

Cache Conscious Star-Join in MapReduce Environments

Guoliang Zhou, Yongli Zhu and Guilan Wang
NCEPU(<http://www.ncepu.edu.cn/>)

Outline

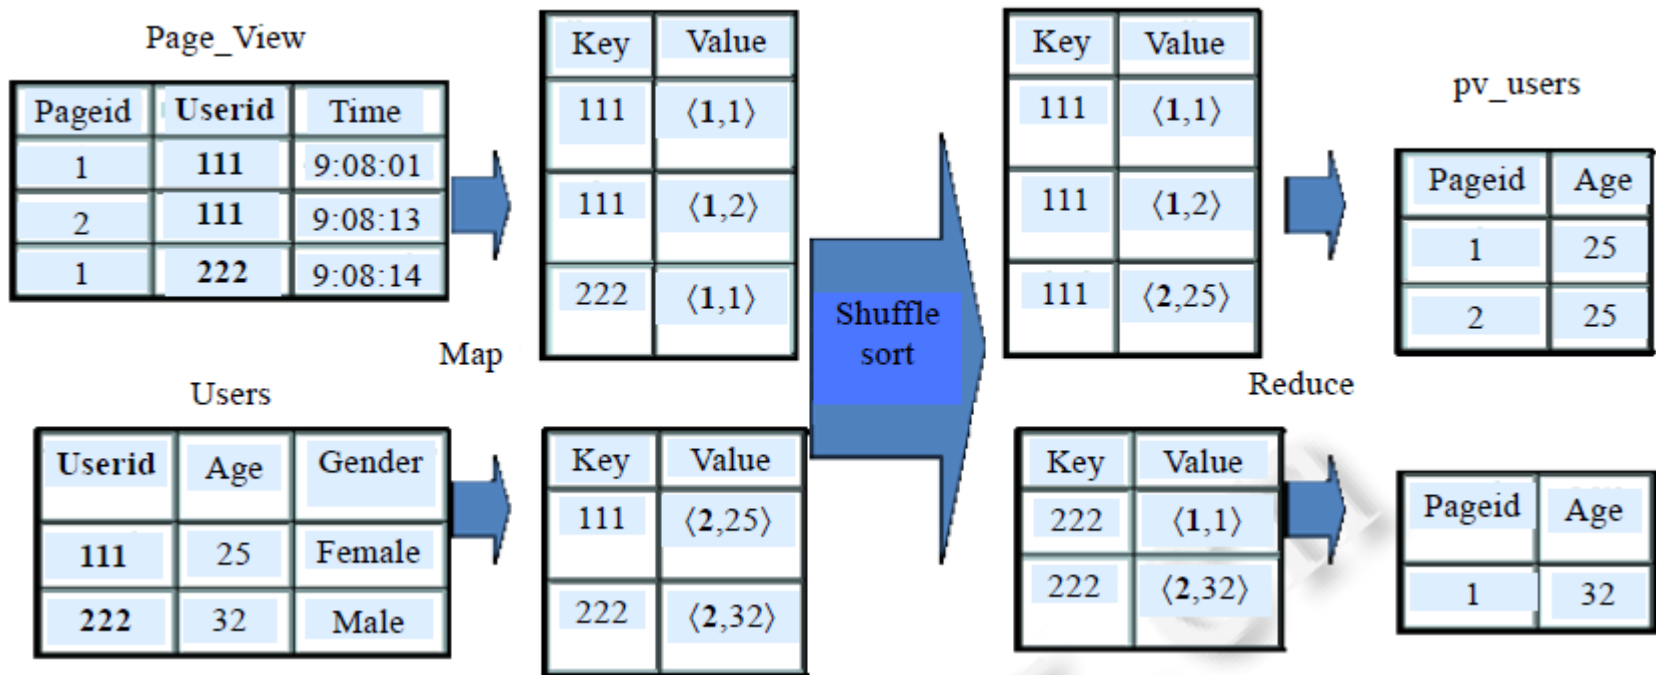
- ▶ INTRODUCTION
- ▶ RELATED WORK
- ▶ MULTI-FRAGMENT-REPLICATION JOIN
- ▶ MAPREDUCE-INVISIBLE JOIN
- ▶ PERFORMANCE EVALUATION
- ▶ CONCLUSIONS

