MapReduce & Hadoop

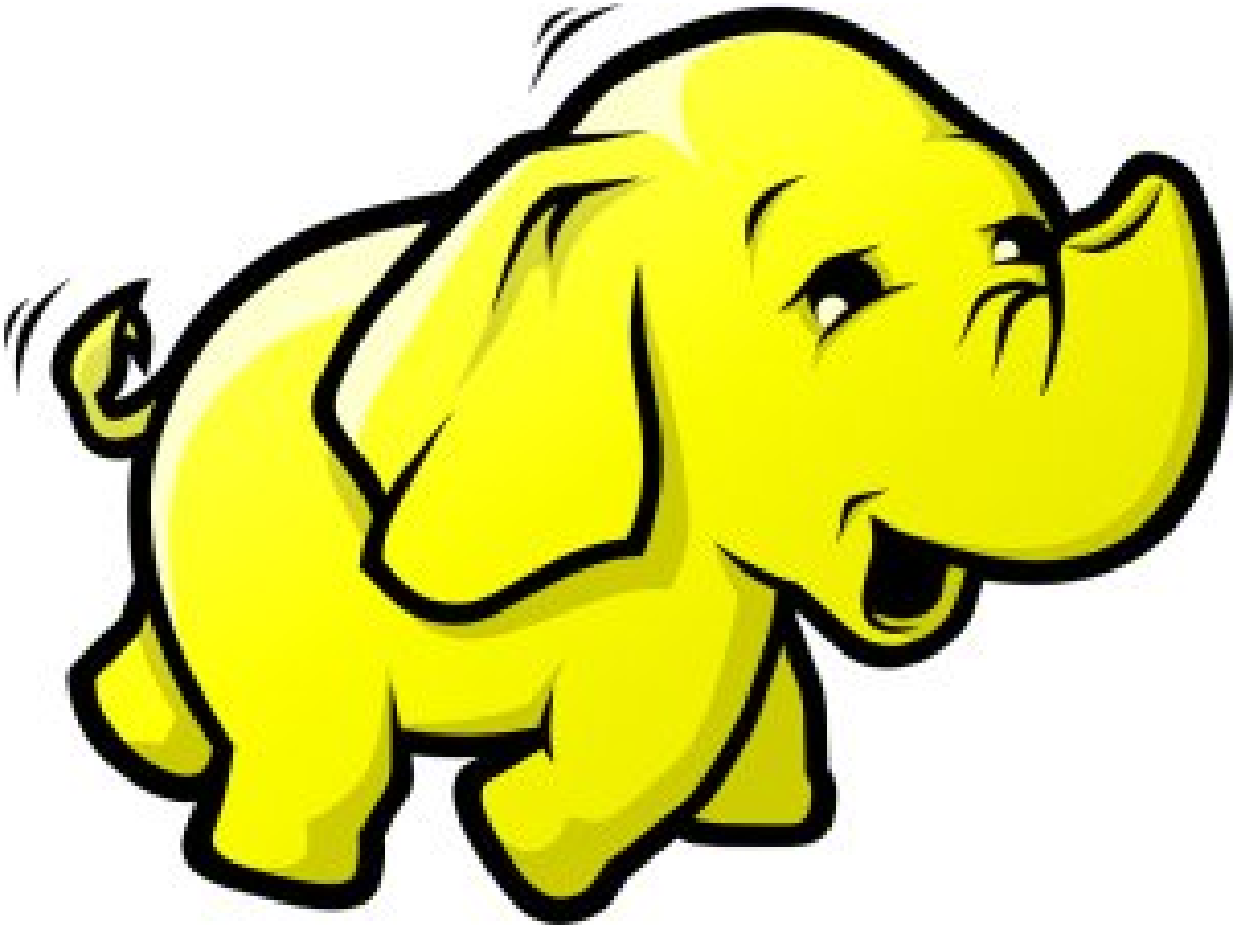

سید مجید عظیمی

Problem

How to process lots of data?

Short answer:

Distribute the processing on many machines

Challenges

- We are not going to care WTF is going on the network.
- Having faulty machine is normal.
- How to store multiple peta bytes of data as input and output

MapReduce

Simple MapReduce job:

