

# Apache Spark

Alexandre da Silva Veith<sup>1</sup> and Marcos Dias de Assunção<sup>1</sup>

<sup>1</sup>Inria Avalon, LIP Laboratory, ENS Lyon, University of Lyon,

Emails: alexandre.veith@ens-lyon.fr, assuncao@acm.org

## 1 Definition

Apache Spark is a cluster computing solution and in-memory processing framework that extends the MapReduce model to support other types of computations such as interactive queries and stream processing [1]. Designed to cover a variety of workloads, Spark introduces an abstraction called Resilient Distributed Datasets (RDDs) that enables running computations in memory in a fault-tolerant manner. RDDs, which are immutable and partitioned collections of records, provide a programming interface for performing operations, such as map, filter and join, over multiple data items. For fault-tolerance purposes, Spark records all transformations carried out to build a dataset, thus forming a *lineage graph*.

## 2 Overview

Spark [2] is an open-source big data framework originally developed at the University of California at Berkeley and later adopted by the Apache Foundation, which has maintained it ever since. Spark was designed to address some of the limitations of the MapReduce model, especially the need for speed processing of large datasets. By using RDDs, purposely designed to store restricted amounts of data in memory, Spark enables performing computations more efficiently than MapReduce, which runs computations on the disk.

Although the project contains multiple components, at its core (Figure 1) Spark is a computing engine that schedules, distributes, and monitors applications comprising multiple tasks across nodes of a computing cluster [3]. For cluster management, Spark supports its native Spark cluster (standalone), Apache YARN [4], or Apache Mesos [5]. At the core also lies the RDD abstraction. RDDs are sets of data items distributed across the cluster nodes and that can be manipulated in parallel. At a higher level,

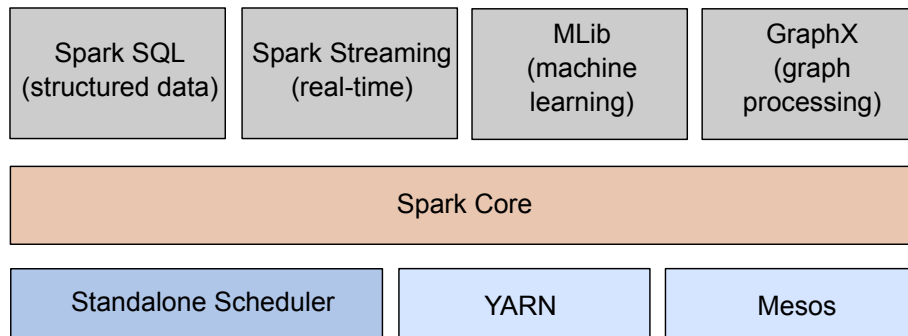


Figure 1: The Apache Spark stack [3].

```

// Create a Scala Spark configuration context
val config = new SparkConf().setAppName("WordCount")
val sc = new SparkContext(config)

// Load the input data
val input = sc.textFile(theInputFile)

// Split it into words
val words = input.flatMap(line => line.split(" "))

// Transform into pairs and count
val counts = words.map(word =>
    (word, 1)).reduceByKey{case (x, y) => x + y}

// Save the word count to a text file
counts.saveAsTextFile(theOutputFile)
  
```

Figure 2: Word count example using Spark's Scala API [3].

it provides support for multiple tightly integrated components for handling various types of workloads such as SQL, streaming and machine learning.

Figure 2 depicts an example of a word count application using Spark's Scala API for manipulating datasets. Spark computes RDDs in a lazy fashion, the first time they are used. Hence, the code in the example is evaluated when the counts are saved to disk, in which moment the results of the computation are required. Spark can also read data from various sources, such as Hadoop Distributed File System (HDFS), Cassandra, OpenStack Swift<sup>1</sup>, and Amazon Simple Storage Service (S3)<sup>2</sup>.

<sup>1</sup><https://wiki.openstack.org/wiki/Swift>

<sup>2</sup><https://aws.amazon.com/s3/>

## 2.1 Spark SQL

Spark SQL [6] is a module for processing structured data<sup>3</sup>. It builds on the RDD abstraction by providing Spark core engine with more information about the structure of the data and the computation being performed. In addition to enabling users to perform SQL queries, Spark SQL provides the Dataset API, which offers Datasets and DataFrames. A Dataset can be built using JVM objects that can then be manipulated using functional transformations. A DataFrame can be built from a large number of sources and is analogous to a table in a relational database; it is a Dataset organised into named columns.

