

NÁSKOK  
DÍKY  
ZNALOSTEM

PROFINIT

# BE0M33BDT Big Data Technologies

## MapReduce

Milan Kratochvíl, Sergii Stamenov, Jan Hučín

December 4th, 2019


# MapReduce

- › Programing model for parallel computations
- › Basic idea:
  - Data are distributed over cluster, I need put them together to groups
  - Mark of a group = **key** (same key = same group)
  - Exchange of data over cluster to get a group together (**shuffle**)
  - After grouping: aggregation (**reduce**)
  - Before grouping: preparation of data and key assignment (**map**)

# MapReduce

- › Programing model for parallel computations
- › Principle:
  - MAP
    - input data  $\langle \text{VALUE1} \rangle$  converts to  $\langle \text{KEY}, \text{VALUE2} \rangle$
  - SHUFFLE & SORT
  - REDUCE
    - data  $\langle \text{KEY}, \text{LIST}(\text{VALUE2}) \rangle$  aggregated to  $\langle \text{KEY}, \text{VALUE3} \rangle$

›  $\langle \text{V1} \rangle \rightarrow \langle \text{K}, \text{V2} \rangle \rightarrow \langle \text{K}, \text{LIST}(\text{V2}) \rangle \rightarrow \langle \text{K}, \text{V3} \rangle$



# Example: number of words by length

**MAP**

Ó, náhlý déšť již zvířil prach a čilá laň teď  
běží s houfcem gazel k úkrytům.

**SHUFFLE  
& SORT**


1: ó  
5: náhlý  
4: déšť  
1: a  
...

**REDUCE**

1: LIST(ó, a, s, k)  
3: LIST(již, laň, teď)  
4: LIST(déšť, čilá, běží)  
5: LIST(náhlý, prach, gazel)  
...

