CS 649 Big Data: Tools and Methods
Spring Semester, 2021

Doc 14 Spark Intro
Mar 9, 2021

Spark

Created at UC Berkeley's AMPLab

2009 Project started
2014 May - 1.0
2016 July - 2.0.2
2017 July - 2.2.0
2020 June - 3.0
2021 March - 3.1.1

Programming interface for
- Java, Python, Scala, R

Interactive shell for
- Python, Scala, R (experimental)

Runs on
- Linux, Mac, Windows

Cluster manager
- Native Spark cluster
- Hadoop YARN
- Apache Mesos

File System
- HDFS
- MapR File System
- Cassandra
- OpenStack Swift
- S3

Pseudo-Distributed Mode
- Single machine
- Uses local file system

Time Line

1991 - Java project started

1995 - Java 1.0 released, Design Patterns book published

2000 - Java 3

2001 - Scala project started

2002 - Nutch started

2004 - Google MapReduce paper

Scala version 1 released

2005 - F# released

2006 - Hadoop split from Nutch

Scala version 2 released

2007 - Clojure released

2009 - Spark project started

2012 - Hadoop 1.0

2014 - Spark 1.0

Hadoop Word Count - Map

```
public class WordCount {

    public static class Map extends Mapper<LongWritable, Text, Text, IntWritable> {
        private final static IntWritable one = new IntWritable(1);
        private Text word = new Text();

        public void map(LongWritable key, Text value, Context context) throws IOException,
                                                                    InterruptedException {
            String line = value.toString();
            StringTokenizer tokenizer = new StringTokenizer(line);
            while (tokenizer.hasMoreTokens()) {
                word.set(tokenizer.nextToken());
                context.write(word, one);
            }
        }
    }
```

Hadoop Word Count - Reduce

```
public static class Reduce extends Reducer<Text, IntWritable, Text, IntWritable> {

    public void reduce(Text key, Iterable<IntWritable> values, Context context)
                                                    throws IOException, InterruptedException {
        int sum = 0;
        for (IntWritable val : values) {
            sum += val.get();
        }
        context.write(key, new IntWritable(sum));
    }
}
```

Hadoop Word Count - Main

```
public static void main(String[] args) throws Exception {
    Configuration conf = new Configuration();

    Job job = new Job(conf, "wordcount");

    job.setOutputKeyClass(Text.class);
    job.setOutputValueClass(IntWritable.class);

    job.setMapperClass(Map.class);
    job.setReducerClass(Reduce.class);

    job.setInputFormatClass(TextInputFormat.class);
    job.setOutputFormatClass(TextOutputFormat.class);

    FileInputFormat.addInputPath(job, new Path(args[0]));
    FileOutputFormat.setOutputPath(job, new Path(args[1]));

    job.waitForCompletion(true);
}
```

Spark Word Count - Python

```
from __future__ import print_function
import sys
from pyspark.sql import SparkSession

if __name__ == "__main__":
    if len(sys.argv) != 2:
        print("Usage: wordcount <file>", file=sys.stderr)
        sys.exit(-1)

spark = SparkSession.builder.appName("PythonWordCount").getOrCreate()

lines = spark.read.text(sys.argv[1]).rdd.map(lambda r: r[0])
counts = lines.flatMap(lambda x: x.split(' ')) \
                  .map(lambda x: (x, 1)) \
                  .reduceByKey(lambda a, b: a + b)
counts.saveAsTextFile("hdfs://...")

spark.stop()
```

Spark Word Count - Java

```
public final class JavaWordCount {
    private static final Pattern SPACE = Pattern.compile(" ");
    public static void main(String[] args) throws Exception {

    if (args.length < 1) {
      System.err.println("Usage: JavaWordCount <file>");
      System.exit(1);
    }

    SparkSession spark = SparkSession.builder().appName("JavaWordCount").getOrCreate();

    JavaRDD<String> lines = spark.read().textFile(args[0]).javaRDD();
    JavaRDD<String> words = lines.flatMap(s -> Arrays.asList(SPACE.split(s)).iterator());
    JavaPairRDD<String, Integer> ones = words.mapToPair(s -> new Tuple2<>(s, 1));
    JavaPairRDD<String, Integer> counts = ones.reduceByKey((i1, i2) -> i1 + i2);

    counts.saveAsTextFile("hdfs://...");
    spark.stop();
  }
}
```

Scala

```
object SparkWordCount {
  def main(args: Array[String]) {
    val spark = SparkSession.builder.appName("Spark Pi").getOrCreate()
    val textFile = sc.textFile("hdfs://...")
    val counts = textFile.flatMap(line => line.split(" "))
                         .map(word => (word, 1))
                        .reduceByKey(_ + _)
    counts.saveAsTextFile("hdfs://...")
    spark.stop()
  }
}
```

Python vs Scala on Spark

Scala is faster that Python

- But that is not so important here
- Most of the computation on Spark is done in Spark

Using Python with Spark

- Python data has to be
  - Converted between Python format and Scala/Java format
  - Sent between Python process and JVM

Installing PySpark using Anaconda

pip install pyspark

make sure using Anaconda pip

Need Java 8 installed

Sample Program

```
import pyspark
import random
sc = pyspark.SparkContext(appName="Pi")
num_samples = 100000
def inside(p):
  x, y = random.random(), random.random()
  return x*x + y*y < 1
count = sc.parallelize(range(0, num_samples)).filter(inside).count()
pi = 4 * count / num_samples
print(pi)
sc.stop()
```

Installing Spark - Java/Scala/Python

http://spark.apache.org/downloads.html

Download Apache Spark™

1. Choose a Spark release: 3.1.1 (Mar 02 2021)
2. Choose a package type: Pre-built for Apache Hadoop 3.2 and later
3. Download Spark: spark-3.1.1-bin-hadoop3.2.tgz
4. Verify this release using the 3.1.1 signatures, checksums and project release KEYS.

Read the Readme.md file

Helps to set your path

https://spark.apache.org/docs/latest/

AWS is using
Spark 3.0.1
Hadoop 3.2.1

Running PySpark from Command Line

```
Al pro 12->pyspark
Python 3.7.4 (default, Aug 13 2019, 15:17:50)
[Clang 4.0.1 (tags/RELEASE_401/final)] :: Anaconda, Inc. on darwin
Type "help", "copyright", "credits" or "license" for more information.
21/03/08 15:13:12 WARN NativeCodeLoader: Unable to load native-hadoop library for your platform
using builtin-java classes where applicable
Using Spark's default log4j profile: org/apache/spark/log4j-defaults.properties
Setting default log level to "WARN".
To adjust logging level use sc.setLogLevel(newLevel). For SparkR, use setLogLevel(newLevel).
Welcome to
      ____              __
     / __/__  ___ _____/ /__
    _\ \/ _ \/ _ `/ __/  '_/
   /__ / .__/\_,_/_/ /_/\_\   version 3.1.1
      /_/

Using Python version 3.7.4 (default, Aug 13 2019 15:17:50)
Spark context Web UI available at http://192.168.1.16:4040
Spark context available as 'sc' (master = local[*], app id = local-1615245193879).
SparkSession available as 'spark'.
```

Having PySpark Run in Jupyter Notebook

Don't need PySpark installed via pip, just regular install

In your shell

export SPARK_HOME path to your spark
export PYSPARK_DRIVER_PYTHON jupyter
export PYSPARK_DRIVER_PYTHON_OPTS 'notebook'

