

TDDE31/732A54-Big Data Analytics Lab Compendium

*Notice: Please make sure you have read the whole lab compendium before you start to work on the server
from NSC.*

April 5th, 2024

1 Description and Aim

In the lab you will work on the Sigma¹ set up which is a HPC cluster from the National Supercomputer Centre (NSC). You are supposed to work with the historical meteorological data from the Swedish Meteorological and Hydrological Institute (SMHI). Specifically, you will work with air temperature readings and precipitation readings from 812 meteorological stations in Sweden². In these exercises, you will work with Spark 3.5.1³.

There are three lab topics which are *BDA1-Spark*, *BDA2-Spark SQL* and *BDA3-Machine Learning with Spark*. After completing the first two labs you will have basic knowledge of the programming environment, and PySpark API. You will work on exercises with Spark and Spark SQL and thus will be able to compare the differences between the two approaches. In the third lab, you are supposed to achieve a machine learning method with Spark. The overview of three labs is that you need to upload data to Hadoop Distributed File System (HDFS), read the data from HDFS in your code and then program with PySpark⁴ to solve the assignments in each lab.

2 SMHI Data

The data includes air temperature and precipitation readings from 812 stations in Sweden. The stations include both currently active stations and historical stations that have been closed down. The latest readings available for active stations are from October 10, 2016. The air temperature and precipitation records are **hourly readings**, however some stations provide only one reading every three hours.

The provided files⁵ are prepared csv files with removed headers. Values are separated with semicolons. Some files are too big to be read using some text editors. Therefore, please use

¹Sigma at: <https://www.nsc.liu.se/systems/sigma/>

²If interested in other readings please check: <http://opendata-catalog.smhi.se/explore/>

³Spark 3.5.1 at: <https://spark.apache.org/docs/3.5.1/>

⁴PySpark 3.5.1 at: <https://spark.apache.org/docs/3.5.1/api/python/reference/index.html>

⁵Within the BDA_demo: <https://www.ida.liu.se/~732A54/lab/data.zip>

either python to read the files or bash commands such as *tail* and *more* to get an overview of a file's content. Provided files:

- temperature-readings.csv - ca 2 GB
- precipitation-readings.csv - ca 660 MB
- stations.csv
- stations-Ostergotland.csv

All these files are available at the following folder at Sigma:

`/software/sse2/tetralith_el9/manual/spark/course-examples/BDA_demo/input_data`

The headers of these files are shown in Table 1, Table 2 and Table 3. If you notice any mistakes in the dataset or have any comments please contact the lab assistants.

Station number	Date	Time	Air temperature (in °C)	Quality ^a
----------------	------	------	-------------------------	----------------------

^a G - controlled and confirmed values, Y - suspected or aggregated values

Table 1: Headers for temperature-readings.csv

Station number	Date	Time	Precipitation (in mm)	Quality ^a
----------------	------	------	-----------------------	----------------------

^a G - controlled and confirmed values, Y - suspected or aggregated values

Table 2: Headers for precipitation-readings.csv

Station number	Station name	Measurement height	Latitude	Longitude	Readings from (date and time)	Readings to (date and time)	Elevation
----------------	--------------	--------------------	----------	-----------	-------------------------------	-----------------------------	-----------

Table 3: Headers for stations.csv and stations-Ostergotland.csv

3 Running PySpark Program on Sigma

3.1 Working on Sigma

The Sigma server is available at ***sigma.nsc.liu.se*** (log in using your NSC accounts). You can use (**Option 1:**) ssh forwarding connection or (**Option 2:**) Thinlinc⁶ connection to log in Sigma. You can also use regular ssh connection without forwarding option, in this case, you have to program locally and use *scp* command to upload your code to Sigma, and then check the history log from text file. **Note:** The forwarding function of ssh supports to run graphics applications (e.g., Atom editor) remotely.

