



Contribution ID: 364

Type: Poster

In-Memory Parallel Computing for Partial Wave Analysis

The traditional partial wave analysis (PWA) algorithm is designed to process data serially which requires a large amount of memory that may exceed the memory capacity of one single node to store runtime data. It is quite necessary to parallelize this algorithm in a distributed data computing framework to improve its performance. Within an existing production-level Hadoop cluster, we implement PWA algorithm on the basis of Spark to process data storing on low-level storage system HDFS. But in this case, sharing data through HDFS or internal data communication mechanism of Spark is extremely inefficient. In order to solve this problem, this paper presents an in-memory parallel computing method for PWA algorithm. With this system, we can easily sharing runtime data in parallel algorithms. We can ensure complete data locality to keep compatibility with the traditional data input/output way and cache most repeated used data in memory to improve the performance, owe to the data management mechanism of Alluxio.

Primary author: Mr WEI, Zhanchen (IHEP, CAS)

Co-authors: Dr HUANG, Qiulan (IHEP, CAS); SUN, Gongxing (INSTITUTE OF HIGH ENERGY PHYSICS); Ms LIU, Xiaoyu (IHEP, CAS)

Presenter: Mr WEI, Zhanchen (IHEP, CAS)

Session Classification: Poster Session

Track Classification: Track 1: Computing Technology for Physics Research