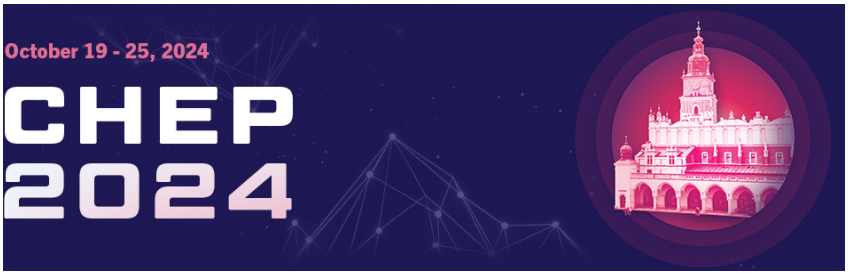

Contribution ID: **326** Type: **Talk**

Heterogeneous Computing and Power Efficiency in HEP

Wednesday 23 October 2024 16:51 (18 minutes)

The Glasgow ScotGrid facility is now a truly heterogeneous site, with over 4k ARM cores representing 20% of our compute nodes, which has enabled large-scale testing by the experiments and more detailed investigations of performance in a production environment. We present here a number of updates and new results related to our efforts to optimise power efficiency for High Energy Physics (HEP) research.

We will show updated benchmark results, including a new figure-of-merit designed to characterise the power usage during the execution of the HEPScore benchmark. Previously, community measurements have used either the average or maximum power, neither of which is a good estimator. We expand our HEP-Score/Watt comparison to include additional machines such as Ampere Altra Q80 and M80, NVidia Grace, and the most recent AMD EPYC chips. We also introduce a Frequency Scan methodology to better characterize performance/watt trade-offs, potentially informing strategies like frequency scaling during peak hours to optimize power efficiency.

In addition, we present a comparison of single-socket versus dual-socket performance, revealing consistent findings that dual-socket configurations exhibit performance degradation compared to two single-socket machines, though of varying magnitudes. Leveraging HEPScore jobs and the 'taskset' command to target specific core configurations, we explore performance variations across core groups within the same socket or across dual sockets. Preliminary results show that same-CPU cores have better performance, confirming the importance of workload optimization strategies, such as fine-tuning the job scheduler to prioritize same-socket core utilization.

Our findings contribute to advancing heterogeneous computing strategies and power efficiency optimizations in HEP, paving the way toward more sustainable hardware solutions.

Primary authors: SIMILI, Emanuele; SIMILI, Emanuele (University of Glasgow (GB))

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

Presenters: SIMILI, Emanuele; SIMILI, Emanuele (University of Glasgow (GB))

Session Classification: Parallel (Track 7)

Track Classification: Track 7 - Computing Infrastructure