

Big Data: Architectures and Data Analytics

February 5, 2020

Student ID _____

First Name _____

Last Name _____

The exam is **closed book** and lasts **2 hours**

Part I

Answer to the following questions. There is only one right answer for each question.

1. (2 points) Consider the input HDFS folder *myFolder* that contains the following two files:
 - NamesItaly.txt
 - the text file NamesItaly.txt contains the following three lines (size: 33 bytes)
Luis,Turin
Luca,Rome
Paolo,Turin
 - NamesFrance.txt
 - the text file NamesFrance.txt contains the following two lines (size: 22 bytes)
Paolo,Nice
Luis,Paris

Suppose that you are using a Hadoop cluster that can potentially run up to 5 instances of the mapper class in parallel. Suppose the HDFS block size is 128MB. Suppose to execute a MapReduce application for Hadoop that analyzes the content of *myFolder*. Suppose the map phase emits overall the following 5 key-value pairs (the key part is a name while the value part is always 1):

("Luis", 1)
("Luca", 1)
("Paolo", 1)
("Paolo", 1)
("Luis", 1)

Suppose the number of instances of the reducer class is set to 4 and suppose the reduce method of the reducer class sums the values associated with each key and emits one pair (name, sum values) for the keys associated with "sum values" greater than or equal to 2. Suppose the following 2 pairs are overall emitted by the reduce phase:

("Paolo", 2)
("Luis", 2)

Considering all the instances of the reducer class, overall, how many times is the **reduce method** invoked?

- a) 2
- b) 3
- c) 4
- d) 5

2. (2 points) Consider the input HDFS folder *myFolder* that contains the following two files:

- Temperatures2019.txt: size=1030MB
- Temperatures2020.txt: size=1018MB

Suppose that you are using a Hadoop cluster that can potentially run up to 10 instances of the mapper class in parallel. Suppose to execute a MapReduce application for Hadoop that computes some statistics about temperatures. This application is based on one single MapReduce job. The input of this Hadoop application is the HDFS folder *myFolder*. The HDFS block size is 1024MB. The number of instances of the reducer class is set to 5. How many mappers are instantiated by Hadoop (i.e., how many instances of the mapper class) when you execute this MapReduce application by specifying the folder *myFolder* as input?

- a) 10
- b) 5
- c) 3
- d) 2

Part II

PoliCars is an international company characterized by several production plants around the world. In each production plant, robots are used to produce some components of the PoliCars's cars. PoliCars computes a set of statistics about its production plants and robots. The analyses are performed by considering the following input data sets/files related to the failures of the PoliCars's robots.

- ProductionPlants.txt

- ProductionPlants.txt is a text file containing the list of production plants of PoliCars. Each line of ProductionPlants.txt is associated with one production plant. The number of production plants is 1000. In some cities, there is more than one production plant.

- Each line of ProductionPlants.txt has the following format

- PlantID, City, Country

where *PlantID* is the production plant identifier while *City* and *Country* are the city and the country in which the production plant is located, respectively.

- For example, the following line

PID1, Turin, Italy

means that the production plant **PID1** is located in **Turin, Italy**

- Robots.txt

- Robots.txt is a text file containing the list of robots managed by PoliCars. One line for each robot of PoliCars is stored in Robots.txt. The number of managed robots is more than 15000.

- Each line of Robots.txt has the following format

- RID, PlantID, IP

where *RID* is the robot identifier, *PlantID* is the identifier of the production plant in which the robot is installed, and *IP* is its IP address.

- For example, the following line

R15, PID1, 130.192.20.21

means that robot **R15** is installed in the production plant **PID1** and the IP address associated with that robot is **130.192.20.21**.

- Failures.txt
 - Failures.txt containing the historical information about the failures of the robots. A new line is inserted in Failures.txt every time there is a failure of a robot. The number of managed robots is more than 15000 and Failures.txt contains the historical data about the Failures of the last 20 years.
 - Each line of Failures.txt has the following format
 - RID,FailureTypeCode,Date,Time

where *RID* is the identifier of the robot that had a failure, *Date* is the date of the failure, *Time* is the time of the failure, and *FailureTypeCode* is the code of the type of Failure.

- For example, the following line

R15,FCODE122,2020/05/01,06:40:51

means that robot **R15** had a failure of type **FCODE122** at **06:40:51** (hour=06, minute=40, second=51) of **May 1, 2020**.

Exercise 1 – MapReduce and Hadoop (7 points)

The managers of PoliCars are interested in performing some analyses about their robots.