Introduction

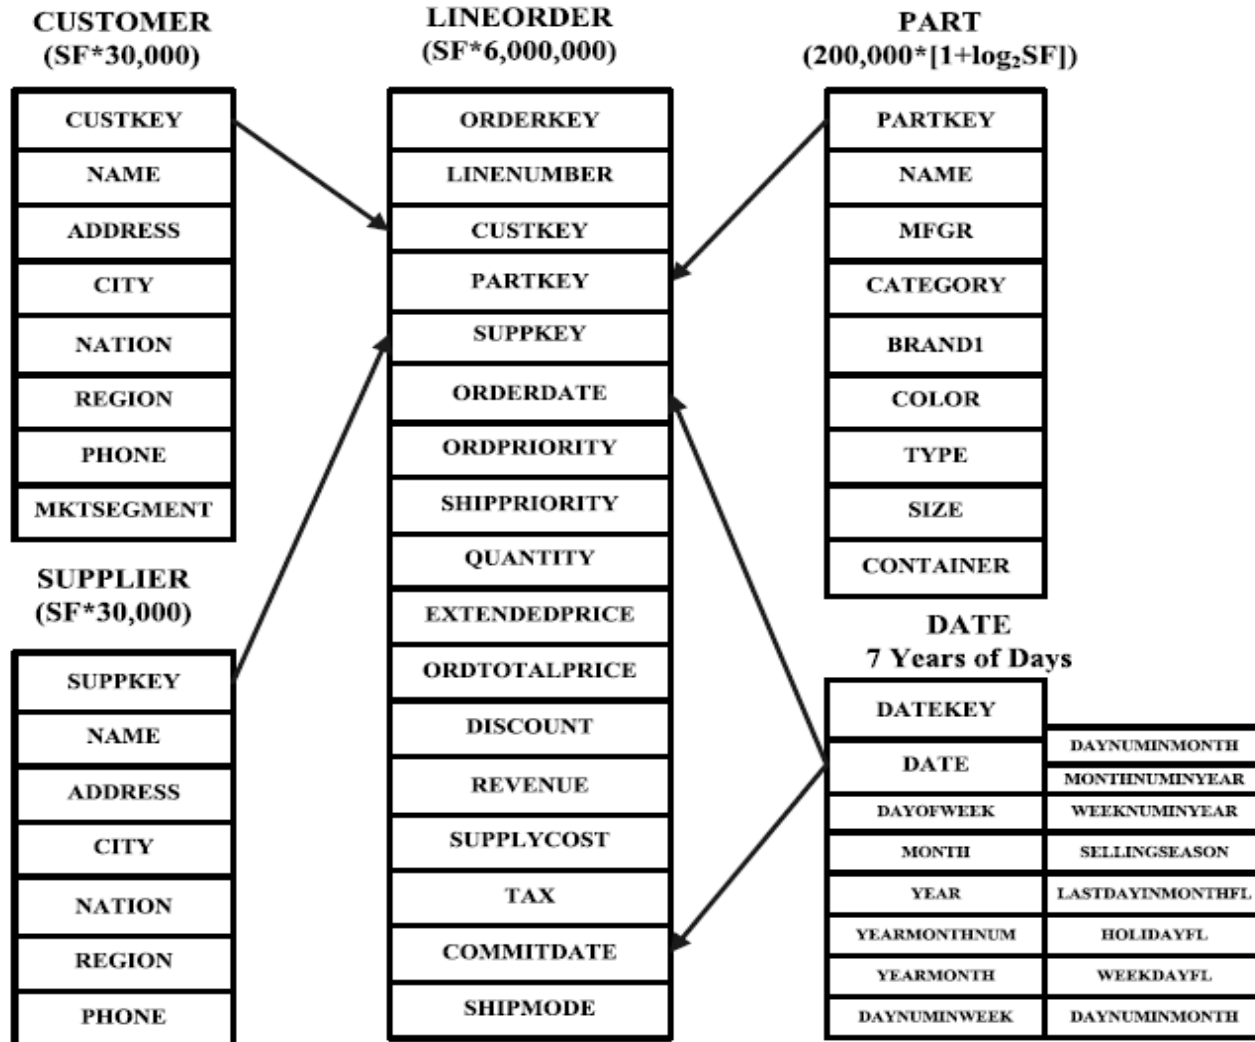
- ▶ Big data MapReduce
- ▶ Hive Pig

- ▶ Star join
- ▶ Cache conscious

Mapreduce join

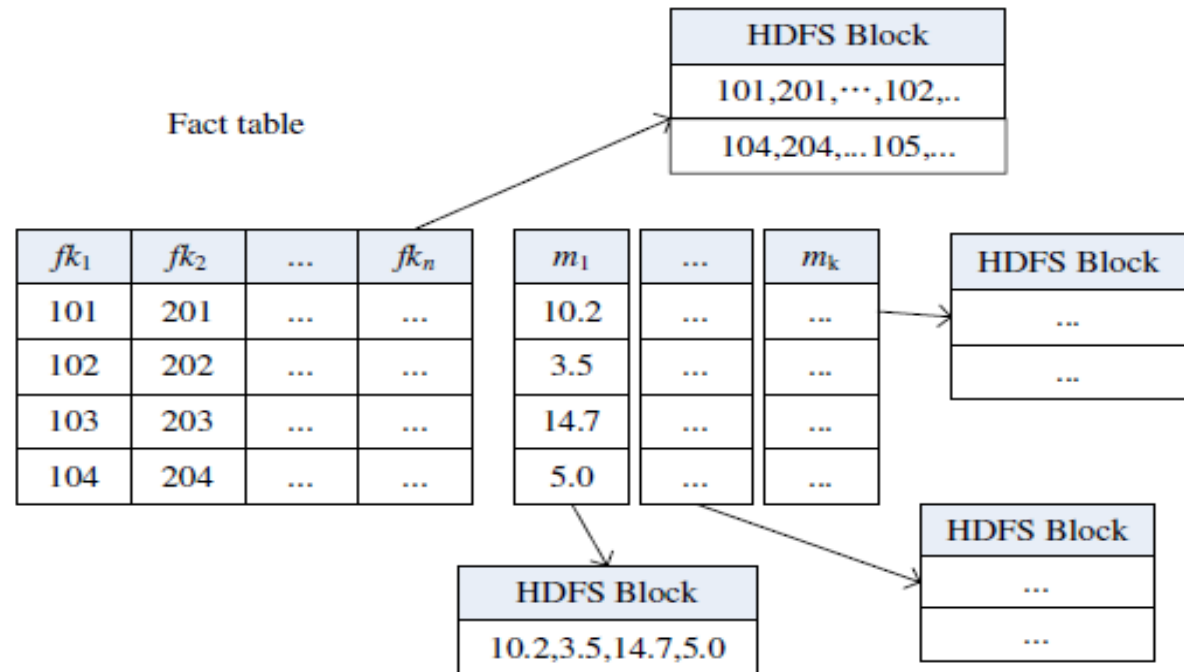


Star join(SSB)



MULTI-FRAGMENT-REPLICATION JOIN

- ▶ Map Join
- ▶ Fact table partition into (column families)



- ▶ Firstly, we cache all dimension tables in local memory for each mapper.
- ▶ Then streaming the value of fact table joins with all dimension tables.
- ▶ That is Multi-Fragment-Replication Join (MFRJ).

