



15-826: Multimedia Databases and Data Mining

Extra: intro to hadoop

C. Faloutsos



Resources

- Software
 - <http://hadoop.apache.org/>
- Map/reduce paper [Dean & Ghemawat]
 - <http://labs.google.com/papers/mapreduce.html>
- Tutorial: see part1, foils #9-20
 - videolectures.net/kdd2010_papadimitriou_sun_yan_lsdm/

15-826

(c) 2012 C. Faloutsos

2



Motivation: Scalability

- Google: > 450,000 processors in clusters of ~2000 processors each [Barroso, Dean, Hölzle, “*Web Search for a Planet: The Google Cluster Architecture*” IEEE Micro 2003]
- Yahoo: 5Pb of data [Fayyad, KDD’07]
- Problem: **machine failures**, on a daily basis
- How to parallelize data mining tasks, then?
- A: map/reduce – hadoop (open-source clone)
 - <http://hadoop.apache.org/>



15-826

(c) 2012 C. Faloutsos

3

 CMU SCS



2' intro to hadoop

- master-slave architecture; n-way replication (default n=3)
- ‘group by’ of SQL (in parallel, fault-tolerant way)
- e.g. find histogram of word frequency
 - compute local histograms
 - then merge into global histogram

```
select course-id, count(*)
from ENROLLMENT
group by course-id
```

15-826 (c) 2012 C. Faloutsos 4

 CMU SCS

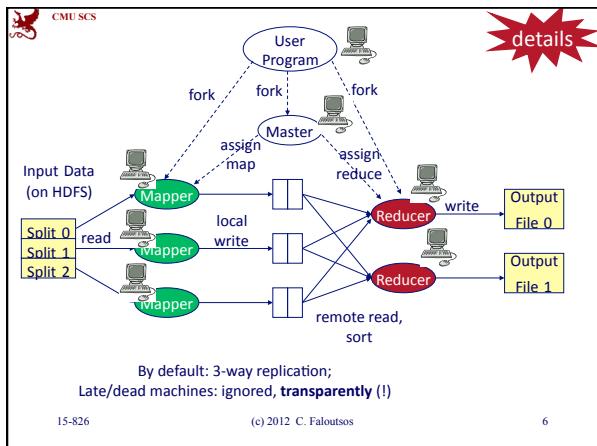


2' intro to hadoop

- master-slave architecture; n-way replication (default n=3)
- ‘group by’ of SQL (in parallel, fault-tolerant way)
- e.g. find histogram of word frequency
 - compute local histograms
 - then merge into global histogram

```
select course-id, count(*) reduce
from ENROLLMENT
group by course-id map
```

15-826 (c) 2012 C. Faloutsos 5



 CMU SCS

More details:

- (thanks to U Kang for the animations)

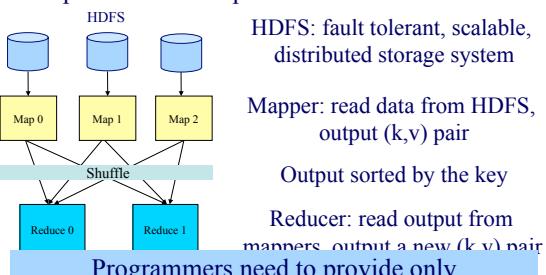


15-826 (c) 2012 C. Faloutsos 7

 CMU SCS

Background: MapReduce

- MapReduce/Hadoop Framework



HDFS: fault tolerant, scalable, distributed storage system

Mapper: read data from HDFS, output (k, v) pair

Output sorted by the key

Reducer: read output from mappers, output a new (k, v) pair

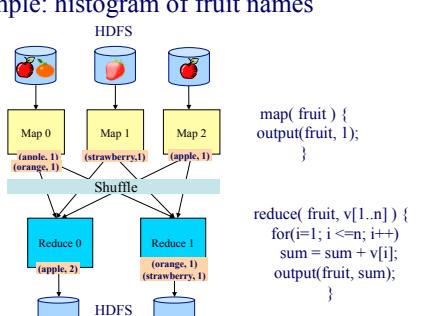
Programmers need to provide only `map()` and `reduce()` functions

15-826 (c) 2012 C. Faloutsos 8

 CMU SCS

Background: MapReduce

- Example: histogram of fruit names



```

map( fruit ) {
    output(fruit, 1);
}

reduce( fruit, v[1..n] ) {
    for(i=1; i <=n; i++)
        sum = sum + v[i];
    output(fruit, sum);
}

```

15-826 (c) 2012 C. Faloutsos 9



Conclusions

- Convenient mechanism to process Tera- and Peta-bytes of data, on >hundreds of machines
- User provides `map()` and `reduce()` functions
- Hadoop handles the rest (eg., machine failures etc)

15-826

(c) 2012 C. Faloutsos

10
