

MapReduce for Efficient SPARQL Queries

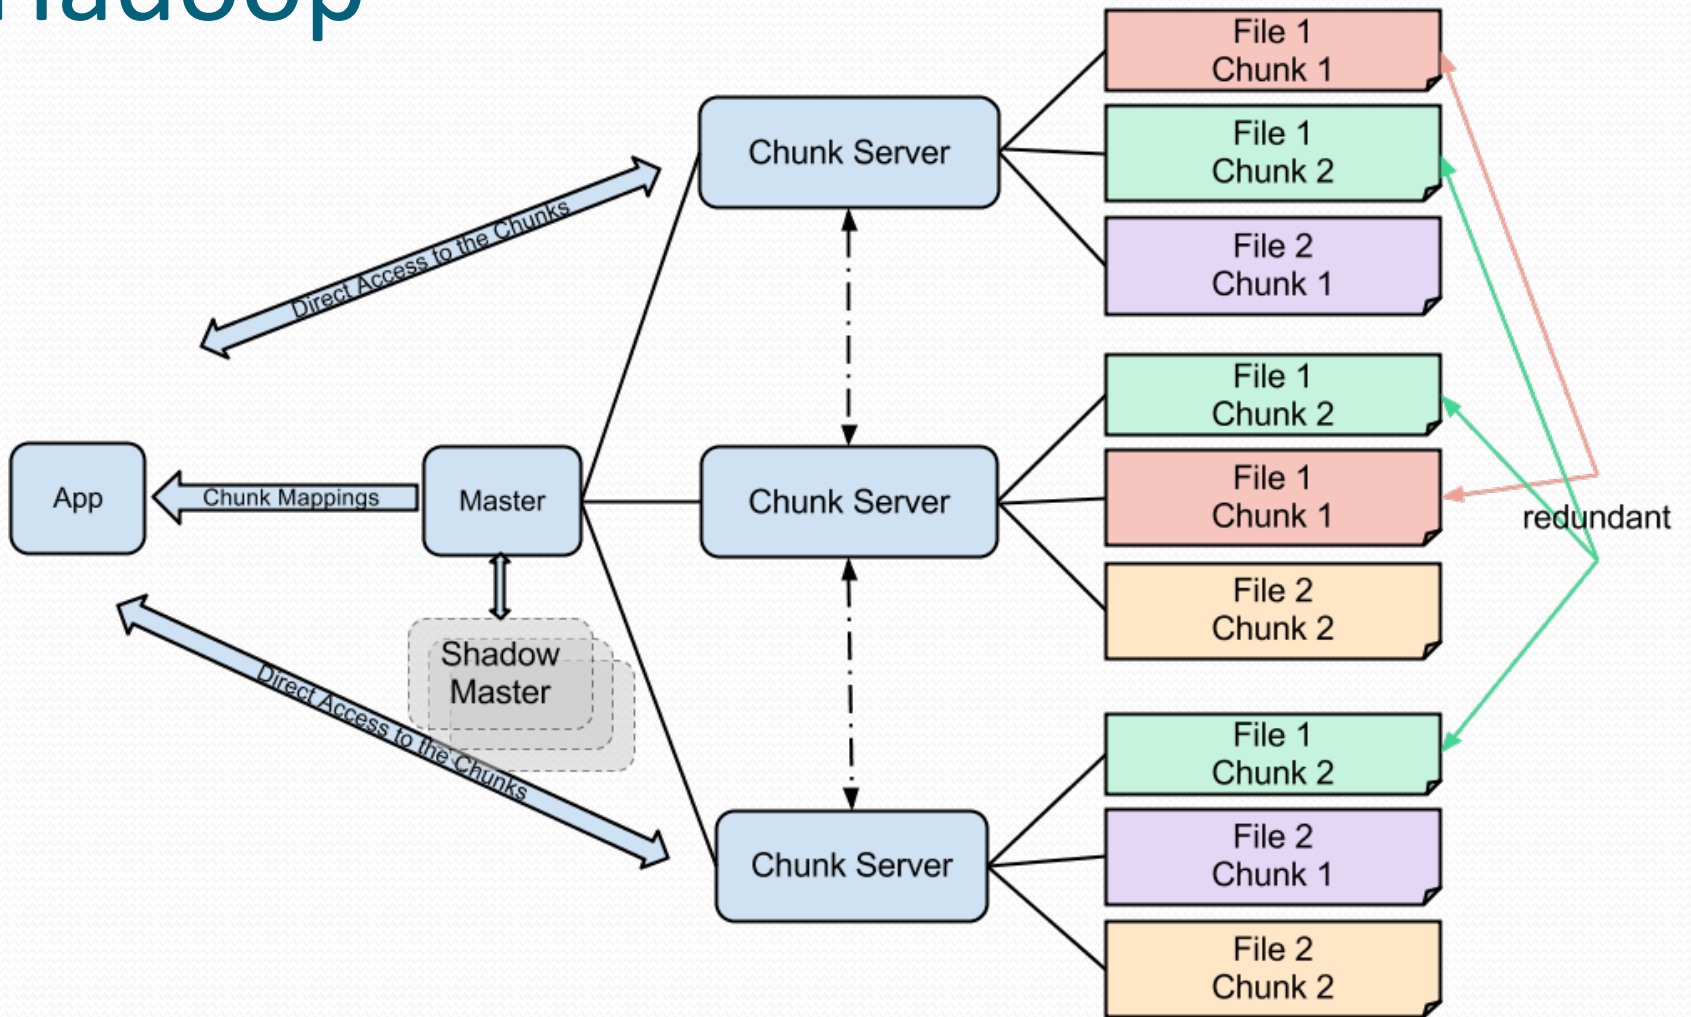
Hassan Jamous

A.M: 783

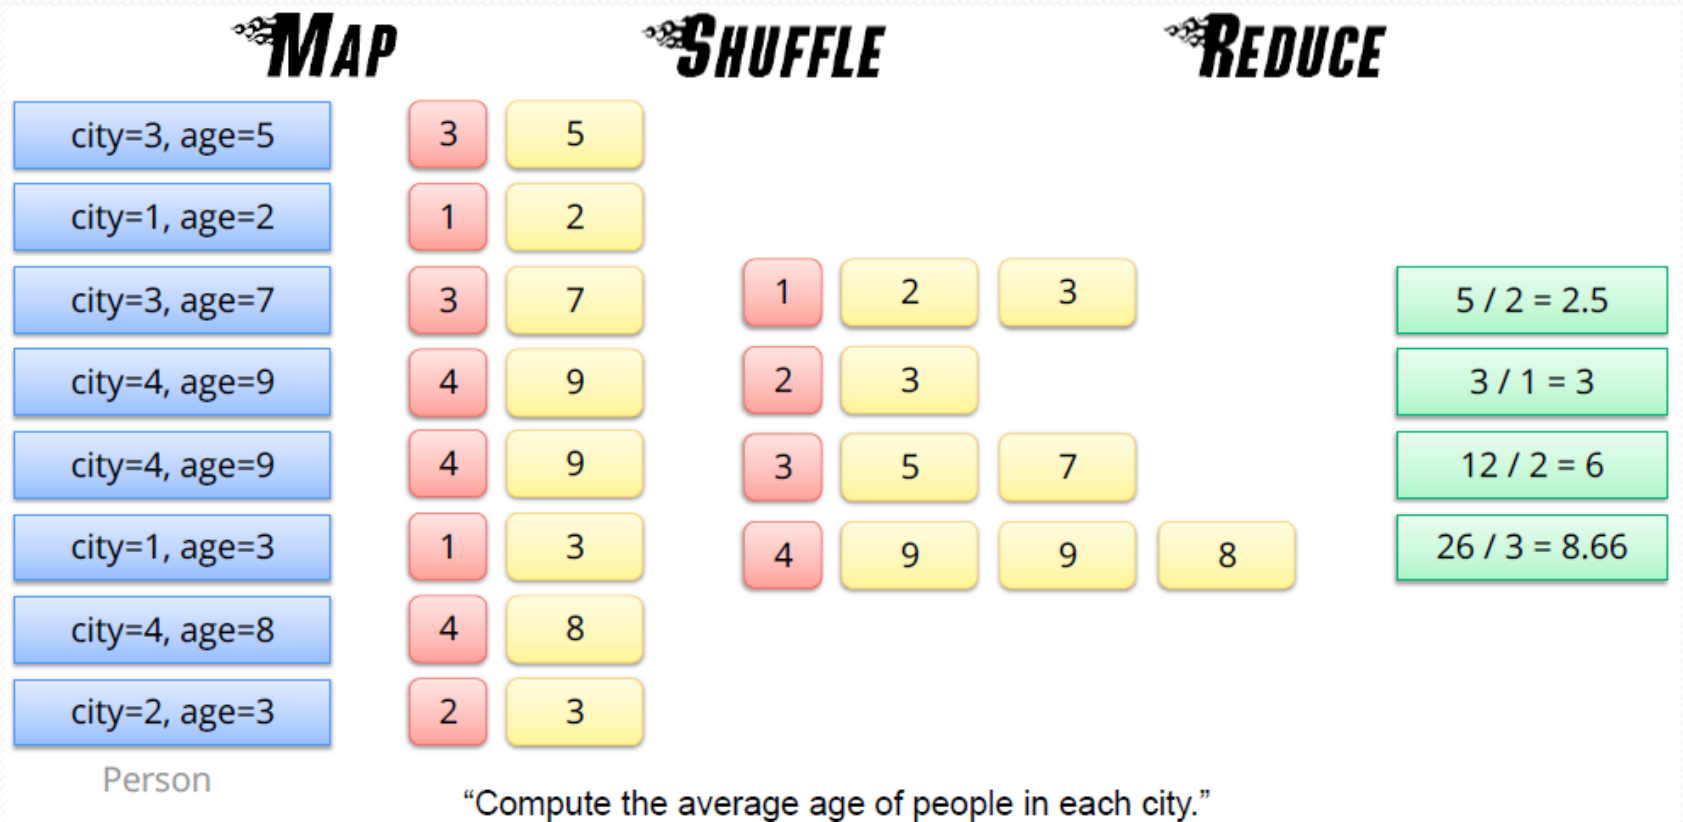
Overview

- Google File System & Apache Hadoop
- MapRedcue
- MapReduce Architecture
- RDF Structure
- Graph Pattern Matching
- SPARQL, Hadoop and MapReduce
- Reduce MapReduce cycles – Pig
- Nested TripleGroup Algebra
- MapReduce Tables Join
- Experiment Results