MAPREDUCE-INVISIBLE JOIN

- ▶ Fact table each column is respectively storage in different data node.
- ▶ Dimension table is partitioned as below

CUSTKEY	NAME	CITY	NATION	REGION	PHONE
Customer#00000001	IVhzIApeRb	MOROCCO	MOROCCO	AFRICA	25-705-443-4055
Customer#00000001	XSTf4,NCwDVaWNe6tE	JORDAN 6	JORDAN	MIDDLE EAST	23-453-414-8560

CUSTKEY	NAME
Customer#00000001	IVhzIApeRb
Customer#00000001	XSTf4,NCwDVaWNe6tE

CUSTKEY	CITY
Customer#00000001	MOROCCO
Customer#00000001	JORDAN 6

CUSTKEY	NATION
Customer#00000001	MOROCCO
Customer#00000001	JORDAN

CUSTKEY	REGION
Customer#00000001	AFRICA
Customer#00000001	MIDDLE EAST

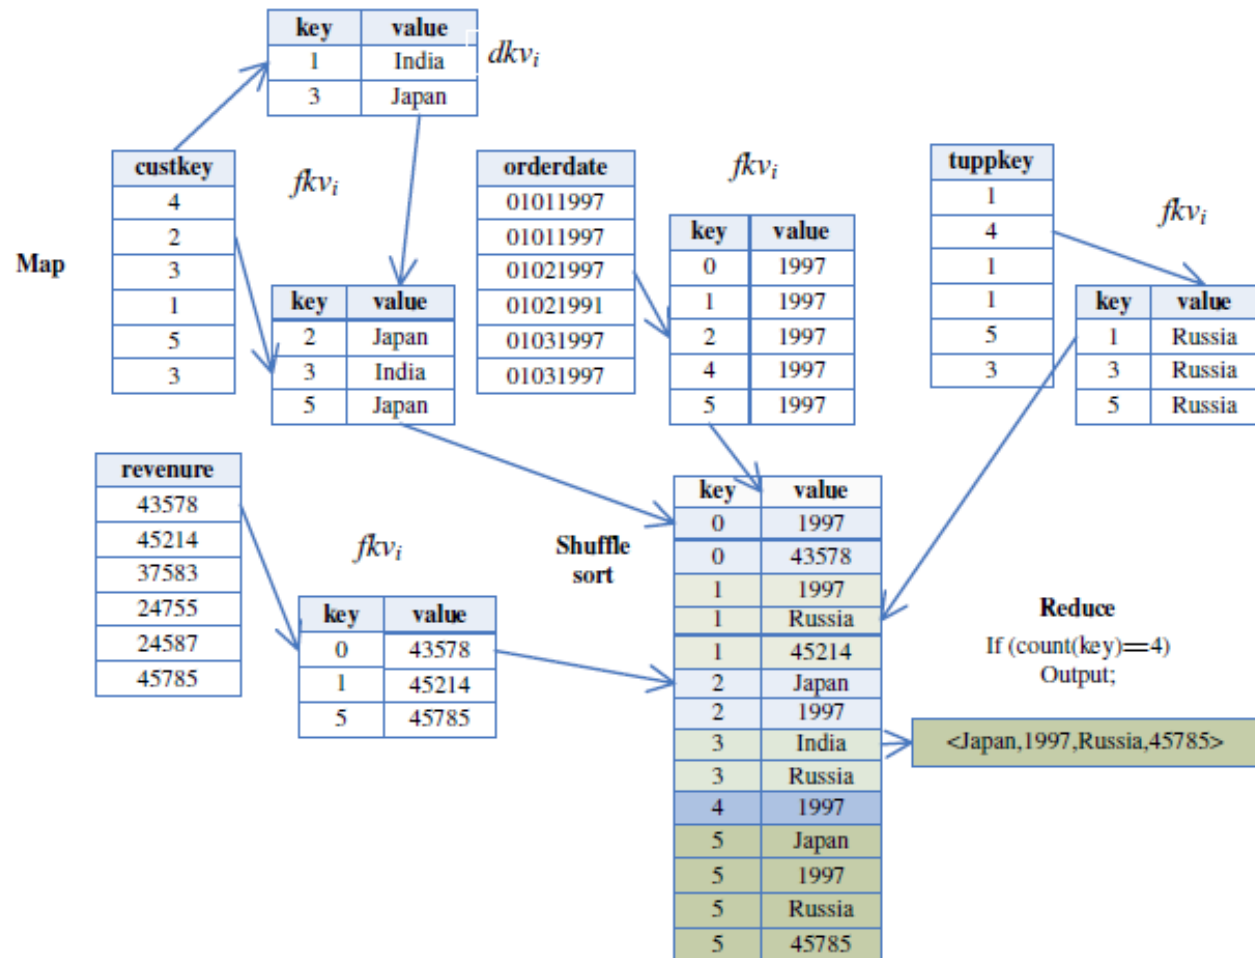
CUSTKEY	PHONE
Customer#00000001	25-705-443-4055
Customer#00000001	23-453-414-8560

