

EPL646 – Advanced Topics in Databases

Advanced Hadoop

<http://www.cs.ucy.ac.cy/~dzeina/courses/epl646/labs/lab.html>



Calculate time

- How to calculate the time

```
long begin = System.currentTimeMillis();  
job.waitForCompletion(true);  
long end = System.currentTimeMillis();  
long second = (end - begin) / 1000;  
System.err.println(job.getJobName() + " takes " +  
second + " seconds");
```

Task1: N-Gram

- Change the code in WordCount so that it counts how many times each set of five consecutive words appears
- You can find the code of WordCount from the solution of the previous lab
- If you don't have the datasets you can download them from the previous lab

Task1: N-Gram

- Function ***map*** will have as input:
 - key = line offset (we can ignore it)
 - value = a whole line from one of the input files
- Function ***map*** will have as output:
 - key = five words
 - value = 1
- Function ***reduce*** will have as input:
 - key = five words
 - value = [a list of number 1]
- The list will as many 1 as there are appearances of the five consecutive words in our data
- Function ***Reduce*** will have as final output:
 - key = five words
 - value = the sum of all 1 (i.e. the same as WordCount)

Task2: Anagram

- An anagram is a word that can be created by the movement of the letters of another word
- E.g.
 - Refills → fillers
 - Relayed → layered
 - Rentals → antlers
 - Rebuild → builder
- You must find the anagrams in a huge input file. How would you do it?

```
public static boolean isAnagram(String first, String second) {
    // Checks that the two inputs are anagrams, by checking they have all the
    // same characters.
    // Left as exercise for the user...
}
```

Task2: Anagram

- Hadoopifying...
 - (input) $\langle k1, v1 \rangle \rightarrow$
 - map $\rightarrow \langle k2, v2 \rangle \rightarrow$
 - combine $\rightarrow \langle k2, v2 \rangle \rightarrow$
 - reduce $\rightarrow \langle k3, v3 \rangle$ (**output**)
- Download
 - /usr/share/dict/words or /usr/dict/words

