

Introduction to Data Management CSE 344

Lecture 23: Pig Latin

Some slides are courtesy of
Alan Gates, Yahoo!Research

Magda Balazinska – CSE344 Fall 2011

1

Outline

- Review of MapReduce
- Introduction to Pig System
- Pig Latin tutorial

Magda Balazinska – CSE344 Fall 2011

2

MapReduce

- Google: paper published 2004
- Open source implementation: Hadoop
- MapReduce = high-level programming model and implementation for large-scale parallel data processing
- Main competing alternative Dryad from Microsoft

3

MapReduce Data Model

- Files !
- A file = a bag of (key, value) pairs
- A MapReduce program:
 - Input: a bag of (input key, value) pairs
 - Output: a bag of (output key, value) pairs

Magda Balazinska – CSE344 Fall 2011

4

Step 1: the MAP Phase

- User provides the MAP function:
 - Input: one (input key, value)
 - Output: a bag of (intermediate key, value) pairs
- System applies map function in parallel to all (input key, value) pairs in the input file

Magda Balazinska – CSE344 Fall 2011

5

Step 2: the REDUCE Phase

- User provides the REDUCE function:
 - Input: intermediate key, and bag of values
 - Output: bag of output values
- System groups all pairs with the same intermediate key, and passes the bag of values to the REDUCE function

Magda Balazinska – CSE344 Fall 2011

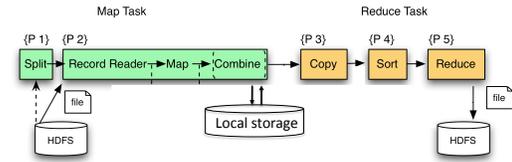
6

Parallel Execution in Cluster

See illustration on white board:

- Data is typically a file in the Google File System
 - HDFS for Hadoop
 - File system partitions file into chunks
 - Each chunk is replicated on k (typically 3) machines
- Each machine can run a few map and reduce tasks simultaneously
- Each map task consumes one chunk
 - Can adjust how much data goes into each map task using "splits"
 - Scheduler tries to schedule map task where its input data is located
- Map output is partitioned across reducers
- Map output is also written locally to disk
- Number of reduce tasks is configurable
- System shuffles data between map and reduce tasks
- Reducers sort-merge data before consuming it

MapReduce Phases

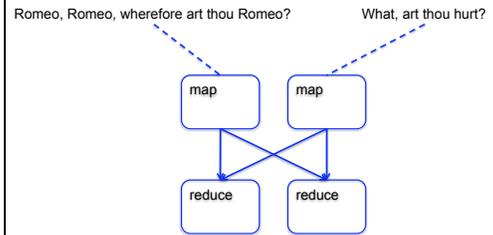


MapReduce

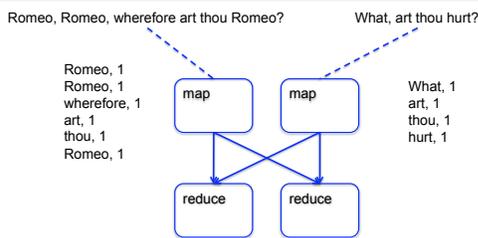
- Computation is moved to the data
- A simple yet powerful *programming* model
 - Map: every record handled individually
 - Shuffle: records collected by key
 - Reduce: key and iterator of all associated values
- User provides:
 - input and output (usually files)
 - map Java function
 - key to aggregate on
 - reduce Java function
- Opportunities for more control: partitioning, sorting, partial aggregations, etc.



