

# Data Intensive Computing – Handout 3

## Training Cluster

You can login to the training cluster *dlrc* via machine `ufallab.ms.mff.cuni.cz` and port 11422, i.e., using `ssh -p 11422 ufallab.ms.mff.cuni.cz`.

All machines in the cluster use same architecture and share `/dlrc_share` directory.

The cluster consists of master `dlrc` and 5 nodes `dlrc-node1` to `dlrc-node5`. Each node has 2 cores and 4GB ram, but four SGE jobs are allowed to run simultaneously (only because otherwise there would be too little slots; the slot number may be even increased later).

## Running Browser with Proxy to the Cluster

To run a browser, which can access master *and* the cluster nodes, run:

```
(chromium --proxy-server=socks://localhost:2020& \  
ssh -ND 2020 -p 11422 ufallab.ms.mff.cuni.cz)
```

## HDFS Distributed Filesystem

The distributed filesystem HDFS is installed on the cluster.

You can browse the filesystem using `dlrc:50070`.

The HDFS filesystem can be manipulated using `hadoop fs` command, which you can shorten to `dfs`. It prints nice usage when executed without arguments, some simple commands are:

- `dfs -ls /`
- `dfs -mkdir /home/straka/test`
- `dfs -put .bashrc /home/straka/test`
- `dfs -get /home/straka/test .`
- `dfs -rm -r /home/straka/test`

Every user has directory `/home/username` which they can write to.

## Wikipedia Data

Compressed and structured Wikipedia data is available in the HDFS:

- `/data/wiki/cs`: Czech Wikipedia data (Sep 2009), 124k articles.
- `/data/wiki/cs-medium`: Subset of Czech Wikipedia data, 5307 articles.
- `/data/wiki/cs-small`: Subset of Czech Wikipedia data, 13 articles.
- `/data/wiki/en`: English Wikipedia data (Sep 2009), 2.9M articles.

All files are so called *sequence files*, every record is a pair (*article name*, *article text*).

## DataFrames API of Apache Spark

DataFrame API is high-level API on top of RDDs.

### Python Example

dataframes.py

```
#!/usr/bin/python

# Create DataFrame from RDD
wiki = sc.textFile("...")
df = sqlContext.createDataFrame(wiki, ['article', 'text'])

# Create DataFrame from Hive table
users = sqlContext.table("users")

# Load from file
val people = sqlContext.read.json(path)

# DataFrame operations
employees
    .join(dept, employees.deptId == dept.id)
    .where(employees.gender == "female")
    .groupBy(dept.id, dept.name)
    .count()

df.select("name").show()
df.select(df.name, df.age + 1).show()
df.filter(df.age > 21).show()
df.groupBy(df.age).count().show()

# Create RDD from DataFrame
employees.map(lambda row: row.name)

# Use SQL
val people = sqlContext.read.json("...")
people.registerTempTable("people")
teenagers = sqlContext.sql("SELECT name FROM people WHERE age >= 13 AND age <= 19")
```

The example is available in /home/straka/examples.

### Tasks

When needing to split given text into words (called *tokens*), use function `wordpunct_tokenize` from `nltk.tokenize` package.

<i>Task</i>	<i>Points</i>	<i>Description</i>
spark_unique_words	2	<p>Create a list of unique words used in the articles using Spark. Convert them to lowercase to ignore case.</p> <p>Use HDFS <code>/data/wiki/{cs,en}</code> as input and either <code>/dlrc_share/data/wiki/tokenizer/{cs,en}_tokenizer</code> or <code>nlTK.tokenize.wordpunct_tokenize</code> as a tokenizer.</p>
anagrams	2	<p>Two words are anagrams if one is a letter permutation of the other (ignoring case).</p> <p>For a given input, find all anagram classes that contain at least <math>A</math> words. Output each anagram class on a separate line.</p> <p>Use HDFS <code>/data/wiki/{cs,en}</code> as input and either <code>/dlrc_share/data/wiki/tokenizer/{cs,en}_tokenizer</code> or <code>nlTK.tokenize.wordpunct_tokenize</code> as a tokenizer.</p>
sort	3	<p>You are given data consisting of (31-bit integer, string data) pairs. These are available in plain text format:</p> <ul style="list-style-type: none"> <li>• <code>/dlrc_share/data/numbers-txt/numbers-small:</code> 3MB</li> <li>• <code>/dlrc_share/data/numbers-txt/numbers-medium:</code> 184MB</li> <li>• <code>/dlrc_share/data/numbers-txt/numbers-large:</code> 916MB</li> </ul> <p>You can assume that the integers are uniformly distributed. Your task is to sort these data, comparing the key numerically and not lexicographically. The lines in the output must be the same as in the input, only in different order.</p> <p>Your solution should work for TBs of data. For that reason, you must use multiple machines. If your job is executed using <math>m</math> machines, the output consists of <math>m</math> files, which when concatenated would produce sorted (key, value) pairs. In other words, each of the output files contains sorted (integer, data) pairs and all keys in one file are either smaller or larger than in other file. Your solution should work for any number of machines specified.</p> <p>Obviously, do not use <code>sort</code> nor <code>sortByKey</code> method in your solution.</p>

