Cache, Accumulators, Broadcast variables

RDDs: Advanced topics

Persistence and Cache

Persistence and Cache

- Spark computes the content of an RDD each time an action is invoked on it
- If the same RDD is used multiple times in an application, Spark recomputes its content every time an action is invoked on the RDD, or on one of its “descendants”
- This is expensive, especially for iterative applications
- We can ask Spark to persist/cache RDDs

Persistence and Cache

- When you ask Spark to persist/cache an RDD, each node stores the content of its partitions in memory and reuses them in other actions on that RDD/dataset (or RDDs derived from it)
  - The first time the content of a persistent/cached RDD is computed in an action, it will be kept in the main memory of the nodes
  - The next actions on the same RDD will read its content from memory
    - i.e., Spark persists/caches the content of the RDD across operations
    - This allows future actions to be much faster (often by more than 10x

Persistence and Cache

- Spark supports several storage levels
  - The storage level is used to specify if the content of the RDD is stored
    - In the main memory of the nodes
    - On the local disks of the nodes
    - Partially in the main memory and partially on disk

Persistence and Cache: Storage levels

Storage Level	Meaning
MEMORY_ONLY	Store RDD as deserialized Java objects in the JVM. If the RDD does not fit in memory, some partitions will not be cached and will be recomputed on the fly each time they're needed. This is the default level.
MEMORY_AND_DISK	Store RDD as deserialized Java objects in the JVM. If the RDD does not fit in memory, store the partitions that don't fit on (local) disk, and read them from there when they're needed.
MEMORY_ONLY_SER	Store RDD as serialized Java objects (one byte array per partition). This is generally more space-efficient than deserialized objects, especially when using a fast serializer, but more CPU-intensive to read.
MEMORY_AND_DISK_SER	Similar to MEMORY_ONLY_SER, but spill partitions that don't fit in memory to disk instead of recomputing them on the fly each time they're needed.

http://spark.apache.org/docs/2.2.0/rdd-programming-guide.html#rdd-persistence

Persistence and Cache: Storage levels

Storage Level	Meaning
DISK_ONLY	Store the RDD partitions only on disk.
MEMORY_ONLY_2, MEMORY_AND_DISK_2, etc.	Same as the levels above, but replicate each partition on two cluster nodes.
OFF_HEAP (experimental)	Similar to MEMORY_ONLY_SER, but store the data in off-heap memory. This requires off-heap memory to be enabled.

http://spark.apache.org/docs/2.2.0/rdd-programming-guide.html#rdd-persistence

Persistence and Cache

- You can mark an RDD to be persisted by using the **JavaRDD<T> persist(StorageLevel level)** method of the **JavaRDD<T>** class
- The parameter of persist can assume the following values
  - StorageLevel.MEMORY_ONLY()
  - StorageLevel.MEMORY_AND_DISK()
  - StorageLevel.MEMORY_ONLY_SER()
  - StorageLevel.MEMORY_AND_DISK_SER()
  - StorageLevel.DISK_ONLY()
  - StorageLevel.NONE()
  - StorageLevel.OFF_HEAP()

Persistence and Cache

  - StorageLevel.MEMORY_ONLY_2()
  - StorageLevel.MEMORY_AND_DISK_2()
  - StorageLevel.MEMORY_ONLY_SER_2()
  - StorageLevel.MEMORY_AND_DISK_SER_2()
- The storage level *_2() replicate each partition on two cluster nodes
  - If one node fails, the other one can be used to perform the actions on the RDD without recomputing the content of the RDD

Persistence and Cache

- You can cache an RDD by using the **JavaRDD<T> cache()** method of the **JavaRDD<T>** class
  - It corresponds to persist the RDD with the storage level `MEMORY_ONLY`
  - i.e., it is equivalent to inRDD.persist(StorageLevel.MEMORY_ONLY())
- Note that both persist and cache return a new JavaRDD
  - Because RDDs are immutable

Persistence and Cache

- The use of the persist/cache mechanism on an RDD provides an advantage if the same RDD is used multiple times
  - i.e., multiples actions are applied on it or on its descendants