1: 4  
3: 3  
4: 3  
5: 3  
6: 1  
7: 2

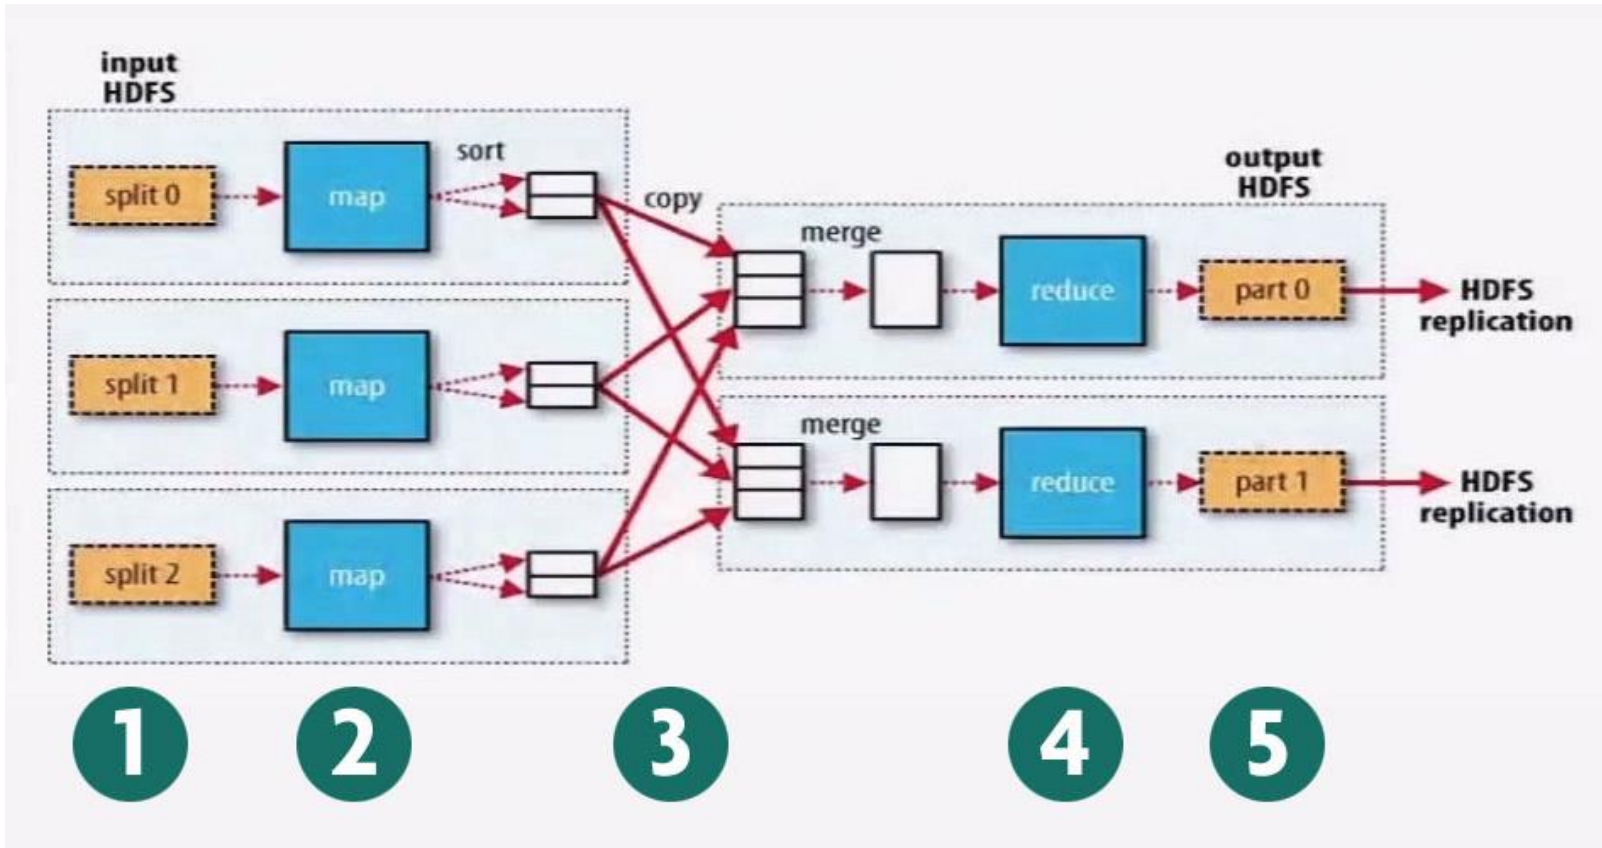
# MapReduce schema



# Stages

- › Planning
  - Do locally as most as you can – exploit HDFS replication
  - Transfer only necessary data over network
- › Mapper saves its result to a local FS (not HDFS)
- › Combine can do a useful aggregation on mapper side
- › Shuffle & sort
  - Transfers all data with same key to same node
  - The most demanding stage of MapReduce (network load)
- › Reduce
  - number of reducers  $\ll$  number of mappers
  - one reducer = one output file to HDFS

# MapReduce data flow



<https://content.pivotal.io/blog/hadoop-101-programming-mapreduce-with-native-libraries-hive-pig-and-cascading>

# Task resolvable by MapReduce

- › „Divide and conquer“
  
- › **Formal requirements on reduce operation**
  
- › ▲ must hold
  - a) asociativity:  $(A \blacktriangle B) \blacktriangle C = A \blacktriangle (B \blacktriangle C)$
  - b) exists  $\diamond$  such that  $A \blacktriangle \diamond = A$
  - c) commutativity:  $(A \blacktriangle B) = (B \blacktriangle A)$
  
- › Why?
  - a) we have no control over order of operations
  - b) node with no data must have no influence on results
  - c) mapper/shuffle can change data order




# Suitable tasks

- › „School tasks“
  - word count
- › Real tasks
  - summing up of partial reports
  - aggregating by items, clients, product types etc.
  - data filtering and validation
  - image processing
  - table joining
  - big data processing (matrix multiplication)

# Unsuitable tasks

- › Median
  - does not hold associativity
- › Average
  - average of partial averages  $\neq$  overall average
- › BUT: Task can be carried out smarter way.