- **Option 1:** when you are at a computer in an SU room at the university,

⁶<https://www.cendio.com/thinlinc/download>

- **ssh -X username@sigma.nsc.liu.se** where username is your NSC username (not the LiU one), -X indicates forwarding function of ssh which is used for running graphics applications remotely.
 - **[username@sigma ~] \$ exit** is used to logout Sigma. If it is hung on, please use ctrl-c to terminate the connection.
 - You can use Emacs editor for coding by running following command first in a terminal:
 - * **emacs &**
 - You can also use VS Code editor for coding by running following two commands first in a terminal:
 - * **/proj/liu-compute-2024-2/shared/VSCode-linux-x64/code --no-sandbox &**
 - **scp LOCAL_FILE_PATH username@sigma.nsc.liu.se:Documents** is used for uploading files from your local machine to Sigma. (**Note:** you are supposed to run *scp* command before you log in Sigma when you want to upload files to Sigma.)
 - You can always first connect to the university server and from there connect to Sigma. Information regarding remote access to LiU's system can be found [here](#).
- **Option 2:** Sigma's Thinlinc server is available same as **sigma.nsc.liu.se**. In this way, you can get a graphical environment on Sigma and given that you work directly on Sigma there is no need to use *ssh* or *scp*. **Note:** Please remember to log out when done working on the labs so that Sigma does not keep open Thinlinc sessions. Also, each pair of students, please use **at most one** Thinlinc connection during lab sessions, due to the limited number of Thinlinc licenses on Sigma.
 - When you are at a computer in an SU room at the university:
 - * You need to open a terminal and run following two lines:
 - * **\$ module load courses/732A54** or **\$ module load courses/TDDE31**
 - * **\$ tlclient**
 - If you use your own computers, you can just download Thinlinc first and then connect directly.
 - As mentioned above, on Sigma, you can use Emacs or VS Code editors for coding by running aforementioned commands in a terminal.

NSC reserves nodes for each lab session of the course, which means other jobs on Sigma will not use these nodes during our lab sessions. Table 4 shows reservation names for our scheduled lab sessions as well as a reservation, **devel**, which you can use at any time (e.g., outside of lab sessions).

Each time when you submit a job, the options '**-A liu-compute-2024-2**', and '**--reservation RESERVATION_NAME**' should be specified.

- **[username@sigma ~] \$ listreservations**, by running this command, you can see all the reservations.
- **[username@sigma ~] \$ sbatch -A liu-compute-2024-2 --reservation liu-bda-2024-04-09 run_yarn_with_historyserver.q**

Table 4: Time and Reservation Name

RESERVATION_NAME	Time
liu-bda-2024-04-09	04-09, 17:00 to 19:15
liu-bda-2024-04-11	04-11, 08:00 to 10:15
liu-bda-2024-04-12	04-12, 15:00 to 17:15
liu-bda-2024-04-16	04-16, 13:00 to 17:15
liu-bda-2024-04-18	04-18, 08:00 to 10:15
liu-bda-2024-04-19	04-19, 15:00 to 17:15
liu-bda-2024-04-23	04-23, 13:00 to 17:15
liu-bda-2024-04-25	04-25, 08:00 to 10:15
liu-bda-2024-04-26	04-26, 15:00 to 17:15
liu-bda-2024-04-30	04-30, 13:00 to 17:15
liu-bda-2024-05-02	05-02, 08:00 to 10:15
liu-bda-2024-05-03	05-03, 15:00 to 17:15
liu-bda-2024-05-14	05-14, 13:00 to 17:15
liu-bda-2024-05-16	05-16, 08:00 to 10:15
liu-bda-2024-05-17	05-17, 15:00 to 17:15
liu-bda-2024-05-21	05-21, 13:00 to 17:15
devel	Anytime

3.2 Submitting Jobs to Hadoop Cluster and Sigma

To submit your pyspark code to the Hadoop cluster, you will use:

- ***spark-submit --deploy-mode cluster --master yarn --num-executors 9 --driver-memory 2g --executor-memory 2g --executor-cores 4 CODE.py*** where CODE.py is the python script in the current folder. In this command, Yarn is used for resource management and the cluster-deploy mode is used.

