Getting Started With Spark using Python

Estimated time needed: **15** minutes



The Python API

Spark is written in Scala, which compiles to Java bytecode, but you can write python code to communicate to the java virtual machine through a library called py4j. Python has the richest API, but it can be somewhat limiting if you need to use a method that is not available, or if you need to write a specialized piece of code. The latency associated with communicating back and forth to the JVM can sometimes cause the code to run slower. An exception to this is the SparkSQL library, which has an execution planning engine that precompiles the queries. Even with this optimization, there are cases where the code may run slower than the native scala version. The general recommendation for PySpark code is to use the "out of the box" methods available as much as possible and avoid overly frequent (iterative) calls to Spark methods. If you need to write high-performance or specialized code, try doing it in scala. But hey, we know Python rules, and the plotting libraries are way better. So, it's up to you!

How to run Jupyter notebook cell ?

Jupyter Notebooks are a powerful way to write and iterate on your Python code for data analysis.The Jupyter Notebook is organized into cells. Each cell can contain code, text, or visual elements. Rather than writing and re-writing an entire program, Jupyter Notebooks allow you to write code in separate blocks (or “cells”) and run each block of code individually. Then, if you need to make a change, you can go back and make an edit and rerun the program again, all in the same window.

To select the cell, simply click on it. The selected cell will be highlighted, indicating that it is ready for execution.Once you have the desired cell selected, you have multiple options to run it:

- Keyboard Shortcut: Press Shift + Enter on your keyboard to run the selected cell. This will execute the code or process the content within the cell and display the output below it.

If the cell has any output, it will be shown just after running.

- Toolbar: Look for the "Run Cell" button in the toolbar at the top of the Theia Lab interface. Clicking on this button will also execute the selected cell.

After running the cell, you can see the output displayed below it. If the cell contains any code that produces visual outputs, such as plots or images, they will be shown in the output area.

To run subsequent cells, repeat the same steps: select the next cell you want to execute and use the appropriate method mentioned above. Remember to execute cells in the correct order if there are dependencies between them. For example,if you have variables defined in a previous cell that are needed in the current cell, make sure to run the preceding cell first.

Objectives

In this lab, we will go over the basics of Apache Spark and PySpark. We will start with creating the SparkContext and SparkSession. We then create an RDD and apply some basic transformations and actions. Finally we demonstrate the basics dataframes and SparkSQL.

After this lab you will be able to:

- Create the SparkContext and SparkSession
- Create an RDD and apply some basic transformations and actions to RDDs
- Demonstrate the use of the basics Dataframes and SparkSQL

Setup

For this lab, we are going to be using Python and Spark (PySpark). These libraries should be installed in your lab environment or in SN Labs.

In [1]:
```
# Installing required packages
!pip install pyspark
!pip install findspark
```

```
Requirement already satisfied: pyspark in /home/jupyterlab/conda/envs/pyth
on/lib/python3.7/site-packages (3.4.4)
Requirement already satisfied: py4j==0.10.9.7 in /home/jupyterlab/conda/en
vs/python/lib/python3.7/site-packages (from pyspark) (0.10.9.7)
Requirement already satisfied: findspark in /home/jupyterlab/conda/envs/py
thon/lib/python3.7/site-packages (2.0.1)
```

In [2]:
```
import findspark
findspark.init()
```

In [3]:
```
# PySpark is the Spark API for Python. In this lab, we use PySpark to ini
from pyspark import SparkContext, SparkConf
from pyspark.sql import SparkSession
```

Exercise 1 - Spark Context and Spark Session

In this exercise, you will create the Spark Context and initialize the Spark session needed for SparkSQL and DataFrames. SparkContext is the entry point for Spark applications and contains functions to create RDDs such as `parallelize()`. SparkSession is needed for SparkSQL and DataFrame operations.

