

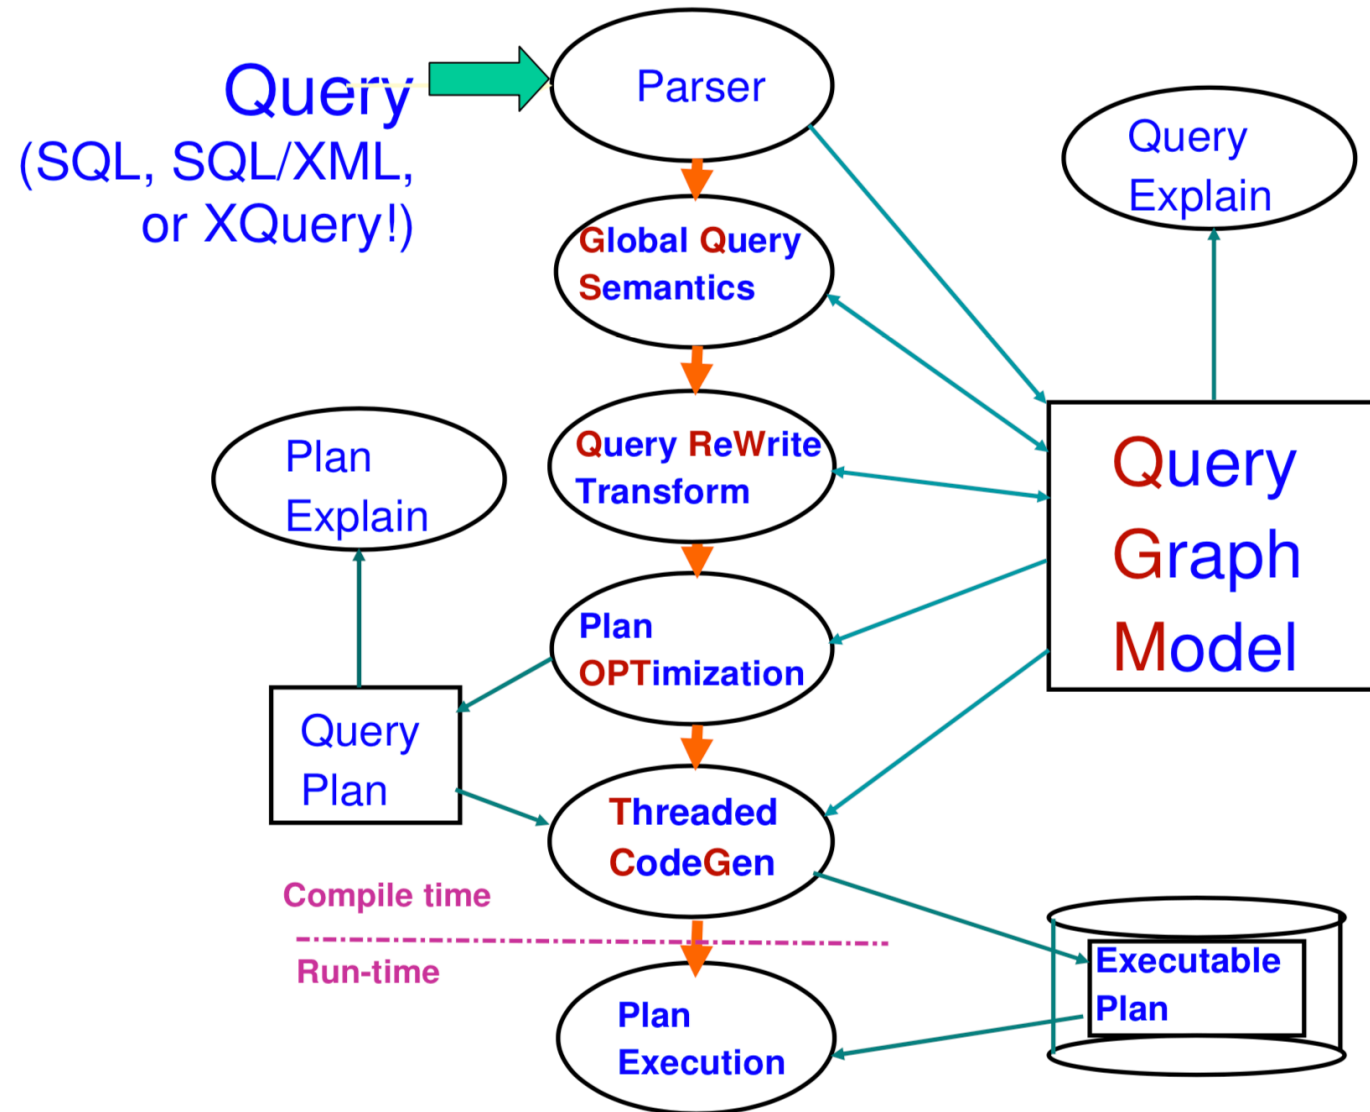
Encapsulation of Parallelism in the Volcano Query Processing System