Google File System & Apache Hadoop

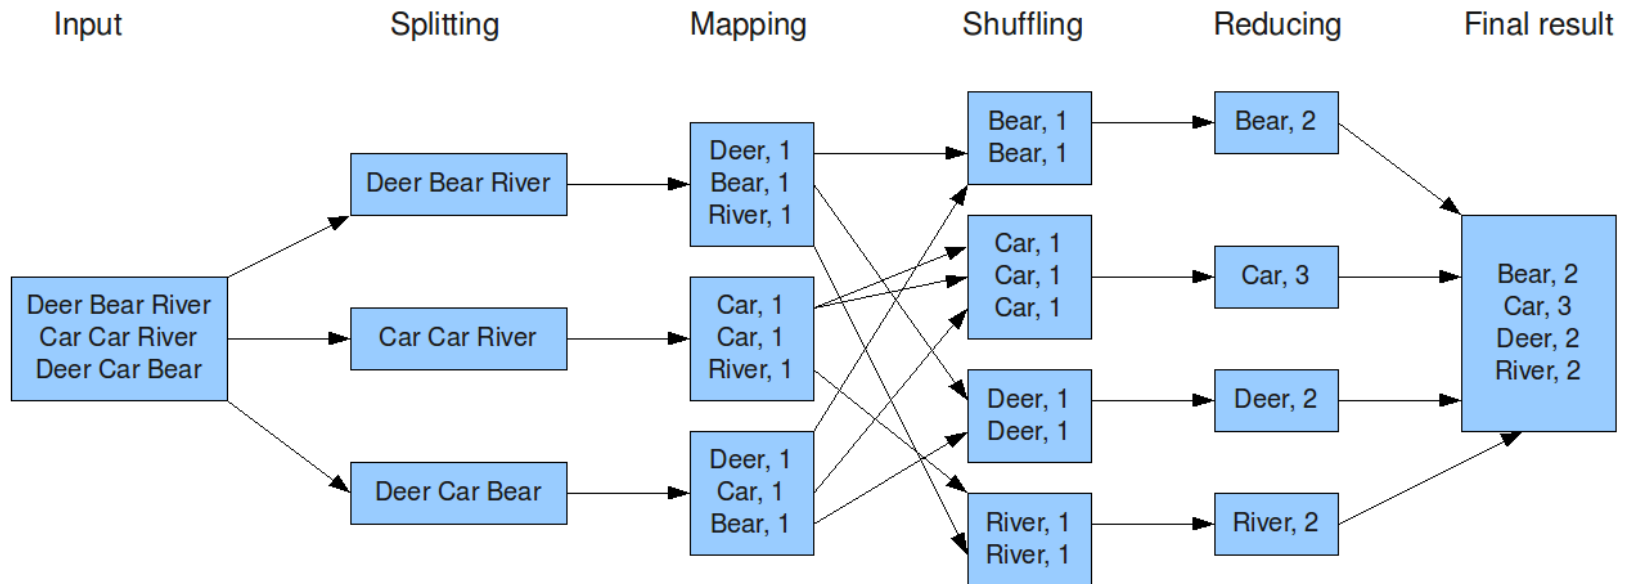


MapReduce

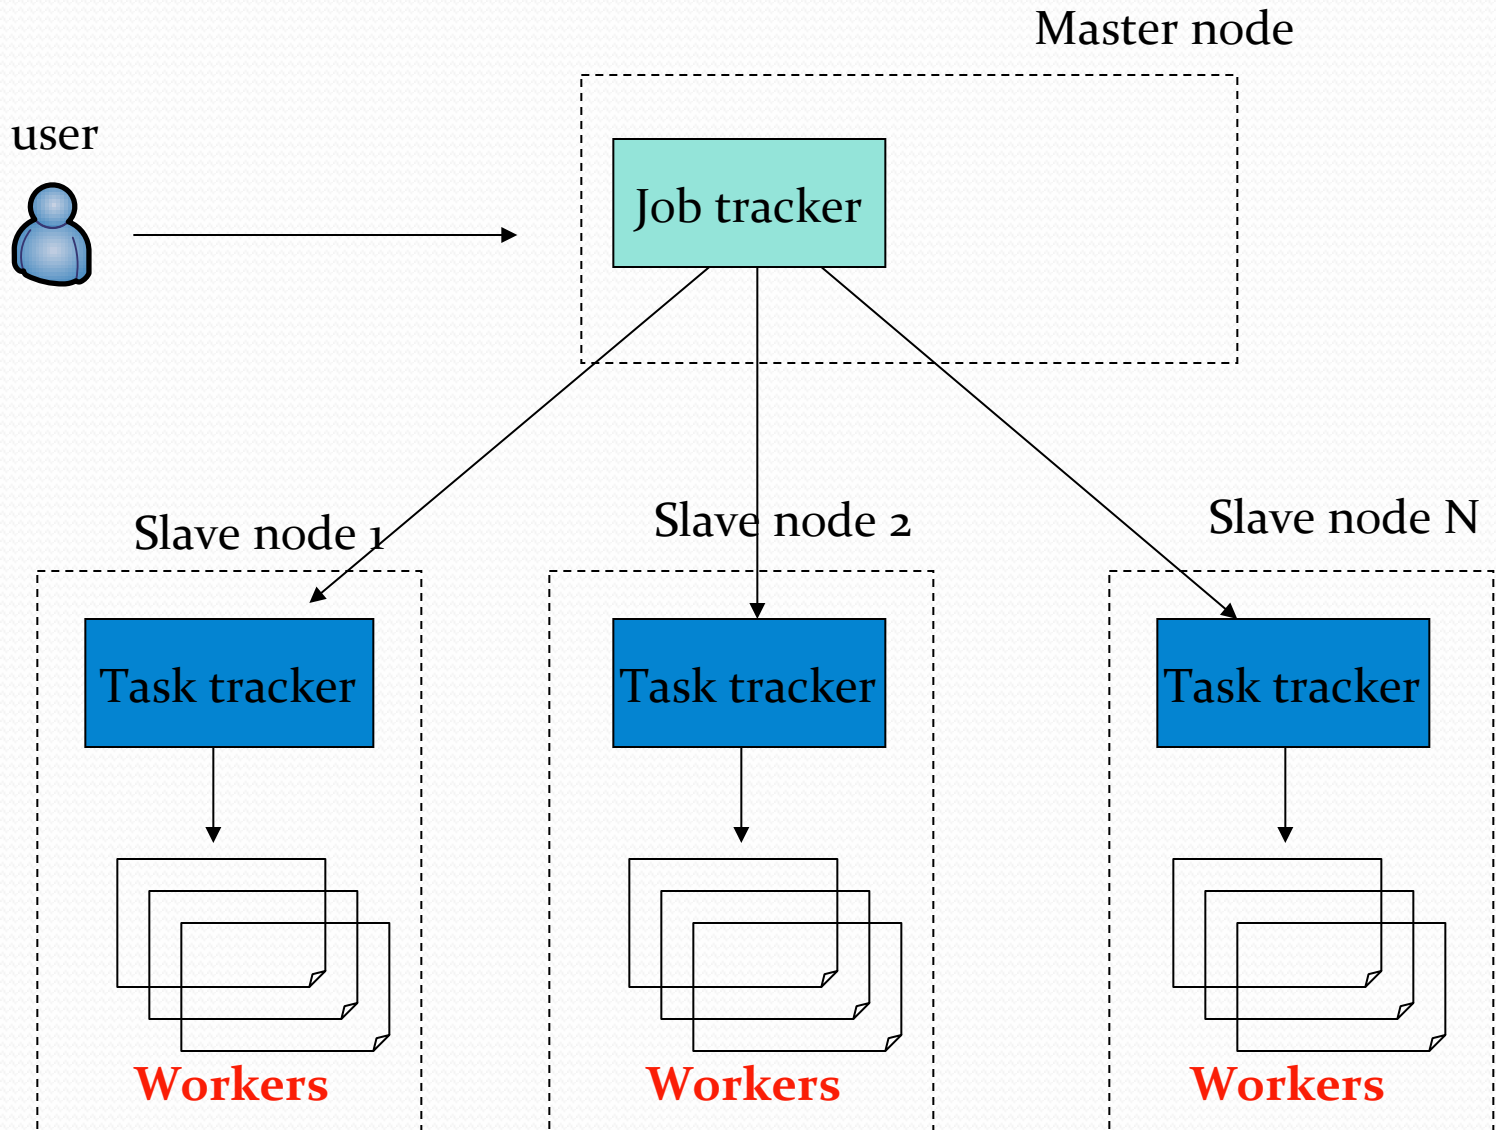


MapReduce

The overall MapReduce word count process



MapReduce Architecture



Comments

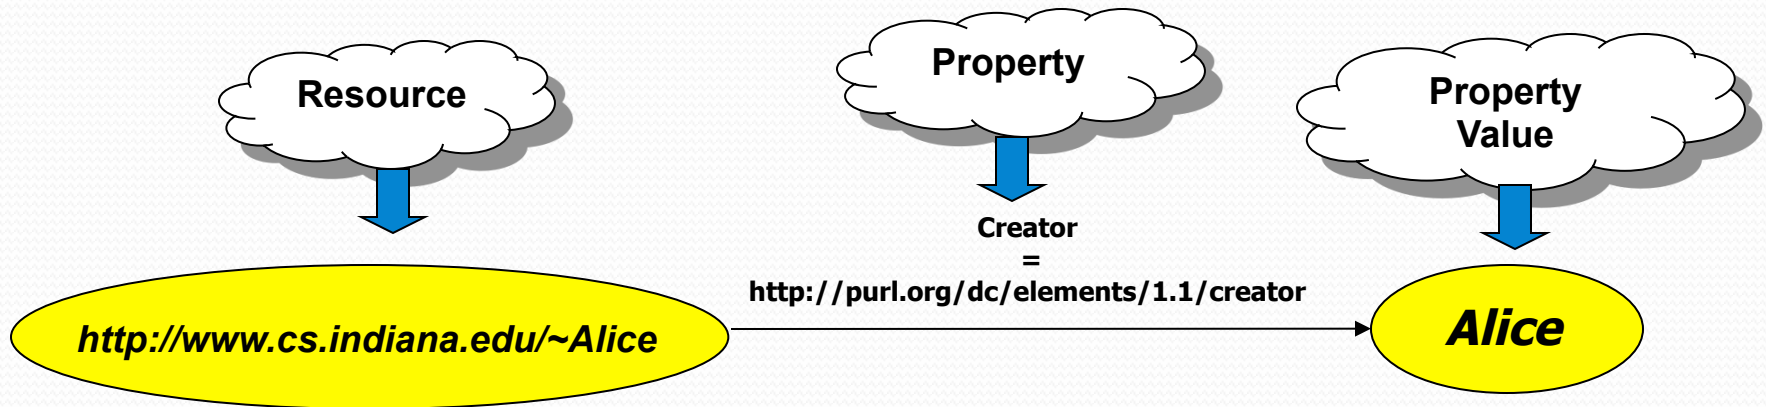
- MapReduce as a computing paradigm is not new in the field of Distributed system.
- Google gives a standard way for this kind of processing.
- The MapReduce main power is in implementation of the JobTracker and TaskTracker .