Persistence and Cache

- The storage levels that store RDDs on disk are useful if and only if
  - The “size” of the RDD is significantly smaller than the size of the input dataset
  - Or the functions that are used to compute the content of the RDD are expensive
  - Otherwise, recomputing a partition may be as fast as reading it from disk

Remove data from cache

- Spark automatically monitors cache usage on each node and drops out old data partitions in a least-recently-used (LRU) fashion
- You can manually remove an RDD from the cache by using the **JavaRDD<T> unpersist()** method of the **JavaRDD<T>** class

Cache: Example

- Create an RDD from a textual file containing a list of words
  - One word for each line
- Print on the standard output
  - The number of lines of the input file
  - The number of distinct words

Cache: Example

```
// Read the content of a textual file
// and cache the associated RDD
JavaRDD<String> inputRDD = sc.textFile("words.txt").cache();

System.out.println("Number of words: "+inputRDD.count());
System.out.println("Number of distinct words: "
                                        +inputRDD.distinct().count());
```

Cache: Example

```
// Read the content of a textual file
// and cache the associated RDD
JavaRDD<String> inputRDD = sc.textFile("words.txt").cache();

System.out.println("Number of wor
System.out.println("Number of dist
                         +inputRDD.distinct().count());
```

The cache method is invoked.
Hence, inputRDD is a “cached” RDD

Cache: Example

```
// Read the content of a textual file
// and cache the associated RDD
JavaRDD<String> inputRDD = sc.textFile("words.txt").cache();

System.out.println("Number of words: "+inputRDD.count());
System.out.println("Number of distinct words: "
                                        +inputRDD.distinct().count());
```

This is the first time an action is invoked on the inputRDD RDD.
The content of the RDD is computed by reading the lines of the words.txt file and the result of the count action is returned.
The content of inputRDD is also stored in the main memory of the nodes of the cluster.

Cache: Example

```
// Read the content of a textual file
// and cache the associated RDD
JavaRDD<String> inputRDD = sc.textFile("words.txt").cache();

System.out.println("Number of words: "+inputRDD.count());
System.out.println("Number of distinct words: "
                        +inputRDD.distinct().count());
```

The content of inputRDD is in the main memory if the nodes of the cluster. Hence the computation of distinct() is performed by reading the data from the main memory and not from the input (HDFS) file words.txt

Accumulators

Accumulators

- When a “function” passed to a Spark operation is executed on a remote cluster node, it works on separate copies of all the variables used in the function
  - These variables are copied to each node of the cluster, and no updates to the variables on the nodes are propagated back to the driver program

Accumulators

- Spark provides a type of shared variables called **accumulators**
- Accumulators are shared variables that are only “added” to through an associative operation and can therefore be efficiently supported in parallel
- They can be used to implement counters (as in MapReduce) or sums

Accumulators

- Accumulators are usually used to compute simple statistics while performing some other actions on the input RDD
  - The avoid using actions like reduce() to compute simple statistics (e.g., count the number of lines with some characteristics)

Accumulators

- The driver defines and initializes accumulators
- The code executed in the worker nodes increases the value of the accumulators
  - i.e., the code in the "functions" associated with transformations or actions
- The final values of the accumulators are returned to the driver node
  - Only the driver node can access the final value of each accumulator
  - The worker nodes cannot access the values of the accumulators
    - They can only add values to them

Accumulators

- Pay attention that the accumulators are increased in the functions associated with transformations or actions
- If you update accumulators inside transformations, pay attention that transformations are lazily evaluated
  - The values of the accumulators are updated only when there is an action that trigger the execution of the transformations

Accumulators

- Pay attention that if the part of the code in which an accumulator is increased is executed multiple times, the value of the accumulator could be wrong

Accumulators

- To be sure that the values of the accumulators are correct you must update them in transformations or actions that are executed only one time in your application
  - Or at least, you must know how many times those transformations or actions will be executed