Huawei Wang

Overview

- ▶ Architecture
- ▶ Bracket Model
- ▶ Operator Model
- ▶ Pros & Cons
- ▶ Comparison

Typical Query Engine Architecture



Similar Systems



- ▶ System R

- ▶ Starburst

Bracket Model

Problem:

- Extensibility
- Large overhead



Operator Model

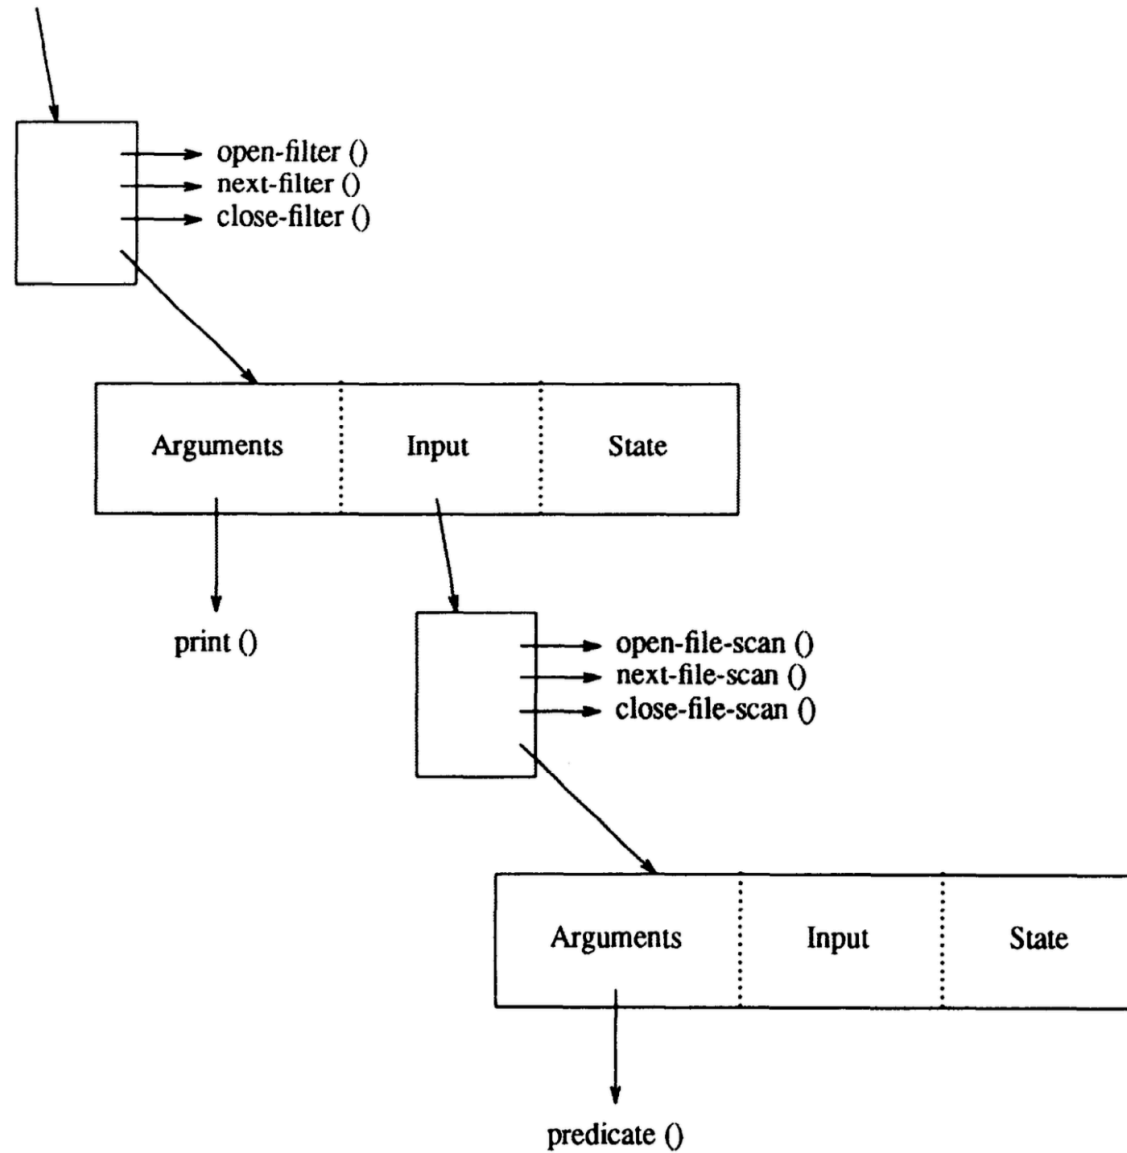
Single process:

- Operator mapping to iterator(Can be applied to big data system)
- Use stream as abstractions for input between operators

Multiple process:

