T 44: Data, AI, Computing 4 (workflow)

Time: Tuesday 16:00–18:00 Location: Geb. 30.34: LTI

T 44.1 Tue 16:00 Geb. 30.34: LTI

Testing ATLAS computing resources with HammerCloud — •Alexander Mario Lory[1], Günter Duckeck[1], Benjamin Rottler[2], Michael Böhler[2], and Otmar Biebel[1] — [1]Ludwig-Maximilians-Universität, München — [2]Albert-Ludwigs-Universität, Freiburg

HammerCloud is a framework for testing and benchmarking distributed resources of the LHC experiments ATLAS and CMS. It is a key component for the operation of the ATLAS distributed computing system, as the test results are used to automatically include or exclude resources from the pool of resources available to users. In this presentation recent use cases and developments within HammerCloud ATLAS are highlighted, including a new feature – the automatic massive recovery mechanism, which efficiently recovers GRID resources on a large scale, addressing massive exclusions resulting from central issues.

T 44.2 Tue 16:15 Geb. 30.34: LTI

Integration of the Goettingen HPC resources Emmy to the WLCG Tier-2 grid computing environment of GoeGrid and performance results of the ATLAS jobs — •Saidev Polisetty, Arnulf Quadt, Daniel Schindler, and Sebastian Wozniewski — II. Physikalisches Institut, Georg-August-Universitaet Goettingen

For the upcoming Run 4 of the LHC, there is a necessity to increase the computing resources for simulation, reconstruction and analysis in terms of storage and computing power. In this context, the German community is preparing to integrate the National High Performance Computing (NHR) resources and make them usable within the WLCG under the FIDIUM project.

At Goettingen campus, there is both, a WLCG Tier-2 site (GoeGrid) and a large HPC EMMY cluster by NHR and the North German Supercomputing Alliance (HLRN). In this project, the integration is done by virtually extending the GoeGrid batch system with containers turning the HPC nodes into virtual nodes with own partitionable job scheduling to run the ATLAS jobs. Their performance is studied for quantitative analysis and optimisation of the environment running the jobs. Results and the performance is discussed in this talk.

T 44.3 Tue 16:30 Geb. 30.34: LTI

The GPU driven journey towards more sustainable HEP computing — Manuel Giffels, Matthias Schnepf, Günter Quast, and •Tim Voigtländer — Karlsruhe Institute of Technology, Karlsruhe, Germany

Energy efficient usage of our hardware resources has become a topic of ever-increasing importance. Both, high operational costs and environmental concerns are good reasons for us to optimize the way we use hardware. In comparison with CPUs, GPUs are a convenient alternative that has the possibility to achieve a significantly higher energy efficiency in many of the HEP relevant workflows, like simulation, reconstruction, or machine learning based end user analyses. This talk aims to shed some light on how efficient such applications run on different hardware, using the example of a number of GPU benchmarks that have been performed on the Karlsruhe Tier 3 cluster TOpAS.

T 44.4 Tue 16:45 Geb. 30.34: LTI

Parallelization and benchmarking of a Jupyter based HEP data analysis with Dask — •Karl Erik Bode, Michael Böhler, and Markus Schumacher — Institute of Physics, Albert-Ludwigs-University Freiburg, Freiburg, Germany

Using the combination of the scientific Python software stack, Jupyter notebooks, and Dask it is possible to scale an interactive HEP data analysis, both on local resources and on a computing cluster.

After vectorisation of the reference Higgs boson to di-photon decay analysis, the required compute time is decreased and it is even possible to analyze larger data sets.

With Dask, the vectorized algorithm can be scaled to utilize all CPU cores of the local machine and at the same time provides data structures to enable analysis of data sets larger than memory.

Only minor changes are required, to port this analysis setup from a laptop to a to a High Throughput Cluster (HTC) or to a High Performance Cluster (HPC).

This contribution introduces the used software stack, specific for scaling the algorithm from single threaded to a mult threaded analysis. Finally we discuss the performance improvement both on a typical laptop as well as on an HPC and HPC cluster.

