



Introduction to Spark

- MapReduce enables big data analytics using large, unreliable clusters
- There are multiple implementation of MapReduce
 - Based on HDFS
 - Based on No-SQL databases
 - Based on cloud

- Can not support Complex/iterative applications efficiently
- Root Cause: Inefficient data sharing
 - Only way to share data across jobs is via stable storage*
- There is room to further improvement with MapReduce
 - iterative algorithms,
 - interactive ad-hoc queries



Slow due to replication and disk I/O,
but necessary for fault tolerance

- In other words, MapReduce lacks efficient primitives for data sharing
- This is where Spark comes into the picture – instead of load the data from disk for every query, why not do In-Memory data sharing?

- Here is how Spark addresses this issue:
Resilient Distributed Datasets (RDDs)
- RDD allows applications to keep working sets in memory for reuse.

RDD Overview

- A new programming model (RDD)
 - parallel/distributed computing
 - in-memory sharing
 - fault-tolerance
- An implementation of RDD:

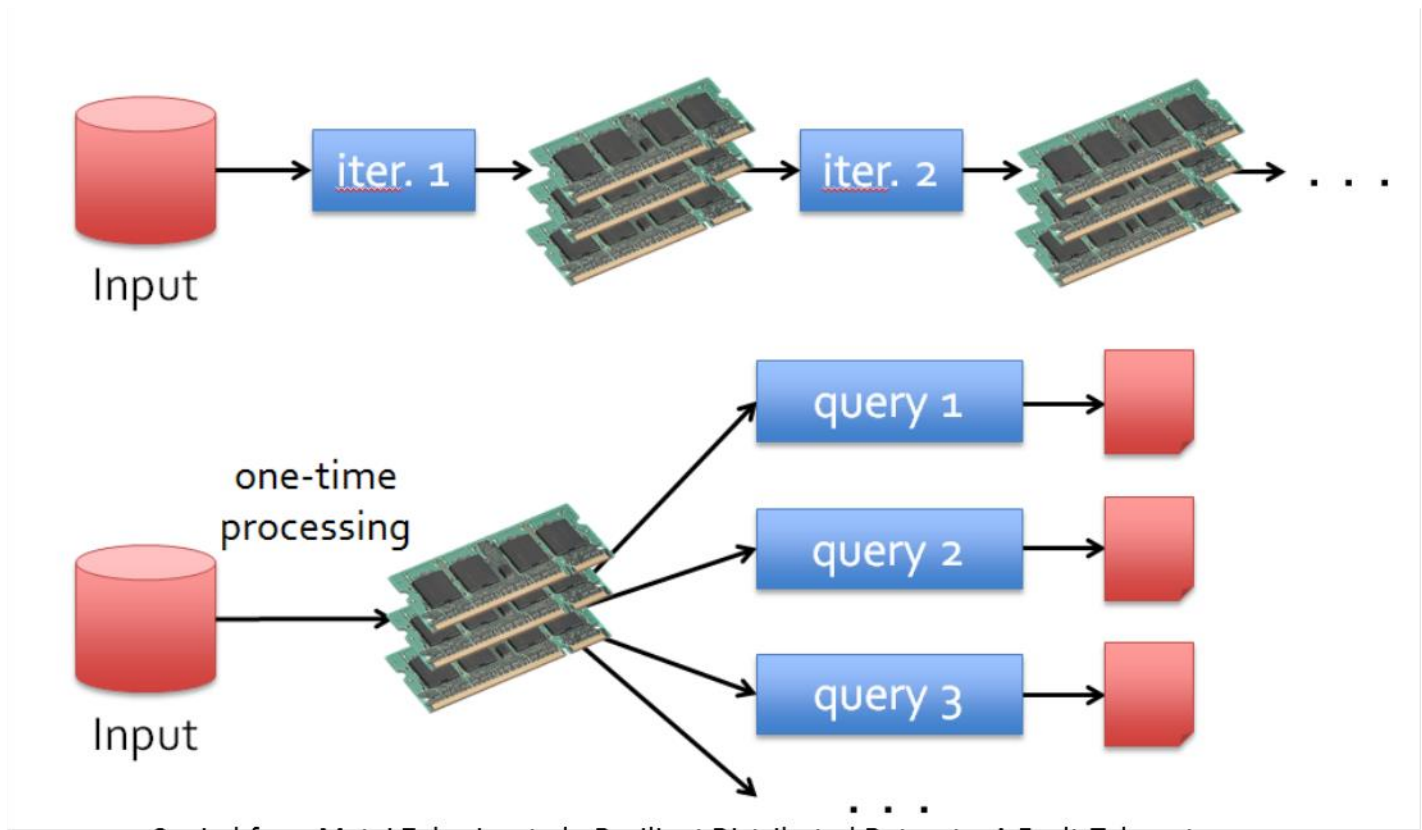


Spark's Philosophy

- Generalize MapReduce
- Richer Programming Model
 - ✓ Fewer Systems to Master
- Better Memory Management
 - ✓ Less data movement leads to better performance for complex analytics



Something like this.....



