


Data Science with Spark

Shelly Garion

IBM Research -- Haifa



Overview – Advanced Data Analysis Tools

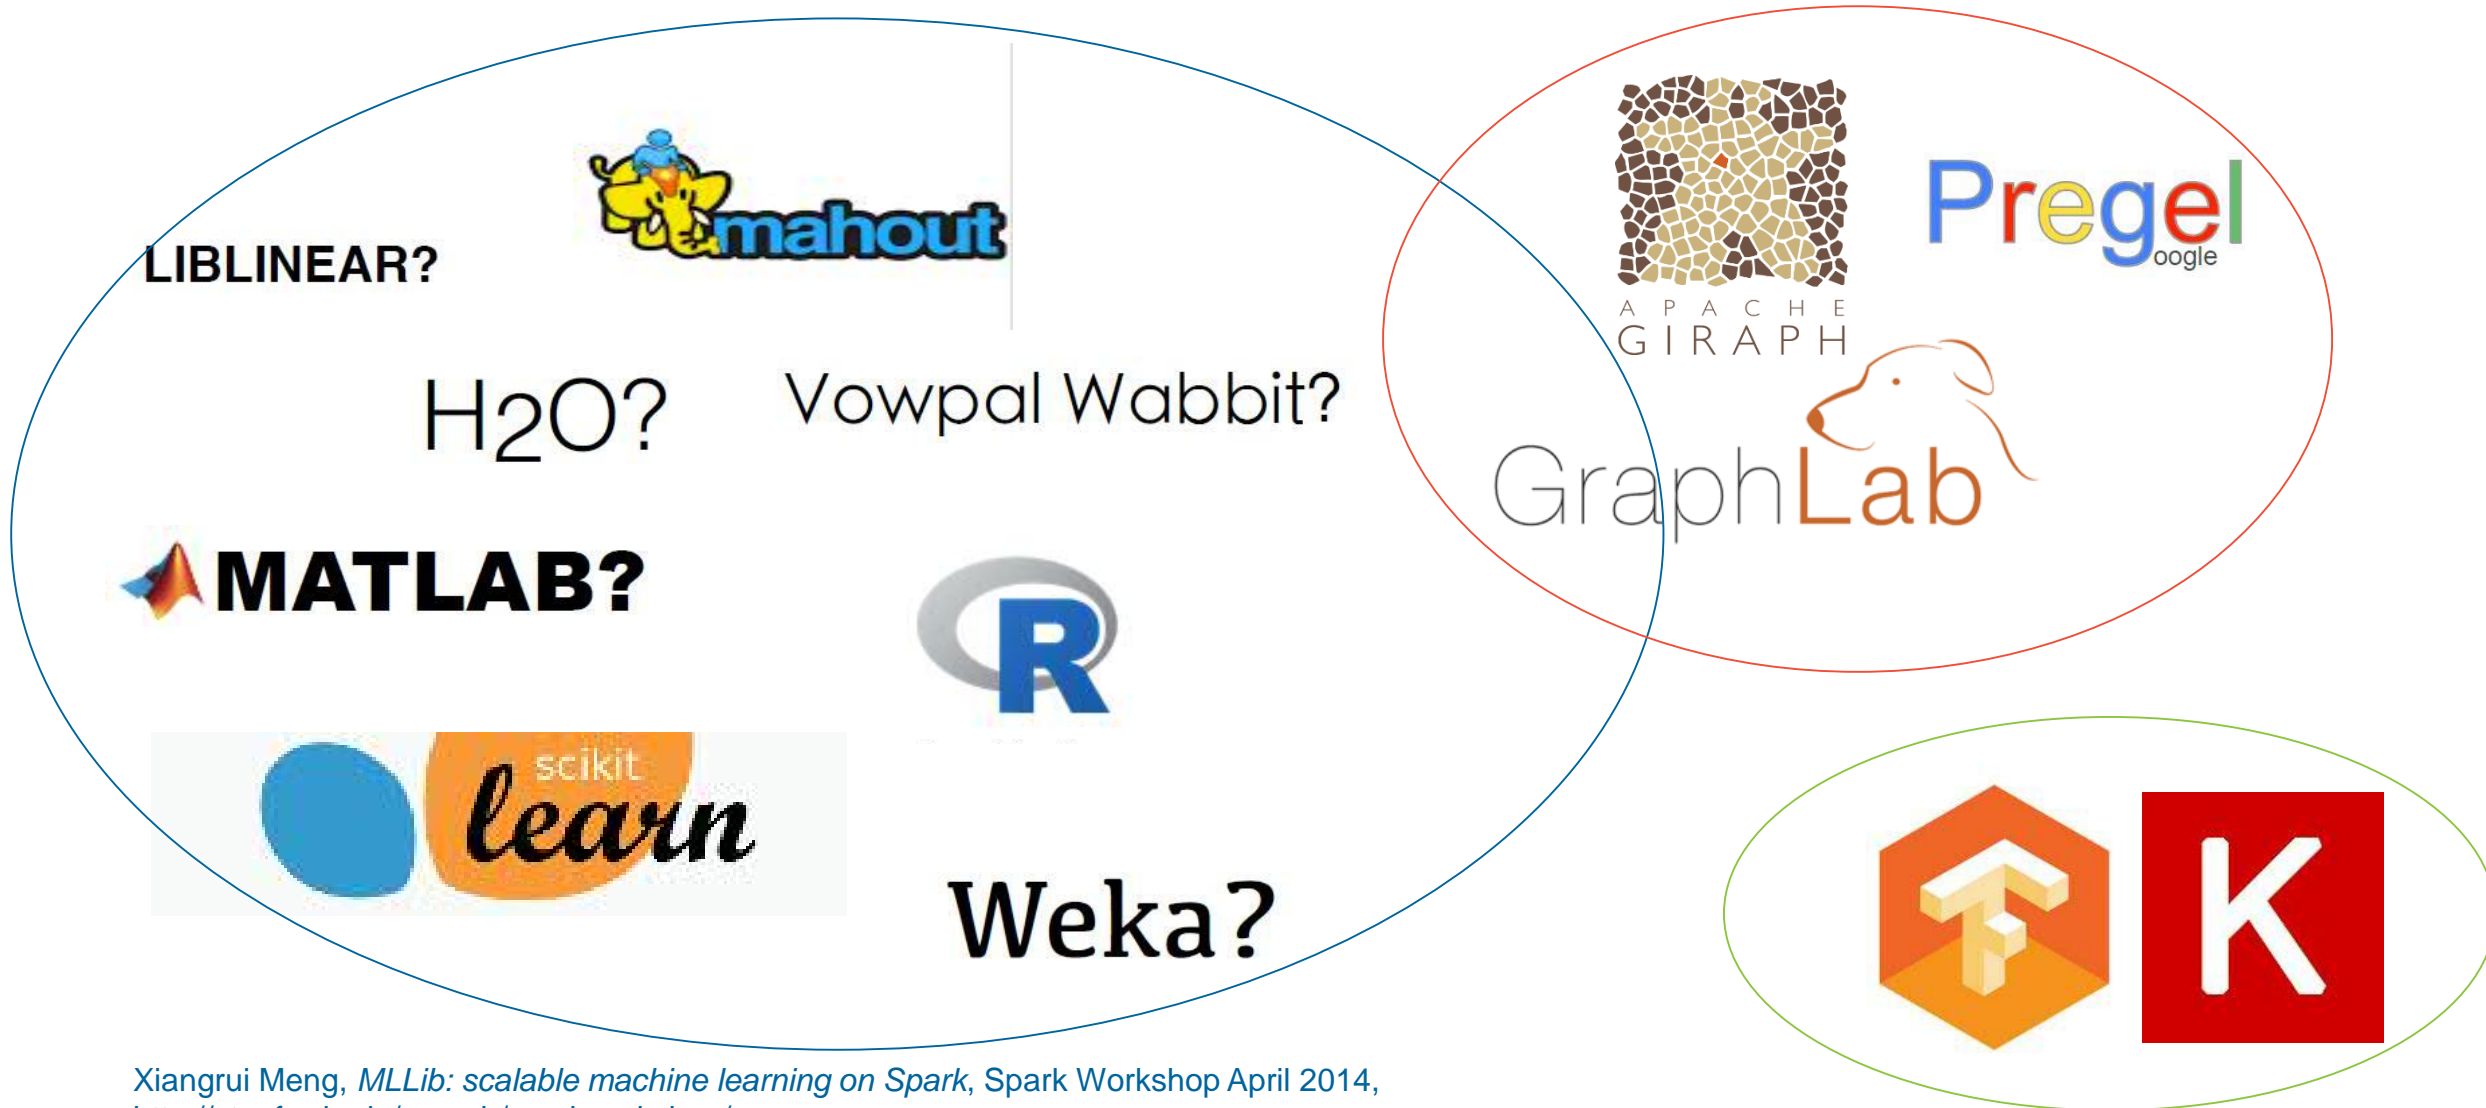