T 44.5 Tue 17:00 Geb. 30.34: LTI

Orchestrated columnar-based analysis with Columnflow — •Mathis Frahm, Johannes Haller, Philip Keicher, Nathan Provoust, Marcel Rieger, Daniel Savoiu, Peter Schleper, Matthias Schröder, and Bogdan Wiederspan — Institut für Experimentalphysik, Universität Hamburg

The large datasets and increasing complexity of modern physics analysis in high energy collider physics pose a major challenge to the analysis workflows. Systems are required that can process large scales of data efficiently, while keeping the execution of the full analysis manageable. In this talk, we present Columnflow, a tool for columnar-based data analysis. Columnflow provides an orchestrated, yet flexible workflow that handles the bookkeeping of results and dependencies automatically. Typical analysis tasks such as propagating systematic uncertainties, machine-learning applications, and statistical inference are transparently included in the workflow. The implemented workflow allows the usage of distributed computing resources and is fully configurable but at the same time accessible to newcomers.

T 44.6 Tue 17:15 Geb. 30.34: LTI

Workflow Management with Snakemake: A Case Study in Tau Physics at the ATLAS Detector — •Luka Vomberg[1], Christian Grefe[2], Philip Bechtle[1], and Klaus Desch[1] — [1]Physikalisches Institut, 53115 Bonn — [2]CERN

Data analyses in particle physics often rely on complex software workflows. Reproducing and reusing such analyses is frequently challenging, as it is difficult to keep track of the many interdependent individual steps. This presentation illustrates, through a case study, how the workflow management tool Snakemake contributes to declaratively uniting all necessary steps into a comprehensive package.

The case study focuses on an analysis measuring the TauID efficiencies in Run 3 of the LHC at the ATLAS detector, which uses the ABCD method to estimate the contribution from misidentified tauleptons. Various ATLAS software packages are integrated into the overall workflow, initially structured to fully reproduce the results of the equivalent Run 2 analysis and then adapted to the requirements of Run 3.

T 44.7 Tue 17:30 Geb. 30.34: LTI

User-oriented sustainable operation of the VISPA Cloud Services — Niclas Eich, Johannes Erdmann, Martin Erdmann, Benjamin Fischer, Paul Gilles, •Tim Hauptreif, and Jan Kelleter — RWTH Aachen University

As the effects of climate change become more imminent, scientists from all disciplines are striving to make their research more sustainable. As computing clusters are an essential part of modern research, they need to be improved in terms of sustainability. Various measures to reduce the carbon footprint of our computing centres include monitoring power consumption, increasing the efficiency of hardware and software, and influencing user behaviour.

We present our efforts to increase the sustainability of the VISPA computing cluster (https://vispa.physik.rwth-aachen.de) by introducing resource-aware scheduling and leveraging user interaction. Maintaining the relatively small VISPA cluster puts us in a unique position where we can freely test different energy saving schedules and get direct feedback from academic researchers and students alike. While scheduling is an ongoing research topic for many, especially large computing clusters, few can adapt directly to user feedback. We are sharing our experience in providing more sustainable computing resources with minimal disruption to our users' research progress.

T 44.8 Tue 17:45 Geb. 30.34: LTI

Search for Hidden Job Failure Risk Factors in ATLAS Job Meta Data — Arnulf Quadt[1], Sebastian Wozniewski[2], and •Kia-Jüng Yang[3] — [1]II. Physikalisches Institut, Georg-August-Universität Göttingen — [2]II. Physikalisches Institut, Georg-August-Universität Göttingen — [3]II. Physikalisches Institut, Georg-August-Universität Göttingen

The ATLAS Detector records over 10,000 TB of data per year and it increases even further with the upcoming upgrades. The Worldwide LHC Computing Grid (WLCG) provides a distributed computing infrastructure to store and process these data. It is crucial, that the WLCG is also reliable, meaning that the failure rate of submitted jobs by the users is low. While some job failures can be clearly traced back to known temporary issues, others seem to happen more randomly due to various more hidden reasons. An investigation of the job failure rates depending on the job attributes may reveal correlations, which might allow for mitigating actions in order to further reduce the number of job failures. This task is supported by the training of a neural network, which helps to identify and investigate correlations in the multi-dimensional space of job attributes.