

# 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 140

Type: oral presentation

## Fine grained event processing on HPCs with the ATLAS Yoda system

*Thursday, April 16, 2015 11:00 AM (15 minutes)*

High performance computing facilities present unique challenges and opportunities for HENP event processing. The massive scale of many HPC systems means that fractionally small utilizations can yield large returns in processing throughput. Parallel applications which can dynamically and efficiently fill any scheduling opportunities the resource presents benefit both the facility (maximal utilization) and the (compute-limited) science. The ATLAS Yoda system provides this capability to HENP-like event processing applications by implementing event-level processing in an MPI-based master-client model that integrates seamlessly with the more broadly scoped ATLAS Event Service. Fine grained, event level work assignments are intelligently dispatched to parallel workers to sustain full utilization on all cores, with outputs streamed off to destination object stores in near real time with similarly fine granularity, such that processing can proceed until termination with full utilization. The system offers the efficiency and scheduling flexibility of preemption without requiring the application actually support or employ checkpointing. We will present the new Yoda system, its motivations, architecture, implementation, and applications in ATLAS data processing at several US HPC centers.

**Primary author:** CALAFIURA, Paolo (Lawrence Berkeley National Lab. (US))

**Co-authors:** OLEYNIK, Danila (Joint Inst. for Nuclear Research (RU)); NILSSON, Paul (Brookhaven National Laboratory (US)); Dr VAN GEMMEREN, Peter (Argonne National Laboratory (US)); PANITKIN, Sergey (Brookhaven National Laboratory (US)); MAENO, Tadashi (Brookhaven National Laboratory (US)); Dr WENAUS, Torre (Brookhaven National Laboratory (US)); TSULAIA, Vakho (Lawrence Berkeley National Lab. (US)); GUAN, Wen (University of Wisconsin (US))

**Presenter:** TSULAIA, Vakho (Lawrence Berkeley National Lab. (US))

**Session Classification:** Track 8 Session

**Track Classification:** Track8: Performance increase and optimization exploiting hardware features