



Contribution ID: 558

Type: **Poster**

## TMPIFile: A New Parallel I/O Solution in ROOT

*Thursday 7 November 2019 16:15 (15 minutes)*

Communication among processes is generating considerable interest in the scientific computing community due to the increasing use of distributed memory systems. In the field of high energy physics (HEP), however, little research has been addressed on this topic. More precisely in ROOT I/O, the de facto standard for data persistence in HEP applications, no such feature is provided. In order to perform efficient and robust cross-node communications, we introduce the TMPIFile functionality into ROOT where Message Passing Interface (MPI) is used to pass data across the entire distributed system. In the case of ATLAS workflows, instead of writing to file, the compressed data for each node is now stored into the sender side of TMPIFile. After a certain amount of data is collected, the TMPIFile sender can automatically package the data into a memory buffer and send it via MPI. On the other end of the communication, a collector receives buffers from multiple senders, merges them into one file, and at last, writes the file into the disk. Multiple collectors can be created to avoid I/O contention. Test results will be shown from runs at NERSC and ALCF.

### Consider for promotion

No

**Authors:** WANG, Yunsong (Lawrence Berkeley National Lab. (US)); Mr BASHYAL, Amit (Oregon State University); CHILDERS, Taylor (Argonne National Laboratory (US)); VAN GEMMEREN, Peter (Argonne National Laboratory (US)); TSULAIA, Vakho (Lawrence Berkeley National Lab. (US)); CANAL, Philippe (Fermi National Accelerator Lab. (US)); FARRELL, Steven (Lawrence Berkeley National Lab (US))

**Presenter:** FARRELL, Steven (Lawrence Berkeley National Lab (US))

**Session Classification:** Posters

**Track Classification:** Track 5 –Software Development