## 2.2 Spark Streaming

Spark Streaming provides a micro-batch based framework for processing data streams. Data can be ingested from systems such as Apache Kafka<sup>4</sup>, Flume or Amazon Kinesis<sup>5</sup>. Under the traditional stream processing approach based on a graph of continuous operators that process tuples as they arrive (*i.e.* the dataflow model), it is arguably difficult to achieve fault tolerance and handle stragglers. As application state is often kept by multiple operators, fault tolerance is achieved either by replicating sections of the processing graph or via upstream backup. The former demands synchronisation of operators via a protocol such as Flux [7] or other transactional protocols [8], whereas the latter, when a node fails, requires parents to replay previously sent messages to rebuild the state.




Figure 3: High-level view of discretised streams.

Spark Streaming uses a high-level abstraction called *discretised stream* or *DStream* [9]. As depicted in Figure 3, DStreams follows a micro-batch approach that organises stream processing as batch computations carried out periodically over small time windows. During a short time interval, DStreams stores the received data, which the cluster resources then use as input dataset for performing parallel computations once the interval elapses. These computations produce new datasets that represent an intermediate state or computation outputs. The intermediate state consists of RDDs that DStreams processes along with the datasets stored during the next interval.

<sup>3</sup><https://spark.apache.org/docs/latest/sql-programming-guide.html>

<sup>4</sup><https://kafka.apache.org/>

<sup>5</sup><https://aws.amazon.com/kinesis/>

In addition to providing a strong unification with batch processing, this model stores the state in memory as RDDs that DStreams can deterministically recompute. This micro-batch approach, however, sacrifices response time as the delay for processing events is dependent on the length of the micro batches.

### 2.3 MLlib

Spark contains a library with common Machine Learning (ML) functionality such as learning algorithms for classification, regression, clustering, and collaborative filtering; and featurisation including feature extraction, transformation, dimensionality reduction, and selection [10]. MLlib also enables the creation of ML pipelines and persistence of algorithms, models and pipelines. These features are designed to scale out across large computing clusters using Spark's core engine.

### 2.4 GraphX

GraphX [11] extends the RDD API by enabling the creation of a multigraph (*i.e.*, the property graph) with arbitrary properties attached to vertices and edges. The library is designed for manipulating graphs, exposing a set of operators (*e.g.*, subgraph, joinVertices), and for carrying out parallel computations. It contains a library of common graph algorithms, such as PageRank and triangle counting.

## 3 Examples of Applications

Spark has been used for several data processing and data science tasks, but the range of applications that it enables is endless. Freeman *et al.* [12], for instance, designed a library called Thunder on top of Spark for large-scale analysis of neural data. Many machine learning and statistical algorithms have been implemented for MLlib, which simplifies the construction of machine learning pipelines. The source code of Spark has also grown substantially since it became an Apache project. Numerous third party libraries and packages have been included for performing tasks in certain domains or for simplifying the use of existing APIs.

Spark provides the functionalities that data scientists need to perform data transformation, processing, and analysis. Data scientists often need to perform ad hoc exploration during which they have to test new algorithms or verify the results in the least amount of time. Spark provides APIs, libraries and shells that allow scientists to perform such tasks while enabling them to test their algorithms on large problem sizes. Once the exploration phase is performed, the solution is productised by engineers who integrate the data analysis tasks into an often more complex business application.

Examples of applications built using Apache Spark include analysis of data from mobile devices [13] and Internet of Things (IoT), web-scale graph analytics, anomaly detection of user behaviour and network traffic for information security [14], real-time machine learning, data stream processing pipelines, engineering workloads, geospatial data processing, to cite just a few. New application scenarios are presented each year during Spark’s Summit<sup>6</sup>, an event that has become a showcase of next-generation big data applications.

## 4 Future Directions of Research

Spark APIs shine both during exploratory work and when engineering a solution deployed in production. Over the years, much effort has been paid towards making APIs easier to use and to optimise; for instance, the introduction of the Dataset API to avoid certain performance degradations that could occur if a user did not properly design the chain of operations executed by the Spark engine. As mentioned earlier, considerable work has also focused on creating new libraries and packages, including for processing live streams. The discretised model employed by Spark’s stream processing API, however, introduces some delays depending on the time length of micro batches.

More recent and emerging application scenarios, such as analysis of vehicular traffic and networks, monitoring of operational infrastructure, wearable assistance [15], and 5G services, require data processing and service response under very short delays. Spark can be used as part of a larger service workflow, but alone it does not provide means to address some of the challenging scenarios that require very short response times. This requires the use of other frameworks such as Apache Storm.

To reduce the latency of applications delivered to users, many service components are also increasingly being deployed at the edges of the Internet [16] under a model commonly called edge computing. Some frameworks are available for processing streams of data using resources at the edge (*e.g.* Apache Edgent<sup>7</sup>), whereas others are emerging. There are also frameworks that aim to provide high-level programming abstractions for creating dataflows (*e.g.* Apache Beam<sup>8</sup>) with underlying execution engines for several processing solutions. There is, however, a lack of unifying solutions on programming models and abstractions for exploring resources from both cloud and edge computing deployments, as well as scheduling and resource management tools for deciding what processing tasks need to be offloaded to the edge.

