

Lab 5

The objective of this laboratory is to start playing around with Apache Spark. We provide you a template in Java, from which you can write your first application.

The workflow for building and executing an Apache Spark application is the same as the one we used with Hadoop: once unzipped the template, import and open the project in your IDE (Eclipse), modify it, and then export a jar. Then, submit this job either locally (i.e., on the PC of LABINF):

```
spark-submit --class it.polito.bigdata.spark.example.SparkDriver
--deploy-mode client --master local SparkProject.jar arguments
```

or on the cluster (you must execute spark-submit on the bigdatalab.polito.it gateway to submit your application on the cluster):

```
spark-submit --class it.polito.bigdata.spark.example.SparkDriver -
-deploy-mode cluster --master yarn SparkProject.jar arguments
```

Notice how in the template we did not, even if possible, configure the Spark master (we did not use the `setMaster(...)` method) and other settings explicitly in the main class of the driver. In this way, you can always submit the same jar on any configuration, either Spark standalone locally or a Yarn cluster. The main advantage of this method is for testing your application: first you can try your application on your machine, on a sample file, and then safely submit the same application for the real work.

1. Problem specification

If you completed Lab 1, you should now have (at least one) huge files with the word frequencies in the amazon food reviews, in the format `word\tfreq`, where *freq* is an integer (a copy of the output of Lab 1 is available in the HDFS shared folder `/data/students/bigdata-01QYD/Lab2/`).

Your task is to write a **Spark** application to filter these results, analyze the filtered data and compute some statistics on them.

The first filter you should implement is the following:

- keep only the lines containing words that start with a prefix (a string) that is passed as a command-line parameter

The result is the set of lines (`word\tfreq`) that satisfy the filtering operation.

Print on the standard output:

- the number of selected lines
- the maximum frequency (*maxfreq*) among the ones of the selected lines (i.e., the maximum value of *freq* in the lines obtained by applying the filter).

In the second part of your application, among the lines selected by the first filter, you have to apply another filter to select only the most frequent words. Specifically, your application must select those lines that contain words with a frequency (*freq*) greater than 80% of the maximum frequency (*maxfreq*) computed before. Hence, implement the following filter:

- keep only the lines with a frequency *freq* greater than $0.8 * maxfreq$.

Finally, perform the following operations on the selected lines (the ones selected by applying both filters):

- count the number of selected lines and print this number on the standard output
- save the selected words (without frequency) in an output folder for inspecting the results

2. Testing the application

1. Run you application locally on the **LABINF PC** by using the following command and select only the lines containing the words starting with "ho" from the local file SampleLocalFile.csv (open a Linux shell on the LABINF PC and execute the following command in the folder containing the jar file associated with your application and SampleLocalFile.csv)

```
spark-submit --class it.polito.bigdata.spark.example.SparkDriver
--deploy-mode client --master local SparkProject.jar
SampleLocalFile.csv LocalOutputFolder "ho"
```

Analyze the content of the local output folder.

2. Run you application on the cluster by using the following command and select only the lines containing the words starting with "ho" from the content of the HDFS folder /data/students/bigdata-01QYD/Lab2/ (Use ssh to open a shell on the gateway bigdatalab.polito.it by using the usual procedure and then execute the following command)

```
spark-submit --class it.polito.bigdata.spark.example.SparkDriver
--deploy-mode cluster --master yarn SparkProject.jar
/data/students/bigdata-01QYD/Lab2/ HDFSOutputFolder "ho"
```

Analyze the content of the HDFS output folder.

How to access logs by using the job history web page

The logs of your applications on the cluster is available at <https://ma1-bigdata.polito.it:18489/>
If there is any error and you need to analyze the content of the error log files follow these steps:

1. Open a new terminal on your local machine (do **not** use the previous terminal connected to bigdatalab.polito.it).

2. On the local terminal command line, type the command
kinit sXXXXXX
(then, you will be asked for your BigDataLab password)
This command is used to obtain a Kerberos ticket (without it you cannot access the log web page)
3. Now you should have access to the spark job history and logs at
<https://ma1-bigdata.polito.it:18489/>
4. Select your job and click on its App ID
You can use the search box to select only your jobs
5. Click on the executor tab and analyze the stderr associated with you application.
Pay attention that there is one log file for the driver and one log file for each executor.
If you application raised an exception, analyze all the stderr files to identify the reason of the error (depending on the error type, the log containing the needed information can be the one associated with the driver or that of an executor)

How to access logs without the web page and the Kerberos ticket

If you are connecting from outside Polito, and hence 1) you cannot request a Kerberos ticket and 2) you cannot access the history web page, you can proceed as follows to retrieve the log files from the command line:

1. Open a shell on the gateway by using ssh
ssh smatricola@bigdatalab.polito.it
2. Execute the following command in the remote shell:
yarn logs -applicationId *application_1521819176307_2195*
The last parameter is the application/job ID. You can retrieve the job ID of your application on the HUE interface: <https://bigdatalab.polito.it:8080/hue/jobbrowser#!jobs>

Bonus questions

Explore the web dashboard of your jobs¹. The history and the dashboard of the jobs can be accessed at the following URL <https://ma1-bigdata.polito.it:18489/>

Note that a Kerberos ticket is needed to access that web page. Use the same step reported above to obtain a Kerberos ticket and access the interface available at <https://ma1-bigdata.polito.it:18489/>

Check out the DAG of your application (under Jobs tab click on the saveAsTextFile at SparkDriver.java link to access the DAG) and the statistics for the application, and then answer the following questions:

- How many executors have been allocated to your job? (The information is available in the Executors tab)
- How was the input data split among them?
- Was your application well balanced among the executors, or were there tasks much slower/heavier/bigger than the others? Why?

Keep in mind these questions for the next labs as well.

¹ For local applications, the web dashboard is killed as soon as your job finishes. On the cluster, it is moved on a history server, so you can access it also after the job has completed.