Questions



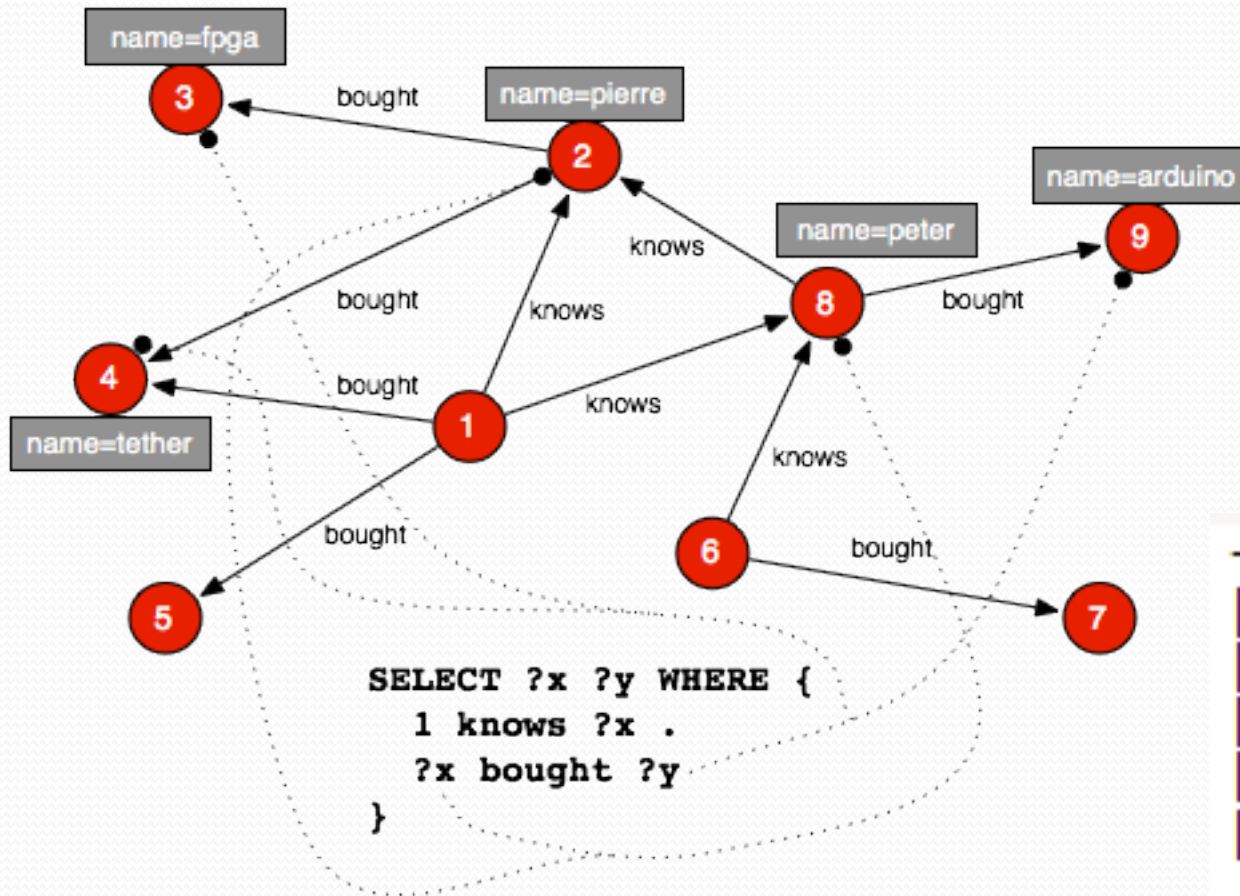
RDF Structure

Alice is the creator of the resource <http://www.cs.indiana.edu/~Alice>.



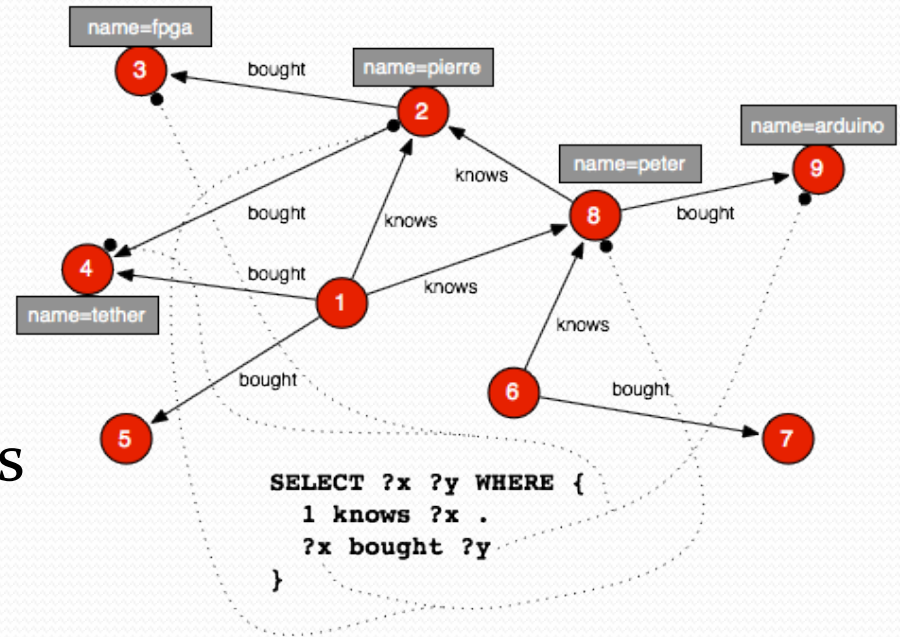
- Can you imagine the number of triples.

Graph Pattern Matching



SPARQL Engine & Join

- Compute first statement then the second one.
- Compute both statements then Join.



1	Bought	4
1	Knows	8
1	Knows	2
...

SPARQL, Hadoop and MapReduce

- There are very large number of RDF Triples.
- We can use Hadoop technology to store this amount of data.
- MapReduce is optimal to query data from Hadoop Store.
- Can we convert SPARQL Query to MapReduce.

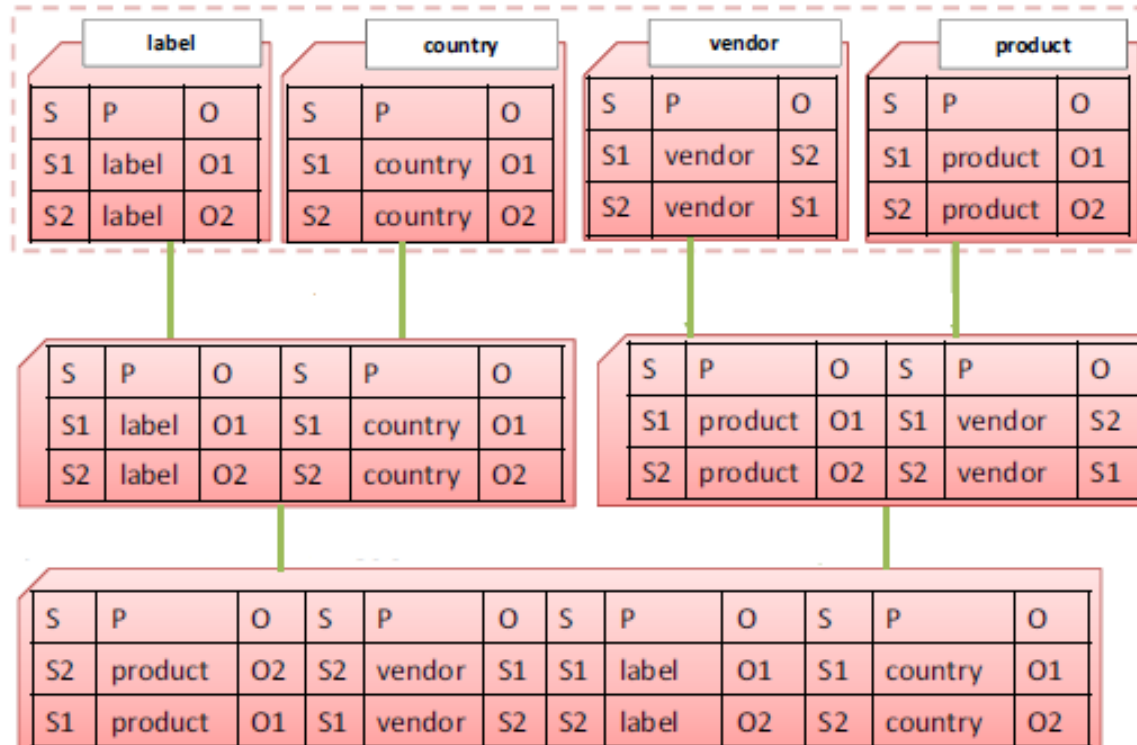
SPARQL & MapReduce - 1

- Retrieve the name of Vendors, the country where the vendors are located, and their product Offers
- One MapReduce request for each statement.
- One MapReduce for each Join
- Total of 7 in this example.

```
SELECT ?vlabel, ?vcountry ,?prod
WHERE { ?v :label ?vlabel.
        ?v :country ?vcountry.
        ?o :vender ?v.
        ?o :product ?prod. }
```

SPARQL & MapReduce - 1

```
SELECT ?vlabel, ?vcountry, ?vprod
WHERE { ?v :label ?vlabel.
       ?v :country ?vcountry.
       ?o :vender ?v.
       ?o :product ?vprod. }
```



