

CS 455: INTRODUCTION TO DISTRIBUTED SYSTEMS

[SPARK]

Spark: It's all about transformation and actions

Transformations

- Wrangle with the data
- Consume, and beget, an RDD
- Flock together ... to form daisy chains

But it is actions

- That trigger evaluations
- Providing them potency
- Revealing their expressive power

Shrideep Pallickara
Computer Science
Colorado State University

COMPUTER SCIENCE DEPARTMENT



1

Topics covered in this lecture

- Resilient Distributed Datasets
- Common Transformations and Actions



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.2

2

RESILIENT DISTRIBUTED DATASET [RDD]

COMPUTER SCIENCE DEPARTMENT



3

Resilient Distributed Dataset (RDD)

- RDD is an **immutable distributed collection** of objects
- Each RDD is split into *multiple partitions*
 - ▣ Maybe computed on different nodes in the cluster
- Can contain any type of Java, Scala, or Python objects
 - ▣ Including user-defined classes



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.4

4

Creation of RDDs

- ① Loading an external dataset
- ② Distributing a collection of objects via the driver program

```
>>> lines = sc.textFile("README.md")
```



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.5

5

Once created, RDDs offer two types of operations

□ Transformations

- Construct a new RDD from a previous one
- E.g.: Filtering data that matches a predicate

□ Actions

- Compute a result based on an RDD
- Return result to the driver program or save it in an external storage system (HDFS)



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.6

6

Some more about RDDs

- Although you can define new RDDs anytime
 - Spark computes them in a **lazy fashion**
 - When?
 - The first time they are used in an *action*
- Loading lazily allows transformations to be performed *before* the action



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.7

7

Lazy loading allows Spark to see the whole chain of transformations

- Allows it to **compute just the data needed** for the result
- Example:

```
lines = sc.textFile("README.md")
pythonLines= lines.filter(lambda line: "Python" in line)
```
- If Spark were to load and store all lines in the file, as soon as we wrote `lines=sc.textFile()`?
 - Would waste a lot of storage space, since we immediately filter out a lot of lines



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.8

8

RDD and actions

- RDDs are **recomputed** (by default) every time you run an action on them
- If you wanted to **reuse** an RDD?
 - Ask Spark to **persist** it using `RDD.persist()`
 - After computing it the first time, Spark will store RDD contents in memory (**partitioned** across cluster machines)
 - Persisted RDD is used in future actions



RDDs: memory residency and immutability implications

- Spark can keep an RDD loaded in-memory on the executor nodes throughout the life of a Spark application for faster access in **repeated computations**
- RDDs are immutable, so **transforming an RDD returns a new RDD** rather than the existing one
- Cross-cutting implications?
 - Lazy evaluation, in-memory storage, and immutability allows Spark to be easy-to-use, fault-tolerant, scalable, and efficient



Every Spark program and shell works as follows

- ① **Create** some input RDD from external data
- ② **Transform** them to define new RDDs using transformations like `filter()`
- ③ Ask Spark to **`persist()`** any intermediate RDDs that needs to be reused
- ④ **Launch actions** such as `count()`, etc. to kickoff a parallel computation
 - ▣ Computing is optimized and executed by Spark



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.11

11

A CLOSER LOOK AT RDD OPERATIONS

COMPUTER SCIENCE DEPARTMENT



COLORADO STATE UNIVERSITY

12

RDDs support two types of operations

- Transformations
 - Operations that **return a new RDD**. E.g.: `filter()`
- Actions
 - Operations that **return a result** to the driver program or write to storage
 - Kicks of a computation. E.g.: `count()`
- Distinguishing aspect?
 - Transformations return RDDs
 - Actions return *some other* data type



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.13

13

Transformations

- Many transformations are **element-wise**
 - Work on only one element at a time
- Some transformations are not element-wise
 - E.g.: We have a logfile, `log.txt`, with several messages, but we only want to select error messages

```
inputRDD = sc.textFile("log.txt")
errorsRDD = inputRDD.filter(lambda x:"error" in x)
```



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.14

14

In our previous example ...