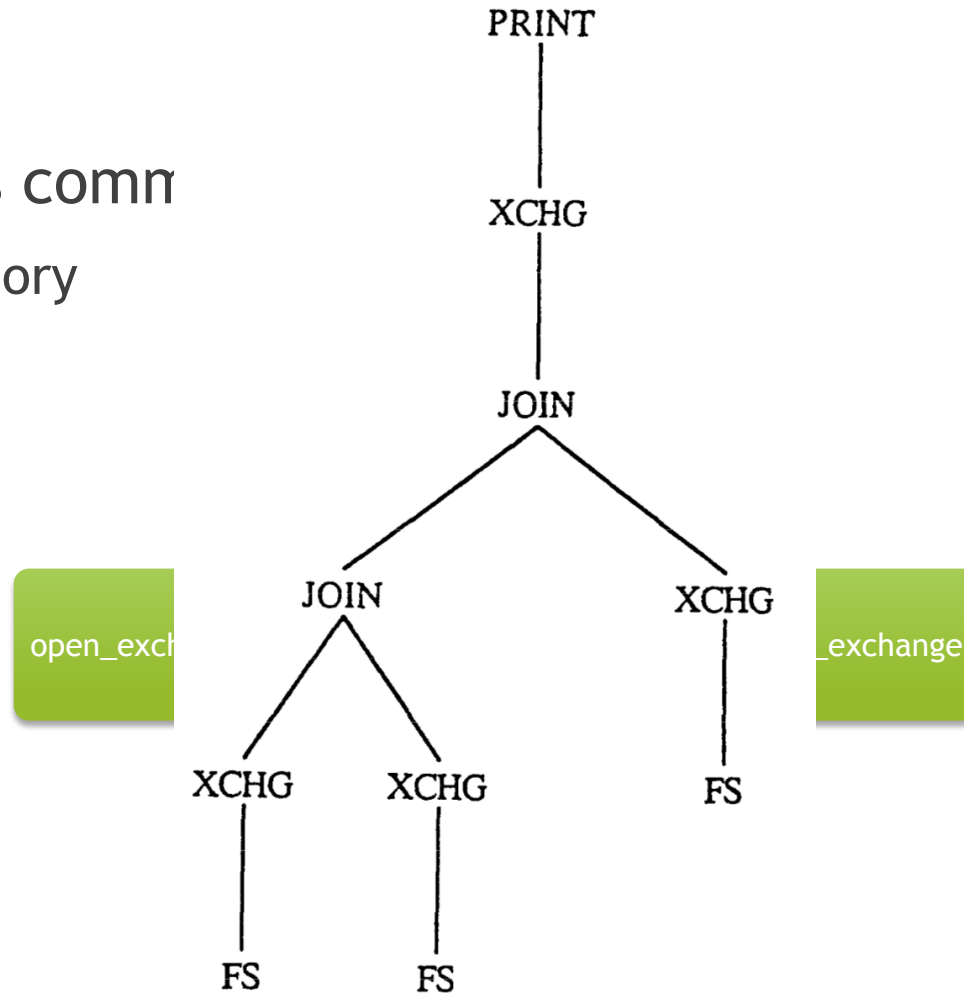
- Introduce exchange operator
 - Vertical parallelism
 - Horizontal parallelism

Operator Model




Vertical Parallelism

- ▶ Inter-process comrn
 - ▶ Shared Memory
 - ▶ Semaphore



Horizontal Parallelism

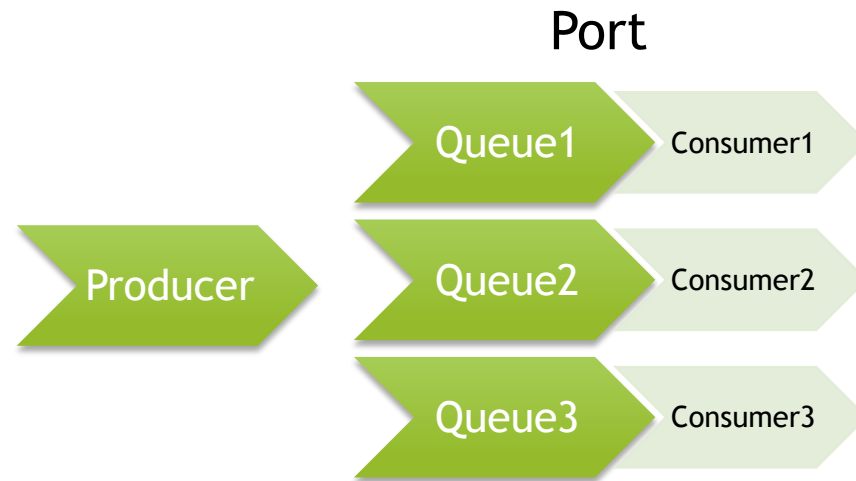
- ▶ Bushy parallelism



- ▶ Intra-operator parallelism

Horizontal Parallelism

How to partition data?



Horizontal Parallelism

- Centralized scheme
- Propagation tree scheme
- Primed process

Pros & Cons

- ▶ More generalized
 - ▶ Algorithm Level
 - ▶ System Level
- ▶ Easy Implementation
- ▶ Heavy weight creating process

Comparison

- ▶ Spark can choose whether to persist RDD
- ▶ Volcano only let intermediate results exist in buffer
- ▶ Volcano is only a query execution engine with 2 key meta operators.