- **Spark MLlib** – large scale machine learning
 - RDD based API
 - DataFrame based API
- **Spark GraphX** – graph-parallel processing

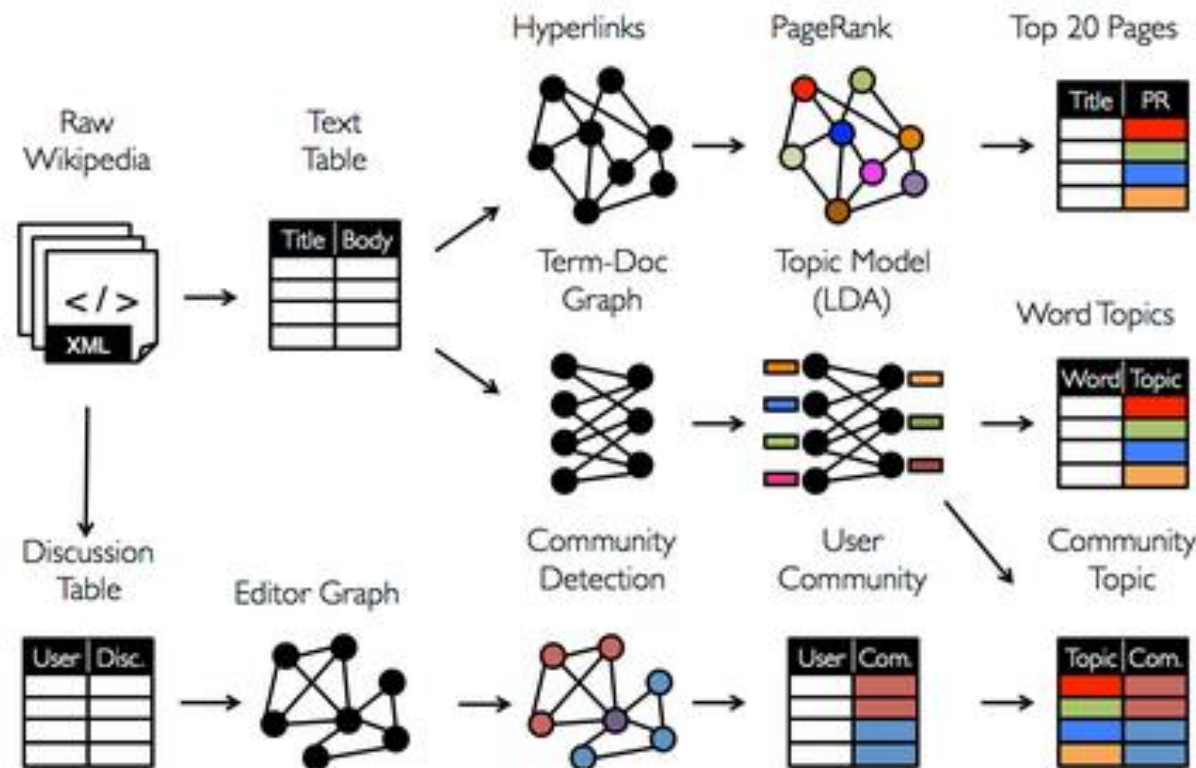
➤ How to clean your data?
➤ How to combine it all?
➤ How to visualize it?