- filter **does not mutate** inputRDD
 - Returns a pointer to an entirely new RDD
 - inputRDD can still be reused later in the program
- We could use inputRDD to search for lines with the word “warning”
 - While we are at it, we will use another transformation, union(), to print number of lines that contained either

```
errorsRDD = inputRDD.filter(lambda x: "error" in x)
warningsRDD = inputRDD.filter(lambda x: "warning" in x)
badlinesRDD = errorsRDD.union(warningsRDD)
```



In our previous example

- Note how union() is different from filter()
 - Operates on 2 RDDs instead of one
- Transformations can actually operate on **any number** of RDDs



RDD Lineage graphs

- As new RDDs are derived from each other using transformations, Spark *tracks dependencies*
 - **Lineage graph**
- Uses lineage graph to
 - Compute each RDD on demand
 - Recover lost data if part of persistent RDD is lost



COLORADO STATE UNIVERSITY

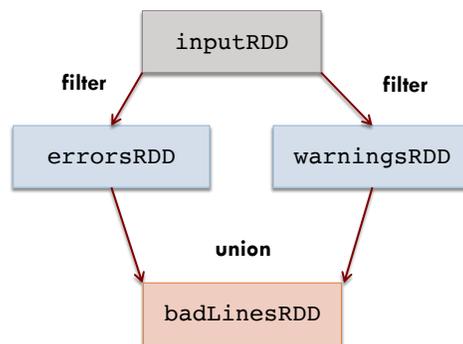
Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.17

17

RDD lineage graph for our example



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.18

18

Actions

- We can create RDDs from each other using transformations
- At some point, we need to actually **do something** with the dataset
 - ▣ Actions
- Forces *evaluations of the transformations* required for the RDD they were called on



Let's try to print information about badLinesRDD

```
print "Input had " + badLinesRDD.count() + "concerning lines"  
print "here are 10 examples:"  
for line in badLinesRDD.take(10)  
    print line
```



RDDs also have a `collect` to retrieve the entire RDD

- Useful if program filters RDD to a very small size and you want to deal locally
 - Your entire dataset must fit in memory on a single machine to use `collect()` on it
 - Should NOT be used on large datasets
- In most cases, RDDs **cannot be** `collect()`ed to the driver
 - Common to write data out to a distributed storage system ... HDFS or S3



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.21

21

Lazy Evaluation

- Transformations on RDDs are **lazily evaluated**
 - Spark will not begin to execute until it sees an action
- Uses this to **reduce the number of passes** it has to take over data by grouping operations together
- What does this mean?
 - When you call a transformation on an RDD (for e.g. `map`) the operation is not immediately performed
 - Spark internally records metadata that operation is requested



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.22

22

How you should think of RDDs

- Rather than thinking of it as containing specific data
 - Best to think of it as **containing instructions on how to compute the data** that we build through transformations
- Loading data into a RDD is lazily evaluated just as transformations are



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.23

23

COMMON TRANSFORMATIONS AND ACTIONS

COMPUTER SCIENCE DEPARTMENT



COLORADO STATE UNIVERSITY

24

Element-wise transformations: `filter()`

- Takes in a function and returns an RDD that only has elements that pass the `filter()` function



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

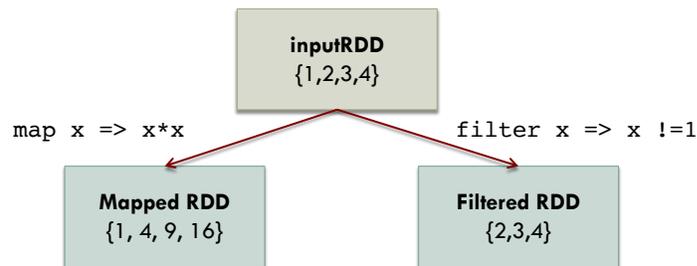
SPARK

L19.25

25

Element-wise transformations: `map()`

- Takes in a function and applies it to each element in the RDD
- Result of the function is the new value of each element in the resulting RDD



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.26

26

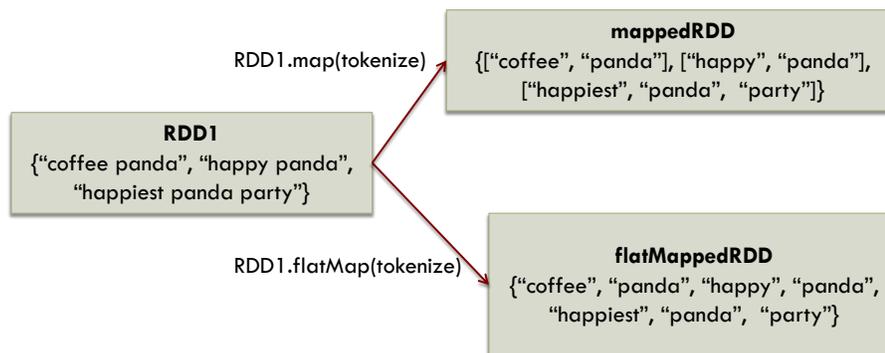
Things that can be done with map ()

- Fetch website associated with each URL in collection to just squaring numbers
- map ()'s return type does not have to be the same as its input type
- Multiple output elements for each input element?
 - Use flatMap ()

```
lines=sc.parallelize(["hello world", "hi"])  
words=lines.flatMap(lambda line: line.split(" "))  
words.first() # returns hello
```