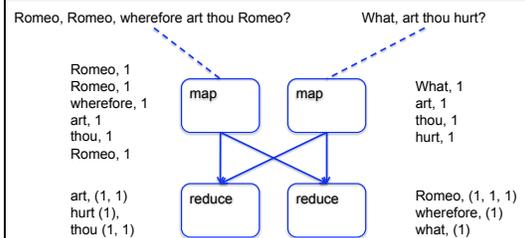
MapReduce Illustrated

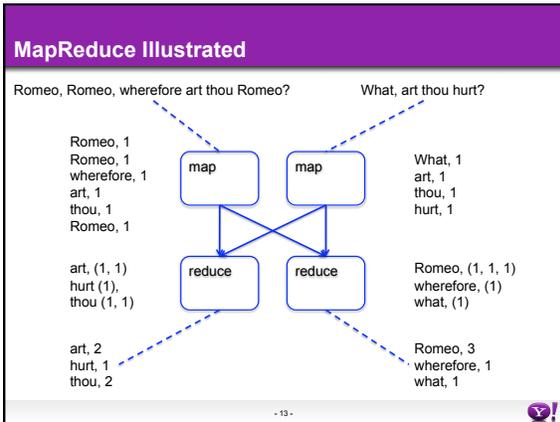


MapReduce Illustrated



MapReduce Illustrated





- ### Making Parallelism Simple
- Sequential reads = good read speeds
 - In large cluster failures are guaranteed; MapReduce handles retries
 - Good fit for batch processing applications that need to touch all your data:
 - data mining
 - model tuning
 - Bad fit for applications that need to find one particular record
 - Bad fit for applications that need to communicate between processes; oriented around independent units of work
- 14 -

MapReduce (MR) tools

Open source MapReduce implementation: 

One MR query language: *Pig Latin*

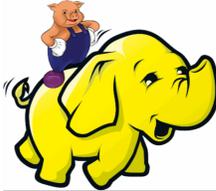
Query engine: 

Graphics taken from: hadoop.apache.org and research.yahoo.com/node/90

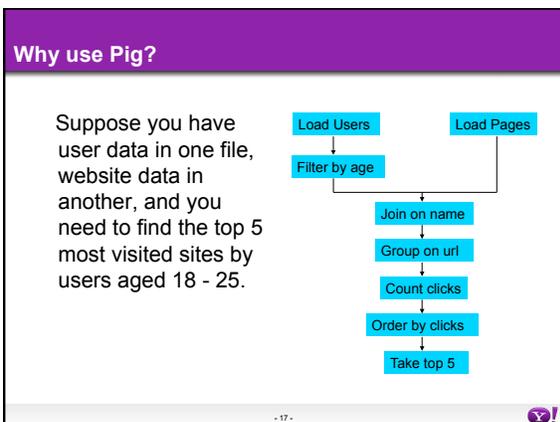
15

What is Pig?

- An engine for executing programs on top of Hadoop
- It provides a language, Pig Latin, to specify these programs
- An Apache open source project
<http://hadoop.apache.org/pig/>



- 16 -



In MapReduce

```

    import java.io.*;
    import java.util.*;
    import org.apache.hadoop.conf.*;
    import org.apache.hadoop.fs.*;
    import org.apache.hadoop.mapreduce.*;
    import org.apache.hadoop.mapreduce.Mapper;
    import org.apache.hadoop.mapreduce.Reducer;
    import org.apache.hadoop.mapreduce.lib.input.*;
    import org.apache.hadoop.mapreduce.lib.output.*;
    import org.apache.hadoop.util.*;

    public class MapReduce {
        public static void main(String[] args) throws Exception {
            Configuration conf = new Configuration();
            Job job = new Job(conf, "MapReduce");
            job.setJarByClass(MapReduce.class);
            job.setMapperClass(MyMapper.class);
            job.setReducerClass(MyReducer.class);
            job.setInputFormatClass(TextInputFormat.class);
            job.setOutputFormatClass(TextOutputFormat.class);
            FileInputFormat.addInputPath(job, new File("input"));
            FileOutputFormat.addOutputPath(job, new File("output"));
            job.waitForCompletion(true);
        }
    }

    class MyMapper extends Mapper<Text, Text> {
        public void map(Text key, Text value, Context context) throws IOException, InterruptedException {
            // Map logic here
        }
    }

    class MyReducer extends Reducer<Text, Text, Text, Text> {
        public void reduce(Text key, Iterable<Text> values, Context context) throws IOException, InterruptedException {
            // Reduce logic here
        }
    }
    
```

170 lines of code, 4 hours to write

- 18 -

In Pig Latin

```

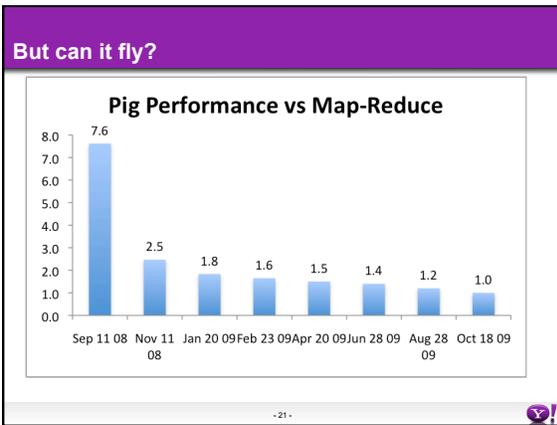
Users = load 'users' as (name, age);
Fltrd = filter Users by
    age >= 18 and age <= 25;
Pages = load 'pages' as (user, url);
Jnd = join Fltrd by name, Pages by user;
Grpd = group Jnd by url;
Smmd = foreach Grpd generate group,
    COUNT(Jnd) as clicks;
Srttd = order Smmd by clicks desc;
Top5 = limit Srttd 5;
store Top5 into 'top5sites';
    
```