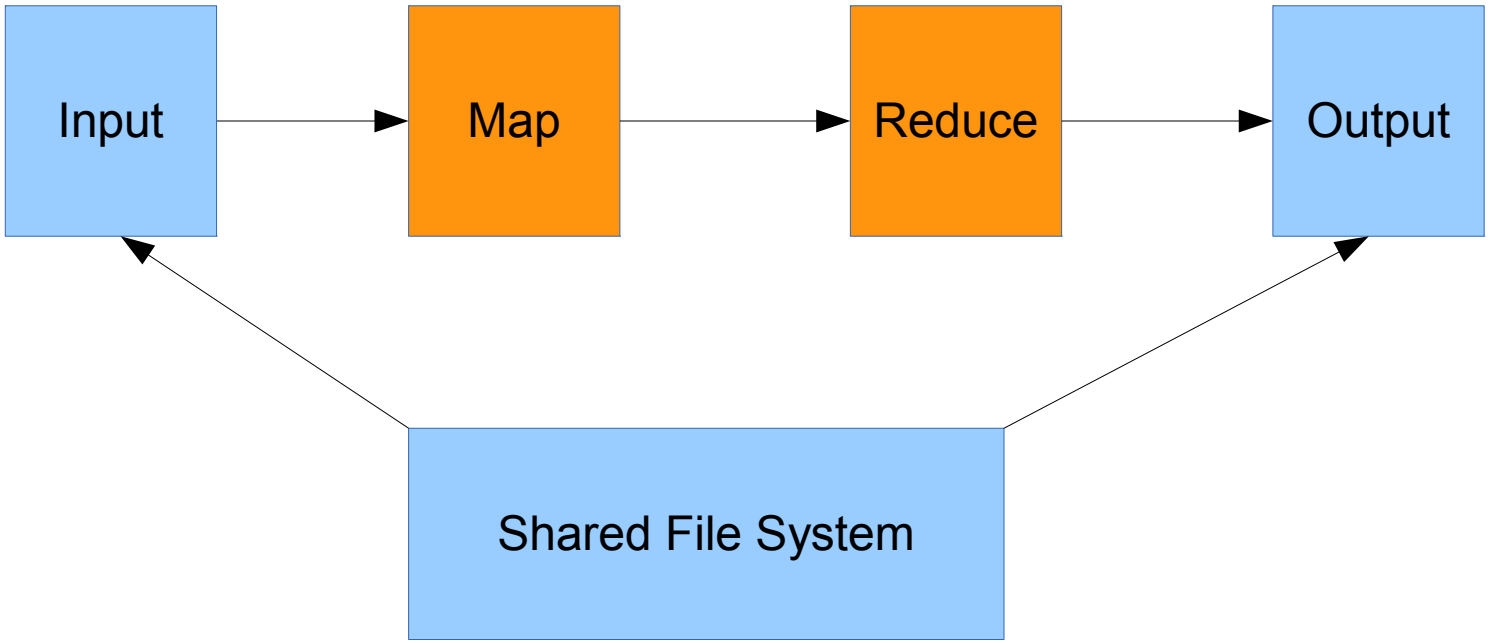

MapReduce

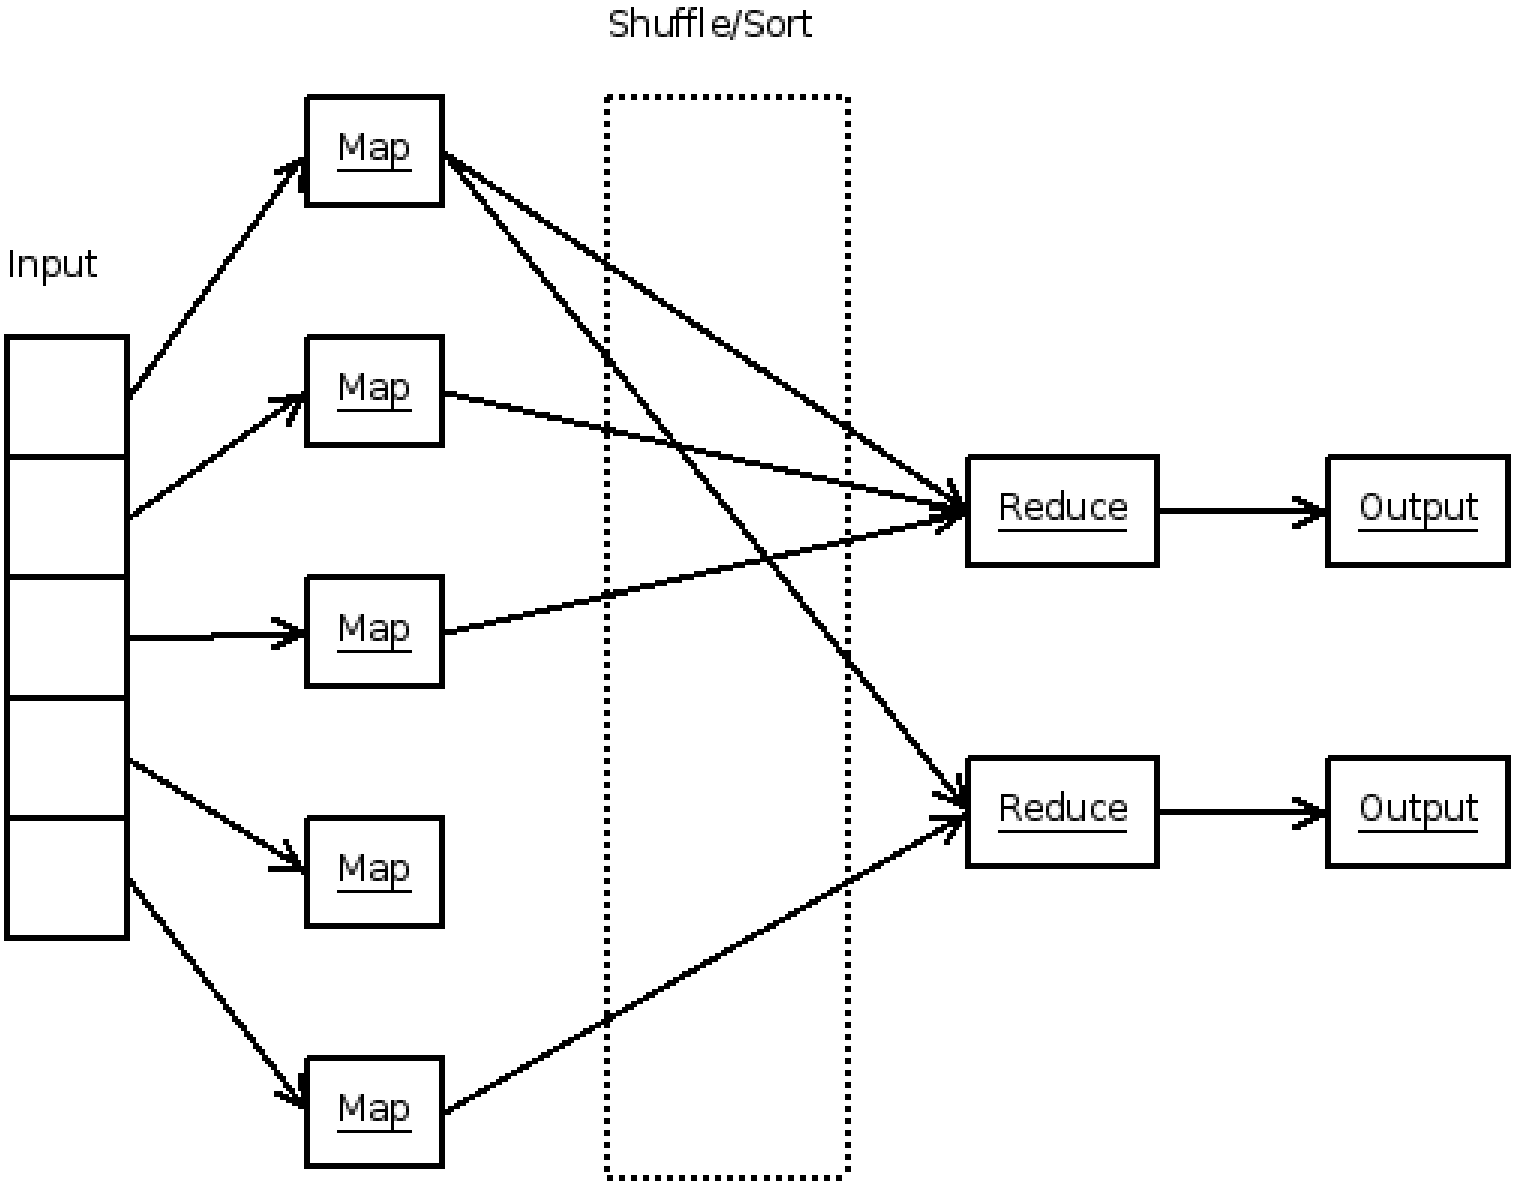
Shuffle/Sort
Input
Map
Map
Map
Map
Map
Reduce
Reduce
Output
Output

What Google does everyday

We have a web crawler that collects a copy of each web page: X1, X2, X3, …., Xn

The job:

Create Reverse Index

Mapper

- Each mapper process one file. Output of Mapper is a list of Key/Value:

  Key: keyword , Value: Web page address

Mapper1		Mapper2		Mapper3		Mapper4	
Key	Value	Key	Value	Key	Value	Key	Value
is	x1	is	x2	is	x3	is	x4
driver	x1	kernel	x2	kernel	x3	kernel	x4
Linux	x1	me	x2	hello	x3	Apache	x4
me	x1	glib	x2	Apache	x3	me	x4

