

Spark

Cluster computing (take 2)

J. Fernando Sánchez, Joaquín Salvachúa, Gabriel Huecas

2016

Universidad Politécnica de Madrid

Schedule for today

Recap

Other Collection Methods

Scala collections provide many other functional methods; for example, Google for “Scala Seq”

Method on Seq[T]	Explanation
<code>map(f: T => U): Seq[U]</code>	Pass each element through f
<code>flatMap(f: T => Seq[U]): Seq[U]</code>	One-to-many map
<code>filter(f: T => Boolean): Seq[T]</code>	Keep elements passing f
<code>exists(f: T => Boolean): Boolean</code>	True if one element passes
<code>forall(f: T => Boolean): Boolean</code>	True if all elements pass
<code>reduce(f: (T, T) => T): T</code>	Merge elements using f
<code>groupBy(f: T => K): Map[K, List[T]]</code>	Group elements by f(element)
<code>sortBy(f: T => K): Seq[T]</code>	Sort elements by f(element)
. . .	

Spark specific methods

Some examples:

- collect: to apply all transformations and get their results
- cache: to save results for later use
- groupByKey: to group items by key
- reduceByKey¹: group and apply reduce in the same step

¹Prefer this to groupbykey+map+reduce. See the slides from our previous class for more information

Spark specific methods

Some examples:

- collect: to apply all transformations and get their results
- cache: to save results for later use
- groupByKey: to group items by key
- reduceByKey¹: group and apply reduce in the same step
- keyBy: to creates [K,V] tuples by applying a function to each value

¹Prefer this to groupbykey+map+reduce. See the slides from our previous class for more information

Spark specific methods

Some examples:

- collect: to apply all transformations and get their results
- cache: to save results for later use
- groupByKey: to group items by key
- reduceByKey¹: group and apply reduce in the same step
- keyBy: to creates [K,V] tuples by applying a function to each value
- ...

¹Prefer this to groupbykey+map+reduce. See the slides from our previous class for more information

Now, I'll start using the terminal.

Now, I'll start using the terminal.
Don't freak out

Apologies in advance.

Advanced Spark configuration

Available Properties

Most of the properties that control internal settings have reasonable default values. Some of the most common options to set are:

Application Properties

Property Name	Default	Meaning
<code>spark.app.name</code>	(none)	The name of your application. This will appear in the UI and in log data.
<code>spark.driver.cores</code>	1	Number of cores to use for the driver process, only in cluster mode.
<code>spark.driver.maxResultSize</code>	1g	Limit of total size of serialized results of all partitions for each Spark action (e.g. collect). Should be at least 1M, or 0 for unlimited. Jobs will be aborted if the total size is above this limit. Having a high limit may cause out-of-memory errors in driver (depends on <code>spark.driver.memory</code> and memory overhead of objects in JVM). Setting a proper limit can protect the driver from out-of-memory errors.
<code>spark.driver.memory</code>	1g	Amount of memory to use for the driver process, i.e. where <code>SparkContext</code> is initialized. (e.g. 1g, 2g). <i>Note:</i> In client mode, this config must not be set through the <code>SparkConf</code> directly in your application, because the driver JVM has already started at that point. Instead, please set this through the <code>--driver-memory</code> command line option or in your default properties file.
<code>spark.executor.memory</code>	1g	Amount of memory to use per executor process (e.g. 2g, 8g).
<code>spark.extraListeners</code>	(none)	A comma-separated list of classes that implement <code>SparkListener</code> ; when initializing <code>SparkContext</code> , instances of these classes will be created and registered with Spark's listener bus. If a class has a single-argument constructor that accepts a <code>SparkConf</code> , that constructor will be called; otherwise, a zero-argument constructor will be called. If no valid constructor can be found, the <code>SparkContext</code> creation will fail with an exception.
<code>spark.local.dir</code>	/tmp	Directory to use for "scratch" space in Spark, including map output files and RDDs that get stored on disk. This should be on a fast, local disk in your system. It can also be a comma-separated list of multiple directories on different disks. <i>NOTE:</i> In Spark 1.0 and later this will be overridden by <code>SPARK_LOCAL_DIRS</code> (Standalone, Mesos) or <code>LOCAL_DIRS</code> (YARN) environment variables set by the cluster manager.
<code>spark.logConf</code>	false	Logs the effective <code>SparkConf</code> as INFO when a <code>SparkContext</code> is started.
<code>spark.master</code>	(none)	The cluster manager to connect to. See the list of allowed master URL's .
<code>spark.submit.deployMode</code>	(none)	The deploy mode of Spark driver program, either "client" or "remote". Which means to launch driver

- Application, Runtime, UI and RDD settings
- Highlights
 - Master URL and Port
 - CPU and memory per worker
 - Parallelism: default partition size

Deploying a cluster with 20+ nodes

Our big data cluster

We will illustrate what we've already covered by deploying a cluster of more than 20 nodes.

Our big data cluster

We will illustrate what we've already covered by deploying a cluster of more than 20 nodes.

Top of the line servers in Azure and Amazon.

Our big data cluster

We will illustrate what we've already covered by deploying a cluster of more than 20 nodes.

