



Contribution ID: 238

Type: Poster

Simulating the ATLAS Distributed Computing Infrastructure to Optimize Workload Allocation Strategies

In large-scale distributed computing systems, workload dispatching and the associated data management are critical factors that determine key metrics such as resource utilization of distributed computing and resilience of scientific workflows. As the Large Hadron Collider (LHC) advances into its high luminosity era, the ATLAS distributed computing infrastructure must improve these metrics to manage exponentially larger data volumes (exceeding ExaBytes) and support the demanding needs of high-energy physics research.

To improve the distributed computing operation, the existing workload allocation strategies can be optimized, or novel strategies can be designed. However, in practice, it is not viable to test new workload allocation strategies on the actual ATLAS distributed computing. To address this, we have developed an agile simulation framework of the ATLAS distributed computing system using the SimGrid toolkit to evaluate and refine workload dispatching strategies for heterogeneous computing infrastructure. Moreover, it is crucial to also address the inherent overhead and potential bottlenecks associated with the management of the large data volumes required by these workloads. Therefore, we extensively analyze the historical remote transfers to understand the root causes of slowdowns, vulnerabilities, and inefficient resource utilization linked to data movement.

To ensure the accuracy and reliability of our framework, we calibrate and validate the ATLAS distributed computing implementation in the simulation framework by testing real workloads from historical ATLAS data. This calibrated simulation framework will serve as the testbed for evaluating custom allocation algorithms and also generate the datasets required to train ML surrogates to enable fast and scalable simulations. In addition, an interactive monitoring interface is being developed to visualize the workload dispatching and the resource utilization. Apart from serving as a platform for testing and executing new strategies that can improve the resilience of the ATLAS distributed computing, our framework is ultimately experiment-agnostic and open sourced, providing an example case that can enable users to configure large-scale distributed computing grids and implement custom workload allocation algorithms through dynamic plugins.

Significance

References

Experiment context, if any

Authors: HSU, Kuan-Chieh (Brookhaven National Laboratory (US)); KHAN, Raees Ahmad (University of Pitts-

burgh (US)); SRI VATSAVAI, Sairam (Brookhaven National Laboratory (US))

Co-authors: HOISIE, Adolfy (Brookhaven National Laboratory (US)); KLIMENTOV, Alexei (Brookhaven National Laboratory (US)); PARK, David (Brookhaven National Laboratory); AKMAN, Fatih Furkan (University of Massachusetts (US)); SUTER, Frederic; Mr KIM, Jaehyung (Carnegie Mellon University); STEELE, John Rembrandt (University of Massachusetts (US)); BOUDREAU, Joseph (University of Pittsburgh); PODHORSZKI, Norbert (Oak Ridge National Laboratory); KILIC, Ozgur Ozan (Brookhaven National Laboratory); NILSSON, Paul (Brookhaven National Laboratory (US)); REN, Ray (Brookhaven National Laboratory (US)); DUTTA, Sankha Baran (Brookhaven National Laboratory (US)); KLASKY, Scott; FENG, Shengyu; YOO, Shinjae; MAENO, Tadashi (Brookhaven National Laboratory (US)); KORCHUGANOVA, Tania; Dr CHOWDHURY, Tasnuva (Brookhaven National Laboratory (US)); MARTINEZ OUTSCHOORN, Verena Ingrid (University of Massachusetts (US)); YANG, Wei (SLAC National Accelerator Laboratory (US)); YANG, Yiming

Presenter: KHAN, Raees Ahmad (University of Pittsburgh (US))

Session Classification: Poster session with coffee break

Track Classification: Track 1: Computing Technology for Physics Research