In path

$SPARK_HOME/bin $SPARK_HOME/sbin

Having PySpark Run in Jupyter Notebook

pyspark

Then will start pyspark in jupyter notebook

In jupyter notebook "spark" will be a SparkSession

Standard Warning

Unable to load native-hadoop library for your platform... using builtin-java classes where applicable

Major Parts of Spark

Spark Core
  Resilient Distributed Dataset (RDD)

Spark SQL
  SQL, csv, json
  Dataframe

Spark Streaming
  Near real-time response

MLib Machine Learning Library
  Statistics, regression, clustering, dimension reduction, feature extraction
  Optimization

GraphX

Original (Old) Spark

Newer Version Spark

Spark

Ecosystem of packages, libraries and systems on top of Spark Core

Unstructured API

Resilient Distributed Datasets (RDD)
Accumulators
Broadcast variables

Old Spark

Structured API

DataFrames
Datasets
Spark SQL

Newer, faster, higher level
Preferred over Unstructured

Basic Architecture

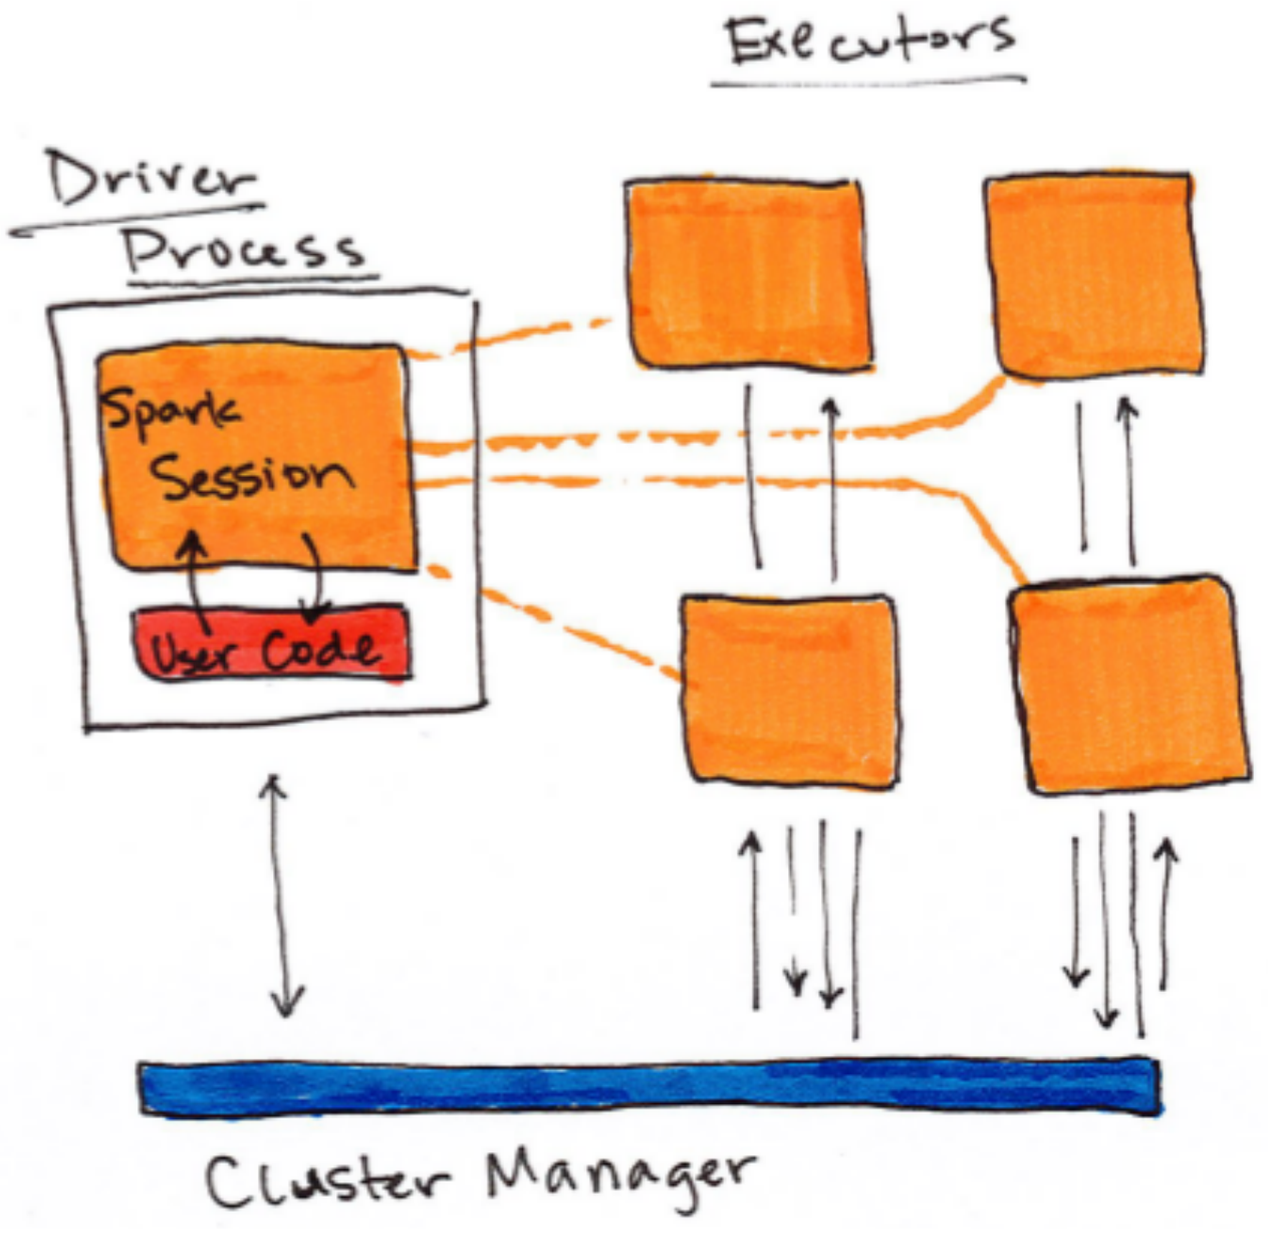

Local Mode

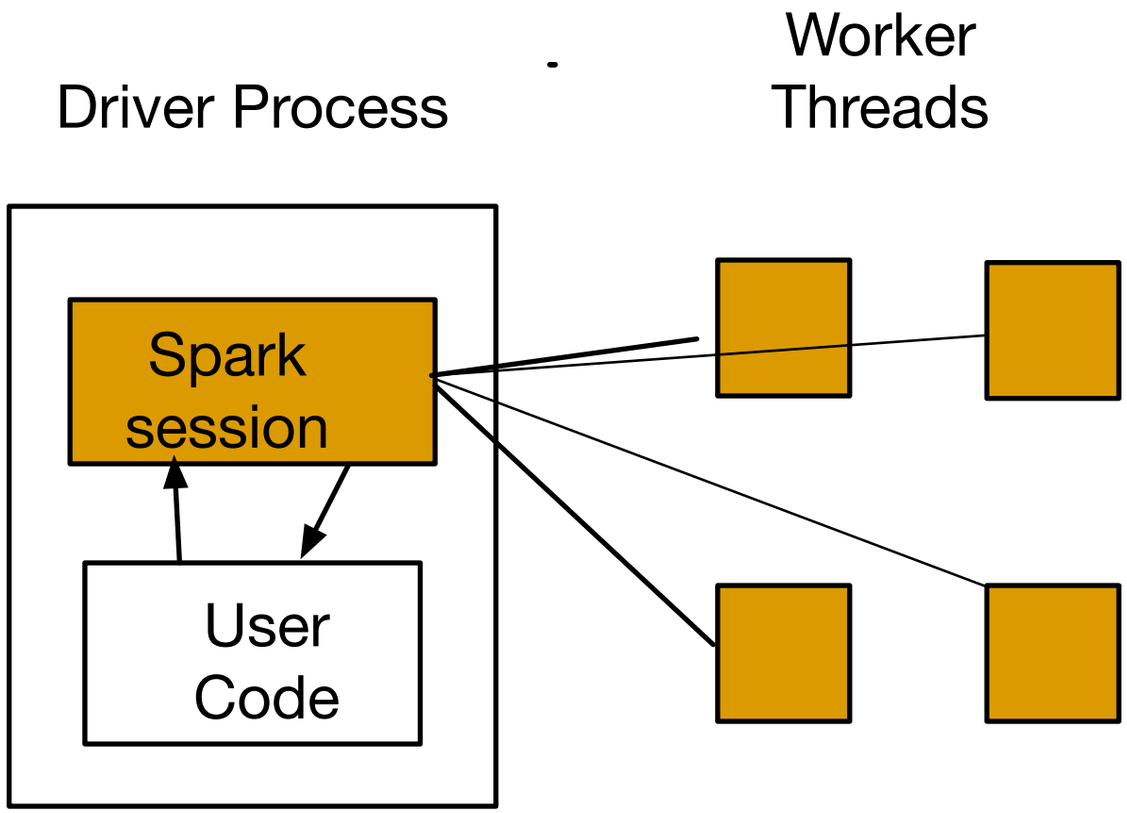

We will start using local mode

Use local mode to

Develop Spark code

SparkContext

Entry point for Unstructured API (Old Spark)

- Connection to Spark cluster
- Runs on master node
