

www.twosigma.com

Huohua 火花 Distributed Time Series Analysis Framework For Spark

Wenbo Zhao

STAC Summit 2016



- Software Engineer @
 TWO SIGMA
- Focus on analytics related tools, libraries and Systems



We view everything as a time series

- Stock market prices
- Temperatures
- Sensor logs

٠

...

Presidential polls



What is a time series?

- A sequence of observations obtained in successive time order
- Our goal is to forecast future values given past observations





Multivariate time series

- We can forecast better by joining multiple time series
- *Temporal join* is a fundamental operation for time series analysis
- Huohua enables fast distributed temporal join of large scale unaligned time series





What is temporal join?

- A particular *join* function defined by a matching criteria over *time*
- Examples of criteria
 - look-backward find the most recent observation in the past
 - look-forward find the closest observation in the future





time	weather			
08:00 AM	60 °F 🦛			
10:00 AM	70 °F 金 🔪	time	weather	corn price
12:00 AM	80 °F (08:00 AM		
		10:00 AM		
time	corn price	12:00 AM		
08:00 AM				
11:00 AM				

















💩 TWO SIGMA

- A single time series may not fit into a single machine
- Forecasting may involve hundreds of time series
- Existing packages don't support temporal join or can't handle large time series
 - MatLab, R, SAS, Pandas
 - Even Spark based solutions fall short
 - PairRDDFunctions, DataFrame/Dataset, spark-ts

Huohua – a new time series library for Spark

- Goal
 - provide a collection of functions to *manipulate* and *analyze* time series at scale
 - group, temporal join, summarize, aggregate ...
- How
 - build a time series aware data structure
 - extending RDD to TimeSeriesRDD
 - optimize using temporal locality
 - reduce shuffling
 - reduce memory pressure by streaming

What is a TimeSeriesRDD in Huohua?

- TimeSeriesRDD extends RDD to represent time series data
 - associates a time range to each partition
 - tracks partitions' time-ranges through operations
 - preserves the temporal order





TimeSeriesRDD- an RDD representing time series

time	temperature	RDD
6:00 AM	60°F	(6:00 AM, 60°F) (6:01 AM, 61°F)
6:01 AM	61°F	
7:00 AM	70°F	(7:00 AM, 70°F) (7:01 AM, 71°F)
7:01 AM	71°F	
8:00 AM	80°F	(8:00 AM, 80°F)
8:01 AM	81°F	(8.01 An; 811)



TimeSeriesRDD- an RDD representing time series







• A group function groups rows with exactly the same timestamps





• A group function groups rows with nearby timestamps







• Groups rows with exactly the same timestamps







• Data is shuffled and materialized







• Data is shuffled and materialized







• Data is shuffled and materialized







Temporal order is not preserved







• Another sort is required







• Another sort is required







Back to correct temporal order







• Back to temporal order



















Benchmark for group

- Running time of *count* after *group*
 - 16 executors (10G memory and 4 cores per executor)
 - data is read from HDFS





- A *temporal join* function is defined by a matching criteria over *time*
- A typical matching criteria has two parameters
 - direction whether it should look-backward or look-forward
 - window how much it should look-backward or look-forward







- A *temporal join* function is defined by a matching criteria over *time*
- A typical matching criteria has two parameters
 - direction whether it should look-backward or look-forward
 - window how much it should look-backward or look-forward







Temporal join with criteria look-back and window of length 1







- Temporal join with criteria look-back and window of length 1
 - How do we do temporal join in TimeSeriesRDD?





- Temporal join with criteria look-back and window of length 1
 - partition *time* space into disjoint intervals





- Temporal join with criteria look-back and window of length 1
 - Build dependency graph for the joined TimeSeriesRDD





- Temporal join with criteria look-back and window 1
 - Join data as streams per partition





- Temporal join with criteria look-back and window 1
 - Join data as streams





- Temporal join with criteria look-back and window 1
 - Join data as streams





- Temporal join with criteria look-back and window 1
 - Join data as streams





Benchmark for temporal join

- Running time of *count* after *temporal join*
 - 16 executors (10G memory and 4 cores per executor)
 - data is read from HDFS



Functions over TimeSeriesRDD

- group functions such as window, intervalization etc.
- temporal joins such as look-forward, look-backward etc.
- summarizers such as average, variance, z-score etc. over
 - windows
 - Intervals
 - cycles





- Not quite yet ...
- https://github.com/twosigma



Future work

- Dataframe / Dataset integration
 - Speed up
 - Richer APIs
- Python bindings
- More summarizers



Key contributors

- Christopher Aycock
- Jonathan Coveney
- Jin Li
- David Medina
- David Palaitis
- Ris Sawyer
- Leif Walsh
- Wenbo Zhao





