Geophysical Research Abstracts
Vol. 19, EGU2017-11755, 2017
EGU General Assembly 2017

SciSpark: In-Memory Map-Reduce for Earth Science Algorithms

Brian Wilson, Rahul Palamuttam, Kim Whitehall, Chris Mattmann, Sujen Shah, Alexander Goodman, Paul Ramirez, and Wayne Burke

JET PROPULSION LABORATORY, PASADENA, United States (Brian.Wilson@jpl.nasa.gov)

We are developing a lightning fast Big Data technology called SciSpark based on Apache Spark under a NASA AIST grant (PI Mattmann). Spark implements the map-reduce paradigm for parallel computing on a cluster, but emphasizes in-memory computation, “spilling” to disk (SSD’s preferred) only as needed, and so outperforms the disk-based Apache Hadoop by 100 to 1000x in memory and by 10x on disk.

SciSpark extends Spark to support Earth Science use in three ways:
• Efficient ingest of N-dimensional geo-located arrays (physical variables) with attributes from netCDF3/4, HDF4/5, and/or OPeNDAP URLS;
• Array operations for dense arrays in scala and Java using the ND4S/ND4J or Breeze libraries, and in python using numpy and scipy;
• Operations to “split” datasets across a Spark cluster by time or space or both.

For example, a decade-long time-series of geo-variables can be split across time to enable parallel “speedups” of analysis by day, month, or season. Similarly, very high-resolution climate grids can be partitioned into spatial tiles for parallel operations across rows, columns, or blocks. In addition, using Spark’s gateway into python, PySpark, one can utilize the entire ecosystem of numpy, scipy, etc. Finally, SciSpark Notebooks provide a modern eNotebook technology in which scala, python, or spark-sql codes are entered into cells in the Notebook and executed on the cluster, with results, plots, or graph visualizations displayed in “live widgets”. A three-session SciSpark Tutorial was taught in July 2016 using a set of 15 example Notebooks illustrating basic use and advanced science cases in both python and scala.

We have exercised SciSpark by implementing three complex Use Cases: discovery and evolution of Mesoscale Convective Complexes (MCCs) in storms, yielding a graph of connected components; PDF Clustering of atmospheric state using parallel K-Means; and statistical “rollups” of geo-variables or model-to-obs. differences (i.e. mean, stddev, skewness, & kurtosis) by month, season, year, and multi-year at interactive speeds (e.g. 15-year statistical rollup of TRMM precipitation in 40 seconds). Geo-variables are ingested and split across the cluster using methods on the sciSparkContext object to create a scientific sRDD (resilient distributed dataset), a partitioned series of sciDataset objects in cluster memory. Each sciDataset contains one or more physical variables, with the coordinates and attributes needed to implement geophysical algorithms such as subsetting, quality masking, regridding, and statistical comparisons.

The presentation will cover the architecture of SciSpark, the design of the scientific RDD (sRDD) data structures for N-dim. arrays, results from the three science Use Cases, example Notebooks, lessons learned from the algorithm implementations, parallel performance metrics, and coming enhancements.