- Used to create RDDs, accumulators, broadcast variables
- Only one SparkContext per JVM
- stop() the current SparkContext before starting another

SparkContext org.apache.spark.SparkContext
  Scala version

JavaSparkContext org.apache.spark.api.java.JavaSparkContext
  Java version

pyspark.SparkContext
  Python version

SparkSession

org.apache.spark.sql.SparkSession
pyspark.sql.SparkSession

Contains a SparkContext

Entry point to use Dataset & DataFrame

Connection to Spark cluster
Runs on master node

Major Data Structures

Resilient Distributed Datasets (RDDs)

Fault-tolerant collection of elements that can be operated on in parallel

Dataset & Dataframes

Fault-tolerant collection of elements that can be operated on in parallel

Rows & Columns

JSON, csv, SQL tables

Part of SparkSQL

Use RDDs as underlying data structure

Partitions

RDD & Dataset

- Divided into partitions
- Each partition is on different machine

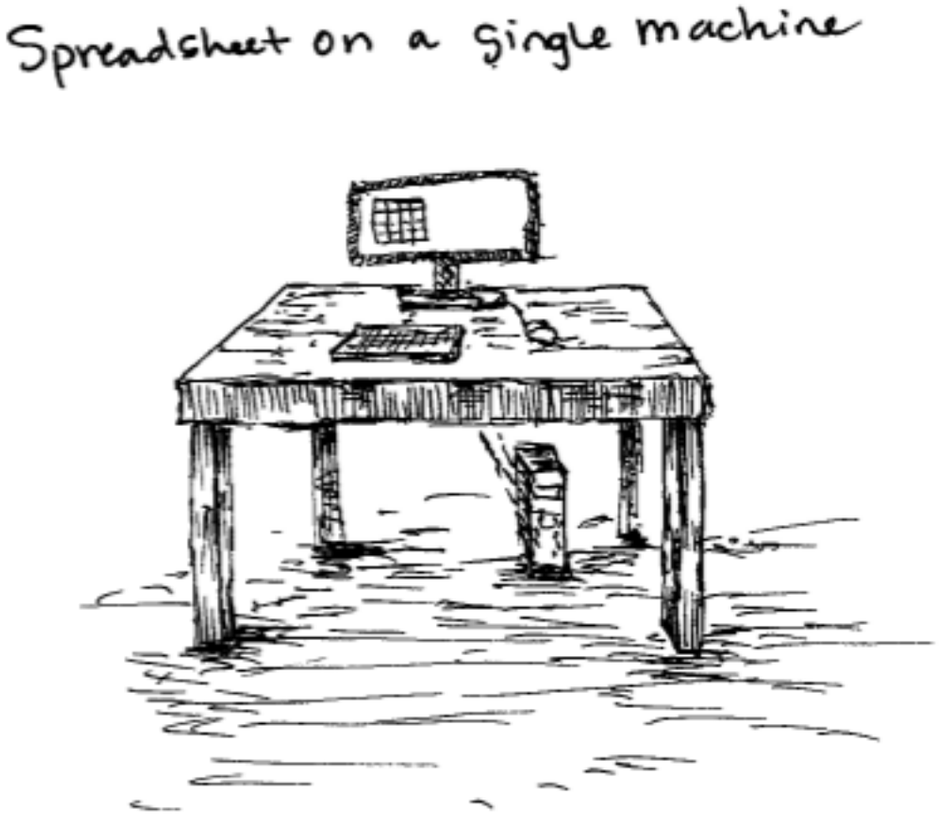

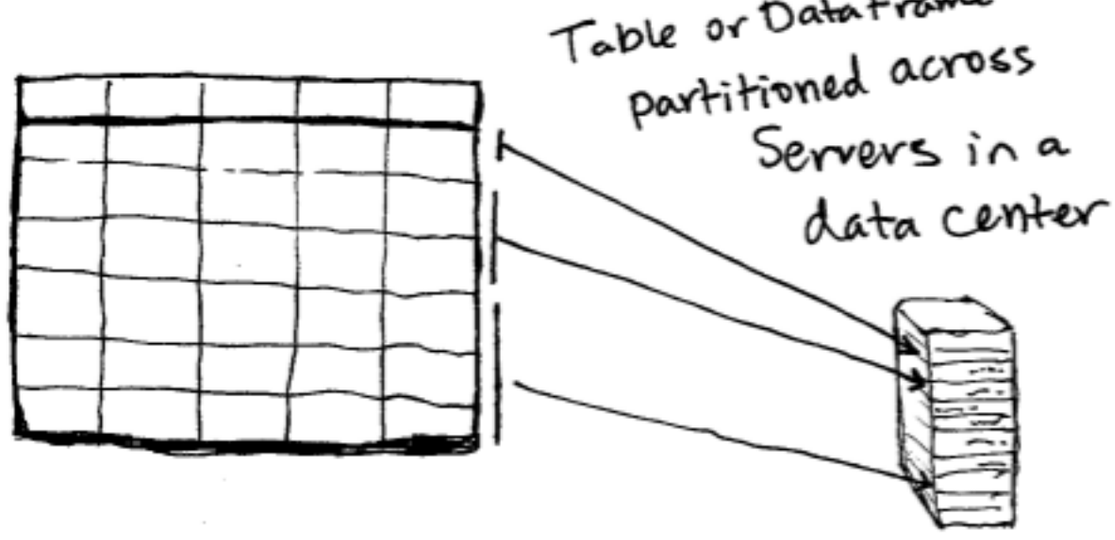

Resilient & Distributed

Distributed

- Partitions on different machines

Resilient

- Each partition can be replicated on multiple machines
- Data structure knows how to reproduce operations