To run your pyspark code on Hadoop cluster, you also need to first submit a job to Sigma. Some scripts in the demo (in next section) are provided to you so you can make the call of your *spark-submit* command. You will use the non-interactive way to submit a job on Sigma. Each time after you submit, the job will enter the scheduling queue. You can use **sbatch** command to submit the job, **squeue** command to monitor the submitted job and may use **scancel** command to cancel a job.

- ***[username@sigma ~] \$ sbatch -A liu-compute-2024-2 --reservation liu-bda-2024-04-09 run_yarn_with_historyserver.q*** (**Note:** Please remember to use '-A liu-compute-2024-2' if you are in more than one project from NSC, it will guarantee to use the allocation reserved on Sigma for our course.)
- ***[username@sigma ~] \$ squeue -u username***
- ***[username@sigma ~] \$ scancel JOB_ID***

Sigma uses Slurm for scheduling. Once you submit a job, the job will be assigned an ID. After the job is finished, you will see a **slurm-JOB_ID.out** file returned, which includes the

output information of the job script. The script `run_yarn_with_historyserver.q` is supposed to contain commands for constructing the Hadoop cluster for your account, commands for interacting with Hadoop Distributed File System (HDFS) and `spark-submit` command for running pyspark code. A detailed example (Figure 3) will be shown in following section which introduces the demo.

Another way to access the logs after the execution, you will need to set the `spark.eventLog.enabled` which is already contained in `'run_yarn_with_historyserver.q'`. After the job finishes, you can access the history server by running:

- `spark_browse_job`

Then a firefox session that points at the history server web interface will be opened. In this way, it only lists the latest finished job on the history server.

3.3 Demo Code

You can find the demo from following folders after you log in Sigma.

- `/software/sse2/tetralith_el9/manual/spark/course-examples/BDA_demo/`

The steps to run the `BDA_demo` are shown as below, and in Figure 1 and Figure 2.

- Step 1: Login Sigma with `'ssh -X'` connection or Thinlinec.
- Step 2: Copy the demo to your home folder on Sigma.
- Step 3: Use `sbatch` to submit the job. Then use `squeue` command to monitor your job. Once the job is finished, check the returned file named `slurm-ID.out`. **Notice:** when you want to run your own code, you need to make some changes in `'run_yarn_with_historyserver.q'` to make sure you create folders on HDFS and copy data to HDFS correctly (see Figure 3).
- Step 4: Check the history log from `slurm-ID.out` or run `spark_browse_job`, or check the output folder.

In `'run_yarn_with_historyserver.q'` as shown in Figure 3, you can see a number of commands that are used to interact with HDFS.