Why Spark MLlib & GraphX?



Combined Analytics of Data



Analyze tabular data with SQL

Analyze graph data using GraphX graph analytics engine

Use same machine learning Infrastructure

Use same solution for streaming data

Machine Learning Algorithms

- Classification

- Logistic regression
- Linear support vector machine (SVM)
- Naïve Bayes
- Decision trees and forests

- Regression

- Generalized linear regression (GLM)

- Recommendation

- Alternating least squares (ALS)

- Clustering

- K-means and Streaming K-means
- Gaussian mixture
- Latent Dirichlet allocation (LDA)

- Dimensionality reduction

- Singular value decomposition (SVD)
- Principal component analysis (PCA)

- Feature extraction & selection

- Word2Vec


- ...

See: <https://spark.apache.org/docs/latest/mllib-guide.html>

Performance of MLlib



- It is built on Apache Spark, a fast and general engine for large-scale data processing.
- Run programs up to 100x faster than Hadoop MapReduce in memory, or 10x faster on disk.

Logistic Regression



Logistic regression in Hadoop and Spark

ALS Results



Performance of MLlib

- Speed-up between MLlib versions




Figure 2: (a) Benchmarking results for ALS. (b) MLlib speedup between versions.

Example: K-Means Clustering (RDD based API)

Goal:

Segment tweets into clusters by geolocation using Spark MLlib K-means clustering

```
1 <longitude>, <latitude>, <timestamp>, <userId>, <tweet message>
2
3 -56.544541,-29.089541,1403918487000,1706271294,Por que ni estamos jugando, son más pajeros e:
4 -69.922686,18.462675,1403918487000,2266363318,Aprenda hablar amigo
5 -118.565107,34.280215,1403918487000,541836358,today a boy told me I'm pretty and he loved me
6 121.039399,14.72272,1403918487000,362868852,@Kringgelss labuyoo. Hahaha
7 -34.875339,-7.158832,1403918487000,285758331,@keithmeneses_ oi td bem? sdds 😊❤️
8 103.766123,1.380696,1403918487000,121042839,Xian Lim on iShine 3 2
```


Example: K-Means Clustering (RDD based API)

To run the k-means algorithm in Spark, we need to first read the csv file

```
1  val sc = new SparkContext("local[4]", "kmeans")
2  // Load and parse the data, we only extract the latitude and longitude of each line
3  val data = sc.textFile(arg)
4  val parsedData = data.map {
5      line =>
6          Vectors.dense(line.split(',').slice(0, 2).map(_.toDouble))
7  }
```

Then we can run the spark kmeans algorithm:

```
1  val iterationCount = 100
2  val clusterCount = 10
3  val model = KMeans.train(parsedData, clusterCount, iterationCount)
```

<https://chimpler.wordpress.com/2014/07/11/segmenting-audience-with-kmeans-and-voronoi-diagram-using-spark-and-mllib/>