Basic Operations

RDDs, Dataframes, Datasets
- Immutable

Transformations
- Create new dataset (RDD) from existing one
- Lazy
- Only done when needed by an action
- Examples
  - map, filter, sample, union, distinct, groupByKey, repartition

Actions
- Return results to driver program
- Examples
  - reduce, collect, count, first, take

Actions & Transformations on DataSet

View the Spark Scala API

org.apache.spark.sql.Dataset

Actions

```
def collect(): Array[T]
```
Returns an array that contains all rows in this Dataset.

```
def collectAsList(): List[T]
```
Returns a Java list that contains all rows in this Dataset.

```
def count(): Long
```
Returns the number of rows in the Dataset.

```
def describe(cols: String*): DataFrame
```
Computes statistics for numeric and string columns, including count, mean, stddev, min, and max.

```
def first(): T
```
Returns the first row.

Typed transformations

```
def alias(alias: Symbol): Dataset[T]
```
(Scala-specific) Returns a new Dataset with an alias set.

```
def alias(alias: String): Dataset[T]
```
Returns a new Dataset with an alias set.

```
def as(alias: Symbol): Dataset[T]
```
(Scala-specific) Returns a new Dataset with an alias set.

DataFrame, DataSet & RDD

What are they

What is the difference

When do use which one

Which languages can use them

DataFrame

Table with rows and Columns

Row

org.apache.spark.sql.Row

Schema

Column labels

Column types

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

Partitioner

Distributes DataFrame among cluster

Plan

Series of transformations to perform on DataFrame

Langauges

Scala, Java, JVM languages, Python, R

Optimized

Spark Catalyst Optimizer

Python DataFrame & Spark DataFrame

They are different

Need Apache Arrow to convert between them

DataSet

Same as DataFrame except for Rows

Programmer defines Row class
- Scala Cas Class
- Java Bean

Difference from DataFrame
- Compiler knows column names and column types in DataSet
- Compile time error checking
- Better data layout

Languages
- Scala, Java, JVM languages

RDD

Table

No information about types

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

No compile time or runtime type checking

Shares same basic operations as DataFrames & DataSets

Far fewer optimizations

No Catalyst Optimizer

No space optimization

Example - Same data

RDD	33.3 MB
DataFrame	7.3 MB

Languages

Java, Scala

Python, R - not recommended

Spark Types

Java Types are not space efficient

“abcd” - 48 bytes

Spark has its own types

Special memory representation of each type

Space efficient

Cache aware

Spark	Scala	Python	Python API
ByteType	Byte	int or long	ByteType()
ShortType	Short	int or long	ShortType()
IntegerType	Int	int or long	IntergerType()
LongType	Long	int or long	LongType()

Structured verses Unstructured

Structured = DataSet, DataFrame

Unstructured = RDD

Typed verses Untyped

Typed = DataSet

Untyped = DataFrame

Some Sample Data

JSON flight Data 2015

United States Bureau of Transportation statistics

The Definitive Guide, Zaharia & Chambers, O'Reilly Media, Inc, 2017-10-??

2015-summary.json

{"ORIGIN_COUNTRY_NAME":"Romania","DEST_COUNTRY_NAME":"United States","count":15}
{"ORIGIN_COUNTRY_NAME":"Croatia","DEST_COUNTRY_NAME":"United States","count":1}
{"ORIGIN_COUNTRY_NAME":"Ireland","DEST_COUNTRY_NAME":"United States","count":344}
{"ORIGIN_COUNTRY_NAME":"United States","DEST_COUNTRY_NAME":"Egypt","count":15}
{"ORIGIN_COUNTRY_NAME":"India","DEST_COUNTRY_NAME":"United States","count":62}
{"ORIGIN_COUNTRY_NAME":"Singapore","DEST_COUNTRY_NAME":"United States","count":1}
{"ORIGIN_COUNTRY_NAME":"Grenada","DEST_COUNTRY_NAME":"United States","count":62}
{"ORIGIN_COUNTRY_NAME":"United States","DEST_COUNTRY_NAME":"Costa Rica","count":588}
{"ORIGIN_COUNTRY_NAME":"United States","DEST_COUNTRY_NAME":"Senegal","count":40}
{"ORIGIN_COUNTRY_NAME":"United States","DEST_COUNTRY_NAME":"Moldova","count":1}

```
jsonFlightFile =
    "/Users/whitney/Courses/696/Fall17/SparkBookData/flight-data/json/2015-summary.json"
```

flightData2015 = spark.read.json(jsonFlightFile)
flightData2015

```
DataFrame[DEST_COUNTRY_NAME: string, ORIGIN_COUNTRY_NAME: string, count: bigint]
```

flightData2015.take(2)

```
[Row(DEST_COUNTRY_NAME='United States', ORIGIN_COUNTRY_NAME='Romania', count=15),
 Row(DEST_COUNTRY_NAME='United States', ORIGIN_COUNTRY_NAME='Croatia', count=1)]
```