Accumulators

- The foreach method is frequently used to update accumulators
- **foreach** is an action
- It applies a function to all elements of an RDD
  - It returns no values (returned data type: void)
  - It is usually used to
    - Print the content of an RDD on the standard output
    - Store the content of an RDD in an external database
    - Update accumulators

Accumulators

- public void foreach(Function<T> **f**)
  - Applies a function **f** to all elements of an RDD
    - i.e., f is invoked one time for each element in the RDD
  - It returns not values (void)
- f can be a function that updates accumulators

Accumulators

- Spark natively supports accumulators of numeric types
  - Long and Double accumulators
- But programmers can add support for new data types

Accumulators

- Accumulators are objects extending **org.apache.spark.util.AccumulatorV2**
  - A Long accumulator can be defined in the driver by invoking the **LongAccumulator longAccumulator ()** method of the **SparkContext** class
    - Pay attention that the scala **SparkContext must be used** instead of the JavaSparkContext
      - org.apache.spark.SparkContext ssc = sc.sc();
  - A Double accumulator can be defined by using the **DoubleAccumulator doubleAccumulator()** method

Accumulators

- The value of an accumulator can be “increased” by using the **void add(T value) method** of the **AccumulatorV2** class
  - Add “value” to the current value of the accumulator
- The final value of an accumulator can be retrieved in the driver program by using the **T value()** method of the **AccumulatorV2** class

Accumulators: Example

- Create an RDD from a textual file containing a list of email addresses
  - One email for each line
- Select the lines containing a valid email and store them in an HDFS file
  - In this example, an email is considered as valid if it contains the @ symbol
- Print also, on the standard output, the number of invalid emails

Accumulators: Example

```
....
// Define an accumulator of type long
final LongAccumulator invalidEmails = sc.sc().longAccumulator();


// Read the content of the input textual file
JavaRDD<String> emailsRDD = sc.textFile("emails.txt");

// Select only valid emails
JavaRDD<String> validEmailsRDD = emailsRDD.filter(line ->
      {           // Increase the accumulator if the line contains an invalid email
                  if (line.contains("@")==false)
                        invalidEmails.add(1);

                  return line.contains("@");
      });
```

Accumulators: Example

```
....
// Define an accumulator of type long
final LongAccumulator invalidEmails = sc.sc().longAccumulator();


// Read th
JavaRDD<String> emailsRDD = sc.textFile("emails.txt");

// Select only valid emails
JavaRDD<String> validEmailsRDD = emailsRDD.filter(line ->
        {           // Increase the accumulator if the line contains an invalid email
                    if (line.contains("@")==false)
                            invalidEmails.add(1);

                    return line.contains("@");
        });
```

Definition of an accumulator of type Long

Accumulators: Example

```
....
// Define an accumulator of type long
final LongAccumulator invalidEmails = sc.sc().longAccumulator();


// Read the content of the input textual file
JavaRDD<String> emailsRDD = sc.textFile("emails.txt");

// Select only valid emails
JavaRDD<String> validEmailsRDD = emailsRDD.filter(line ->
    {       // Increase the accumulator if the line contains an invalid email
            if (line.contains("@")==false)
                invalidEmails.add(1);

            return line.contains("@");
    });
```

The call method increases also the value of invalidEmails that is an accumulator

Accumulators: Example

```
// Store valid emails in the output file
validEmailsRDD.saveAsTextFile(outputPath);

// Print the number of invalid emails
System.out.println("Invalid emails: "+invalidEmails.value());
```

Accumulators: Example

```
// Store valid emails in the output file
validEmailsRDD.saveAsTextFile(outputPath);

// Print the number of invalid emails
System.out.println("Invalid emails: "+invalidEmails.value());
```

Accumulators: Example

```
// Store valid emails in the output file
validEmailsRDD.saveAsTextFile(outputPath);

// Print the number of invalid emails
System.out.println("Invalid emails: "+invalidEmails.value());
```