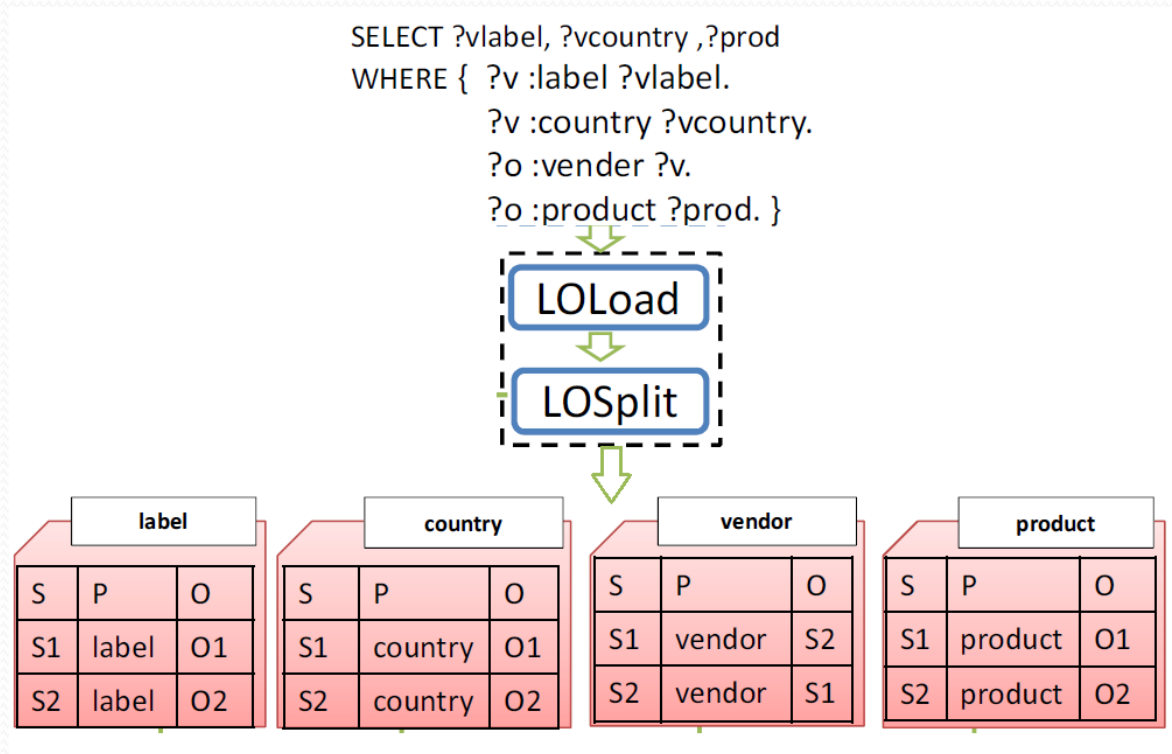
MapReduce cycles and mid results

- The previous solution is very simple
- There is a problem in the number of MapReduce cycles.
- There is a problem in storing the mid results
- we should improve the number of MapReduce cycles and the size of mid results.

Reduce MapReduce cycles - Pig

- Apache Pig has a LOLoad and LOSplit operations.
- allow to load the first set of data in one cycle.

- Reduce the number from 4 to 1.



Comments

- `SELECT ?prod`
Where { `?v :lable 'any value' . ?o :vender ?v`
`?o :product ?prod`}
- `?o :vender ?v` , will have too many mid results.
- We have to choose between running on parallel or reduce the mid results.

Questions



Nested TripleGroup Algebra

- A TripleGroup tg is a relation of triples t_1, t_2, \dots, t_k , whose schema is defined as (S, P, O) .
- Any two triples $t_i, t_j \in tg$ have overlapping components i.e. $t_i [col_i] = t_j [col_j]$ where col_i, col_j refer to subject or object component.
- In our example we have 2 triple groups.

```
SELECT ?vlabel, ?vcountry, ?prod
WHERE { ?v :label ?vlabel.
        ?v :country ?vcountry.
        ?o :vender ?v.
        ?o :product ?prod. }
```

Nested TripleGroup Algebra – MapReduce cycles.

- NTGA provide provide a GroupBy operation which find all triple groups. (we have 2 groups)
- We can get all groups in one MapReduce cycle.
- In this example we have only 2 MapReduce Cycles.

```
SELECT ?vlabel, ?vcountry ,?prod
WHERE { ?v :label ?vlabel.
        ?v :country ?vcountry.
        ?o :vender ?v.
        ?o :product ?prod. }
```

s1	:label	o1
	:country	o2
s2	:product	o3
	:vender	s1

s2	:product	o3		
	:vender	s1	:label	o1
			:country	o2

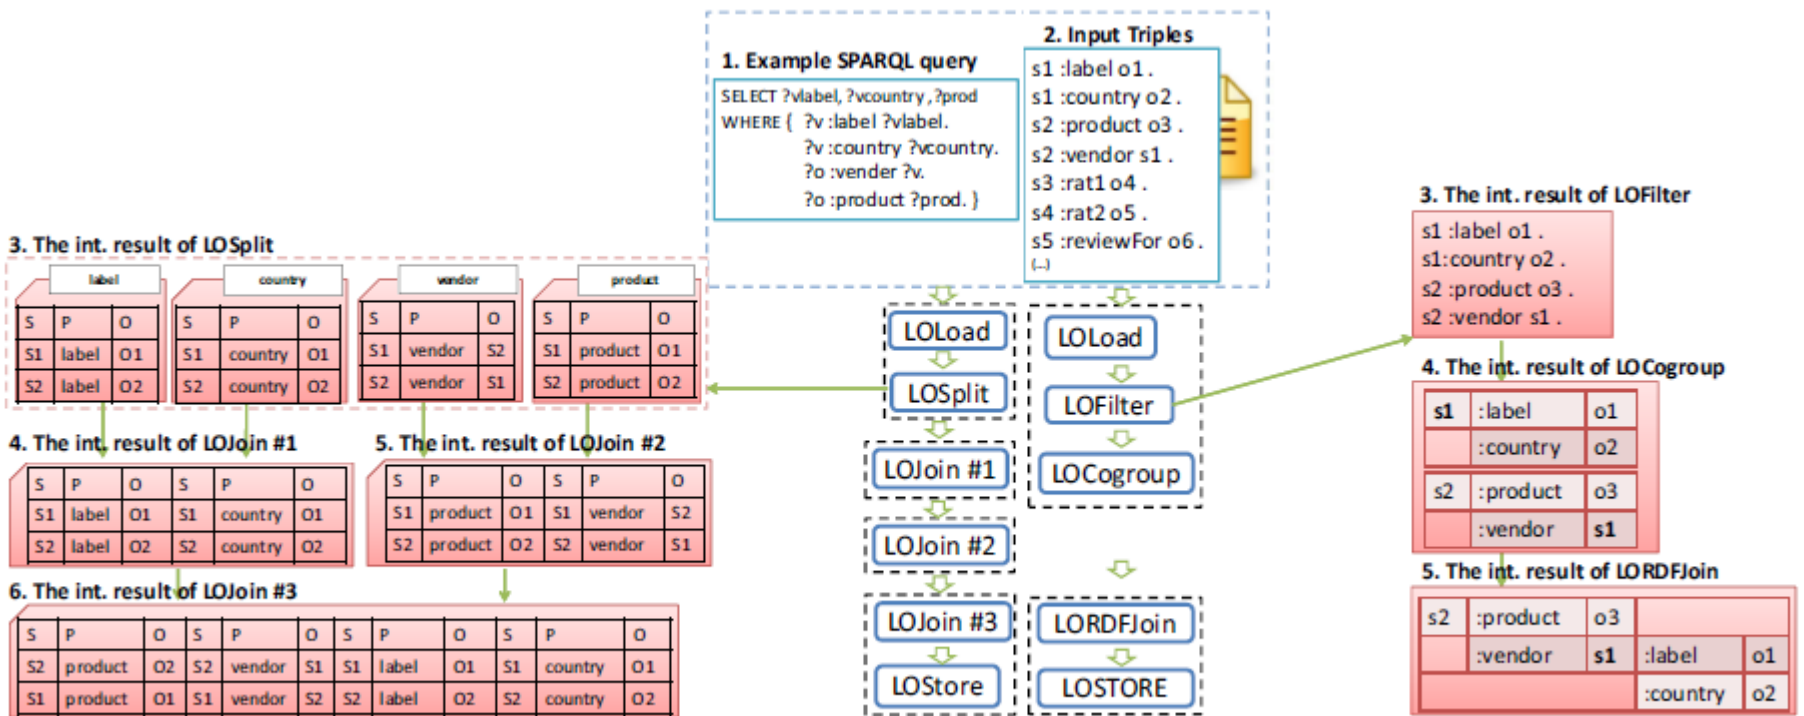
Test Example

```
SELECT ?hpage, ?price, ?rat1
WHERE { ?v homepage ?hpage.
        ?v country ?vcountry.
        ?o vendor ?v .
        ?o price ?price .
        ?o product ?prod .
        ?r revFor ?prod .
        ?r reviewer ?rev .
        ?r rating1 ?rat1 .}
```

RAPID+

- RAPID+ is an extended Apache Pig system that integrates the NTGA data model and algebra
- It provides operations like: LOLoad, LOFilter, LOCogroup, LORDFJoin.