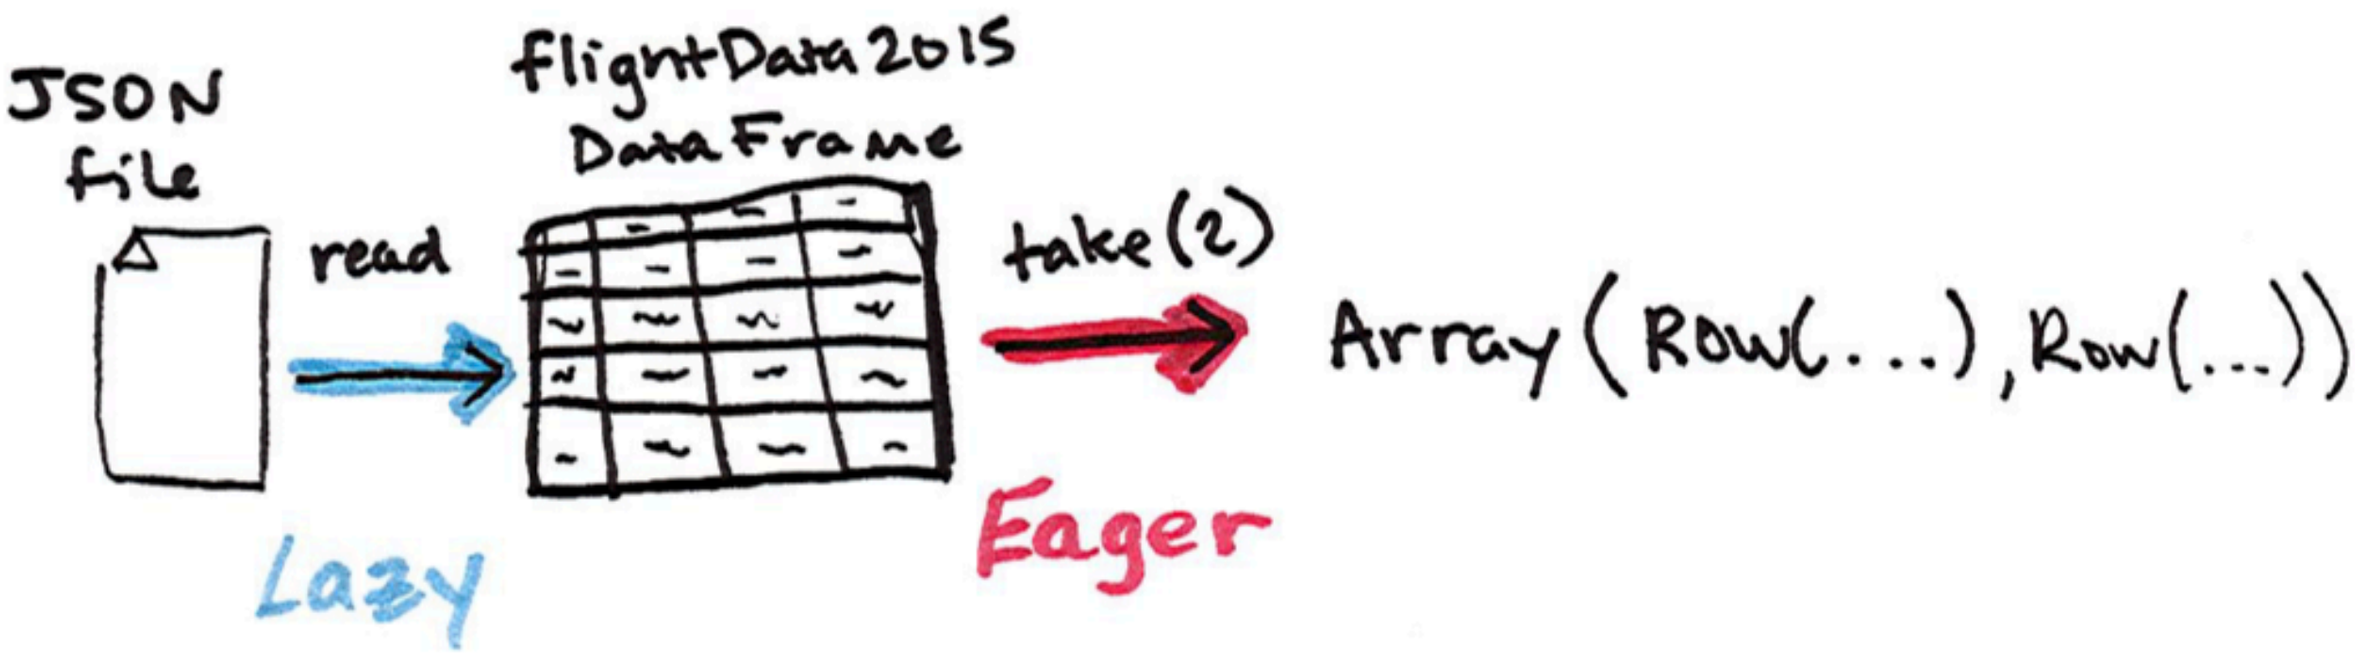

Explain - Spark Plan

flightData2015.explain()

```
== Physical Plan ==
*(1) FileScan json [DEST_COUNTRY_NAME#6,ORIGIN_COUNTRY_NAME#7,count#8L]
  Batched: false, Format: JSON,
  Location: InMemoryFileIndex[file:/Users/whitney/Courses/696/Fall17/SparkBookData/
flight-data/json/2015-summ...,
  PartitionFilters: [],
  PushedFilters: [],
  ReadSchema: struct<DEST_COUNTRY_NAME:string,ORIGIN_COUNTRY_NAME:string,count:bigint>
```