Copied from Matei Zaharia, et.al., Resilient Distributed Datasets, A Fault-Tolerant Abstraction for In-Memory Cluster Computing, amplab, UC Berkeley

But memory is not reliable, how about FT?

- Spark's solution: Resilient Distributed Datasets (RDDs)
 - Read only, partitioned collections of objects
 - A distributed *immutable* array
 - Created through parallel transformations on data in stable storage
 - Can be cached for efficient reuse.

- Log the operations that generates the RDD
- Automatically rebuilt on (partial) failure

Spark API

- Spark provides API to support the following languages: Scala, Java, Python, R
- Data Operations
 - Transformations; lazily create RDDs
`wc = dataset.flatMap(tokenize).reduceByKey(add)`
 - Actions; execute compilation
`wc.collect()`

Abstraction: Dataflow Operators

| | | |
|-------------------------|----------------------|----------------------|
| • map | • reduce | • sample |
| • filter | • count | • take |
| • groupBy | • fold | • first |
| • sort | • reduceByKey | • partitionBy |
| • union | • groupByKey | • mapWith |
| • join | • cogroup | • pipe |
| • leftOuterJoin | • Cross | • Save |
| • rightouterjoin | • Zip | • ... |

Spark Example: Log Mining

- Load error messages from a log into memory and run interactive queries

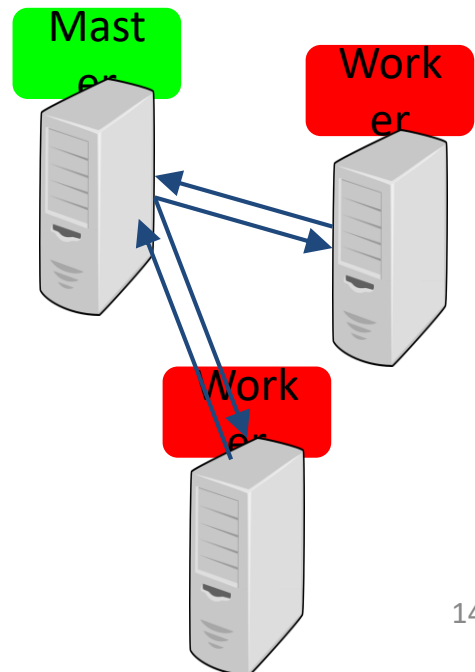
```
lines = spark.textFile("hdfs://...")
errors = lines.filter(_.startsWith("ERROR"))
errors.persist()
errors.filter(_.contains("Foo")).count()
errors.filter(_.contains("Bar")).count()
```

base RDD

transformation

action

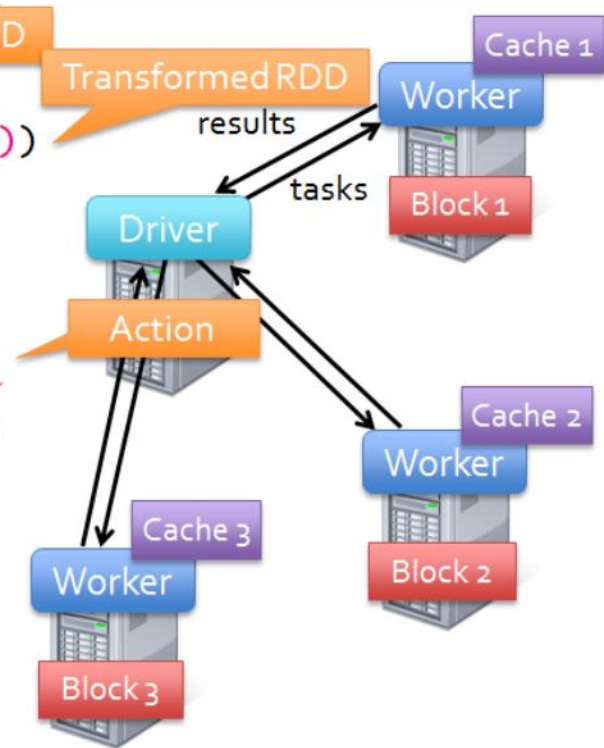
Result: full-text search on 1TB data in 5-7sec
vs. 170sec with on-disk data!