<i>Task</i>	<i>Points</i>	<i>Description</i>
<code>nonuniform_sort</code>	4	<p>Improve the <code>sort</code> task to handle nonuniform data. You can use the following exponentially distributed data:</p> <ul style="list-style-type: none"> <li>• <code>/dlrc_share/data/numbers-txt/nonuniform-small:</code> 3MB</li> <li>• <code>/dlrc_share/data/numbers-txt/nonuniform-medium:</code> 160MB</li> <li>• <code>/dlrc_share/data/numbers-txt/nonuniform-large:</code> 797MB</li> </ul> <p>Assume we want to produce <math>m</math> output files. One of the solutions is the following:</p> <ul style="list-style-type: none"> <li>• Go through the data and sample only a small fraction of the keys. As there are not so many of them, they can fit in one reducer.</li> <li>• Find best <math>m - 1</math> separators using the sampled data.</li> <li>• Run the second pass using the computed separators.</li> </ul>
<code>spark_inverted_index</code>	2	<p>Compute inverted index in Spark – for every lowercased word from the articles, compute <i>(article name, ascending positions of occurrences as word indices)</i> pairs.</p> <p>The output should be a file with one word on a line in the following format:</p> <pre>word \t article name \t space separated occurrences \t article name \t space separated occurrences ...</pre> <p>You will get 2 additional points if the articles will be numbered using consecutive integers. In that case, the output is ascending <i>(article id, ascending positions of occurrences as word indices)</i> pairs, together with a file containing list of articles representing this mapping (the article on line <math>i</math> is the article with id <math>i</math>).</p> <p>Use HDFS <code>/data/wiki/{cs,en}</code> as input and either <code>/dlrc_share/data/wiki/tokenizer/{cs,en}_tokenizer</code> or <code>nltk.tokenize.wordpunct_tokenize</code> as a tokenizer.</p>
<code>no_references</code>	3	<p>An article <math>A</math> is said to reference article <math>B</math>, if it contains <math>B</math> as a token (ignoring case).</p> <p>Run a Spark job which for each article <math>B</math> counts how many references for article <math>B</math> there exist in the whole wiki (summing references in a single article).</p> <p>Use HDFS <code>/data/wiki/{cs,en}</code> as input and either <code>/dlrc_share/data/wiki/tokenizer/{cs,en}_tokenizer</code> or <code>nltk.tokenize.wordpunct_tokenize</code> as a tokenizer.</p>

<i>Task</i>	<i>Points</i>	<i>Description</i>																
spark_wordsim_index	3	<p>In order to implement word similarity search, compute for each form with at least three occurrences all <i>contexts</i> in which it occurs, including their number of occurrences. List the contexts in ascending order.</p> <p>Given <math>N</math> (either 1, 2, 3), the <i>context</i> of a form occurrence is <math>N</math> forms preceding this occurrence and <math>N</math> forms following this occurrence (ignore sentence boundaries, use empty words when article boundaries are reached).</p> <p>The output should be a file with one form on a line in the following format:  form \t context \t counts \t context \t counts ...</p> <p>Use HDFS /data/wiki/{cs,en} as input and either /dlrc_share/data/wiki/tokenizer/{cs,en}_tokenizer or nltk.tokenize.wordpunct_tokenize as a tokenizer.</p>																
spark_wordsim_find	3	<p>Let <math>S</math> be given natural number. Using the index created in <code>spark_wordsim_index</code>, find for each form <math>S</math> most similar forms. List only forms with non-zero similarity. The similarity of two forms is computed using <i>cosine similarity</i> as <math>\frac{C_A \cdot C_B}{ C_A  \cdot  C_B }</math>, where <math>C_F</math> is a vector of occurrences of form <math>F</math> contexts.</p> <p>The output should be a file with one form on a line in the following format:  lemma \t most similar lemma \t cosine similarity  \t 2. most similar lemma \t cosine similarity ...</p>																
kmeans	6	<p>Implement K-means clustering algorithm as described on <a href="http://en.wikipedia.org/wiki/K-means_clustering#Standard_algorithm">http://en.wikipedia.org/wiki/K-means_clustering#Standard_algorithm</a>.</p> <p>The user specifies number of iterations and the program run specified number of K-means clustering algorithm iterations. You can use the following training data. Each line contains space separated coordinates of one points. The coordinates in one input naturally have the same dimension.</p> <table border="1"> <thead> <tr> <th><i>Path in /dlrc_share</i></th> <th><i>Points</i></th> <th><i>Dimension</i></th> <th><i>Clusters</i></th> </tr> </thead> <tbody> <tr> <td>/data/points-txt/small</td> <td>10000</td> <td>50</td> <td>50</td> </tr> <tr> <td>/data/points-txt/medium</td> <td>100000</td> <td>100</td> <td>100</td> </tr> <tr> <td>/data/points-txt/large</td> <td>500000</td> <td>200</td> <td>200</td> </tr> </tbody> </table> <p>You will get 2 additional points if the algorithm stops when none of the centroid positions change more than given <math>\varepsilon</math>.</p>	<i>Path in /dlrc_share</i>	<i>Points</i>	<i>Dimension</i>	<i>Clusters</i>	/data/points-txt/small	10000	50	50	/data/points-txt/medium	100000	100	100	/data/points-txt/large	500000	200	200
<i>Path in /dlrc_share</i>	<i>Points</i>	<i>Dimension</i>	<i>Clusters</i>															
/data/points-txt/small	10000	50	50															
/data/points-txt/medium	100000	100	100															
/data/points-txt/large	500000	200	200															