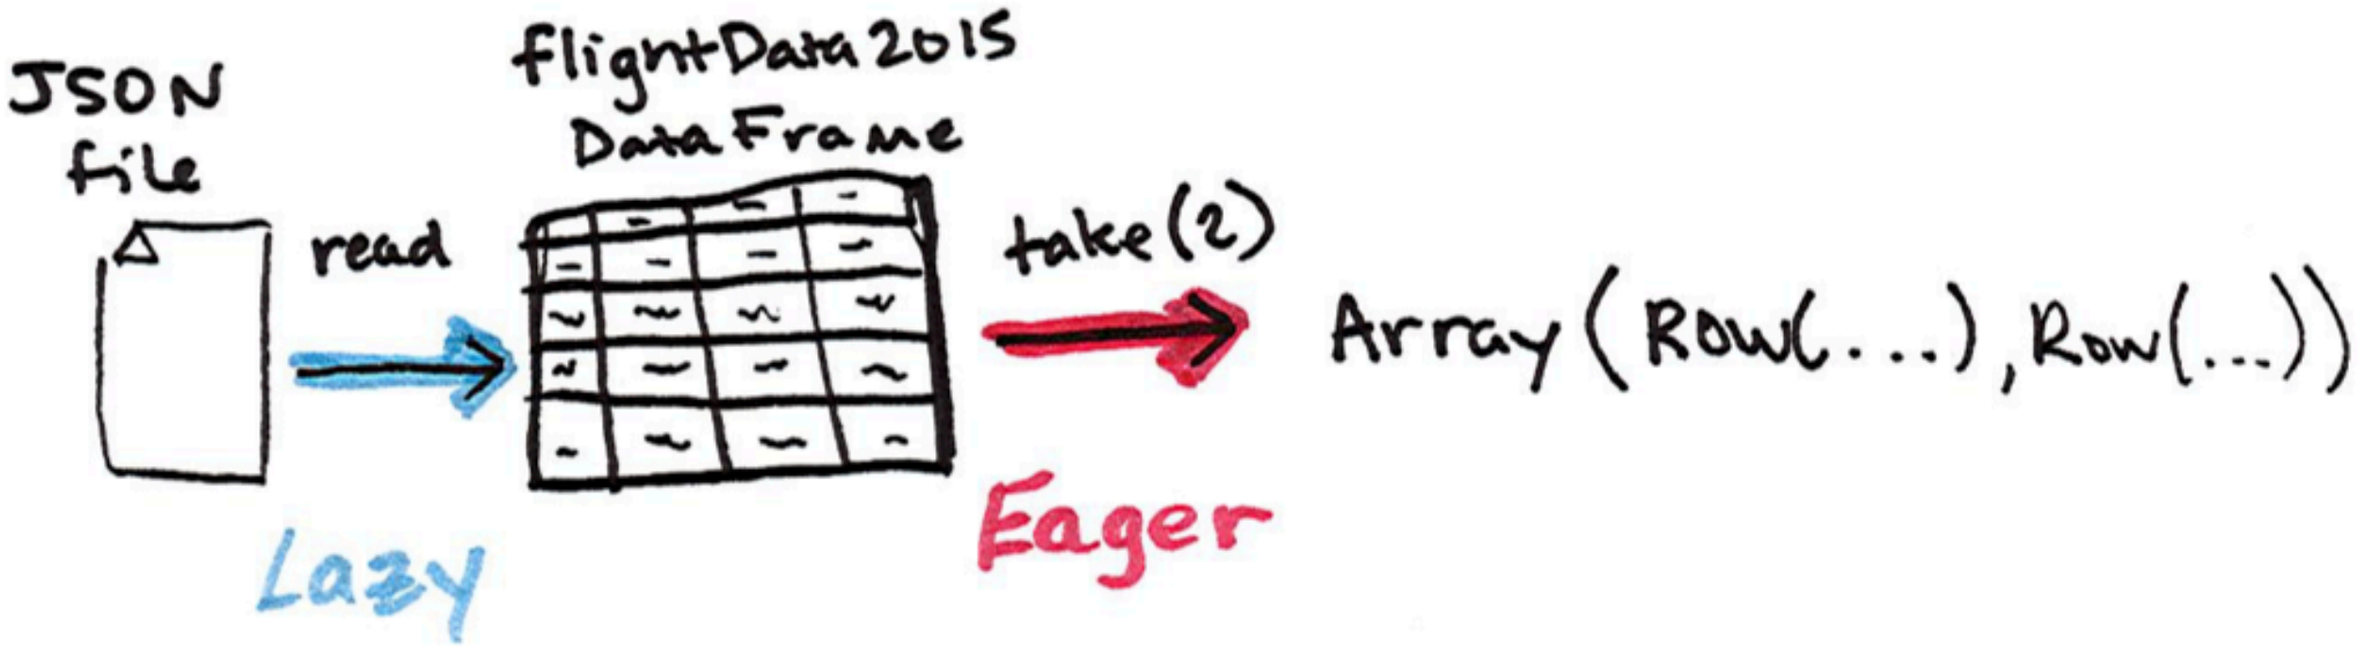

sortedFlightData2015 = flightData2015.sort("count")

sortedFlightData2015: org.apache.spark.sql.Dataset[org.apache.spark.sql.Row] = [DEST_COUNTRY_NAME: string, ORIGIN_COUNTRY_NAME: string ... 1 more field]

sortedFlightData2015.take(2)

```
[Row(DEST_COUNTRY_NAME='United States', ORIGIN_COUNTRY_NAME='Singapore', count=1),
 Row(DEST_COUNTRY_NAME='Moldova', ORIGIN_COUNTRY_NAME='United States', count=1)]
```

sortedFlightData2015.show(3)

```
+-----------------+-------------------+-----+
|DEST_COUNTRY_NAME|ORIGIN_COUNTRY_NAME|count|
+-----------------+-------------------+-----+
|          Moldova|      United States|    1|
|    United States|          Singapore|    1|
|    United States|            Croatia|    1|
+-----------------+-------------------+-----+
```

sortedFlightData2015.explain()

```
== Physical Plan ==
*(2) Sort [count#23L ASC NULLS FIRST], true, 0
+- Exchange rangepartitioning(count#23L ASC NULLS FIRST, 200)
   +- *(1) FileScan json [DEST_COUNTRY_NAME#21,ORIGIN_COUNTRY_NAME#22,count#23
            Batched: false, Format: JSON,
            Location: InMemoryFileIndex[file:/Users/whitney/Courses/696/Fall17/SparkBookData/fli
data/json/2015-summ...,
            PartitionFilters: [],
            PushedFilters: [],
      ReadSchema:
          struct<DEST_COUNTRY_NAME:string,ORIGIN_COUNTRY_NAME:string,count:b
```