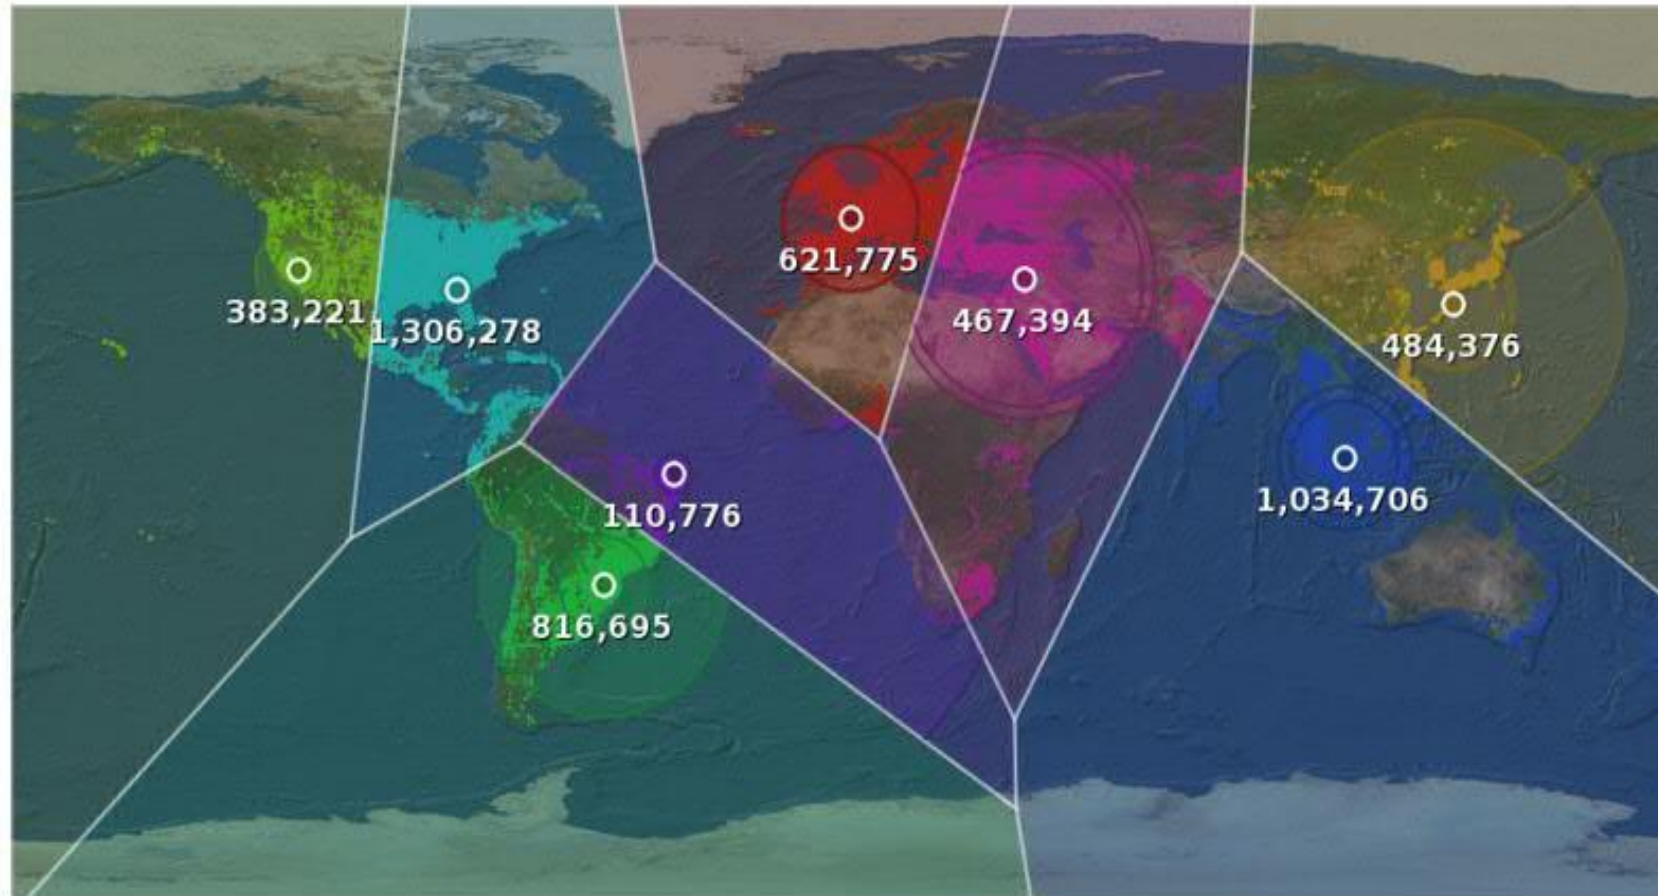
Example: K-Means Clustering (RDD based API)