9 lines of code, 15 minutes to write

Pig System Overview

```

A = LOAD 'file1' AS (sid,pid,mass,px:double);
B = LOAD 'file2' AS (sid,pid,mass,px:double);
C = FILTER A BY px < 1.0;
D = JOIN C BY sid,
    B BY sid;
STORE g INTO 'output.txt';
    
```



Pig Latin Mini-Tutorial

Based entirely on *Pig Latin: A not-so-foreign language for data processing*, by Olston, Reed, Srivastava, Kumar, and Tomkins, 2008

- ### Pig-Latin Overview
- Data model = loosely typed *nested relations*
 - Query model = Relational Algebra (less SQL)
 - Execution model:
 - Option 1: run locally (pig -x local)
 - Option 2: run in a Hadoop cluster
 - Option 1 is good for debugging

Pages(url, category, pagerank)

Example: In SQL...

```

SELECT category, AVG(pagerank)
FROM Pages
WHERE pagerank > 0.2
GROUP By category
HAVING COUNT(*) > 106
    
```

Pages(url, category, pagerank)

Example: ...in Pig-Latin

```

pages = LOAD 'pages-file.txt' as (url, category, pagerank);
good_pages = FILTER pages BY pagerank > 0.2;
groups = GROUP good_pages BY category;
big_groups = FILTER groups
              BY COUNT(good_pages) > 106;
result = FOREACH big_groups GENERATE
          group, AVG(good_pages.pagerank);
    
```

Use "describe" to check schema

For the last statement, can also use:
 result = FOREACH big_groups GENERATE \$0, AVG(\$1.pagerank);

Magda Balazinska -- CSE344 Fall 2011 25

Types in Pig-Latin

- Atomic: string or number, e.g. 'Alice' or 55
- Tuple: ('Alice', 55, 'salesperson')
- Bag: (('Alice', 55, 'salesperson'), ('Betty', 44, 'manager'), ...)
- Maps: we will not use these

Magda Balazinska -- CSE344 Fall 2011 26

Types in Pig-Latin

Bags can be nested !

- {{('a', {1,4,3}), ('c', {}), ('d', {2,2,5,3,2})}

Tuple components can be referenced by name or by number

- \$0, \$1, \$2, ...
- url, category, pagerank, ...

Magda Balazinska -- CSE344 Fall 2011 27

$$t = ('alice', \{ ('lakers', 1), ('iPod', 2) \}, ['age' \rightarrow 20])$$

Let fields of tuple t be called f1, f2, f3

Expression Type	Example	Value for t
Constant	'bob'	Independent of t
Field by position	\$0	'alice'
Field by name	f3	'age' → 20
Projection	f2.\$0	{ ('lakers'), ('iPod') }
Map Lookup	f3#'age'	20
Function Evaluation	SUM(f2.\$1)	1 + 2 = 3
Conditional Expression	f3#'age' > 18? 'adult': 'minor'	'adult'
Flattening	FLATTEN(f2)	'lakers', 1 'iPod', 2

Loading data

- Input data = FILES !
- LOAD command parses an input file
- Parser provided by user
 - In HW6: myudfs.jar, function RDFSplit3
 - Example.pig shows you how to use it

Magda Balazinska -- CSE344 Fall 2011 29

Loading data

```

pages = LOAD 'pages-file.txt'
        USING myLoad( )
        AS (url, category, pagerank)
    
```

```

pages = LOAD 'pages-file.txt'
        USING PigStorage(',')
        AS (url, category, pagerank)
    
```

Magda Balazinska -- CSE344 Fall 2011 30

FOREACH

```
pages_projected =
  FOREACH pages
  GENERATE url, pageRank
```

Note: you cannot write joins in FOREACH
FOREACH pages, users = error

Magda Balazinska -- CSE344 Fall 2011

31

Loading Data in HW6

There is no parser for RDF data, hence:

```
-- example.pig = in your starter code
register s3n://uw-cse344-code/myudfs.jar

raw = LOAD 's3n://uw-cse344-test/cse344-test-file'
      USING TextLoader as (line:chararray);

ntriples = FOREACH raw
           GENERATE FLATTEN(myudfs.RDFSplit3(line))
           AS (subject:chararray,
              predicate:chararray,
              object:chararray);
```