Conceptual Plan

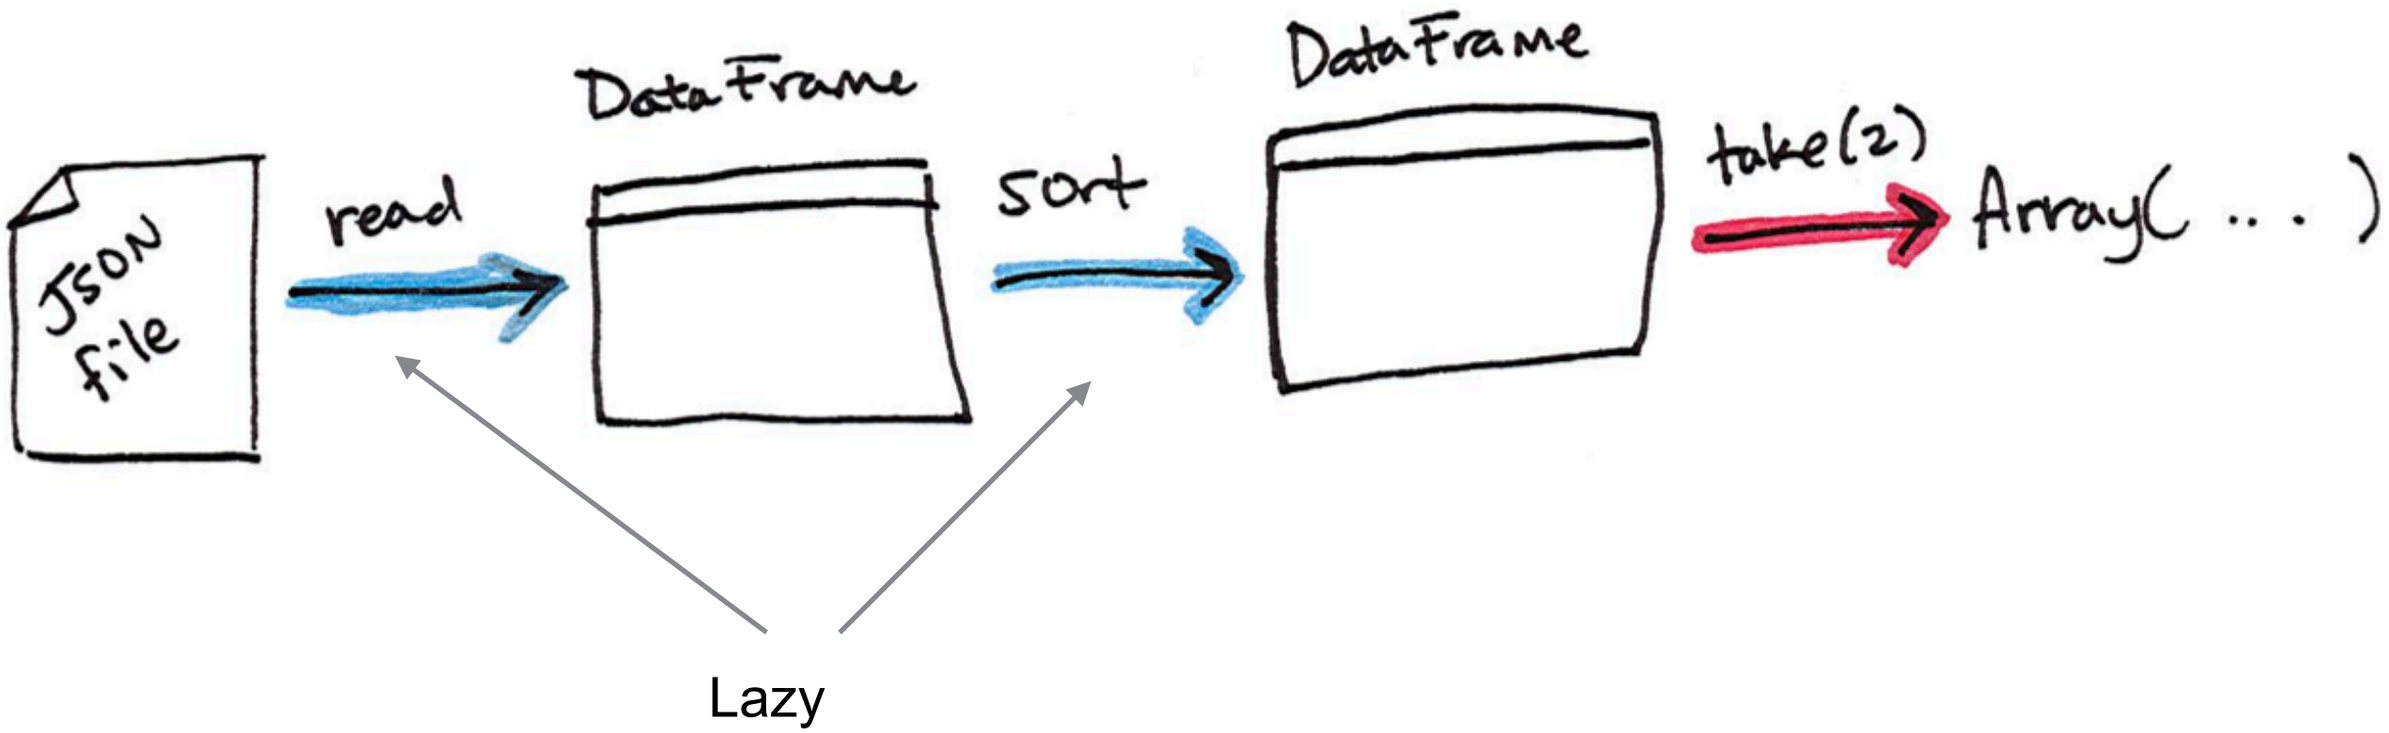

Spark stores the plan in case it needs to recompute the result

Schema

sortedFlightData2015.schema

```
StructType(
   List(
      StructField(DEST_COUNTRY_NAME,StringType,true),
      StructField(ORIGIN_COUNTRY_NAME,StringType,true),
      StructField(count,LongType,true)))
```

StructField

name The name of this field.

dataType The data type of this field.

nullable Indicates if values of this field can be null values.

Reading CSV

people.csv

```
name,age
Andy,30
Justin,19
Michael,
```

```
peopleFile =
    "/Users/whitney/Courses/696/Fall17/SparkExamples/people.csv"

reader = spark.read
reader: org.apache.spark.sql.DataFrameReader = org.apache.spark.sql.DataFrameReader@288

reader.option("header",True)
reader.option("inferSchema",True)

df = reader.csv(peopleFile)
df.show()
+-------+----+
|   name| age|
+-------+----+
|   Andy|  30|
| Justin|  19|
|Michael|null|
+-------+----+
```

Reading CSV

df.printSchema()

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

Some CSV options

encoding
sep (erator)
header
inferSchema
ignoreLeadingWhiteSpace
nullValue
dateFormat
timeStampFormat

mode
  PERMISSIVE - sets record field on corrupt record
  DROPMALFORMED - ignores whole corrupt records
  FAILFAST - throw exception on corrupt record

We can select columns

people.csv

```
name,age
Andy,30
Justin,19
Michael,
```

```
names = df.select("name")
names.show()
```

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

We can select columns

people.csv

```
name,age
Andy,30
Justin,19
Michael,
```

```
from pyspark.sql.functions import col

names = df.select(col("name"))
names.show()
+-------+
|   name|
+-------+
|   Andy|
| Justin|
|Michael|
+-------+
```

Column Operations

older = df.select(col("name"), col("age") + 1)
older.show()

```
+-------+---------+
|   name|(age + 1)|
+-------+---------+
|   Andy|       31|
| Justin|       20|
|Michael|     null|
+-------+---------+
```

older.printSchema()
root
 |-- name: string (nullable = true)
 |-- (age + 1): integer (nullable = true)

Does not work in Python
scala> **val older = df.select($"name", $"age" + 1)**

Column Operations - Java vs Scala

Java or Scala

df.select(col("name"), col("age").plus(1))

Scala Only

df.select($"name", $"age" + 1)

df.select(‘name, 'age + 1)

Python

df.select(col("name"), col("age") + 1)

val adult = older.filter(col("age") > 21)

scala> **adult.show**

```
+----+----------+
|name|(age + 1)|
+----+----------+
|Andy|        31|
+----+----------+
```

scala> **adult.explain**

```
== Physical Plan ==
*Project [name#104, (age#105 + 1) AS (age + 1)#123]
+- *Filter (isnotnull(age#105) && (age#105 > 21))
   +- *FileScan csv [name#104,age#105]
        Batched: false, Format: CSV,
        Location: InMemoryFileIndex[file:/Users/whitney/Courses/696/Fall17/SparkExamples/people.cs
        PartitionFilters: [],
        PushedFilters: [IsNotNull(age), GreaterThan(age,21)],
        ReadSchema: struct<name:string,age:int>
```

df.groupBy("age").count().show()

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

Saving DataFrames

df.write.format("json").save("people.json")

Produces a directory: people.json

Contents:

_SUCCESS (0 Byte file)

part-00000-71516d50-2bcc-4830-ad61-554d1c107f51-c000.json

```
{"name":"Andy","age":30}
{"name":"Justin","age":19}
{"name":"Michael"}
```

Formats

json, parquet, jdbc, orc, libsvm, csv, text

Using SQL

df.createOrReplaceTempView("people")

sqlExample = spark.sql("SELECT * FROM people")

sqlExample.show()

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

show

action

- No return value
- Only prints out value
- So can not use the result

What happens on cluster?

- Actions return value to master node
- But often run in batch mode

Collect - Returns a result

sqlExample

DataFrame[name: string, age: int]

data = sqlExample.collect()

data

[Row(name='Andy', age=30),
 Row(name='Justin', age=19),
 Row(name='Michael', age=None)]

data[0]

Row(name='Andy', age=30)

data[0][0]

'Andy'

