

Sustainable Computing in High-Energy Physics: HEPScore and the HEP Benchmark Suite

Thursday, 28 March 2024 11:40 (20 minutes)

In the contemporary era, where scientific progress meets the imperative of responsible resource utilization, the need for innovative tools is paramount. This paper explores the pivotal role of a new benchmark for High-Energy Physics (HEP), HEPscore, and the HEP Benchmark Suite in steering the HEP community toward sustainable computing practices. As exemplified by projects like the Large Hadron Collider (LHC), HEP demands massive computational resources, making it imperative to strike a balance between scientific excellence and environmental responsibility.

The HEP Benchmark Suite, a toolkit for benchmark orchestration, shares the center stage in this paradigm shift alongside HEPscore. Developed to manage benchmarks from a single application, the Suite characterizes the performance of individual and clustered heterogeneous hardware. Its modular design enables users to configure and run a diverse range of benchmarks, including HEP SPEC06, SPEC 2017, HEPscore23 (HS23), and DB12. Distributed under the GNU General Public License v3, the Suite is open-source, inviting the community to scrutinize and enhance its code.

Modularity extends to the suite's add-ons, enhancing configurability. Existing plugins are able to retrieve hardware and software metadata, offering deeper insights into benchmark running conditions. The prime example of it is the actively developed Energy Plugin, which reports benchmark consumption. This information is critical for the hardware procurement process, as it makes it possible to deliver maximal performance at the lowest energy consumption, reducing at the same time the energy cost.

While plugins may be configured with any compatible benchmark, in the field of High-Energy Physics, HEPscore 23 appears to provide an accurate performance score, taking into account the latest technological advances. The benchmark is composed of seven underlying workloads, containerized applications based on the current applications being used at CERN and the World-Wide LHC Computing Grid. A thorough study led to this final configuration as the best compromise between runtime and accuracy when representing the present applications of the HEP community.

A central database, following the containerized nature of benchmarks, further augments the Suite's capabilities. Reports generated by the Suite can be published to a message broker, facilitating storage in CERN's central benchmarking OpenSearch database. With over 90,000 entries for reference, users can easily query and compare results. The graphical user interface atop the database enables users to create visualizations and dashboards as well.

This paper delves into the architectural design and capabilities of HEPscore and the HEP Benchmark Suite, showcasing their potential to revolutionize benchmarking practices within the HEP community, in particular in terms of energy consumption. Drawing on experiences from worldwide deployments across the Worldwide LHC Computing Grid, we present evidence of the tools' effectiveness in gathering and analyzing measurements globally.

As the HEP community embraces these tools, we anticipate a transformative shift towards a future where scientific progress harmonizes with environmental responsibility, setting a precedent for sustainable computing in the domain of big science.

Primary author: MENÉNDEZ BORGE, Gonzalo (CERN)

Co-author: Mr GIORDANO, Domenico (CERN)

Presenter: MENÉNDEZ BORGE, Gonzalo (CERN)

Session Classification: Physics & Engineering Application

Track Classification: Track 1: Physics and Engineering Applications