T 14: Data, AI, Computing, Electronics II

Time: Monday 16:15–18:15 Location: KH 02.014

T 14.1 Mon 16:15 KH 02.014

Towards a sustainable ET Computing Center — •Stefan Krischer and Achim Stahl — RWTH Aachen University, III. Physikalisches Institut B

For the Einstein Telescope, a dedicated computing center will be required to support detector control, low-latency data processing, and parts of the offline analysis. In preparation, the infrastructure is being explored with a strong focus on sustainability. To assess feasibility, a prototype computing facility is being developed with its own renewable energy system capable of supplying 100% locally generated power. In this setup, computing workloads and hardware adapt their power consumption to the instantaneous energy availability, reducing the required energy-storage capacity. We present the concept and initial insights from the prototype development as a basis for a more sustainable computing center for the Einstein Telescope.

T 14.2 Mon 16:30 KH 02.014

Recent Developments in HEP Computing in Karlsruhe — Giacomo De Pietro, Patrick Ecker, Nils Faltermann, Emelie Fuchs, Johannes Gaessler, Manuel Giffels, Artur Gottmann, Jan Kieseler, Max Kühn, Yannis Kluegl, Günter Quast, Raquel Quishpe, Matthias Schnepf, Nikita Shadskiy, Lars Sowa, •Tim Voigtlaender, and Ralf Florian Von Cube — Karlsruhe Institute of Technology (KIT)

The computing contribution of the Karlsruhe Institute of Technology (KIT) to the WLCG pledge is provided through the Tier-1 center GridKa, together with opportunistic resources from the HoreKa HPC and the local Tier-3 cluster. These resources, together with the development activities of the Karlsruhe computing group, form a tightly integrated component that addresses a broad spectrum of high-energy physics (HEP) workloads and operational needs.

This talk outlines the current status and ongoing developments in HEP computing activities at KIT. This includes the status of the GPU integration into the WLCG from HoreKa and the Karlsruhe Tier-3 cluster, highlighting the first official workflows by LHC experiments on these systems. These early tests provide a baseline for future large-scale GPU-accelerated workflows. In addition, the ongoing effort to sustainably use legacy hardware during periods of renewable-energy surplus is presented as part of a broader move towards an energy-adaptive computing strategy. Finally, planned storage and compute extensions at GridKa Tier-1 center are presented, aimed at meeting future requirements and ensuring compliance with pledged resources.

T 14.3 Mon 16:45 KH 02.014

Integration of NHR resources into the ATLAS computing workflow in Freiburg — •Dirk Sammel, Michael Böhler, and Markus Schumacher — Albert-Ludwigs-Universität Freiburg

At many sites of the Worldwide LHC Computing Grid (WLCG), high-performance computing (HPC) resources are already integrated (or soon will be) into the computing workflows of various experiments. The WLCG Tier-2 cluster ATLAS-BFG in Freiburg has been extended by incorporating resources from the NHR cluster HoreKa at KIT.

To integrate these computing resources into the Freiburg Tier-2 cluster in a transparent and efficient manner, a container-based approach was adopted using the meta-scheduler COBalD/TARDIS. TARDIS launches so-called drones on the HPC system, providing additional compute capacity to the Tier-2 cluster. To differentiate between these augmented resources and the WLCG resources in Freiburg, the AUDITOR accounting ecosystem is used.

This presentation provides an overview of the current status and initial experiences with integrating HPC resources into the Freiburg Tier-2 cluster ATLAS-BFG.

T 14.4 Mon 17:00 KH 02.014

AUDITOR: Tackling HL-LHC accounting challenges and estimating environmental impact of jobs — •Raghuvar Vijayakumar, Michael Böhler, Dirk Sammel, and Markus Schumacher — Universität Freiburg

Distributed computing infrastructures are shared by multiple research communities, particularly within High Energy Physics (HEP), where precise and transparent resource accounting is critical. To meet these demands, we developed AUDITOR (AccoUnting DatahandlIng Toolbox for Opportunistic Resources), a flexible, modular, and extensible accounting ecosystem designed for heterogeneous computing clusters.

AUDITOR captures, processes, and analyzes usage metrics through specialized collectors from HTCondor, Kubernetes, and Slurm batch systems, storing all data in a PostgreSQL database. Its plugin-based architecture allows integration with external tools with Rust and Python clients. Existing plugins include the APEL plugin, which publishes accounting data to the European Grid Initiative (EGI).

