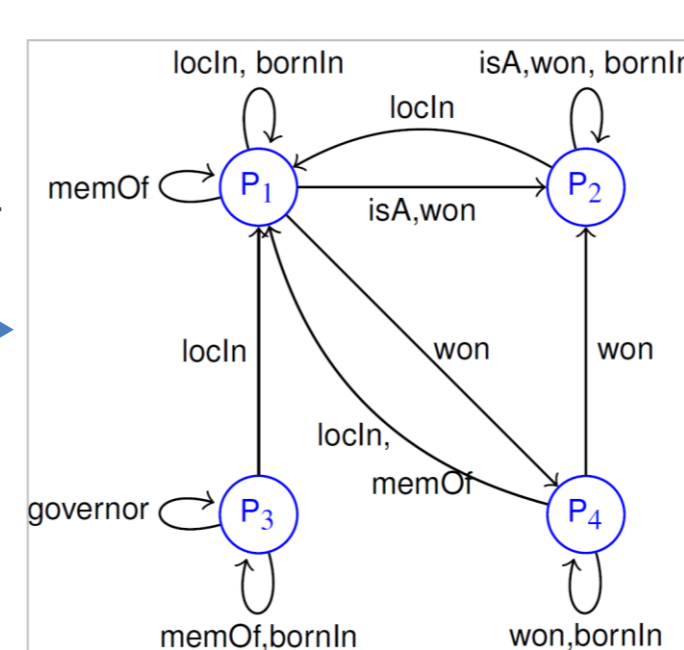
Optimal summary size

$$|V_S| := \sqrt{\frac{\lambda|E_D|}{dn}}$$

λ - tuning parameter

Stage 1: Join-ahead pruning (of dangling triples)

Summary Graph



SPARQL Query

```
SELECT ?person ?city ?prize WHERE{
  ?person <bornIn> ?city .
  ?city <locatedIn> USA .
  ?person <won> ?prize .
  ?prize <hasName> ?name }
```

Graph Exploration

Bindings

```
?person: P1, P2, P4
?city: P1, P2, P4
?prize: P2, P4
```

Stage 2: Distributed & asynchronous query execution