Jar file configuration

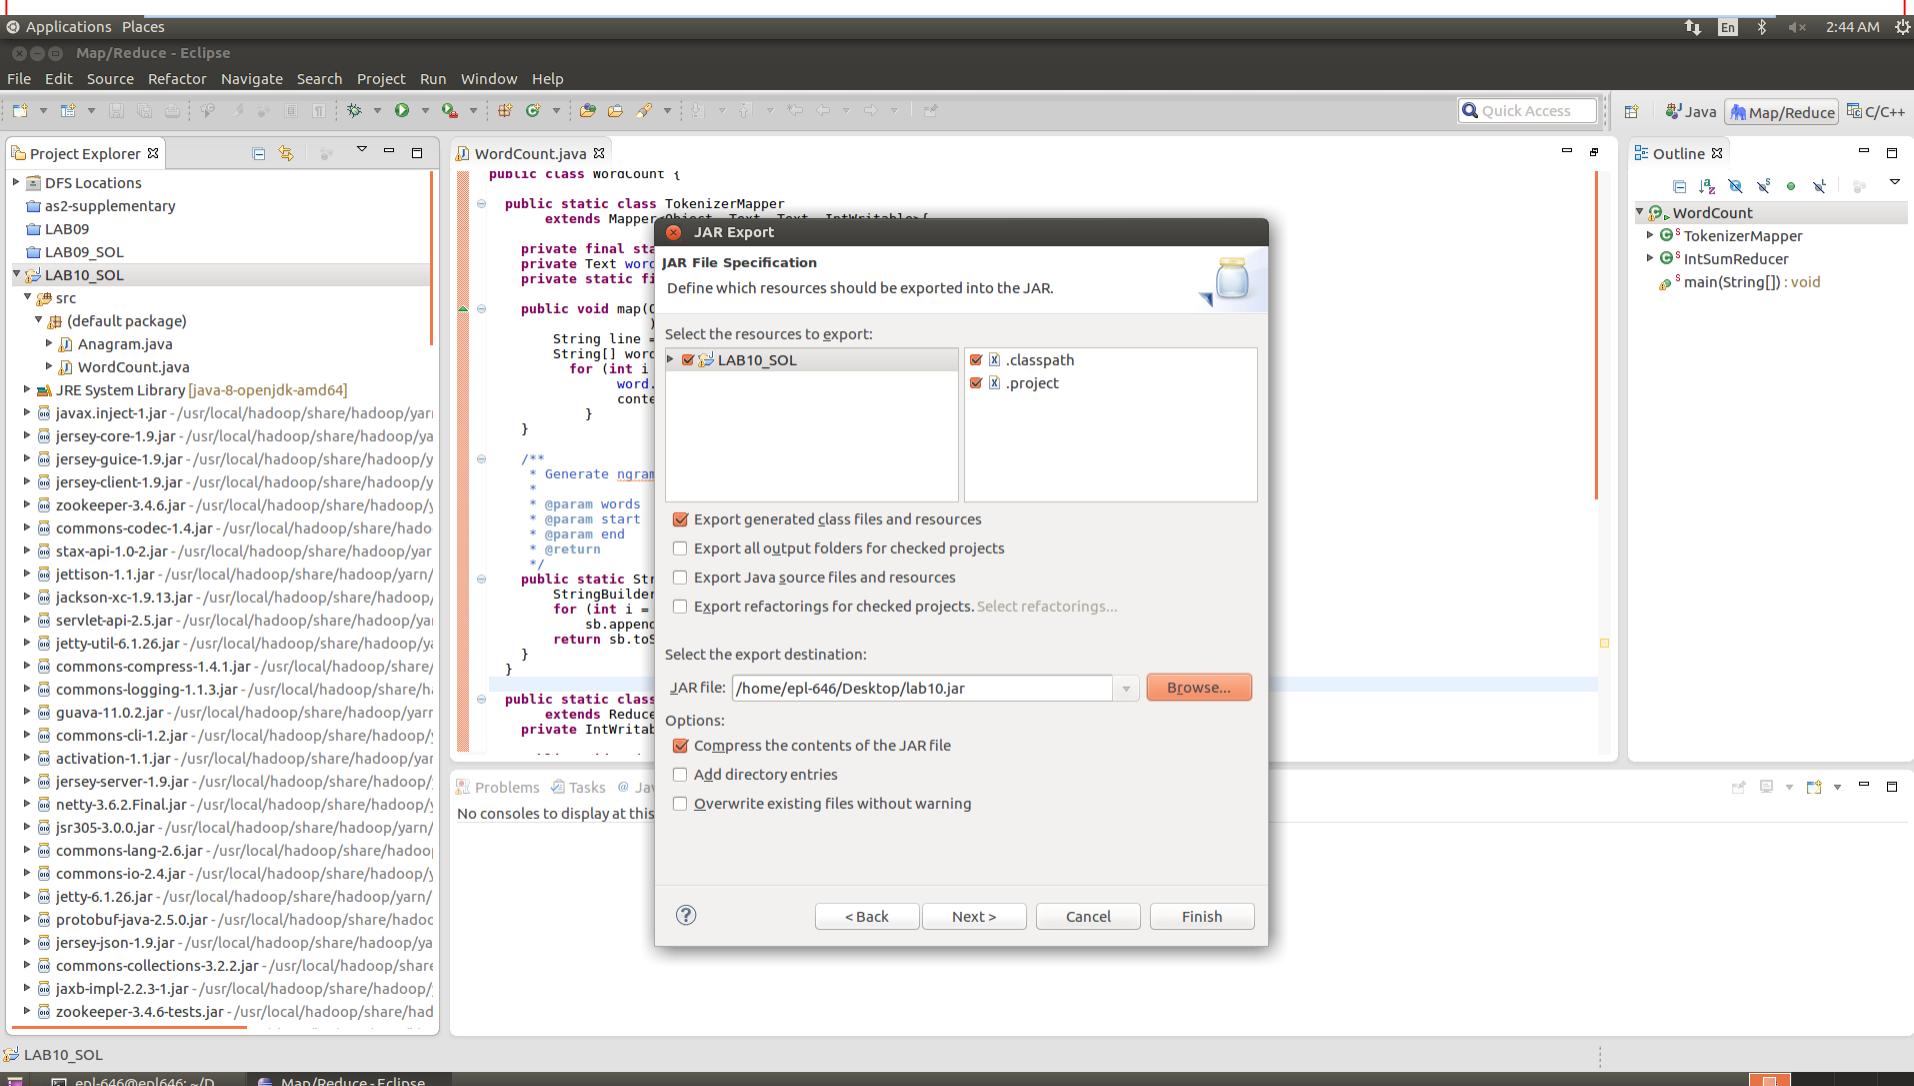
You need to set the jar by class parameter:

```
Configuration conf = new Configuration();
Job job = Job.getInstance(conf, "word count");
job.setJarByClass(WordCount.class);
```