Apache Pig vs RAPID+



(a) With Pig Latin Operators.

(b) With NTGA Operators.

RAPID+ Filter Example

- RAPID+ has TripleStorage operation
- This operation specially for RDF
- It filters data during the load operation.
- This reduce execution time and mid results data size.

```
SELECT
?vlabel ?prod ?price
where {
?v :label ?vlabel.
?v :country ?vcountry.

?o :vender ?v.
?o :price ?itsprice.
?o :delivDays ?days.
?o :product ?prod.

?r :reviewFor ?prod.
?r :rating1 ?rat1.

FILTER (?days <= 2 &&
?vcountry = "US" &&
?rat1 >= 4)}
```

Comments

- Generating Map and Reduce function is more complicated.
- Executing the GROUPBY MapReduce function will take more main memory.
- Filtering data during the load process is great improvement.

Questions



MapReduce Tables Join.

- Normal database can do join magically by defining primary key and foreign key
- There are no relation between tables in Hadoop.
- In order to do Join between tables you should implement this manually in your Map and Reduce functions

Reduce Side Join.

- Customer Information:
CustomerID, NAME, Phone

```
1, Stephanie Leung, 555-555-5555
2, Edward Kim, 123-456-7890
3, Jose Madriz, 281-330-8004
4, David Stork, 408-555-0000
```

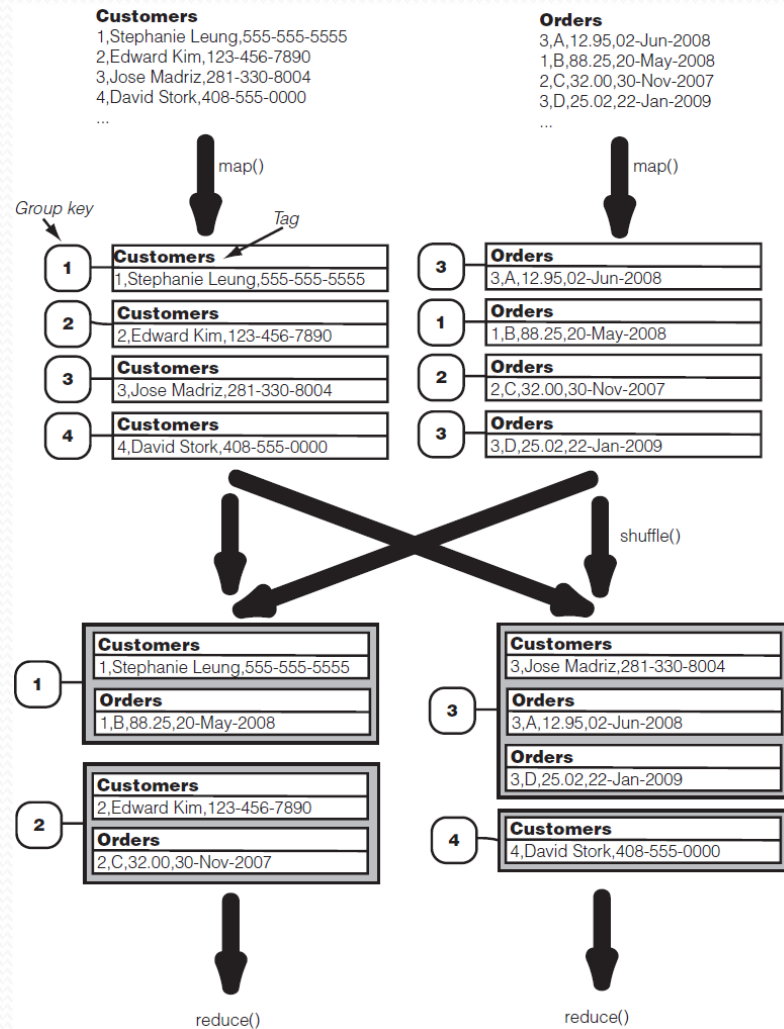
- Orders Information:
CustomerID, OrderID, Price,
Purchase Date.

```
3, A, 12.95, 02-Jun-2008
1, B, 88.25, 20-May-2008
2, C, 32.00, 30-Nov-2007
3, D, 25.02, 22-Jan-2009
```

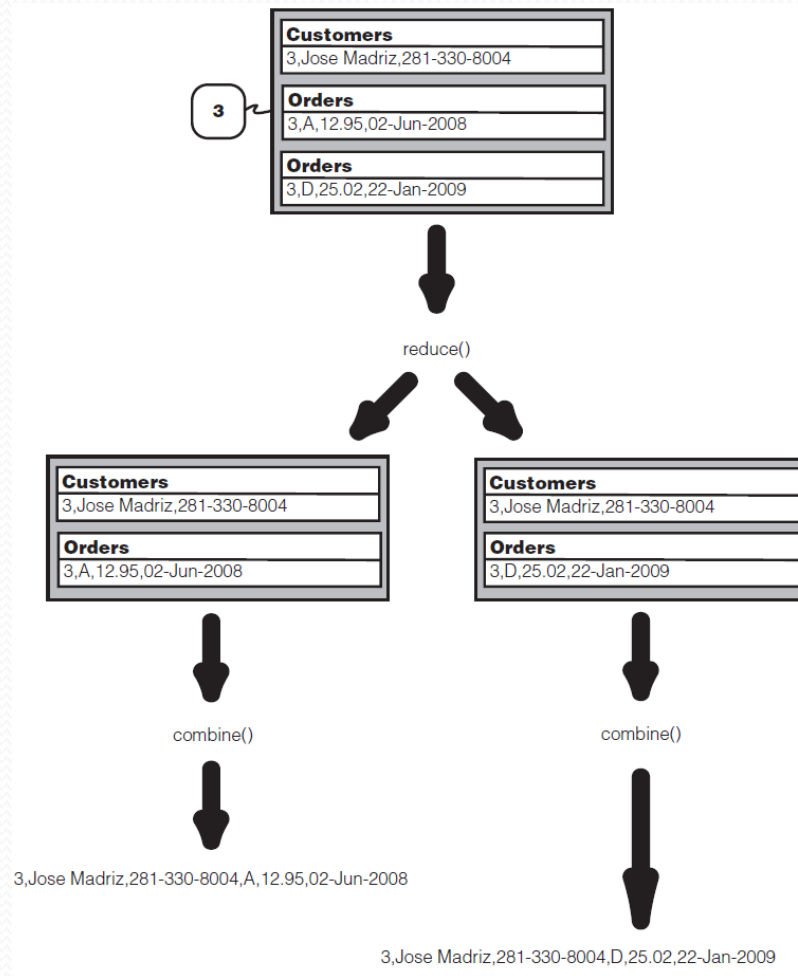
- Join based on
CustomerID

```
1, Stephanie Leung, 555-555-5555, B, 88.25, 20-May-2008
2, Edward Kim, 123-456-7890, C, 32.00, 30-Nov-2007
3, Jose Madriz, 281-330-8004, A, 12.95, 02-Jun-2008
3, Jose Madriz, 281-330-8004, D, 25.02, 22-Jan-2009
```

Reduce Side Join.



Reduce Side Join.



Reduce Side Join.