Difference between map and flatMap



Pseudo set operations

- RDDs support many of the operations of mathematical sets such as union, intersection, etc.
 - Even when the RDDs themselves are not properly sets



COLORADO STATE UNIVERSITY

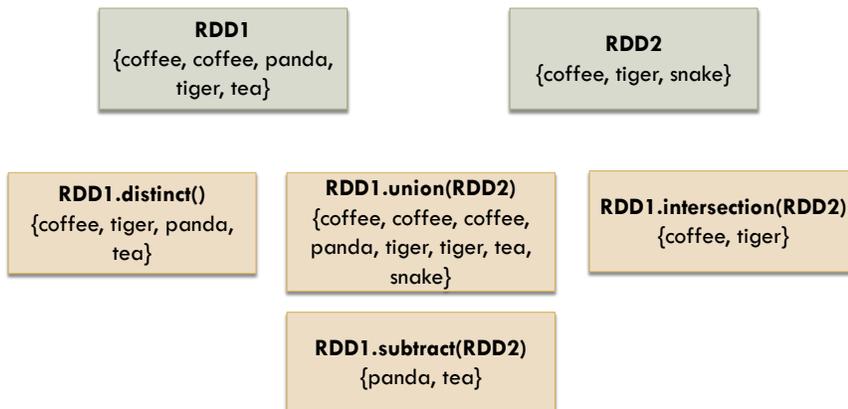
Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.29

29

Some simple set operations



COLORADO STATE UNIVERSITY

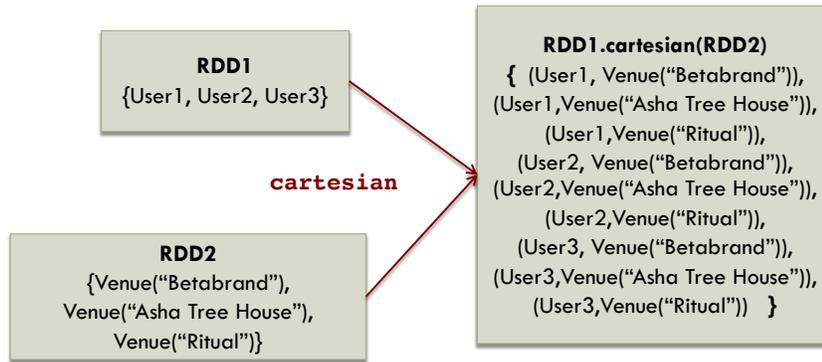
Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.30

30

Cartesian product between two RDDs



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.31

31

COMMON ACTIONS

COMPUTER SCIENCE DEPARTMENT



COLORADO STATE UNIVERSITY

32

Actions on Basic RDDs

- `reduce()`
 - Takes a function that operates on two elements in the RDD; returns an element of the same type
 - E.g. of such an operation? `+` sums the RDD
- ```
sum = rdd.reduce((x,y) => x + y)
```
- `fold()` takes a function with the same signature as `reduce()`, but also takes a “zero value” for initial call
    - “Zero value” is the **identity element** for initial call
    - E.g., 0 for `+`, 1 for `*`, empty list for concatenation



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.33

33

## Both `fold()` and `reduce()` require return type of same type as the RDD elements

- The `aggregate()` removes that constraint
  - For e.g. when computing a running average, maintain both the count so far and the number of elements



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.34

34

## EXAMPLES: BASIC ACTIONS ON RDDs

COMPUTER SCIENCE DEPARTMENT



COLORADO STATE UNIVERSITY

35

### Examples: Basic actions on RDDs

[1/7]

- Our RDD contains {1, 2, 3, 3}
- **collect()**
  - Return all elements from the RDD
  - Invocation: `rdd.collect()`
  - Result: {1, 2, 3, 3}



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.36

36

## Examples: Basic actions on RDDs

[2/7]

- Our RDD contains {1, 2, 3, 3}
- **count()**
  - Number of elements in the RDD
  - Invocation: `rdd.count()`
  - Result: 4



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.37

37

## Examples: Basic actions on RDDs

[3/7]

- Our RDD contains {1, 2, 3, 3}
- **countByValue()**
  - Number of times each element occurs in the RDD
  - Invocation: `rdd.countByValue()`
  - Result: `{ (1,1), (2,1), (3,2) }`



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.38

38

## Examples: Basic actions on RDDs

[4/7]

- Our RDD contains {1, 2, 3, 3}
- **take(num)**
  - Return num elements from the RDD
  - Invocation: `rdd.take(2)`
  - Result: {1, 2}



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.39

39

## Examples: Basic actions on RDDs

[5/7]

- Our RDD contains {1, 2, 3, 3}
- **reduce(func)**
  - Combine the elements of the RDD together in parallel
  - Invocation: `rdd.reduce( (x,y) => x + y )`
  - Result: 9