FILTER

Remove all queries from Web bots:

```
good_pages = FILTER pages BY pageRank > 0.8;
```

Magda Balazinska -- CSE344 Fall 2011

33

JOIN

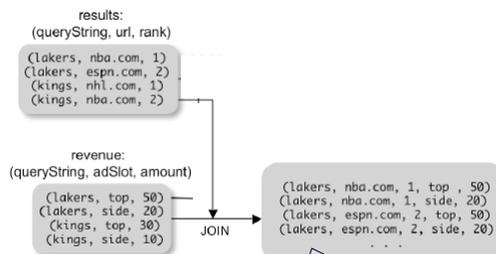
```
results:  {(queryString, url, position)}
revenue:  {(queryString, adSlot, amount)}
```

```
join_result = JOIN results BY queryString
               revenue BY queryString
```

```
join_result : {(queryString, url, position, queryString, adSlot, amount)}
```

Magda Balazinska -- CSE344 Fall 2011

34



Magda Balazinska -- CSE344 Fall 2011

35

GROUP BY

```
revenue:  {(queryString, adSlot, amount)}
```

```
grp_revenue = GROUP revenue BY queryString
query_revenues =
  FOREACH grp_revenue
  GENERATE group as queryString,
           SUM(revenue.amount) AS totalRevenue
```

```
grp_revenue: {(group, {(queryString, adSlot, amount)}}}
query_revenues: {(queryString, totalRevenue)}
```

36

Simple MapReduce

input : {(field1, field2, field3, . . .)}

```
map_result = FOREACH input
    GENERATE FLATTEN(map(*))
key_groups = GROUP map_result BY $0
output = FOREACH key_groups
    GENERATE reduce(*)
```

map_result : {(a1, a2, a3, . . .)}
 key_groups : {(group, {(a1, a2, a3, . . .)}}} 37

Co-Group

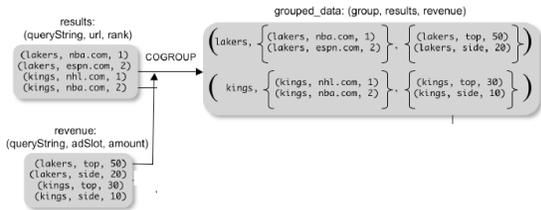
results: {(queryString, url, position)}
 revenue: {(queryString, adSlot, amount)}

```
grouped_data =
    COGROUP results BY queryString,
    revenue BY queryString;
```

grouped_data: {(group, results: {(queryString, url, position)},
 revenue: {(queryString, adSlot, amount)}}}

What is the output type in general ? 38

Co-Group



Is this an inner join or an outer join ?

Co-Group

grouped_data: {(group, results: {(queryString, url, position)},
 revenue: {(queryString, adSlot, amount)}}}

```
url_revenues = FOREACH grouped_data
    GENERATE
    FLATTEN(distributeRevenue(results, revenue));
```

distributeRevenue is a UDF that accepts search results and revenue information for a query string at a time, and outputs a bag of urls and the revenue attributed to them.

Co-Group v.s. Join

grouped_data: {(group, results: {(queryString, url, position)},
 revenue: {(queryString, adSlot, amount)}}}

```
grouped_data = COGROUP results BY queryString,
    revenue BY queryString;
join_result = FOREACH grouped_data
    GENERATE FLATTEN(results),
    FLATTEN(revenue);
```

Result is the same as JOIN

Asking for Output: STORE

In example.pig

```
STORE count_by_object_ordered
    IN '/user/hadoop/example-results'
    USING PigStorage();
```

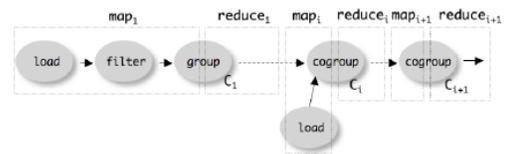
Implementation

- Over Hadoop !
- Parse query:
 - Everything between LOAD and STORE → one logical plan
- Logical plan → sequence of MapReduce jobs
- All statements between two (CO)GROUPs → one MapReduce job

Magda Balazinska – CSE344 Fall 2011

43

Implementation



Magda Balazinska – CSE344 Fall 2011

44