- Joining doesn't take place until the reduce phase
- We shuffle all data across the network first, and in many situations we drop the majority of this data during the joining process.
- It would be more efficient if we eliminate the unnecessary data right in the map phase

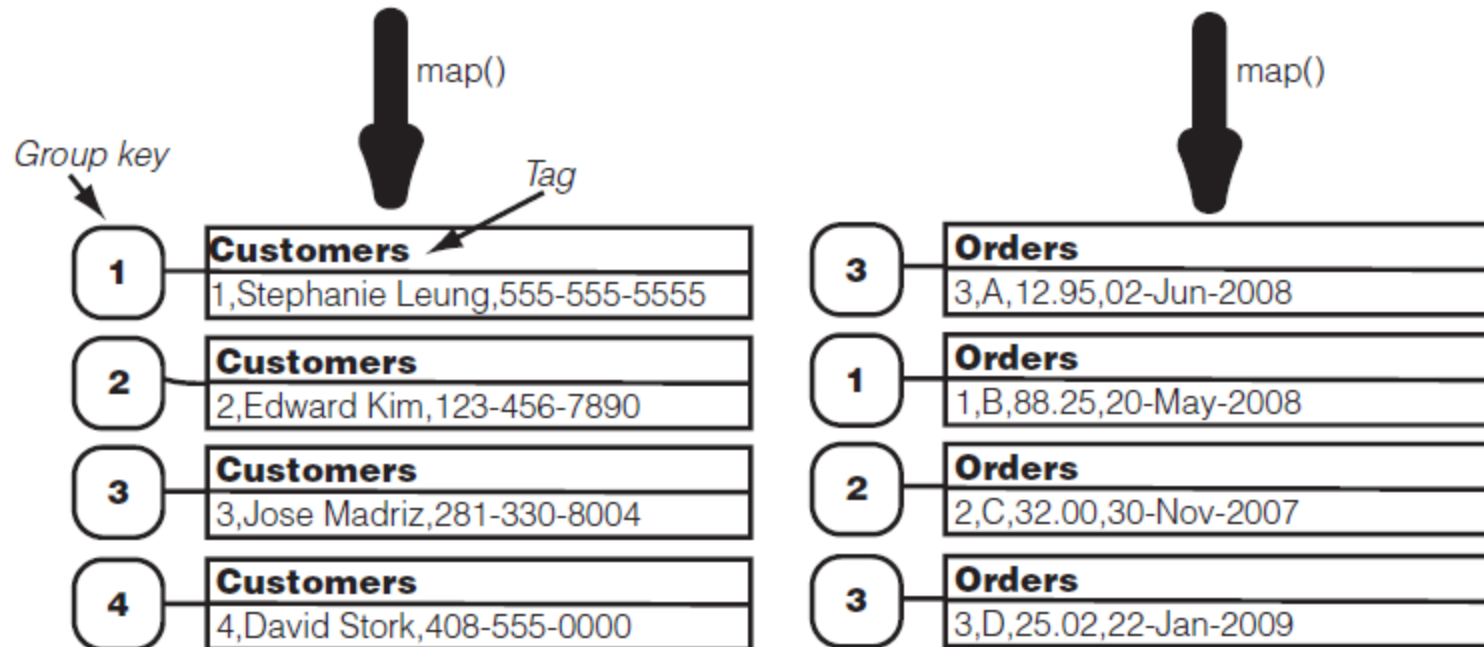
Map Side Join – Multiple Mappers Problem

Customers

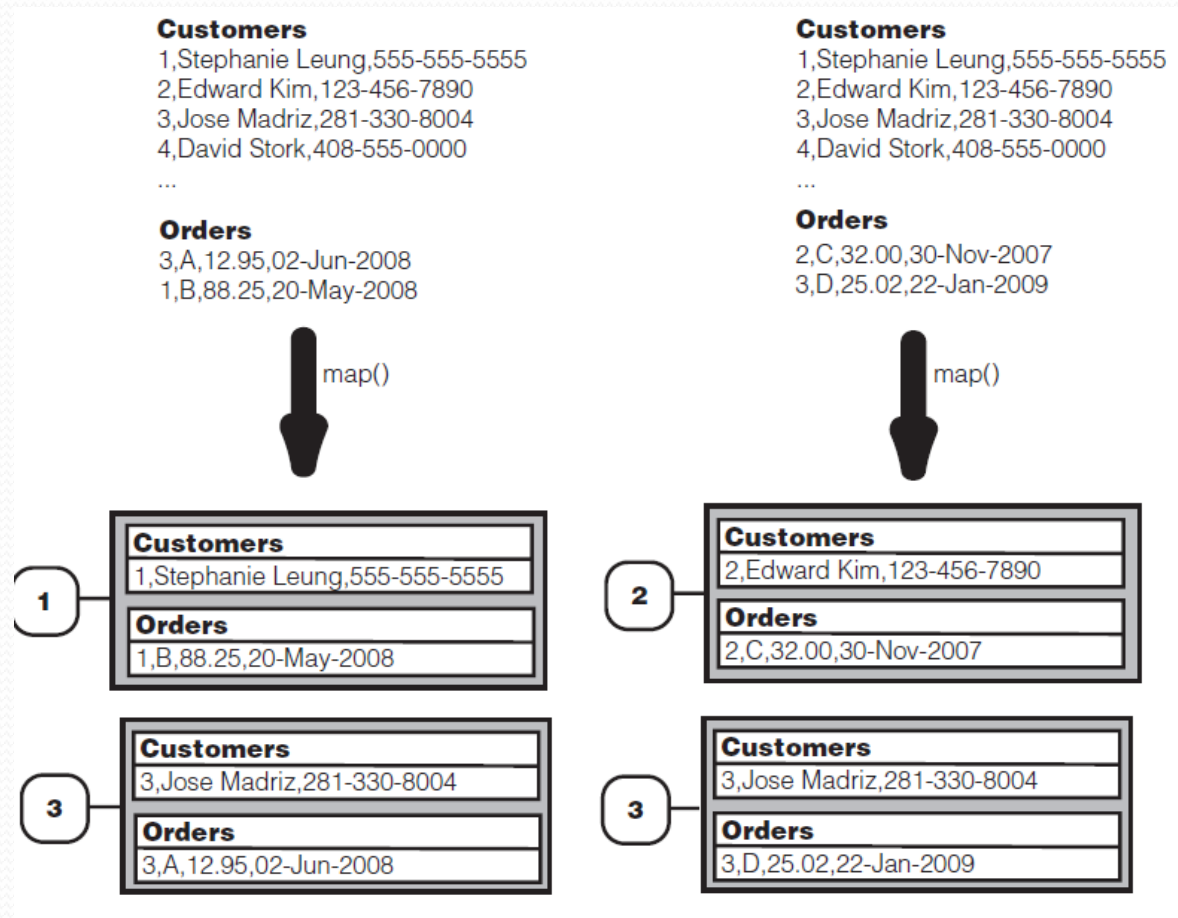
1,Stephanie Leung,555-555-5555
2,Edward Kim,123-456-7890
3,Jose Madriz,281-330-8004
4,David Stork,408-555-0000
...