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.40

40

## Examples: Basic actions on RDDs

[6/7]

- Our RDD contains {1, 2, 3, 3}
- **aggregate(zeroValue)(seqOp, combOp)**
  - Similar to `reduce()` but used to return a different type
  - Invocation:
    - `rdd.aggregate ( (0,0)`  
`((x,y) => (x._1 + y, x._2 + 1),`  
`(x,y) => (x._1 + y._1, x._2 + y._2))`
  - Result: (9, 4)



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.41

41

## Examples: Basic actions on RDDs

[7/7]

- Our RDD contains {1, 2, 3, 3}
- **foreach(func)**
  - Apply the provided function to each element of the RDD
  - Invocation: `rdd.foreach(func)`
  - Result: Nothing



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.42

42

## PERSISTENCE (CACHING)

COMPUTER SCIENCE DEPARTMENT



43

### Why persistence?

- Spark RDDs are lazily evaluated, and we may sometimes wish to use the same RDD multiple times
  - Naively, Spark will **recompute RDD and all of its dependencies** each time we call an action on the RDD
    - Super expensive for iterative algorithms
- To avoid recomputing RDD multiple times?
  - Ask Spark to **persist** the data
  - The nodes that compute the RDD, store the partitions
  - E.g.: `result.persist(StorageLevel.DISK_ONLY)`



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.44

44

## Coping with failures

- If a node that has data persisted on it fails?
  - ▣ Spark recomputes lost partitions of data when needed
- Also, replicate data on multiple nodes
  - ▣ To handle node failures without slowdowns



## Persistence Levels for Spark

| Level               | Space Used | CPU time | In Memory | On disk | Comments                                                                                              |
|---------------------|------------|----------|-----------|---------|-------------------------------------------------------------------------------------------------------|
| MEMORY_ONLY         | High       | Low      | Y         | N       |                                                                                                       |
| MEMORY_ONLY_SER     | Low        | High     | Y         | N       |                                                                                                       |
| MEMORY_AND_DISK     | High       | Medium   | Some      | Some    | Spills to disk if there is too much data to fit in memory                                             |
| MEMORY_AND_DISK_SER | Low        | High     | Some      | Some    | Spills to disk if there is too much data to fit in memory. Stores serialized representation in memory |
| DISK_ONLY           | Low        | High     | N         | Y       |                                                                                                       |



## What if you attempt to cache too much data that does not fit in memory?

- Spark will **evict old partitions** using a Least Recently Used Cache policy
  - For memory only storage partitions, it will be recomputed the next time they are accessed
  - For memory\_and\_disk ones? Write them out to disk
- RDDs also come with a method, `unpersist()`
  - Manually remove data elements from the cache



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.47

47

## WORKING WITH KEY/VALUE PAIRS

COMPUTER SCIENCE DEPARTMENT



COLORADO STATE UNIVERSITY

48

## RDDs of key/value pairs

- Key/value RDDs are commonly used to perform aggregations
  - Might have to do ETL (Extract, Transform, and Load) to get data into key/value formats
- Advanced feature to control layout of pair RDDs across nodes
  - **Partitioning**



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.49

49

## RDDs containing key/value pairs

- Are called **pair RDDs**
- Useful *building block* in many programs
  - Expose operations that allow actions on each key in parallel or regroup data across network
  - `reduceByKey()` to aggregate data separately for each key
  - `join()` to merge two RDDs together by grouping elements of the same key



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.50

50



# PAIR RDDs

COMPUTER SCIENCE DEPARTMENT



COLORADO STATE UNIVERSITY

51

## Pair RDDs

- RDDs that contain **key/value pairs**
- Expose partitions that allow you to act on each key in parallel or regroup data across the network



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.52

52

## Creating Pair RDDs

- `pairs=lines.map(lambda x: (x.split(" ")[0], x))`
  - Creates a pairRDD using the first word as the key
- Java does not have a built-in tuple type
  - `scala.Tuple2` class
    - `new Tuple2(elem1, elem2)`



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.53

53

## The contents of this slide-set are based on the following references

- *Learning Spark: Lightning-Fast Big Data Analysis. 1st Edition. Holden Karau, Andy Konwinski, Patrick Wendell, and Matei Zaharia. O'Reilly. 2015. ISBN-13: 978-1449358624. [Chapters 1-4]*
- Karau, Holden; Warren, Rachel. *High Performance Spark: Best Practices for Scaling and Optimizing Apache Spark.* O'Reilly Media. 2017. ISBN-13: 978-1491943205. [Chapter 2]



COLORADO STATE UNIVERSITY

Professor: SHRIDEEP PALLICKARA  
COMPUTER SCIENCE DEPARTMENT

SPARK

L19.54

54