

HadoopDB

Blending Map/Reduce and PDBMS into one Hybrid System

Qing Zheng

15799 – Adv Topics in DB Systems

Nov.4th 2013

Scale-out OLAP System

- Scale to 100~1000 nodes
- High-Performance/High-Efficiency
- Fault-Tolerance/Heterogeneous H.W.
- User-Friendly Query Interface
- Easy to Install and Tune/Open Source

Big Data

Unstructured
Solution

Hadoop



Structured
Solution

PDBMS

Large-Scale Parallelism
A Large Cluster of Commodity Servers
Share-nothing Architecture

PDBMS

- Scale to 100 nodes
- Optimizing Reads over Writes
 - Compression >> Scan/Aggregation
 - Column-Based >> Selection
 - Index/Locality >> Join
 - Push Model/Query Planner >> Efficient
- SQL >> Flexible, Compact, High-Level

Hadoop

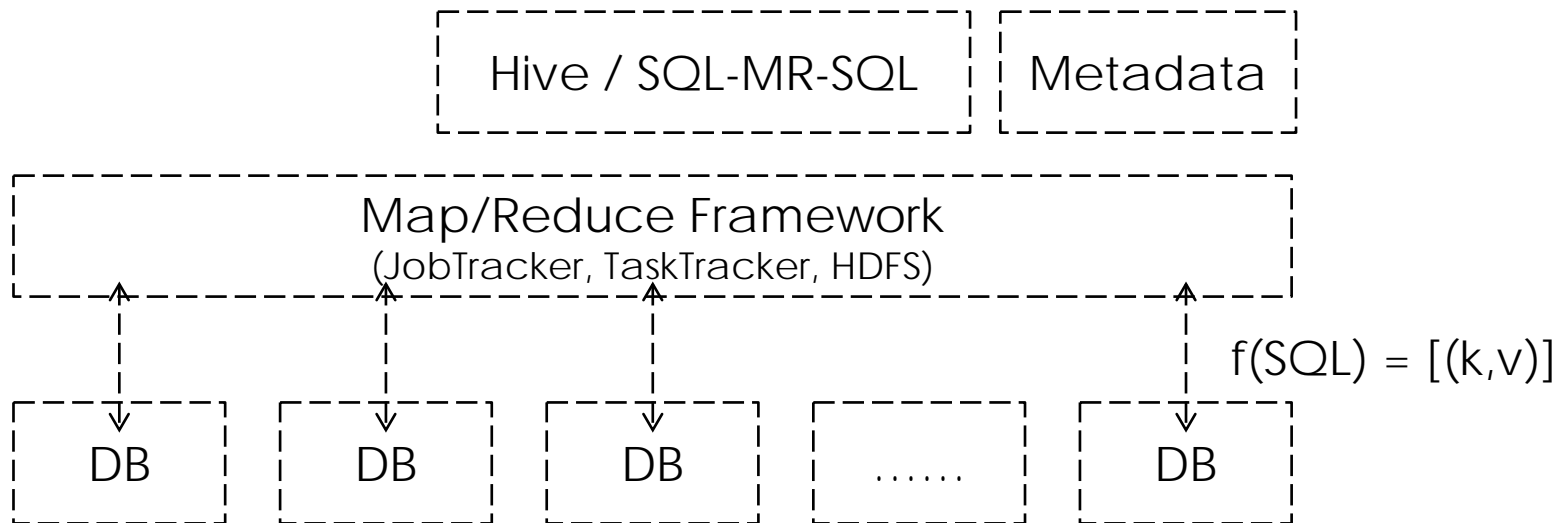
- Scale to 1000 nodes
- Optimizing Writes over Reads
 - Schema-free >> Fast Bulk Loading
 - Schema-ignorance >> Brute-Force Model
- Dynamic Runtime Scheduling
 - Check-pointing >> Fault-Tolerance
 - Speculative Exe. >> Straggler Immunity
 - Fine-grained Partitioning >> Fast Recovery