Q3.1

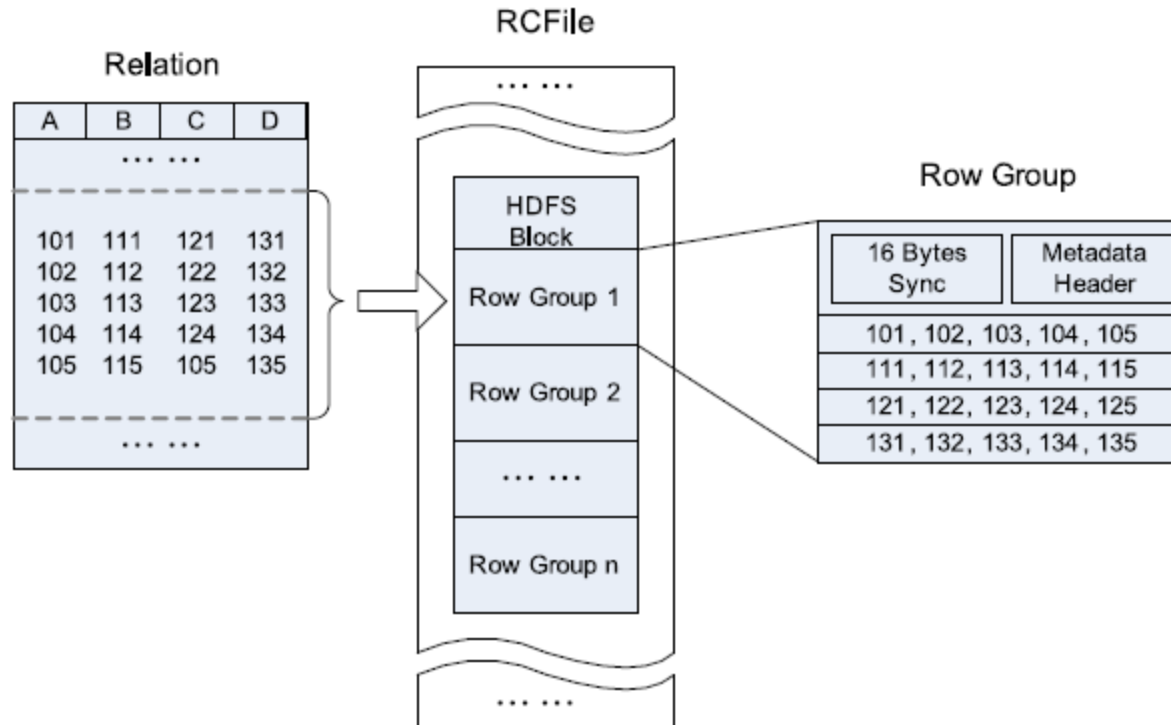
- ▶ SELECT c nation, s nation, d year, lo revenue
- ▶ FROM customer
- ▶ JOIN lineorder ON lo custkey = c custkey
- ▶ JOIN supplier ON lo suppkey = s suppkey
- ▶ JOIN ddate ON lo orderdate = d datekey
- ▶ WHERE c region = 'ASIA'
- ▶ AND s region = 'ASIA'
- ▶ AND d year ≥ 1992 and d year ≤ 1997
- ▶ AND lo revenue > 40000 ;