From the model we can get the cluster centers and group the tweets by cluster:

```
1  val clusterCenters = model.clusterCenters map (_.toArray)
2
3  val cost = model.computeCost(parsedData)
4  println("Cost: " + cost)
5
6  val tweetsByGoup = data
7    .map {_.split(',').slice(0, 2).map(_.toDouble)}
8    .groupBy{rdd => model.predict(Vectors.dense(rdd))}
9    .collect()
10 sc.stop()
```

<https://chimpler.wordpress.com/2014/07/11/segmenting-audience-with-kmeans-and-voronoi-diagram-using-spark-and-mllib/>

Example: K-Means Clustering (RDD based API)



<https://chimpler.wordpress.com/2014/07/11/segmenting-audience-with-kmeans-and-voronoi-diagram-using-spark-and-mllib/>

Machine Learning Pipeline with Spark MLlib

Data pre-processing


Feature extraction

Model fitting

Model training

Validation

Model prediction




Spark MLlib Pipeline (DataFrame based API)

```
// create pipeline
tok = Tokenizer(in="text", out="words")
tf = HashingTF(in="words", out="features")
lr = LogisticRegression(maxIter=10, regParam=0.01)
pipeline = Pipeline(stages=[tok, tf, lr])
```

```
// train pipeline
df = sqlCtx.table("training")
model = pipeline.fit(df)

// make predictions
df = sqlCtx.read.json("/path/to/test")
model.transform(df)
  .select("id", "text", "prediction")
```



Spark MLlib Pipeline (DataFrame based API)

- **DataFrame:**

- Use DataFrame from Spark SQL as ML dataset
- Can have different columns storing text, feature vectors, true labels, and predictions

- **Transformer:**

- A Transformer implements a method `transform()`
- Algorithm that transforms one DataFrame to another DataFrame
 - Feature transformers (e.g., OneHotEncoder)
 - Trained ML models (e.g., LogisticRegressionModel)