# How to compute an average



- › Combiner
  - average on map side
- › Mapper output
  - SPLIT1:40.3333
  - SPLIT2:214.25
- › Reducer output
  - AVG:127.2917



- › Combiner
  - sum and count on map side
- › Mapper output
  - SPLIT1:121:3
  - SPLIT2:857:4
- › Reducer output
  - AVG:139.714

# MapReduce key features

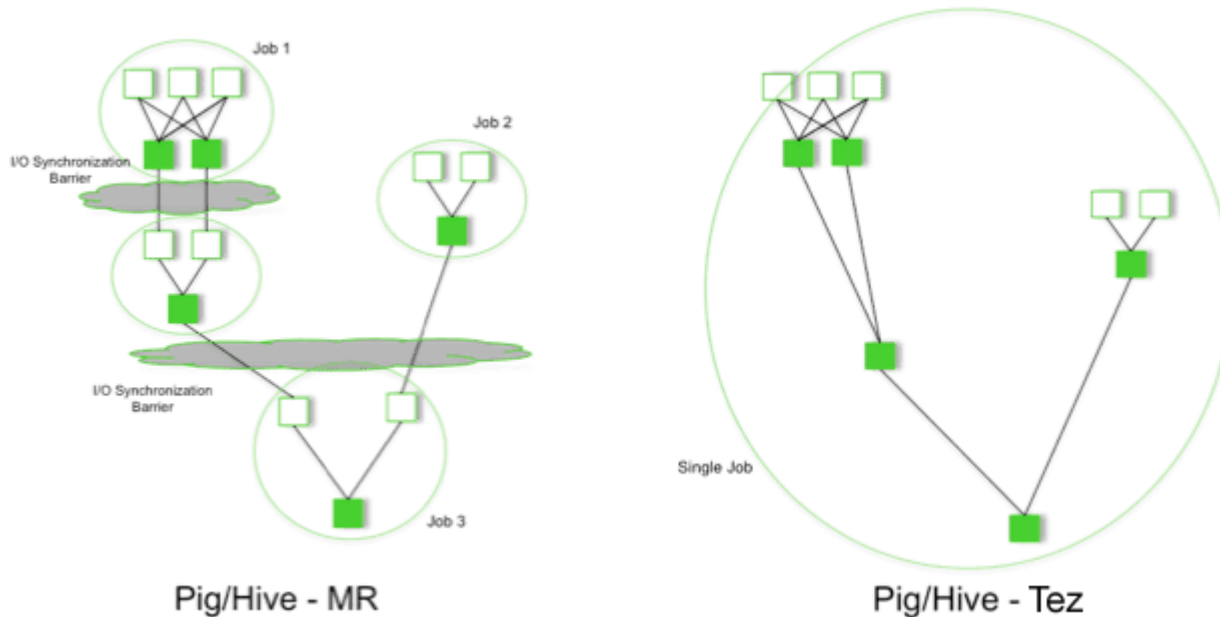
- › Stable, tested, reliable
  - „other components failing, MapReduce keeps running“
- › Very small memory requirements
  - majority of operations is done on disks
- › Very frequent disk usage
  - HDFS: only starting input and final output
  - local FS: intermediate results
  - many I/O operations
- › Rather slow compared to other components
  - long initializing

## When to use MapReduce

- › Task can run in parallel, „good“ reduce operation
- › Processing big big data (petabytes)
- › We have plenty of time
- › Batch processing (not interactive)
- › Lack of memory

# Variations on Hadoop

- › Apache Tez
  - „optimized“ MapReduce
  - similar principle but combining of successive steps
  - less I/O operations



# Variations on Hadoop

- › Apache Spark
  - similar principle but intermediate results are kept in memory
  - enables interactive work
  
- › ... next lecture

# Thanks for your attention

**PROFINIT**

NÁSKOK DÍKY ZNALOSTEM

Profinit EU, s.r.o.

Tychonova 2, 160 00 Praha 6 | Telefon + 420 224 316 016



Web

[www.profinit.eu](http://www.profinit.eu)



LinkedIn

[linkedin.com/company/profinit](https://linkedin.com/company/profinit)



Twitter

[twitter.com/Profinit\\_EU](https://twitter.com/Profinit_EU)



Facebook

[facebook.com/Profinit.EU](https://facebook.com/Profinit.EU)



Youtube

[Profinit EU](https://youtube.com/Profinit_EU)