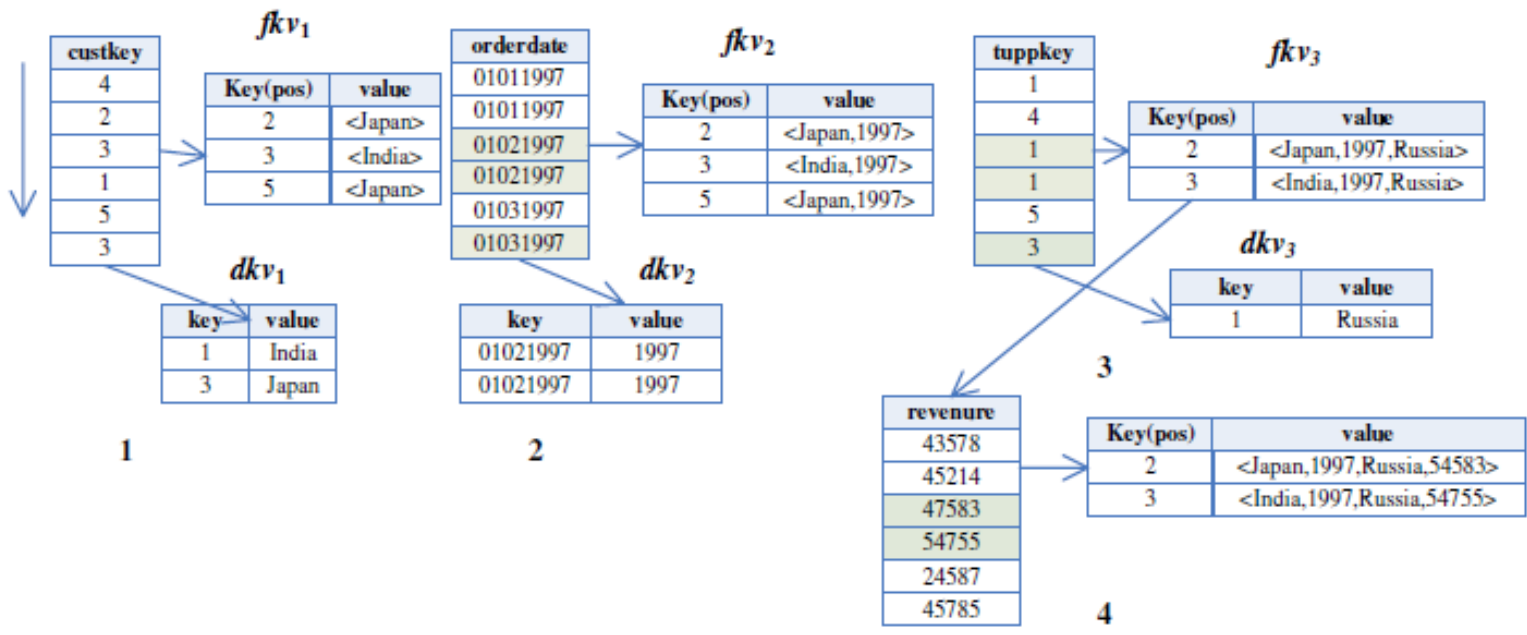
MAPREDUCE-INVISIBLE JOIN

- ▶ Map
- ▶ Reduce



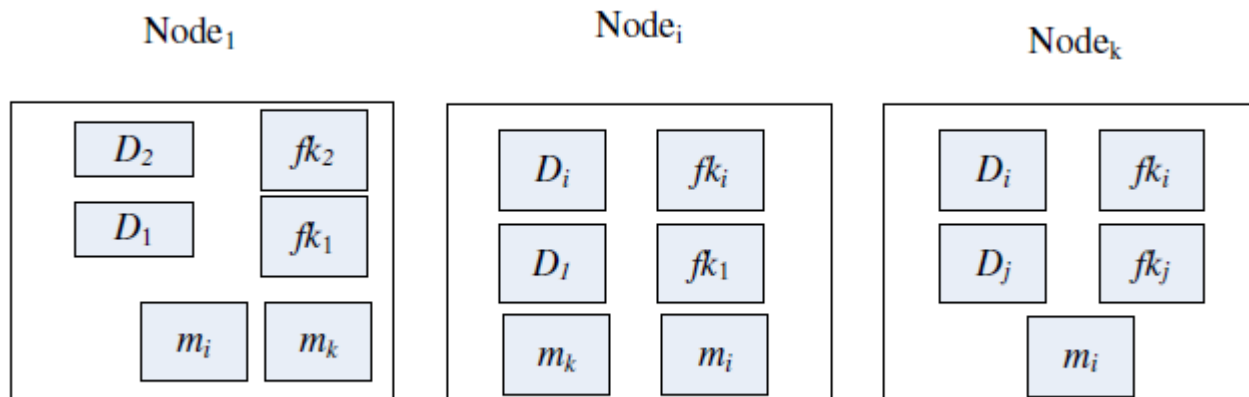
MapReduce-Invisible Join on RCFile







Dimension tables and fact table storage with co-location

- ▶ Dimension table and related foreign key column are co-location storage



Execution plan of MRIJ (big dimension table)