Design a single application, based on MapReduce and Hadoop, and write the corresponding Java code, to address the following point:

1. *Number of robots associated with the failure type FCODE122 at least one time.* The application computes the number of robots associated with the failure type FCODE122 at least one time. Store the computed value in the output HDFS folder (the output will contain one single line associated with the computed value).

Example.

For instance, suppose that Failures.txt contains only the following lines:

```
R1,FCODE122,2020/05/01,06:40:51
R1,FCODE36,2020/05/02,06:40:51
R1,FCODE122,2020/05/03,06:40:51
R2,FCODE122,2020/05/04,06:40:51
R2,FCODE11,2020/05/05,06:40:51
R3,FCODE54,2020/05/14,06:40:51
```

In this case the number of robots associated with the failure type FCODE122 at least one time is two and hence one single line containing the value 2 is stored in the output folder.

Suppose that the input is Failures.txt and it has been already set and also the name of the output folder has been already set.

- Write only the content of the Mapper and Reducer classes (map and reduce methods. setup and cleanup if needed). The content of the Driver must not be reported.
- Use the next two specific multiple-choice questions to specify the number of instances of the reducer class for each job.
- If your application is based on two jobs, specify which methods are associated with the first job and which are associated with the second job.
- If you need personalized classes report for each of them:
 - name of the class
 - attributes/fields of the class (data type and name)
 - personalized methods (if any), e.g, the content of the toString() method if you override it
 - do not report get and set methods. I suppose they are "automatically defined"

Exercise 1.1 - Number of instances of the reducer - Job 1 - MapReduce and Hadoop
(0.5 points)

Select the number of instances of the reducer class of the first Job

- (a) 0
- (b) exactly 1
- (c) any number ≥ 1

Exercise 1.2 - Number of instances of the reducer - Job 2 - MapReduce and Hadoop
(0.5 points)

Select the number of instances of the reducer class of the second Job

- (a) One single job is needed for this MapReduce application
- (b) 0
- (c) exactly 1
- (d) any number ≥ 1

Exercise 2 – Spark and RDDs (19 points)

The managers of PoliCars are interested in performing some analyses related to their production plants and robots.

The managers of PoliCars asked you to develop one single application to address all the analyses they are interested in. The application has five arguments: the three input files ProductionPlants.txt, Robots.txt, Failures.txt and two output folders, “outPart1/” and “outPart2/”, which are associated with the outputs of the following Points 1 and 2, respectively.

Specifically, design a single application, based on Spark, and write the corresponding code, to address the following points:

1. *Production plants with at least one robot with at least 50 failures in year 2020.* The application considers only the failures occurred in year 2020 and selects the identifiers (PlantIDs) of the production plants with at least one robot associated with at least 50 failures in year 2020. The PlantIDs of the selected production plants are stored in the first HDFS output folder (one PlantID per line).
2. *For each production plant compute the number of robots with at least one failure in year 2020.* The application considers only the failures occurred in year 2020 and computes for each production plant the number of its robots associated with at least one failure in year 2020. The application stores in the second HDFS output folder, for each production plant, its PlantID and the computed number of robots with at least one failure in year 2020. **Pay attention** that the output folder must contain also one line for each of the production plants for which all robots had no failures in year 2020. For those production plants the associated output line is (PlantID, 0).

Examples Point 2

- Suppose that in the production plant PID1 there are three robots identified by the identifiers R1, R2, and R3, respectively. Suppose that the number of failures in year 2020 of R1 is 10, the number of failures in year 2020 of R2 is 15, and the number of failures in year 2020 of R3 is 3. In this case, in the second HDFS output folder, the output line associated with the production plant PID1 will be
 - PID1,3
- Suppose that in the production plant PID2 there are two robots identified by the identifiers R4 and R5, respectively. Suppose that the number of failures in year 2020 of R4 is 0 and the number of failures in year 2020 of R5 is also 0. In this case, in the second HDFS output folder, the output line associated with the production plant PID2 will be
 - PID2,0

Predefined Template

...

```
/* Suppose all the needed imports are already set */
```

...

```
public class SparkDriver {  
  
    public static void main(String[] args) {  
        String inProdPlants;  
        String inRobots;  
        String inProdPlants;  
        String outputPathPart1;  
        String outputPathPart2;  
  
        inProdPlants = "ProductionPlants.txt";  
        inRobots = "Robots.txt";  
        inFailures = "Failures.txt";  
  
        outputPathPart1 = "outPart1/";  
        outputPathPart2 = "outPart2/";  
  
        // Create a configuration object and set the name of the application  
        SparkConf conf = new SparkConf().setAppName("Spark Exam - Exercise #2");  
  
        // Create a Spark Context object  
        JavaSparkContext sc = new JavaSparkContext(conf);  
  
        /* Write your code here */  
  
    }  
}
```