- `hadoop fs -mkdir <FOLDER_NAME>` -make a folder on HDFS
- `hadoop fs -mkdir -p <FOLDER_NAME> <FOLDER_NAME>` -make multiple folders
- `hadoop fs -test -d <FOLDER_NAME>` -if the path is a directory, return 0
- `hadoop fs -rm -r <FOLDER_NAME>` -deletes the directory and any content under it recursively
- `hadoop fs -cat <FOLDER_ON_HDFS> [local]` -copy HDFS path to stdout
- `hadoop fs -copyFromLocal <localsrc> ... <dst>` -copy single src, or multiple srcs from local Sigma to HDFS
- `hadoop fs -copyToLocal <dst> ... <localsrc>` -copy single src, or multiple srcs from HDFS to local Sigma

```

(py312) huali50@mac01048 ~ % ssh -X x_huali@sigma.nsc.liu.se
(x_huali@sigma.nsc.liu.se) Password:
(x_huali@sigma.nsc.liu.se) Verification code:
Last login: Tue Jan 16 15:32:45 2024 from 2001:6b0:17:fc08:bccc:1abe:3bd4:be6
Welcome to NSC and Sigma!

**** Project storage directories available to you:
/proj/liu-compute-2023-24/users/x_huali
/proj/theophys/users/x_huali
/proj/liu-compute-2024-2/users/x_huali

**** Documentation and getting help:
https://www.nsc.liu.se/support/systems/sigma-getting-started/
https://www.nsc.liu.se/support

**** Useful commands
To see your active projects and CPU time usage: projinfo
To see available disk storage and usage: snicquota
To see your last jobs: lastjobs
Login to compute node to check running job: jobsh

To tweak job priorities, extend timelimits and reserve nodes: see
https://www.nsc.liu.se/support/batch-jobs/boost-tools/

(Run "nsc-mute-login" to not show this ir

[x_huali@sigma ~]$ ls
2_java_wordcount_1.0 Desktop MapReduce
4_pyspark_wordcount Documents Music
732A54-test Downloads Pictures templates lab-test-2023

[x_huali@sigma ~]$ cp -r /software/sse2/tetralith_el9/manual/spark/course-examples/BDA_demo/ .
[x_huali@sigma ~]$ ls
2_java_wordcount_1.0 BDA_demo Downloads
4_pyspark_wordcount Desktop MapReduceLab
732A54-test Documents Music

[x_huali@sigma ~]$ cd BDA_demo/
[x_huali@sigma BDA_demo]$ ls
demo.py input_data run_local.q run_local. story
server.q

[x_huali@sigma BDA_demo]$ listreservations
Reservations available to user:x_huali / project(s):liu-compute-2024-2,liu-compute-2023-24,liu-2019-26
-26
devel from 2024-03-21T08:00:00 to 2
liu-bda-2024-04-09 from 2024-04-09T: 2)
liu-bda-2024-04-11 from 2024-04-11T08:00:00 to 2024-04-11T10:15:00 (project:liu-compute-2024-2)
liu-bda-2024-04-12 from 2024-04-12T15:00:00 to 2024-04-12T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-04-16 from 2024-04-16T13:00:00 to 2024-04-16T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-04-18 from 2024-04-18T08:00:00 to 2024-04-18T10:15:00 (project:liu-compute-2024-2)
liu-bda-2024-04-19 from 2024-04-19T15:00:00 to 2024-04-19T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-04-23 from 2024-04-23T15:00:00 to 2024-04-23T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-04-25 from 2024-04-25T15:00:00 to 2024-04-25T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-04-26 from 2024-04-26T15:00:00 to 2024-04-26T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-04-30 from 2024-04-30T15:00:00 to 2024-04-30T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-05-02 from 2024-05-02T15:00:00 to 2024-05-02T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-05-03 from 2024-05-03T15:00:00 to 2024-05-03T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-05-14 from 2024-05-14T15:00:00 to 2024-05-14T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-05-16 from 2024-05-16T08:00:00 to 2024-05-16T10:15:00 (project:liu-compute-2024-2)
liu-bda-2024-05-17 from 2024-05-17T15:00:00 to 2024-05-17T17:15:00 (project:liu-compute-2024-2)
liu-bda-2024-05-21 from 2024-05-21T13:00:00 to 2024-05-21T17:15:00 (project:liu-compute-2024-2)

Note: set one of the above as default by running:
userreservation RESERVATIONNAME
Or without the userreservation alias:
source /software/sse2/generic/manual/ssertools/v1.9.5/bin/hpc_userreservation RESERVATIONNAME

```

Step 1

Check your current home folder, you may have an empty one

Step 2: Copy the demo code's folder to your current folder

Enter the folder and list the content

Check reservations before submitting jobs

These are reservations for lab sessions; devel reservation can be used at any time; When you submit a job, specify the reservation name (Step 3).