Recent developments include a Role-Based Access Control system and an archival subsystem that now periodically exports historical data to Parquet files. The new utilization report plugin provides summaries of job counts, HEPScore performance, power consumption, and estimated CO_2 footprint. Looking ahead, our accounting system is being designed to integrate environmental attributes directly into job-level analytics to measure the carbon impact of computing workloads. In this talk, we will present our new features and highlight how to meet the present and future challenges of HL-LHC accounting.

T 14.5 Mon 17:15 KH 02.014

Theia - A General Photon Propagation Framework for GPUs — •Tobias Kerscher — Technical University of Munich

In neutrino telescopes, understanding how light produced by neutrino interactions via Cherenkov radiation propagates through the detector volume is essential. Due to its complexity this is usually done using Monte Carlo simulations. Using popular tools like GEANT4, however, this often turns out to be the most time consuming step in the overall simulation pipeline and thus becomes a bottleneck. Theia aims to solve this by not only running on GPUs, but also utilizing hardware accelerated ray tracing for fast and accurate intersections with the detector geometry. Since it uses an open standard GPU API, it is not limited to CUDA and NVIDIA hardware. Besides the classical photon tracking simulation, it can also solve for the radiance field producing the underlying expected signal distribution the tracking implicitly samples from. Theia was designed with expendability in mind, making minimal assumptions about the simulated environment. This allows it to be used in various scenarios not limited to neutrino telescopes.

T 14.6 Mon 17:30 KH 02.014

Background Jobs for Efficient Resource Utilisation on HPC Systems — Inga Lakomiec[1], •Ughur Mammadzada[1], Saidev Polisetty[1], Arnulf Quadt[1], Rodney Walker[2], and Sebastian Wozniewski[1] — [1]Georg August University of Göttingen — [2]Ludwig Maximilian University of Munich

The upcoming High Luminosity Large Hadron Collider (HL-LHC) era will greatly increase the computational and storage demands of high energy physics beyond what the current Worldwide LHC Computing Grid (WLCG) can provide. In Germany, university based Tier-2 sites are therefore transitioning their workloads to national HPC facilities at Nationales Hochleistungsrechnen (NHR). In this context, the GoeGrid Tier-2 site at the University of Göttingen is migrating its computing activities to the local NHR cluster EMMY.

To integrate HPC resources into the WLCG, virtual worker nodes are created. They must stop accepting WLCG jobs well before the time limit set by the HPC system, which prevents premature termination but leaves a significant fraction of the allocated CPU time unused near the end of the lifecycle. As the HL-LHC will increase the data volume by a factor of ten, such resource loss becomes increasingly inefficient. The heterogeneous WLCG-HPC environment also requires robust monitoring.

This work investigates opportunistic background jobs at GoeGrid and EMMY to recover this otherwise wasted capacity, and reviews the monitoring infrastructure and the concept of virtual worker nodes.

T 14.7 Mon 17:45 KH 02.014

Assessing Power-Saving Potential Through ACPI Sleep States in Data Center Rack Servers at DESY — •Sandro Grizzo[1], Dwayne Spiteri[2], Kilian Schwarz[3], Martin Gasthuber[4], Konrad Kockler[5], and Jan Hartmann[6] — [1]sandro.grizzo@desy.de — [2]dwayne.spiteri@desy.de — [3]kilian.schwarz@desy.de — [4]martin.gasthuber@desy.de — [5]konrad.kockler@desy.de — [6]jan.hartmann@desy.de

This report evaluates the feasibility of using ACPI sleep states to re-

duce power consumption in rack servers at the DESY data center. It outlines the ACPI power-state model, with emphasis on the S5 *soft*off* state and the Linux-specific suspend-to-idle mechanism, as these are the only sleep states typically supported on server-grade hardware. Measurements obtained via IPMI interfaces on several server models show potential power savings above 80% in S5 relative to idle mode. Suspend-to-idle provides a substantially lower reduction, but also lower wake-latencies in return. The findings indicate that using S5 for idle servers could offer considerable power-saving potential in data centers where low wake-latency is not a primary requirement.

T 14.8 Mon 18:00 KH 02.014

Optimisation of GPU Accelerated Algorithms for Track Finding in Particle Physics — •Petr Fiedler, André Sopczak, and Pavel Tvrdik — CTU in Prague

Track finding in particle physics has been an increasing challenge over the past decades because the intensities of collisions in state-of-the-art particle colliders have increased enormously. One of the track reconstruction steps is the track finding. For this, track seeds are determined and the tracks are successively reconstructed by adding more measurements. We explain the main ideas of the state-of-the-art GPU-accelerated implementation of the track finding algorithm and describe several optimisations focusing on an early elimination of fake tracks.