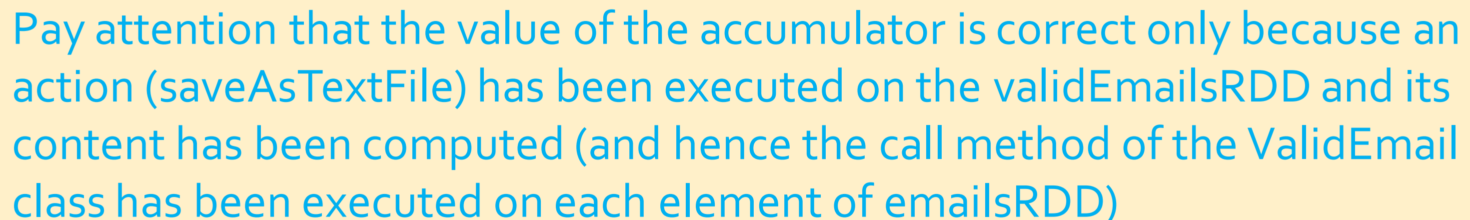

Accumulators: Example with foreach

- Create an RDD from a textual file containing a list of email addresses
  - One email for each line
- Select the lines containing a valid email and store them in an HDFS file
  - In this example, an email is considered as valid if it contains the @ symbol
- Print also, on the standard output, the number of invalid emails

Accumulators: Example with foreach

```
....
// Define an accumulator of type long
final LongAccumulator invalidEmails = sc.sc().longAccumulator();


// Read the content of the input textual file
JavaRDD<String> emailsRDD = sc.textFile("emails.txt");

// Select only valid emails
JavaRDD<String> validEmailsRDD = emailsRDD.filter(line ->
                                        line.contains("@"));
```

Accumulators: Example with foreach

```
// Store valid emails in the output file
validEmailsRDD.saveAsTextFile(outputPath);

// Compute the number of invalid emails
emailsRDD.foreach(line -> {
                    if (line.contains("@") == false)
                          invalidEmails.add(1);
                 });

// Print the number of invalid emails
System.out.println("Invalid emails: "+invalidEmails.value());
```

Accumulators: Example with foreach

```
// Store valid emails in the output file
validEmailsRDD.saveAsTextFile(outputPath);

// Compute the number of invalid emails
emailsRDD.foreach(line -> {
                    if (line.contains("@") == false)
                        invalidEmails.add(1);
                });

// Print the number of invalid emails
System.out.println("Invalid emails: "+invalidEmails.value());
```

Update the accumulator

Personalized accumulators

- Programmers can define accumulators based on new data types (different from Long and Double)
- To define a new accumulator data type of type T, the programmer must define a class extending the abstract class **org.apache.spark.util.AccumulatorV2<T,T>**
  - Several methods must be implemented
    - abstract void add(T value)
    - abstract T value()
    - abstract AccumulatorV2<T,T> copy()
    - ..

Personalized accumulators

- Then, a new accumulator of the new type can be instantiated and “registered” in the context of our application

```
.....
MyAcculumator myAcc = new MyAccumulator();
sc.sc().register(myAcc, "MyNewAcc");
```

Broadcast variables

Broadcast variables

- Spark supports also broadcast variables
- A broadcast variable is a read-only (medium/large) shared variable
  - That is instantiated in the driver
  - And it is sent to all worker nodes that use it in one or more Spark actions

Broadcast variables

- A copy each "standard" variable is sent to all the tasks executing a Spark action using that variable
  - i.e., the variable is sent "num. tasks" times
- A broadcast variable is sent only one time to each executor using it in at least one Spark action (i.e., in at least one of its tasks)
  - Each executor can run multiples tasks using that variable and the broadcast variable is sent only one time
  - Hence, the amount of data sent on the network is limited by using broadcast variables instead of "standard" variables

Broadcast variables

- Broadcast variables are usually used to share (large) lookup-tables

Broadcast variables

- Broadcast variables are objects of type **Broadcast<T>**
- A broadcast variable of type T is defined in the driver by using the **Broadcast<T> broadcast(T value)** method of the **JavaSparkContext** class
- The value of a broadcast variable of type T is retrieved (usually in transformations) by using the **T value()** method of the **Broadcast<T>** class