Simple Example

```
lines = spark.textFile("hdfs://...")
errors = lines.filter(_.startsWith("ERROR"))
messages = errors.map(_.split('\t')(2))
cachedMsgs = messages.cache()

cachedMsgs.filter(_.contains("foo")).count
cachedMsgs.filter(_.contains("bar")).count
. . .
```



Simple Yet Powerful

WordCount Implementation: Hadoop vs. Spark

```
1 public class WordCount {
2
3     public static class TokenizerMapper
4         extends Mapper<Object, Text, Text, IntWritable>{
5
6         private final static IntWritable one = new IntWritable(1);
7         private Text word = new Text();
8
9         public void map(Object key, Text value, Context context
10             ) throws IOException, InterruptedException {
11             StringTokenizer itr = new StringTokenizer(value.toString());
12             while (itr.hasMoreTokens()) {
13                 word.set(itr.nextToken());
14                 context.write(word, one);
15             }
16         }
17     }
18
19     public static class IntSumReducer
20         extends Reducer<Text, IntWritable, Text, IntWritable> {
21         private IntWritable result = new IntWritable();
22
23         public void reduce(Text key, Iterable<IntWritable> values,
24             Context context
25             ) throws IOException, InterruptedException {
26             int sum = 0;
27             for (IntWritable val : values) {
28                 sum += val.get();
29             }
30             result.set(sum);
31             context.write(key, result);
32         }
33     }
34
35     public static void main(String[] args) throws Exception {
36         Configuration conf = new Configuration();
37         Job job = Job.getInstance(conf, "word count");
38         job.setJarByClass(WordCount.class);
39         job.setMapperClass(TokenizerMapper.class);
40         job.setCombinerClass(IntSumReducer.class);
41         job.setReducerClass(IntSumReducer.class);
42         job.setOutputKeyClass(Text.class);
43         job.setOutputValueClass(IntWritable.class);
44         FileInputFormat.addInputPath(job, new Path(args[0]));
45         FileOutputFormat.setOutputPath(job, new Path(args[1]));
46         System.exit(job.waitForCompletion(true) ? 0 : 1);
47     }
48 }
```

```
1 val textFile = sc.textFile("hdfs://...")
2 val counts = textFile.flatMap(line => line.split(" ")).map(word => (word, 1)).reduceByKey(_ + _)
3 counts.saveAsTextFile("hdfs://...")
```



Pregel: iterative graph processing, 200 LoC using Spark

HaLoop: iterative MapReduce, 200 LoC using Spark

WordCount

- `val counts = sc.textFile("hdfs://...").flatMap(line=>line.split("")).map(word=>(word, 1L)).reduceByKey(_+_)`

What is Iterative Algorithm?

data = input data

$w = \langle \text{target vector} \rangle$ -(Shared data structure)

At each iteration,

Do something to item in data:

Update (w) -(Update shared data structure)

```
val data = spark.textFile(...).map(readPoint).cache()

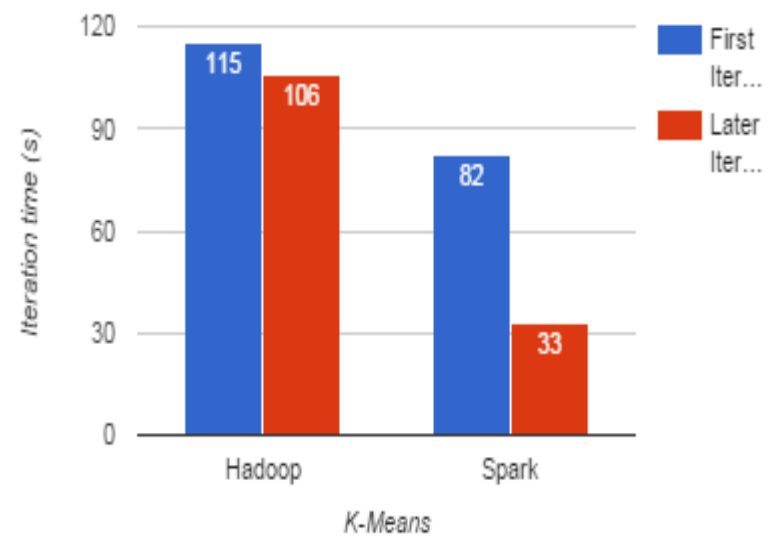

var w = Vector.random(D)

for (i <- 1 to ITERATIONS) {
  val gradient = data.map(p =>
    (1 / (1 + exp(-p.y*(w dot p.x))) - 1) * p.y * p.x
  ).reduce(_ + _)
  w -= gradient
}

println("Final w: " + w)
```