Shuffle + Sort

Apache: [x3, x4]

driver: [x1]

glib: [x2]

hello: [x3]

is: [x1, x2, x3, x4]

kernel: [x2, x3, x4]

Linux: [x1]

me: [x1, x2, x4]

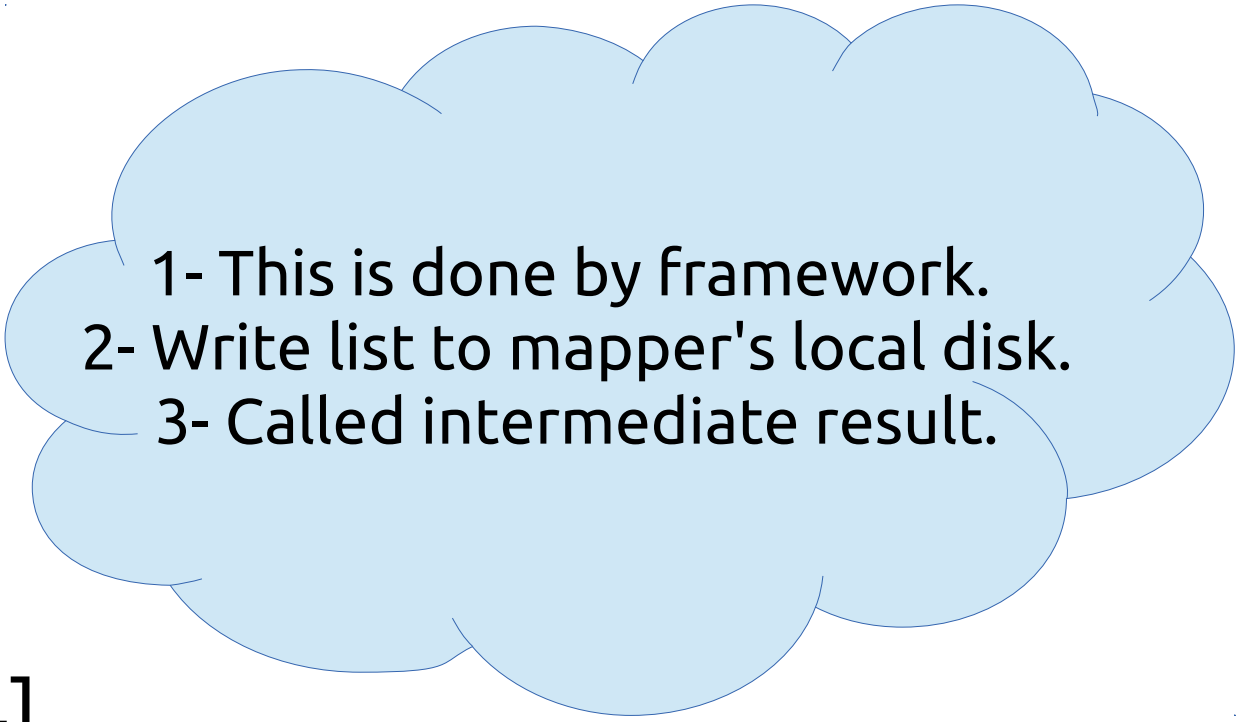

Reducer

- Each partition will be sent to a separate reducer
- Output of each reducer will be written to shared file system based on Key/Value

What can we do in Reducer:

- Do nothing
- Filter some web pages
- Remove duplicates

What is Hadoop?

Best known MapReduce implementation written in Java

What Hadoop provides

- Different input readers(XML, CSV, …)
- Different output writers(XML, CSV, …)
- Distributed replicated file system called HDFS
- Built in support for Avro
- Compression(Gzip, Bzip2, ...)
- Counters
- Built in mechanism for detecting failed nodes

What Hadoop provides

- Support streaming for using in shell pipeline
- Different partitioning algorithms
- Different job schedulers.
- …

Hadoop Ecosystem

- Scoop: Tools to transfer raw data from/to RDBMS
- HBase: Hadoop columnar NoSQL database
- Avro: Fast serialization framework(Java native serialization SUCKS)
- Pig: SQL like language for MapReduce
- Hive: Data warehouse solution based on Hadoop

Hadoop Ecosystem

- Zookeeper: Cluster manager
- Flume: Log aggregation framework
- R: Data analytical language
- Oozie: Workflow scheduler system

Learning resources

- Hadoop: The Definitive Guide
- Hadoop Operations
- Hadoop in Action
- Hadoop in Practice
- HBase: The Definitive Guide
- HBase in Action
- MapReduce Design Patterns
- Programming Hive
- Programming Pig

Question & Answer

?