Broadcast variables: Example

- Create an RDD from a textual file containing a dictionary of pairs (word, integer value)
  - One pair for each line
  - Suppose the content of this file is large but can be stored in main-memory
- Create an RDD from a textual file containing a set of words
  - A sentence (set of words) for each line
- "Transform" the content of the second file mapping each word to an integer based on the dictionary contained in the first file
  - Store the result in an HDFS file

Broadcast variables: Example

- First file (dictionary)

  java 1

  spark 2

  test 3

- Second file (the text to transform)

  java spark

  spark test java

- Output file

  1 2

  2 3 1

Broadcast variables: Example

```
…
// Read the content of the dictionary from the first file and
// map each line to a pair (word, integer value)
JavaPairRDD<String, Integer> dictionaryRDD =
	sc.textFile("dictionary.txt").mapToPair(line ->
	{
		String[] fields = line.split(" ");

		String word=fields[0];
		Integer intWord=Integer.parseInt(fields[1]);

		return new Tuple2<String, Integer>(word, intWord) ;
	});
```

Broadcast variables: Example

```
// Create a local HashMap object that will be used to store the
// mapping word -> integer
HashMap<String, Integer> dictionary=new HashMap<String, Integer>();

// Create a broadcast variable based on the content of dictionaryRDD
// Pay attention that a broadcast variable can be instantiated only
// by passing as parameter a local java variable and not an RDD.
// Hence, the collect method is used to retrieve the content of the
// RDD and store it in the dictionary HashMap<String, Integer> variable
for (Tuple2<String, Integer> pair: dictionaryRDD.collect()) {
	dictionary.put(pair._1(), pair._2());
}

final Broadcast<HashMap<String, Integer>> dictionaryBroadcast =
						sc.broadcast(dictionary);
```

Broadcast variables: Example

```
// Create a local HashMap object that will be used to store the
// mapping word -> integer
HashMap<String, Integer> dictionary=new HashMap<String, Integer>();

// Create a broadcast variable based on the content of dictionaryRDD
// Pay attention that a broadcast variable can be instantiated only
// by passing as parameter a local java variable and not an RDD.
// Hence, the collect method is used to retrieve the content of the
// RDD and store it in the dictionary HashMap<String, Integer> variable
for (Tuple2<String, Integer> pair: dictionaryRDD.collect()) {
	dictionary.put(pair._1(), pair._2());
}

final Broadcast<HashMap<String, Integer>> dictionaryBroadcast =
						sc.broadcast(dictionary);
```

Define a broadcast variable

Broadcast variables: Example

```
// Read the content of the second file
JavaRDD<String> textRDD = sc.textFile("document.txt");
```

Broadcast variables: Example

```
// Map each word in textRDD to the corresponding integer and concatenate
// them
JavaRDD<String> mappedTextRDD=
        textRDD.map(line ->
        {
                String transformedLine=new String("");
                Integer intValue;
                // map each word to the corresponding integer
                String[] words=line.split(" ");
                for (int i=0; i<words.length; i++) {
                        intValue=dictionaryBroadcast.value().get(words[i]);
                        transformedLine=transformedLine.concat(intValue+" ");
                }
                return transformedLine;
        });
```

Broadcast variables: Example

```
// Map each word in textRDD to the corresponding integer and concatenate
// them
JavaRDD<String> mappedTextRDD=
    textRDD.map(line ->
    {
            String transformedLine=new String("");
            Integer intValue;
            // map each word to the corresponding integer
            String[] words=line.split(" ");
            for (int i=0; i<words.length; i++) {
                intValue=dictionaryBroadcast.value().get(words[i]);
                transformedLine=transformedLine.concat(intValue+" ");
            }
            return transformedLine;
    });
```

Retrieve the content of the broadcast variable and use it

Broadcast variables: Example

```
// Store the result in an HDFS file
mappedTextRDD.saveAsTextFile(outputPath);
```