- **Estimator:**

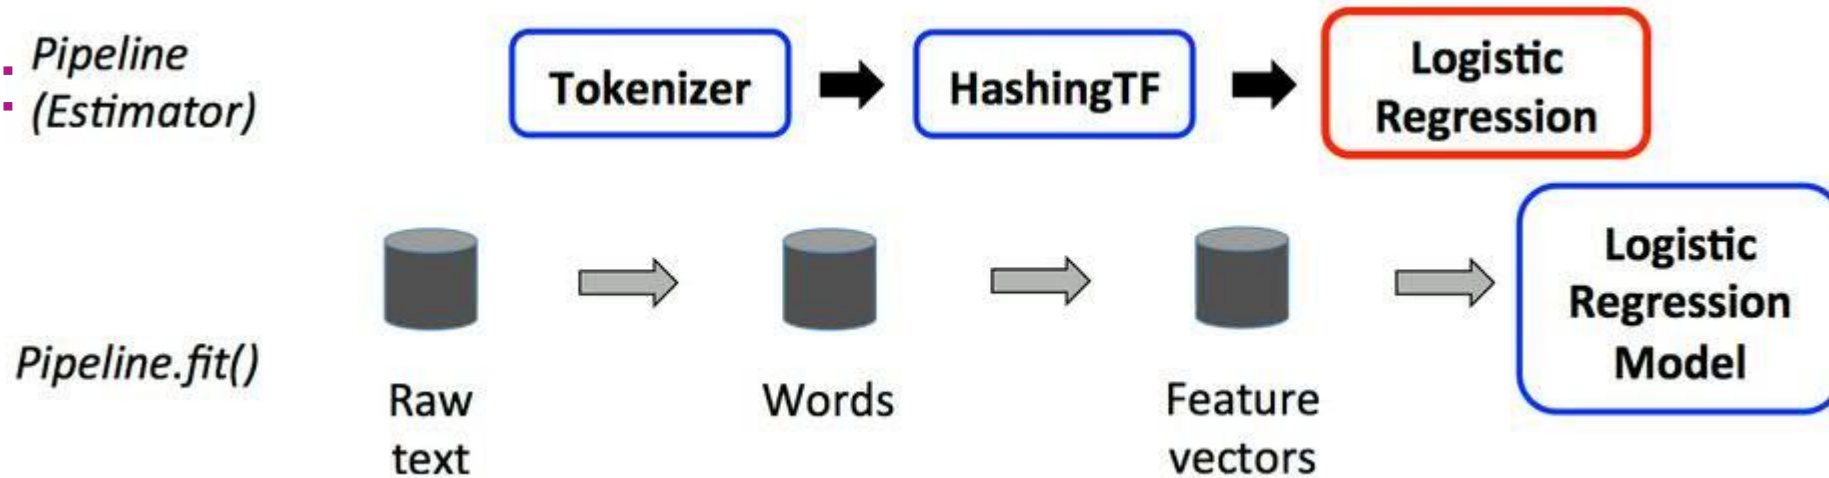
- An Estimator implements a method `fit()`
- Algorithm which can be fit on a DataFrame to produce a transformer
 - ML algorithms which trains on a DataFrame and produces a model (e.g., LogisticRegression)

- **Pipeline:**


- Chains multiple Transformers and Estimators together to specify an ML workflow

Machine Learning Pipeline with Spark MLlib

Learning: *Pipeline (Estimator)*





Model: *PipelineModel (Transformer)*



Spark GraphX

Key idea

- Graphs are essential to analytics (e.g. social networks)
- Tables & Graphs are composable views of the same physical data
- Each view has its own operators that exploit the semantics of the view to achieve efficient execution
- Graph algorithms are based on Pregel API




