Contribution ID: **96** Type: **Oral**

Power Efficiency in HEP (x86 vs. arm)

Tuesday, 25 October 2022 15:30 (20 minutes)

The power consumption of computing is coming under intense scrutiny worldwide, driven both by concerns about the carbon footprint, and by rapidly rising energy costs.
ARM chips, widely used in mobile devices due to their power efficiency, are not currently in widespread use as capacity hardware on the Worldwide LHC Computing Grid.
However, the LHC experiments are increasingly able to compile their workloads on the ARM architecture to take advantage of various HPC facilities (e.g., ATLAS, CMS).

The work described in this paper attempts to compare the energy consumption of various workloads on two almost identical machines, one with an arm64 CPU and the other with a standard AMD x86_64 CPU, operating in identical conditions.
This builds on our initial study of two rather dissimilar machines, located at different UK Universities, which produced some interesting, but at times contradictory, results, showing the need to control the comparison more closely.

The set of benchmarks used include CPU intensive, memory intensive, and I/O bound tasks, ranging from simple scripts, through compiled C programs, to typical HEP workloads (full ATLAS simulations).
We also plan to test the most recent HEPscore containerized jobs, which are actively being developed to match LHC Run3 conditions and can already target different architectures.

The results compare both the power consumption and execution time of the same workload on the two different architectures (arm64 and x86_64).
This will help inform Grid sites whether there are any scenarios where power efficiency can be improved for LHC computing by deploying ARM-based hardware.

Significance

Power efficiency is an important, often overlooked, feature of computing in High-Energy Physics.
A close comparison in efficiency and throughput of HEP workloads among different architectures provides actual data to guide Grid sites in the choice of new hardware.

This builds and largely improves on our previous study by taking advantage of brand new hardware and an extended set of benchmarks.

References

https://indico.cern.ch/event/1128343/contributions/4787174/

Experiment context, if any

LHC, WLCG, ATLAS

Primary author: SIMILI, Emanuele

Co-authors: BRITTON, David (University of Glasgow (GB)); STEWART, Gordon; SKIPSEY, Samuel Cadellin

Presenter: SIMILI, Emanuele

Session Classification: Track 1: Computing Technology for Physics Research

Track Classification: Track 1: Computing Technology for Physics Research