MR + DBMS = HadoopDB

- Scale to 1000 nodes
- SQL-based Query Interface
- Optimizing Reads over Writes
- Dynamic Runtime Scheduling
- Open Source/Un-expensive to own

HadoopDB Architecture

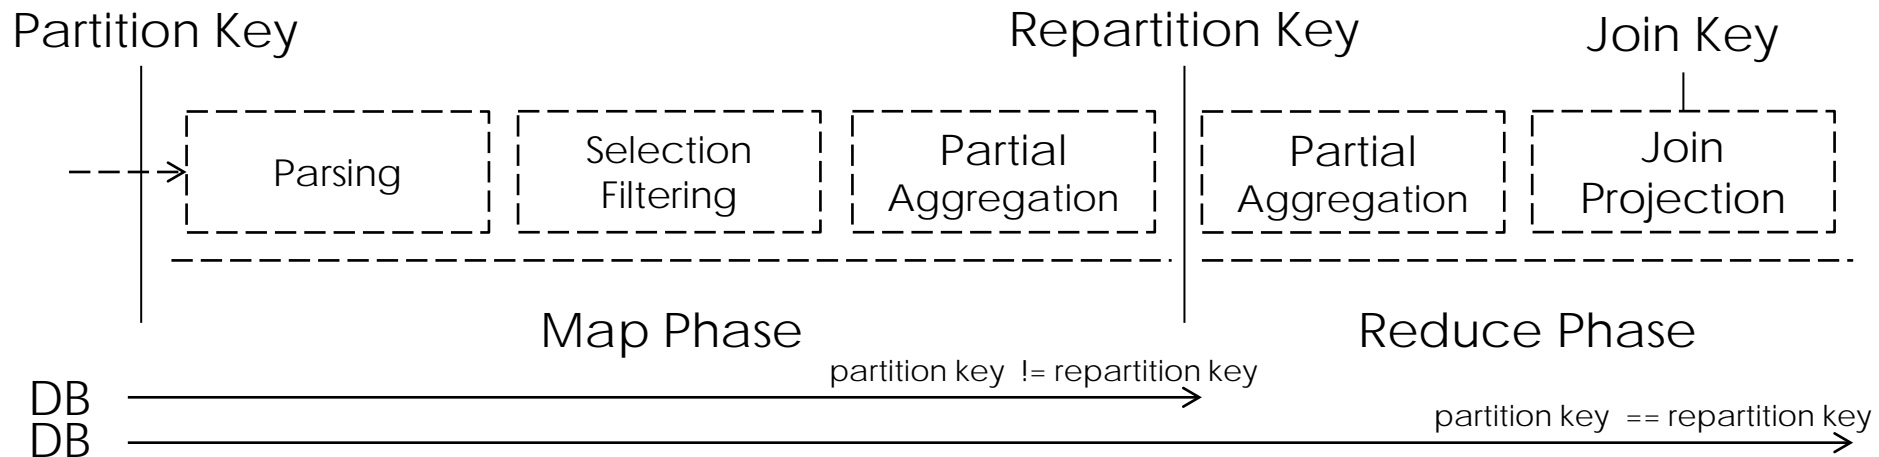
- Static View



connect multiple single-node database systems using Hadoop as the task coordinator and network communication layer

HadoopDB Architecture

- Dynamic View



parts of the incoming query are pushed down and executed inside the database layer

the rest of the query is processed in a more generic MapReduce framework

Evaluation

- Load – No
- Grep – No
- Selection – No, could be Yes
- Aggregation – No, could be Yes
- Join - Yes
- Fault-Tolerance – Yes
- Query Interface – Yes

compared against an ideal system

Bottom Line

- Unstructured Data
 - HadoopDB is no worse than Hadoop
- Structured Data
 - HadoopDB is close to PDBMS

Conclusion

- Hybrid System
 - Performance/SQL – PDBMS
 - Scalability/Resilience – Hadoop
 - Open Source