Figure 1: Login Sigma and Copy the demo code (Steps 1 and 2).

```

[x_huali@sigma BDA_demo]$ sbatch -A liu-compute-2024-2 --reservation devel run_yarn_with_historyserver.q
Submitted batch job 3760375
[x_huali@sigma BDA_demo]$ squeue -u x_huali
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
3760375 sigma run_yarn x_huali R 0:06 2 n[1153,1179]
[x_huali@sigma BDA_demo]$ squeue -u x_huali
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
[x_huali@sigma BDA_demo]$ ls
demo.py output run_local_with_historyserver.q run_yarn_with_historyserver.q spark
input_data run_local.q run_yarn.q slurm-3760375.out
[x_huali@sigma BDA_demo]$ cd output/
[x_huali@sigma output]$ ls
_SUCCESS part-00000 part-00001

```

Step 3

Step 4: check outputs

This indicates that the program runs successfully (without syntax errors).

Figure 2: Submit a job and check the output (Steps 3 and 4).

```

#!/bin/bash
#SBATCH --time=10:00
#SBATCH --nodes=2
#SBATCH --exclusive

echo "START AT: $(date)"

module load spark/3.5.1-hadoop-3.3.6-hpc1-bdist

# Cleanup and start from scratch
rm -rf spark

# Startup hadoop filesystem and yarn
hadoop_setup

echo "Prepare output and input directories and files..."
# The following command will make folders on your home folder on HDFS, the input and output folders should be corresponding to the parameter you give to textFile and saveAsTextFile functions in the code
hadoop fs -mkdir -p "BDA" "BDA/input"
hadoop fs -test -d "BDA/output"
if [ "$?" == "0" ]; then
    hadoop fs -rm -r "BDA/output"
fi

hadoop fs -copyFromLocal ./input_data/temperature-readings-small.csv "BDA/input/"
# Remove the comment when you need specific file below
#hadoop fs -copyFromLocal ./input_data/temperature-readings.csv "BDA/input/"
#hadoop fs -copyFromLocal ./input_data/precipitation-readings.csv "BDA/input/"
#hadoop fs -copyFromLocal ./input_data/stations.csv "BDA/input/"
#hadoop fs -copyFromLocal ./input_data/stations-Ostergotland.csv "BDA/input/"

# Run your program
echo "Running Your program..."
exec 5>&1
APPLICATION_ID=$(spark-submit --conf spark.eventLog.enabled=true --deploy-mode cluster --master yarn --num-executors 9 --driver-memory 2g --executor-memory 2g --executor-cores 4 demo.py 2>&1 | tee >(cat ->&5) | awk '!found && /INFO.*Yarn.*Submitted application/ {tmp=gensub(/^.*Submitted application (.*)$/,"\\1","g");print tmp; found=1}')

echo "===== FINAL OUTPUT ====="
hadoop fs -cat "BDA/output"/*
echo "===== "
echo "Applicaton id: $APPLICATION_ID"
echo "===== "
echo "===== stderr ====="
echo "===== "
yarn logs -applicationId "$APPLICATION_ID" | awk -F: '/^LogType/ {if($2=="stderr") {output=1} else {output=0}} output==1 {print}'
echo "===== "
echo "===== result ====="
echo "===== "
yarn logs -applicationId "$APPLICATION_ID" | awk -F: '/^LogType/ {if($2=="stdout") {output=1} else {output=0}} output==1 {print}' | grep -v "WARN\\|INFO"

rm -rf output
hadoop fs -copyToLocal 'BDA/output' ./
hadoop_stop

echo "END AT: $(date)"

```

Please read these comment lines, you need to change these hadoop commands.

Here you need to replace with your own code.

Figure 3: 'run_yarn_with_historyserver.q' for running the code

4 Hand In

You are supposed to use GitLab⁷ to submit your report and code. For each lab, please submit the code and a report that contains your results (a snippet of the results is enough if the results contain many rows) and answers to the questions. In cases where a plot of your results is asked, you can include the figure directly in the report. You can use a tool of your preference to produce the plots (R, Excel, matplotlib in Python, etc.). Comment each step in your code to provide a clear picture of your reasoning when solving the problem.

⁷<https://gitlab.liu.se/olaha93/bigdata>