Top of the line servers in Azure and Amazon.

Backed up by BlueMix's Spark infrastructure.



This is embarrassing...

This is embarrassing...
We'll have to make do with our
brains and paper.

This is embarrassing...
We'll have to make do with our
brains and paper.
That's probably how Amazon
Mechanical Turk was born.

Input

We have a log of users that know or ask about a topic. It looks like this:

```
Alice knows Scala  
Bob asks about Scala  
Caroline asks about Java  
Don knows about Scala  
...
```

In other words, lines have this format:

```
{student} {action} {topic}
```

Your task

Answer the questions:

- How many questions were asked about each topic?

Your task

Answer the questions:

- How many questions were asked about each topic?
- How many times did each student ask about each topic?

Your task

Answer the questions:

- How many questions were asked about each topic?
- How many times did each student ask about each topic?
- (Harder) Is there any topic with questions that no other student knows about?


Your task

Answer the questions:


- How many questions were asked about each topic?
- How many times did each student ask about each topic?
- (Harder) Is there any topic with questions that no other student knows about?
- (Even harder) Pair students that know about a topic with students that don't

Different orchestrators

Spark architecture



Spark architecture



Cluster managers

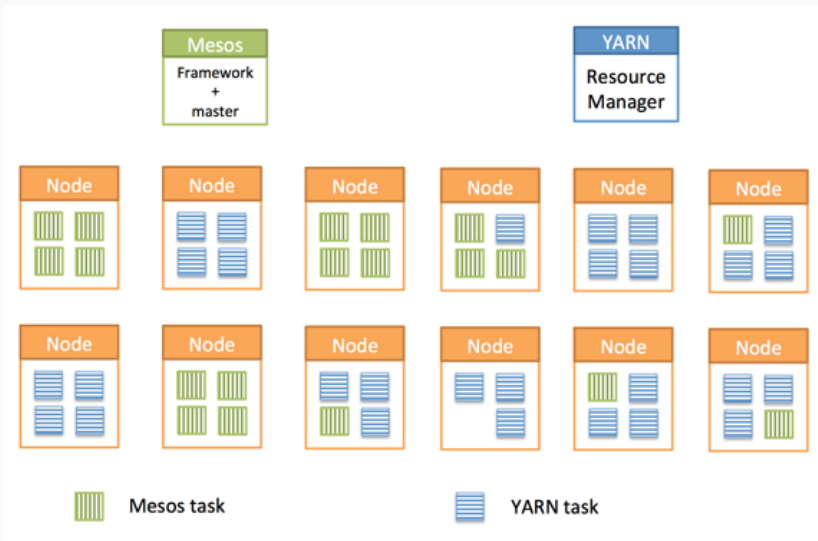
Cluster managers (either Spark's own standalone cluster manager, Mesos or YARN), which allocate resources across applications. Once connected, Spark acquires executors on nodes in the cluster, which are processes that run computations and store data for your application.

Apache Mesos abstracts CPU, memory, storage, and other compute resources away from machines (physical or virtual), enabling fault-tolerant and elastic distributed systems to easily be built and run effectively.


[Download Mesos](#)

[Mesos 1.1.0 Changelog](#)

Cluster managers



Spark architecture




Once again, Spark has astoundingly good documentation

<http://spark.apache.org/docs/latest/running-on-mesos.html>

Spark Ecosystem


Ecosystem



graphx	[SPARK-16345][DOCUMENTATION][EXAMPLES][GRAPHX] Extract graphx program...	5 mon
ml	[SPARK-13770][DOCUMENTATION][ML] Document the ML feature Interaction	21 d
mllib	[SPARK-14300][DOCS][MLLIB] Scala MLlib examples code merge and clean up	a mo
pythonconverters	[SPARK-14744][EXAMPLES] Clean up examples packaging, remove outdated ...	7 mon
sql	[SPARK-18410][STREAMING] Add structured kafka example	13 d
streaming	[SPARK-15208][WIP][CORE][STREAMING][DOCS] Update Spark examples with ...	6 mon
BroadcastTest.scala	[SPARK-16403][EXAMPLES] Cleanup to remove unused imports, consistent ...	5 mon
DFSReadWriteTest.scala	[SPARK-15773][CORE][EXAMPLE] Avoid creating local variable `sc` in ex...	6 mon
DriverSubmissionTest.scala	[SPARK-14444][BUILD] Add a new scalastyle `NoScalaDoc` to prevent Sca...	8 mon
ExceptionHandlingTest.scala	[SPARK-15773][CORE][EXAMPLE] Avoid creating local variable `sc` in ex...	6 mon



Spark Streaming ii



Demo!

We'll show a demo of a modified Spark Streaming task.

Here is a quick “video” of the demo:

<https://asciinema.org/a/6kp49z5m3hq9vja9r7x4rjq26>

All the code and instructions are available in our repository (in several branches):

<https://github.com/balkian/docker-spark>

Acknowledgements and useful links

- Spark programming guide
- Databricks introducing apache spark datasets
- Data Analytics with Hadoop: In-Memory Computing with Spark
- Understanding RDD operations, transformations and actions
- Spark Streaming programming guide