(a) Query 3.1




(b) Query 4.1

Figure 7: Original Query



(a) Query 3.1



(b) Query 4.1

Figure 8: No where conditions

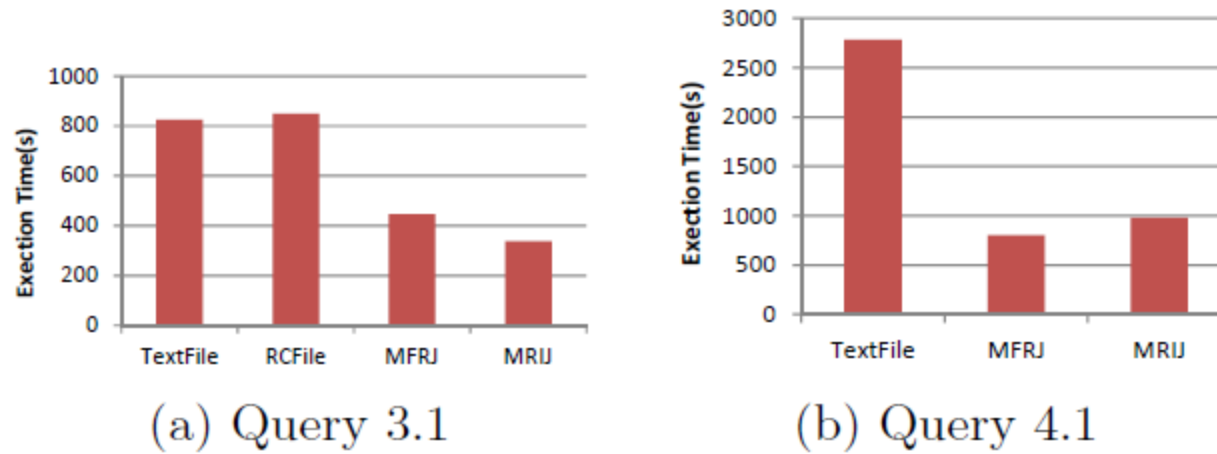


Figure 9: No measure columns

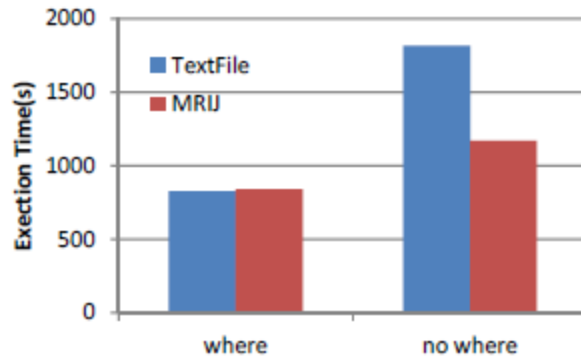


Figure 10: MRJ with big dimension tables

- ▶ data compression in column store
- ▶ Modern processor features utilizing

▶ Thanks

▶ Q & A