Task 1: Creating the spark session and context

In [4]:
```
# Creating a spark context class
sc = SparkContext()

# Creating a spark session
spark = SparkSession \
    .builder \
    .appName("Python Spark DataFrames basic example") \
    .config("spark.some.config.option", "some-value") \
    .getOrCreate()
```

```
25/10/12 09:09:52 WARN util.NativeCodeLoader: Unable to load native-hadoop
library for your platform... using builtin-java classes where applicable
Setting default log level to "WARN".
To adjust logging level use sc.setLogLevel(newLevel). For SparkR, use setL
ogLevel(newLevel).
25/10/12 09:09:54 WARN util.Utils: Service 'SparkUI' could not bind on por
t 4040. Attempting port 4041.
```

Task 2: Initialize Spark session

To work with dataframes we just need to verify that the spark session instance has been created.

In [5]:
```
if 'spark' in locals() and isinstance(spark, SparkSession):
    print("SparkSession is active and ready to use.")
else:
    print("SparkSession is not active. Please create a SparkSession.")
```

```
SparkSession is active and ready to use.
```

Exercise 2: RDDs

In this exercise we work with Resilient Distributed Datasets (RDDs). RDDs are Spark's primitive data abstraction and we use concepts from functional programming to create and manipulate RDDs.

Task 1: Create an RDD.

For demonstration purposes, we create an RDD here by calling `sc.parallelize()`
We create an RDD which has integers from 1 to 30.

In [6]:
```
data = range(1,30)
# print first element of iterator
print(data[0])
len(data)
xrangeRDD = sc.parallelize(data, 4)

# this will let us know that we created an RDD
xrangeRDD
```

```
1
```

Out[6]:
```
PythonRDD[1] at RDD at PythonRDD.scala:53
```

Task 2: Transformations

A transformation is an operation on an RDD that results in a new RDD. The transformed RDD is generated rapidly because the new RDD is lazily evaluated, which means that the calculation is not carried out when the new RDD is generated. The RDD will contain a series of transformations, or computation instructions, that will only be carried out when an action is called. In this transformation, we reduce each element in the RDD by 1. Note the use of the lambda function. We also then filter the RDD to only contain elements <10.

In [7]:
```
subRDD = xrangeRDD.map(lambda x: x-1)
filteredRDD = subRDD.filter(lambda x : x<10)
```

Task 3: Actions

A transformation returns a result to the driver. We now apply the `collect()` action to get the output from the transformation.

In [8]:
```
print(filteredRDD.collect())
filteredRDD.count()
```

```
[0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
```

Out[8]:
```
10
```

Task 4: Caching Data

This simple example shows how to create an RDD and cache it. Notice the **10x speed improvement**! If you wish to see the actual computation time, browse to the Spark UI...it's at host:4040. You'll see that the second calculation took much less time!

In [9]:
```
import time
```

```
test = sc.parallelize(range(1,50000),4)
test.cache()

t1 = time.time()
# first count will trigger evaluation of count *and* cache
count1 = test.count()
dt1 = time.time() - t1
print("dt1: ", dt1)


t2 = time.time()
# second count operates on cached data only
count2 = test.count()
dt2 = time.time() - t2
print("dt2: ", dt2)

#test.count()
```

```
dt1:  0.9789566993713379
dt2:  0.2933692932128906
```

Exercise 3: DataFrames and SparkSQL

In order to work with the extremely powerful SQL engine in Apache Spark, you will need a Spark Session. We have created that in the first Exercise, let us verify that spark session is still active.

In [10]:
```
spark
```

Out[10]:

SparkSession - in-memory

SparkContext

Spark UI

Version	`v2.4.3`
Master	`local[*]`
AppName	`pyspark-shell`

Task 1: Create Your First DataFrame!

You can create a structured data set (much like a database table) in Spark. Once you have done that, you can then use powerful SQL tools to query and join your dataframes.

In [11]:
```
# Download the data first into a local `people.json` file
!curl https://cf-courses-data.s3.us.cloud-object-storage.appdomain.cloud/
```

```
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Cu
rrent
                                 Dload  Upload   Total   Spent    Left  Sp
eed
100    73  100    73    0     0    701      0 --:--:-- --:--:-- --:--:--
708
```

In [12]:
```
# Read the dataset into a spark dataframe using the `read.json()` functio
df = spark.read.json("people.json").cache()
```

In [13]:
```
# Print the dataframe as well as the data schema
df.show()
df.printSchema()
```

```
+----+-------+
| age|   name|
+----+-------+
|null|Michael|
|  30|   Andy|
|  19| Justin|
+----+-------+

root
 |-- age: long (nullable = true)
 |-- name: string (nullable = true)
```

In [14]:
```
# Register the DataFrame as a SQL temporary view
df.createTempView("people")
```

Task 2: Explore the data using DataFrame functions and SparkSQL

In this section, we explore the datasets using functions both from dataframes as well as corresponding SQL queries using sparksql. Note the different ways to achieve the same task!