Orders

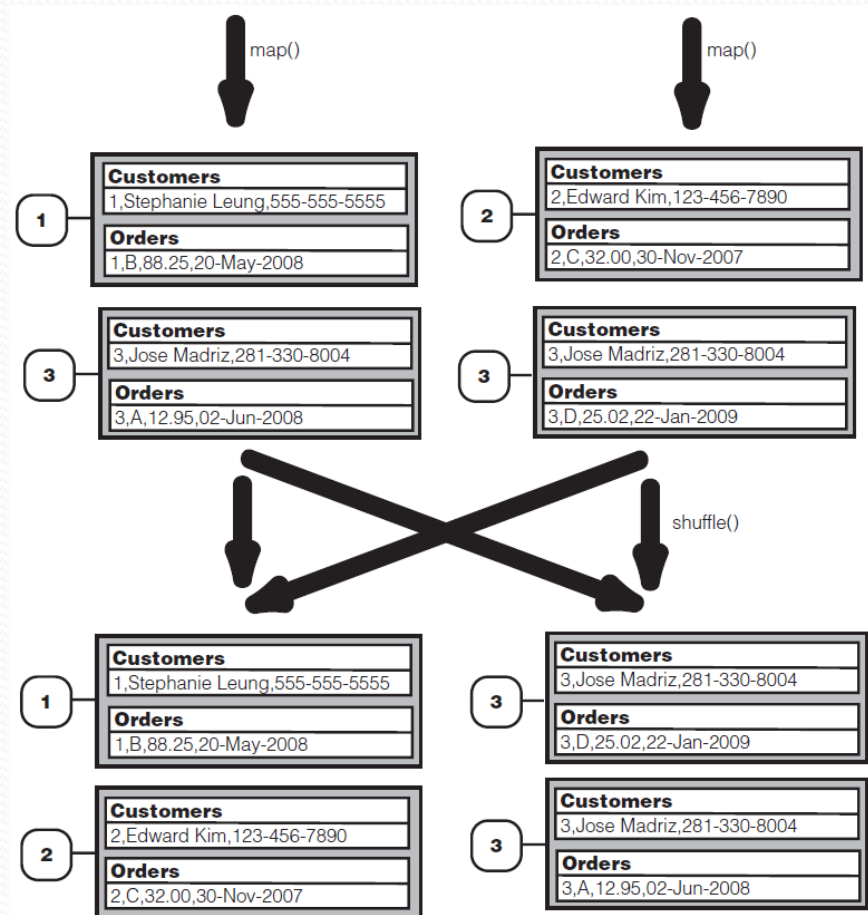
3,A,12.95,02-Jun-2008
1,B,88.25,20-May-2008
2,C,32.00,30-Nov-2007
3,D,25.02,22-Jan-2009
...



Map Side Join



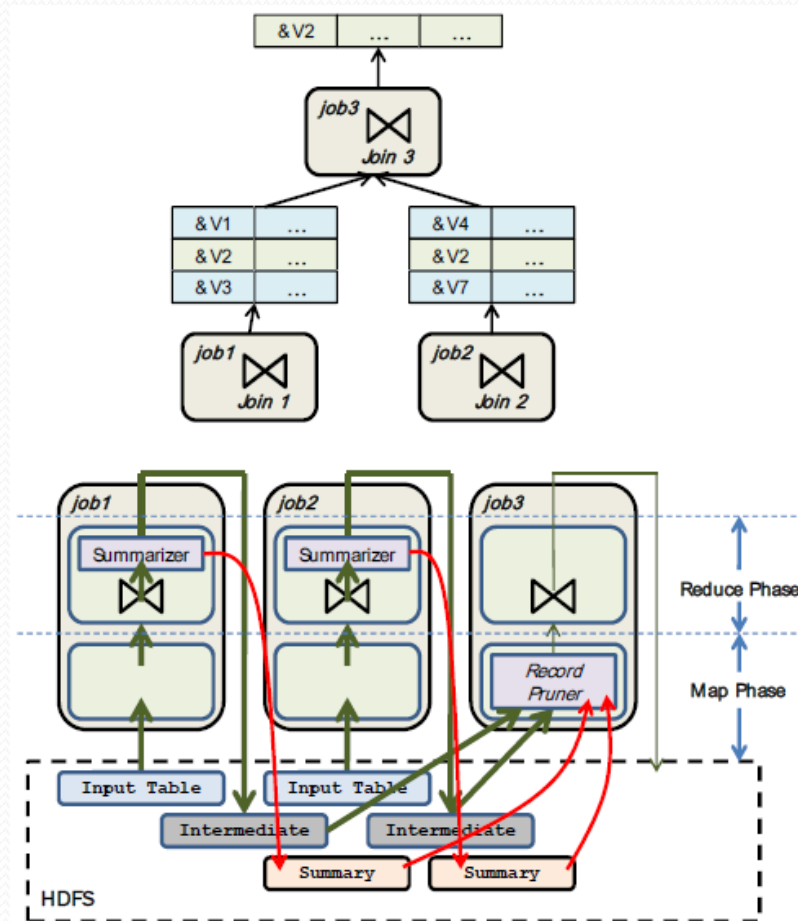
Map Side Join



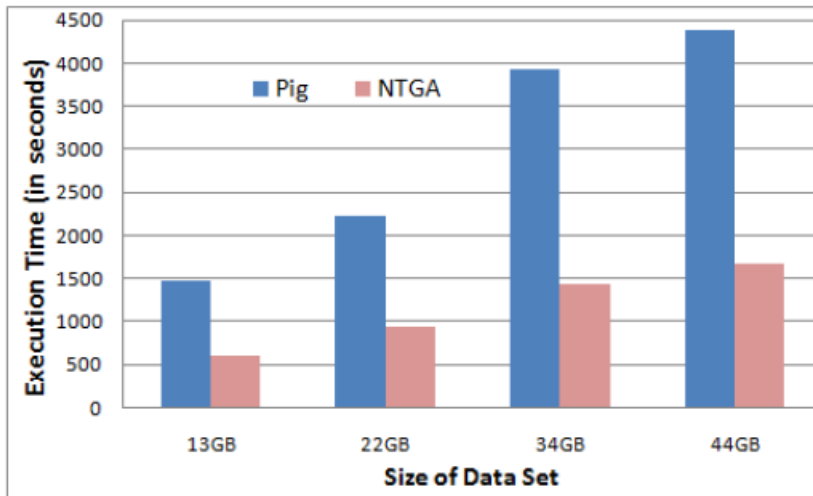
Map Side Join

- Map side join is good when the size of the table is small so it can fit in the memory
- You can use Distributed cache technology in Hadoop to load the table in Distributed RAM.

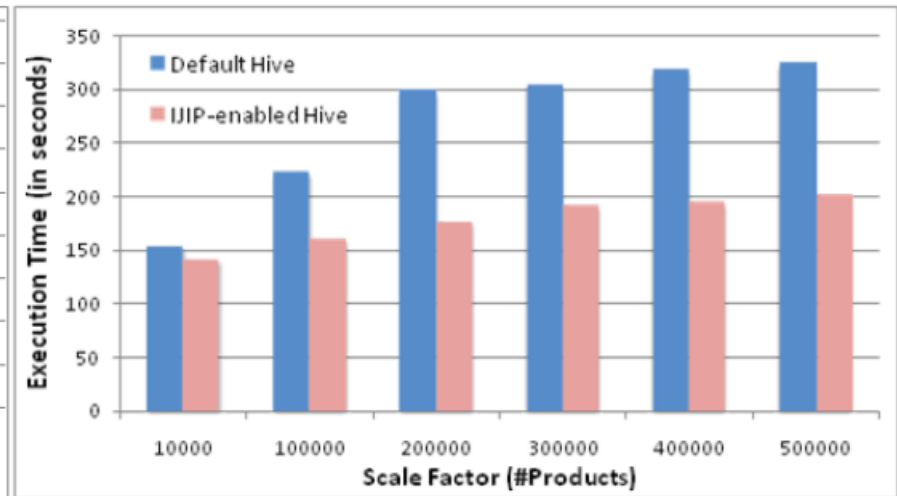
Hadoop - Information Passing



Experiment Results



(a)



(b)

Comments

- Both of papers talked about using Hadoop only. There should be something about using a structure above Hadoop like Hbase or Lucence.
- Nothing mentioned about how to store data in Hadoop. Saving the same kind of triples near each other will reduce the MapReduce execution time.

Bibliography

- **The Google File System** Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung.
- **MapReduce: Simplified Data Processing on Large Clusters** Jeffrey Dean and Sanjay Ghemawat
- **From SPARQL to MapReduce: The Journey Using a Nested TripleGroup Algebra** HyeongSik Kim, Padmashree Ravindra, Kemafor Anyanwu
- **Efficient Processing of RDF Graph Pattern Matching on MapReduce Platforms** Padmashree Ravindra, Seokyong Hong, HyeongSik Kim, Kemafor Anyanwu

Thank You