---

<sup>6</sup><https://spark-summit.org/>

<sup>7</sup><https://edgent.apache.org/>

<sup>8</sup><https://beam.apache.org/>

## 5 Cross References

- Cloud computing for big data analysis.
- Streaming microservices.
- Hadoop.
- In-Memory Processing.
- Advances in MapReduce frameworks.
- Large-scale graph processing.
- Definition of Data Streams.
- Micro-batching vs true streaming.
- Graph processing frameworks.

## References

- [1] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, I. Stoica, Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing, in: 9th USENIX Conference on Networked Systems Design and Implementation, NSDI'12, USENIX Association, Berkeley, USA, 2012, pp. 2–2.
- [2] M. Zaharia, R. S. Xin, P. Wendell, T. Das, M. Armbrust, A. Dave, X. Meng, J. Rosen, S. Venkataraman, M. J. Franklin, A. Ghodsi, J. Gonzalez, S. Shenker, I. Stoica, Apache spark: A unified engine for big data processing, *Commun. ACM* 59 (11) (2016) 56–65.
- [3] H. Karau, A. Konwinski, P. Wendell, M. Zaharia, *Learning Spark: Lightning-Fast Big Data Analysis*, O'Reilly Media, Inc., 2015.
- [4] V. K. Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, B. Saha, C. Curino, O. O'Malley, S. Radia, B. Reed, E. Baldeschwieler, Apache Hadoop YARN: Yet another resource negotiator, in: 4th Annual Symposium on Cloud Computing, SOCC '13, ACM, New York, USA, 2013, pp. 5:1–5:16. [doi:10.1145/2523616.2523633](https://doi.org/10.1145/2523616.2523633).
- [5] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. D. Joseph, R. Katz, S. Shenker, I. Stoica, Mesos: a platform for fine-grained resource sharing in the data center., in: NSDI, Vol. 11, 2011, pp. 22–22.

- [6] M. Armbrust, R. S. Xin, C. Lian, Y. Huai, D. Liu, J. K. Bradley, X. Meng, T. Kaftan, M. J. Franklin, A. Ghodsi, M. Zaharia, Spark sql: Relational data processing in spark, in: Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, SIGMOD '15, ACM, New York, USA, 2015, pp. 1383–1394.
- [7] M. A. Shah, J. M. Hellerstein, S. Chandrasekaran, M. J. Franklin, Flux: An adaptive partitioning operator for continuous query systems, in: 19th International Conference on Data Engineering (ICDE 2003), IEEE Computer Society, 2003, pp. 25–36.
- [8] Y. Wu, K. L. Tan, ChronoStream: Elastic stateful stream computation in the cloud, in: 2015 IEEE 31st International Conference on Data Engineering, 2015, pp. 723–734.
- [9] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker, I. Stoica, Discretized streams: Fault-tolerant streaming computation at scale, in: 24th ACM Symposium on Operating Systems Principles, SOSP '13, ACM, New York, USA, 2013, pp. 423–438.
- [10] X. Meng, J. Bradley, B. Yavuz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen, et al., Mllib: machine learning in apache spark, The Journal of Machine Learning Research 17 (1) (2016) 1235–1241.
- [11] J. E. Gonzalez, R. S. Xin, A. Dave, D. Crankshaw, M. J. Franklin, I. Stoica, Graphx: Graph processing in a distributed dataflow framework., in: OSDI, Vol. 14, 2014, pp. 599–613.
- [12] J. Freeman, N. Vladimirov, T. Kawashima, Y. Mu, N. J. Sofroniew, D. V. Bennett, J. Rosen, C.-T. Yang, L. L. Looger, M. B. Ahrens, Mapping brain activity at scale with cluster computing, Nature Methods 11 (9) (2014) 941–950.
- [13] M. A. Alsheikh, D. Niyato, S. Lin, H.-P. Tan, Z. Han, Mobile big data analytics using deep learning and Apache Spark, IEEE Network 30 (3) (2016) 22–29.
- [14] S. Ryza, U. Laserson, S. Owen, J. Wills, Advanced Analytics with Spark: Patterns for Learning from Data at Scale, O'Reilly Media, Inc., 2017.
- [15] K. Ha, Z. Chen, W. Hu, W. Richter, P. Pillai, M. Satyanarayanan, [Towards wearable cognitive assistance](#), in: 12th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys '14, ACM, New York, USA, 2014, pp. 68–81. [doi:10.1145/2594368.2594383](https://doi.org/10.1145/2594368.2594383).  
URL <http://doi.acm.org/10.1145/2594368.2594383>

- [16] Y. C. Hu, M. Patel, D. Sabella, N. Sprecher, V. Young, Mobile edge computing – a key technology towards 5G, ETSI White Paper 11 (11) (2015) 1–16.