Fewer Triangles
Weaker Community




More Triangles
Stronger Community

Spark GraphX

Main components

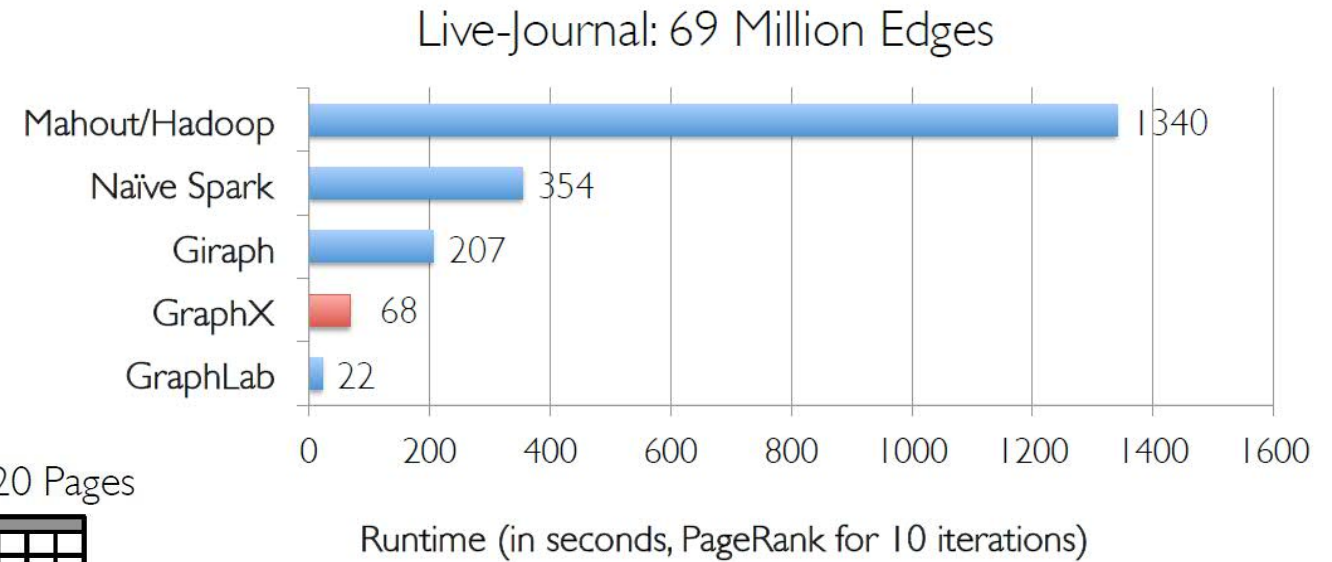



- **VertexRDD** maps IDs to vertex content
- **EdgeRDD** are of the form (ID1, ID2, ET)
- **Triplets** are a combination of Vertex & Edge RDDs



```
def Graph(vertices: Table[ (Id, V) ],
          edges: Table[ (Id, Id, E) ])
// Table Views -----
def vertices: Table[ (Id, V) ]
def edges: Table[ (Id, Id, E) ]
def triplets: Table [ ((Id, V), (Id, V), E)]
```


Performance of GraphX



Example - PageRank




Popular algorithm originally introduced by Google



Sergei Brin and Lawrence Page, *"The anatomy of a large-scale hypertextual Web search engine"*, Computer Networks and ISDN Systems. (1998) 30: 107–117.

Example - PageRank

Popular algorithm originally introduced by Google



PageRank Algorithm

- Start each page with a rank of 1
- On each iteration:


$$A. \text{ contrib} = \frac{\text{curRank}}{|\text{neighbors}|}$$

$$B. \text{ curRank} = 0.15 + 0.85 \sum_{\text{neighbors}} \text{contrib}_i$$

Sergei Brin and Lawrence Page, ["The anatomy of a large-scale hypertextual Web search engine"](#), Computer Networks and ISDN Systems. (1998) 30: 107–117.

Example: PageRank Spark GraphX

Popular algorithm originally introduced by Google



```
// get people with top-k pageranks
def findTopPageRank(allPeople: RDD[String], links: RDD[(String, String, Double)], k: Int) = {
  val versRDD = allPeople.map(p => (uid(p), p))
  val edgesRDD = links.map{ case (l, r, score) => Edge(uid(l), uid(r), score) }


  val g = Graph(versRDD, edgesRDD).cache
  val ranks = g.pageRank(0.001)

  ranks.vertices.top(k)(Ordering.by( . 2)).map(p => (fromUid(p._1), p._2))
}
```

Example: PageRank

How is it implemented in Pregel?

Popular algorithm originally introduced by Google



```
def PageRank(v: Id, msgs: List[Double]) {  
  // Compute the message sum  
  var msgSum = 0  
  for (m <- msgs) { msgSum += m }  
  // Update the PageRank  
  PR(v) = 0.15 + 0.85 * msgSum  
  // Broadcast messages with new PR  
  for (j <- OutNbrs(v)) {  
    msg = PR(v) / NumLinks(v)  
    send_msg(to=j, msg)  
  }  
  // Check for termination  
  if (converged(PR(v))) voteToHalt(v)  
}
```

Reynold S. Xin, Daniel Crankshaw, Ankur Dave, Joseph E. Gonzalez, Michael J. Franklin, Ion Stoica.
[GraphX: Unifying Data-Parallel and Graph-Parallel Analytics. OSDI 2014.](#) October 2014.