In [15]:
```
# Select and show basic data columns

df.select("name").show()
df.select(df["name"]).show()
spark.sql("SELECT name FROM people").show()
```

```
+-------+
|   name|
+-------+
|Michael|
|   Andy|
| Justin|
+-------+

+-------+
|   name|
+-------+
|Michael|
|   Andy|
| Justin|
+-------+

+-------+
|   name|
+-------+
|Michael|
|   Andy|
| Justin|
+-------+
```

In [16]:
```
# Perform basic filtering

df.filter(df["age"] > 21).show()
spark.sql("SELECT age, name FROM people WHERE age > 21").show()
```

```
+---+----+
|age|name|
+---+----+
| 30|Andy|
+---+----+

+---+----+
|age|name|
+---+----+
| 30|Andy|
+---+----+
```

In [17]:
```
# Perfom basic aggregation of data

df.groupBy("age").count().show()
spark.sql("SELECT age, COUNT(age) as count FROM people GROUP BY age").sho
```

```
+----+-----+
| age|count|
+----+-----+
|  19|    1|
|null|    1|
|  30|    1|
+----+-----+

[Stage 30:=====================================>                 (51 + 10)
/ 75]
```

```
+----+-----+
| age|count|
+----+-----+
|  19|    1|
|null|    0|
|  30|    1|
+----+-----+
```

Question 1 - RDDs

Create an RDD with integers from 1-50. Apply a transformation to multiply every number by 2, resulting in an RDD that contains the first 50 even numbers.

In [18]:
```
# starter code
# numbers = range(1, 50)
# numbers_RDD = ...
# even_numbers_RDD = numbers_RDD.map(lambda x: ..)
```

In [20]:
```
# Code block for learners to answer
numbers = range(1, 50)
numbers_RDD = sc.parallelize(numbers)
even_numbers_RDD = numbers_RDD.map(lambda x: x * 2)
print( even_numbers_RDD.collect())
```

```
[2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32, 34, 36, 38, 4
0, 42, 44, 46, 48, 50, 52, 54, 56, 58, 60, 62, 64, 66, 68, 70, 72, 74, 76,
78, 80, 82, 84, 86, 88, 90, 92, 94, 96, 98]
```

Double-click **here** for the solution.

Question 2 - DataFrames and SparkSQL

Similar to the `people.json` file, now read the `people2.json` file into the notebook, load it into a dataframe and apply SQL operations to determine the average age in our people2 file.

In [ ]:
```
# starter code
# !curl https://cf-courses-data.s3.us.cloud-object-storage.appdomain.clou
# df = spark.read...
# df.createTempView..
# spark.sql("SELECT ...")
```

In [21]:
```
# Code block for learners to answer
!curl https://cf-courses-data.s3.us.cloud-object-storage.appdomain.cloud/
df = spark.read.json("people2.json").cache()
df.createTempView("people2")
result = spark.sql("SELECT AVG(age) from people2")
result.show()
```

```
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Cu
rrent
                                 Dload  Upload   Total   Spent    Left  Sp
eed
100   136  100   136    0     0   1478      0 --:--:-- --:--:-- --:--:--
1462
+--------+
|avg(age)|
+--------+
|    24.8|
+--------+
```

Double-click **here** for a hint.

Double-click **here** for the solution.

Question 3 - SparkSession

Close the SparkSession we created for this notebook

In [22]:
```
# Code block for learners to answer
spark.stop()
```

Double-click **here** for the solution.

Authors

Karthik Muthuraman

Other Contributors

Jerome Nilmeier

© IBM Corporation. All rights reserved.