Planning

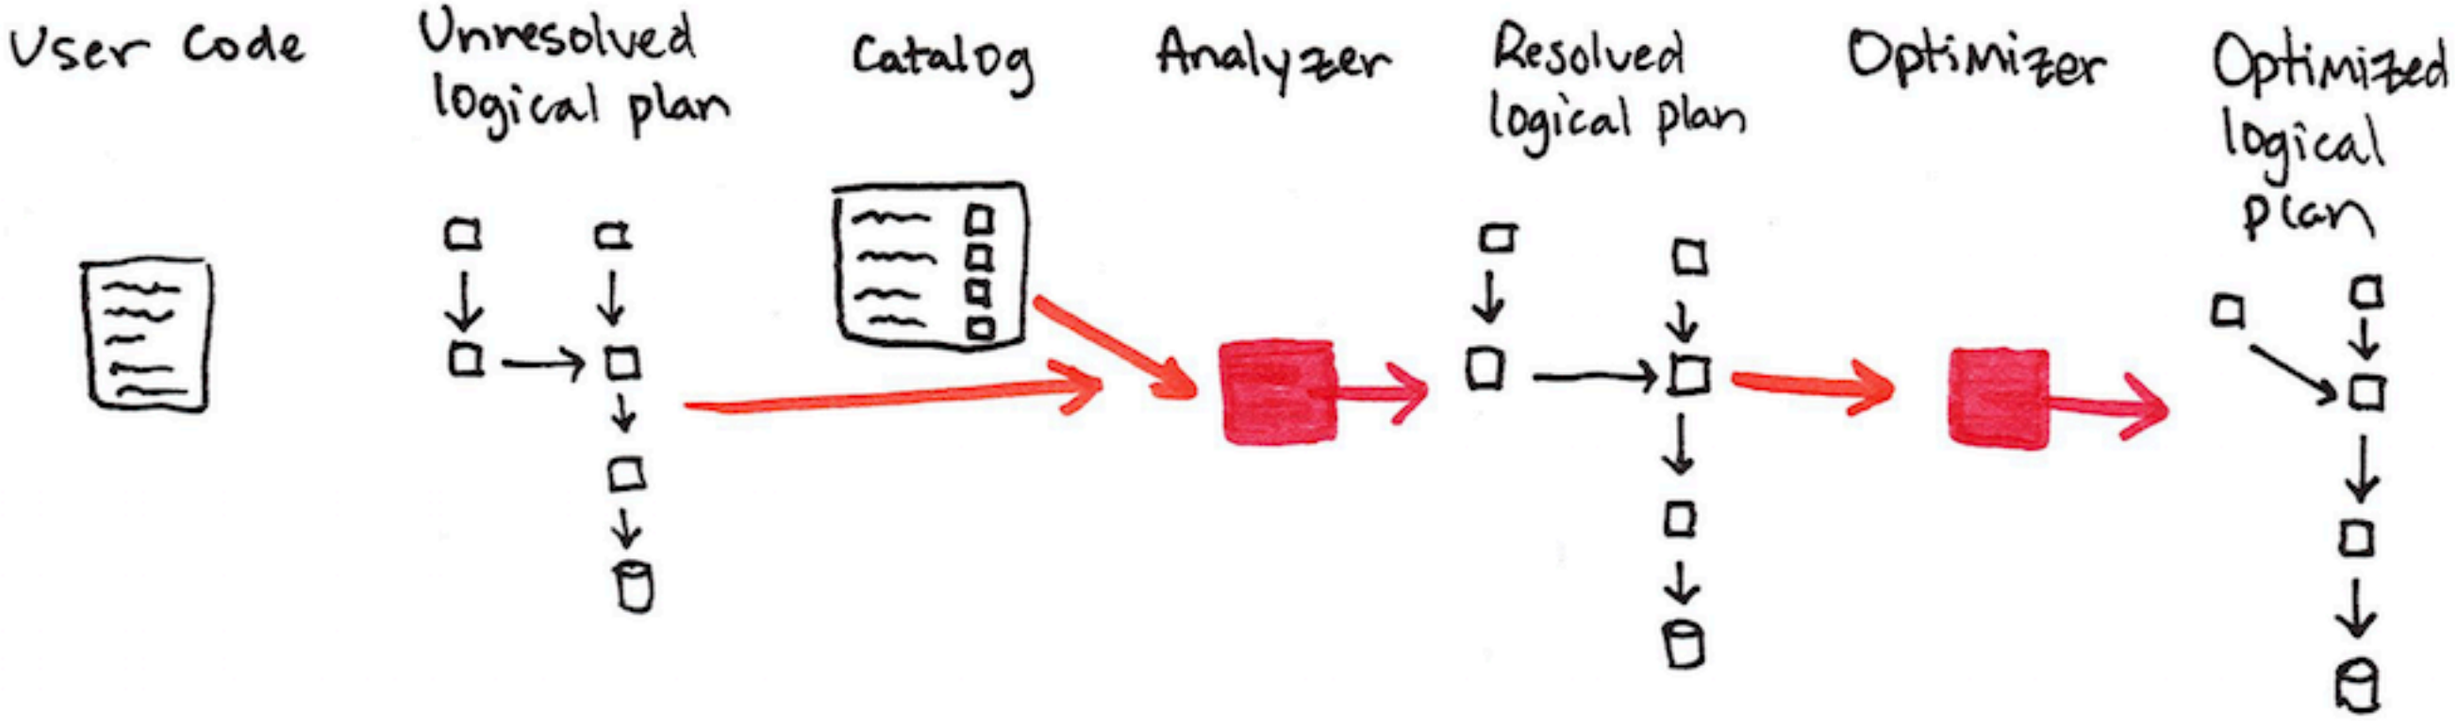
User Code
Unresolved logical plan
Catalog
Analyzer
Resolved logical plan
Optimizer
Optimized logical plan

Schema

famousPeople.json

```
{"first": "Roger", "last":"Rabbit","age":29}
{ "last":"Modi","age":67}
{"first": "Sachin", "last":"Tendulkar","age":44}
```

```
people = spark.read.format("json").load("famousPeople.json")
people.show()
```

```
+---+------+---------+
|age| first|     last|
+---+------+---------+
| 29| Roger|   Rabbit|
| 67|  null|     Modi|
| 44|Sachin|Tendulkar|
+---+------+---------+
```

Schema

{"first": "Roger", "last":"Rabbit","age":29}
{ "last":"Modi","age":67}
{"first": "Sachin", "last":"Tendulkar","age":44}

people = spark.read.format("json").load("famousPeople.json")
people.schema

StructType(
    StructField(age,**LongType**,true),
    StructField(first,StringType,true),
    StructField(last,StringType,true))

Schema

{"first": "Roger", "last":"Rabbit","age":29}
{ "last":"Modi","age":67}
{"first": "Sachin", "last":"Tendulkar","age":44}

```
from pyspark.sql.types import StructField, StructType, StringType, IntegerType

manualSchema = StructType([StructField("first", StringType(), True),
                           StructField("last", StringType(), False),
                           StructField("age", IntegerType(), False)
                          ])

people2 = spark.read.format("json").schema(manualSchema).load("famousPeople.json")
people2.show()
```

```
+------+---------+---+
| first|     last|age|
+------+---------+---+
| Roger|   Rabbit| 29|
|  null|     Modi| 67|
|Sachin|Tendulkar| 44|
+------+---------+---+
```

Schema

{"first": "Roger", "last":"Rabbit","age":29}
{ "last":"Modi","age":67}
{"first": "Sachin", "last":"Tendulkar","age":44}

```
manualSchema2 = StructType([StructField("first", StringType(), True),
                            StructField("last", StringType(), False),
                            StructField("address", StringType(), False)
                            ])

people3 = spark.read.format("json").schema(manualSchema2).load("famousPeople.json")
people3.show()
```

```
+------+---------+-------+
| first|     last|address|
+------+---------+-------+
| Roger|   Rabbit|   null|
|  null|     Modi|   null|
|Sachin|Tendulkar|   null|
+------+---------+-------+
```

Time

```
dwellFile = "/Users/whitney/Courses/696/Fall17/datasets/data/dwell-times.tsv"

reader = spark.read
reader.option("header", True)
reader.option("inferSchema", True)
reader.option("sep", "\t")

dwellDf = reader.csv(dwellFile)
dwellDf.printSchema()
dwellDf.show(3)
```

dwell-times.tsv

date	dwell-time
2015-01-01T00:03:43Z	74

```
root
 |-- date: timestamp (nullable = true)
 |-- dwell-time: integer (nullable = true)

+-------------------+----------+
|               date|dwell-time|
+-------------------+----------+
|2014-12-31 16:03:43|        74|
|2014-12-31 16:32:12|       109|
|2014-12-31 17:52:18|        88|
+-------------------+----------+
```