Evaluation

10 iterations on 100GB data using 25-100 machines



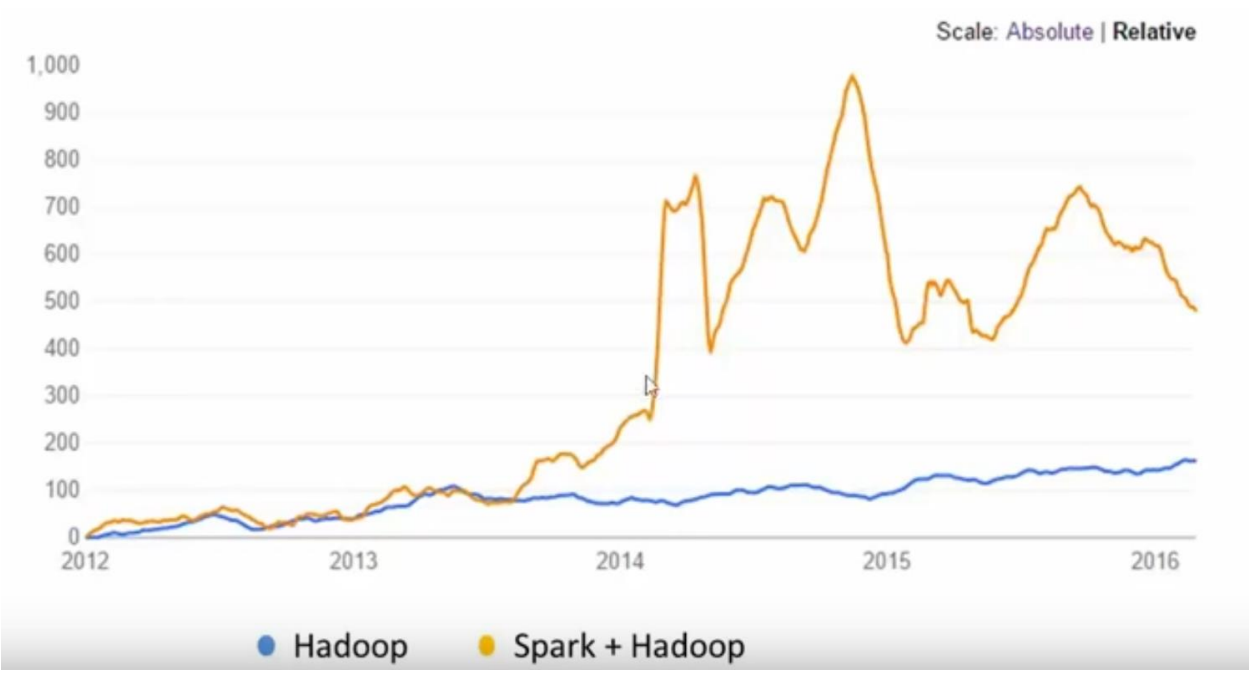
Spark Ecosystem



- Spark Core: is the execution engine for the spark platform. It provides distributed in-memory computing capabilities
- Spark SQL: is an engine for Hadoop Hive that enables unmodified Hadoop Hive queries to run up to 100x faster on existing deployments and data
- Spark Streaming: is an engine that enables powerful interactive and analytical applications on streaming data
- MLlib: is a scalable machine learning library
- GraphX is a graph computation engine built on top of Spark

Hadoop Vs Spark

- Spark is a computing framework that can be deployed upon Hadoop.
- You can view Hadoop as an operating system for a distributed computing cluster, while Spark is an application running on the system to provide in-memory analytics functions.





Pivotal™



STRATIO

ORACLE

guavus

cloudera



bluedata

MicroStrategy®



Typesafe

elasticsearch.



pentaho®



ADATAO
DATA INTELLIGENCE FOR AI



tresata



Atigeo™

NUBE



ZOOMDATA®
DATA in Motion

Qlik Q®



Alpine

DIYOTTA

talend*

 **tableau** *MicroStrategy* Qlik 

elasticsearch.  pentaho **talend**

 **tresata**  TRIFACTA **SKYTREE**
THE MACHINE LEARNING COMPANY

Alpine  **at**scale **looker** 
virdata

 **FAiM**DATA  **ADATAO**
DATA INTELLIGENCE FOR ALL **DIYOTTA**

 **OOM**DATA  **platfora**  **APERVI**  **NUBE**

 **Atigeo**  日志易
rizhiyi.com **ZALONI**  **Typesafe**