Else you will get:

java.lang.RuntimeException: java.lang.ClassNotFoundException:

Export the .jar file



Run the jar file

The screenshot shows the Hadoop Web UI interface. On the left, there's a sidebar with 'Cluster Metrics' and 'Scheduler Metrics' sections. The main area is titled 'All Applications' and lists one application entry:

ID	User	Name	Application Type	Queue	StartTime	FinishTime	State	FinalStatus	Progress	Tracking UI	Blacklisted Nodes
application_1510706565009_0001	epl-646	word count	MAPREDUCE	default	Wed Nov 15 02:45:27 +0200 2017	Wed Nov 15 02:46:04 +0200 2017	FINISHED	SUCCEEDED	History	N/A	

Below the application list is a terminal window showing the command-line output of a Hadoop job:

```
epl-646@epl646:~/Desktop$ hadoop jar lab10.jar Anagram
17/11/15 02:45:23 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-Java classes where applicable
17/11/15 02:45:25 INFO client.RMProxy: Connecting to ResourceManager at /0.0.0.0:8032
17/11/15 02:45:25 WARN mapreduce.JobResourceUploader: Hadoop command-line option parsing not performed. Implement the Tool interface and execute your application via ToolRunner to run it from the command-line.
17/11/15 02:45:26 INFO InputFormat: Total input paths to process : 1
17/11/15 02:45:27 INFO mapreduce.JobSubmitter: number of splits:1
17/11/15 02:45:27 INFO mapreduce.JobSubmitter: Submitting tokens for job: job_150706565009_0001
17/11/15 02:45:28 INFO Impl.YarnClientImpl: Submitted application application_150706565009_0001
17/11/15 02:45:28 INFO mapreduce.Job: The url to track the job: http://epl646:88/proxy/application_150706565009_0001/
17/11/15 02:45:28 INFO mapreduce.Job: Running job: job_150706565009_0001
17/11/15 02:45:41 INFO mapreduce.Job: Job job_150706565009_0001 running in uber mode : false
17/11/15 02:45:41 INFO mapreduce.Job: map 0% reduce 0%
17/11/15 02:45:53 INFO mapreduce.Job: map 100% reduce 0%
17/11/15 02:46:04 INFO mapreduce.Job: map 100% reduce 100%
17/11/15 02:46:06 INFO mapreduce.Job: Job job_150706565009_0001 completed successfully
17/11/15 02:46:06 INFO mapreduce.Job: Counters:
File System Counters
FILE: Number of bytes read=462972
```

- If the node is unhealthy you may need to execute the following command and then restart hadoop:

```
chown -r epl-646:epl-646 /app/hadoop/
```

Questions?

<http://www.cs.ucy.ac.cy/~dzeina/courses/epl646/labs/lab.html>

