# ARMing up for HEP

Emanuele Simili, Gordon Stewart, Samuel Skipsey, Dwayne Spiteri, Albert Borbely, David Britton

University of Glasgow

**Abstract:** We present a case for ARM chips as an alternative to standard x86 at WLCG sites to help reduce power consumption [1]. New measurements are shown on the performance and energy consumption of two machines (one ARM [2] and one x86), that were otherwise similar in specification and cost. These new results include the energy-efficiency and speed of single- and multi-thread jobs; the effect of hyper-threading; and an initial look at clock throttling as a way of shaping power-load. **We observed significantly lower power consumption and often slightly better performance on the ARM machine** and, noting the increased availability of ARM software builds from all LHC experiments and beyond, we plan to install a 1920-core ARM cluster at our WLCG Tier2 site at Glasgow in the summer of 2023. This will enable testing, physics-validation, and eventually an ARM production environment that will inform and influence other WLCG sites in the UK and worldwide.

## Methodology

Compared two almost identical servers (**ARM** & **x86**) of similar price:

**AMD48c / AMD96ht: Single AMD EPYC 7003 series (GIGABYTE)**
CPU: x86 AMD EPYC 7643 48C/96HT @ 2.3GHz (TDP 225W)
RAM: 256GB (16 x 16GB) DDR4 3200MHz
HDD: 3.84TB Samsung PM9A3 M.2 (2280)

**ARM80c: Single socket Ampere Altra (GIGABYTE)**
CPU: ARM Q80-30 80C @ 3GHz (TDP 210W)
RAM: 256GB (16 x 16GB) DDR4 3200MHz
HDD: 3.84TB Samsung PM9A3 M.2 (2280)

Power measurement collected using IPMItools [3]
Readings were validated with external plugs.

Results agree within **~4%**

(the difference has opposite sign for the two servers, strengthening our conclusion about the power efficiency of **ARM**)

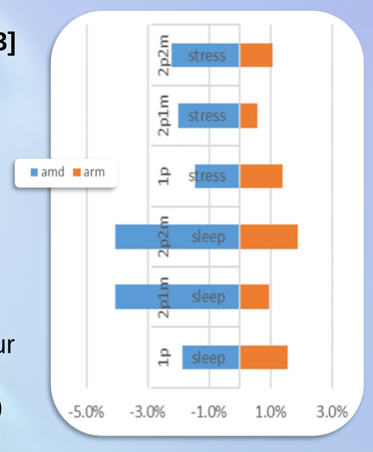


Comparison extended to a dual **x86** compute node, representative of the Glasgow site (the node is part of a 2U/4-node chassis, but has a similar price/node as the machines above):

**2xAMD32c / 2xAMD64ht: Dual Socket AMD EPYC 7513 series Processors (DELL)**
CPU: 2 * x86 AMD EPYC 7513, 32C/64HT @ 2.6GHz (TDP 200W)
RAM: 512GB (16 x 32GB) DDR4 3200MHz
HDD: 3.84TB SSD SATA Read Intensive

## Score & Power

Monitor power usage while running the latest HEP-Score [4], in order to benchmark performance and power efficiency of **ARM** vs. **x86**:

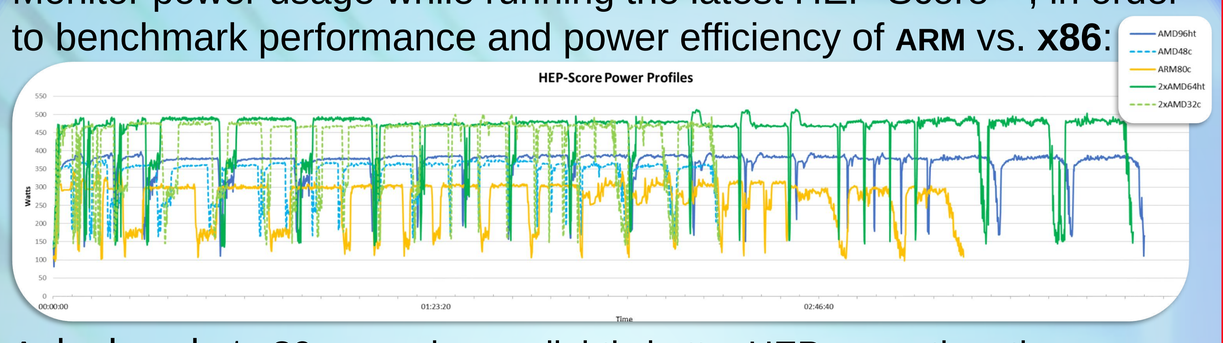


A dual-socket **x86** server has a slightly better HEP-score than the **ARM**, but at a higher energy consumption.

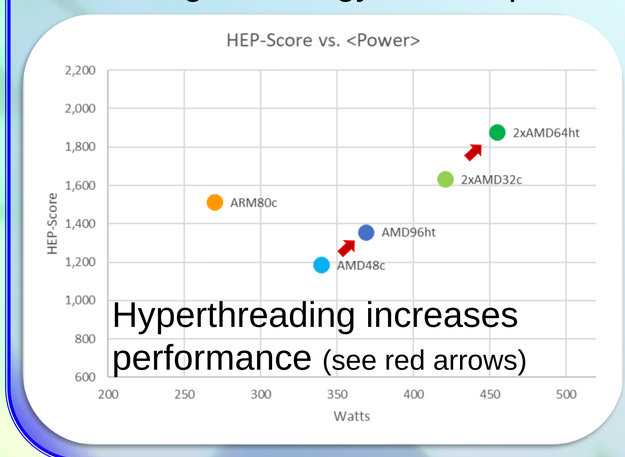


Hyperthreading increases performance (see red arrows)

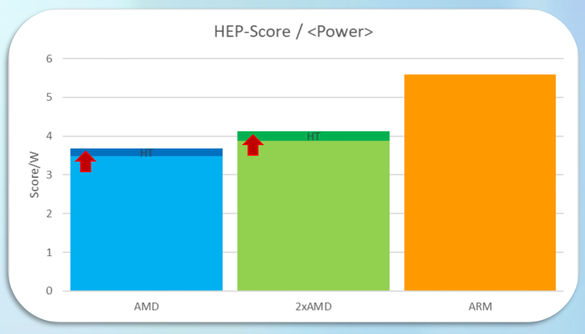


When energy is part of the rating, the **ARM** outperforms the **x86**s.

## Thread-Scan

Increase the number of HEP-Score containerized jobs executed simultaneously until the CPU is saturated:

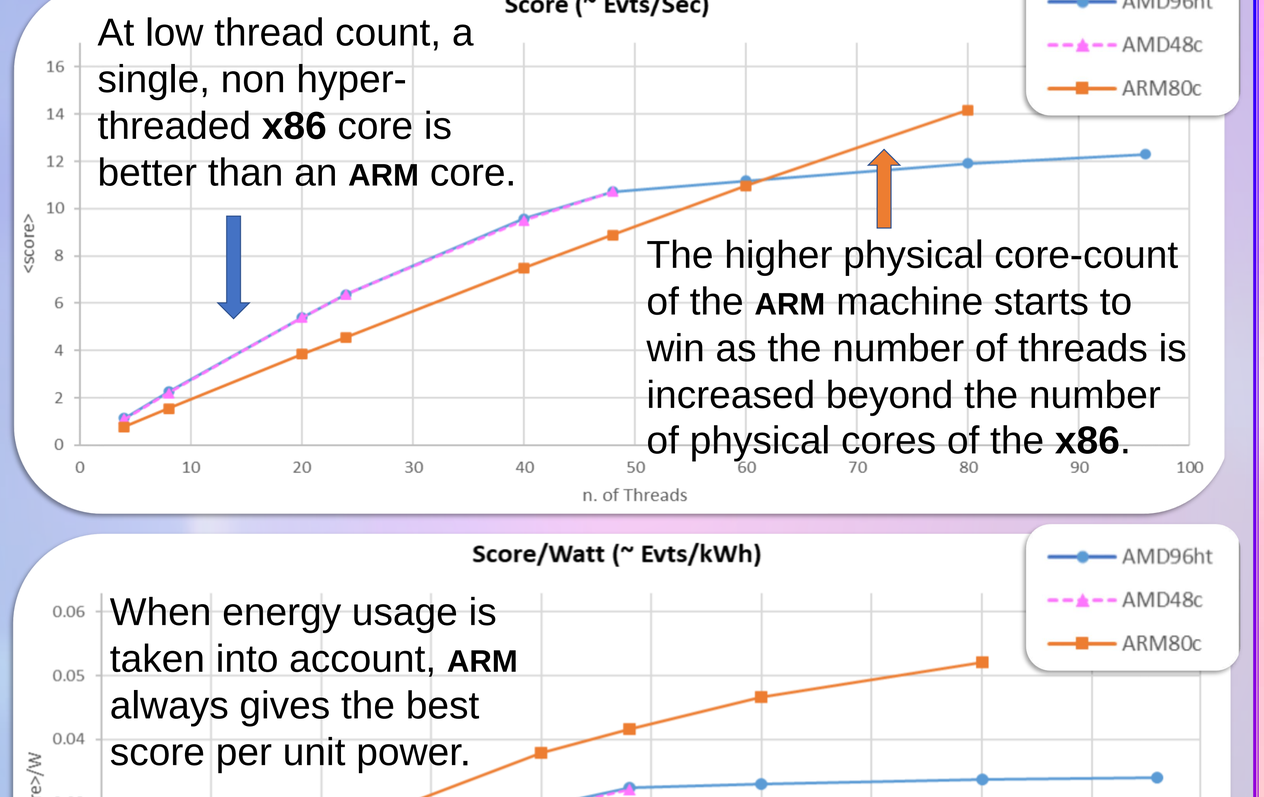


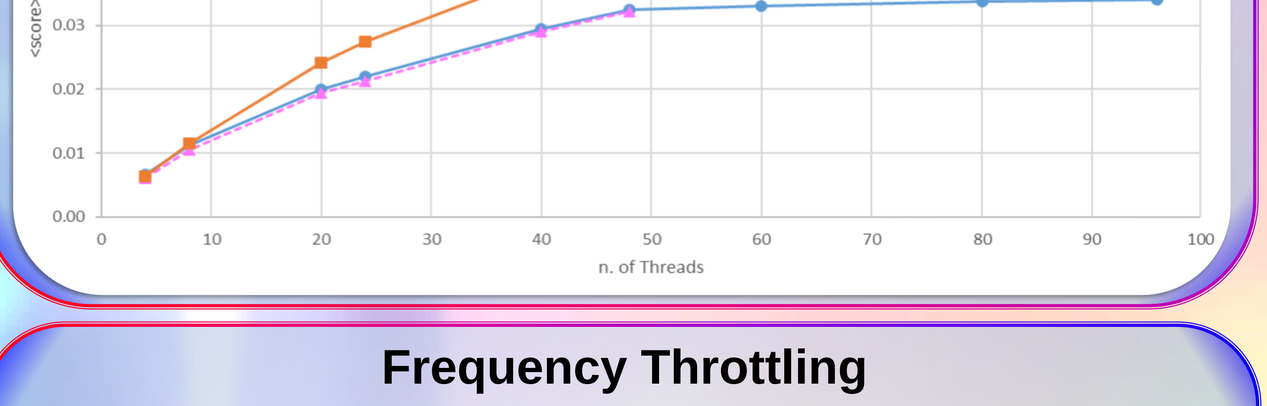


## Frequency Throttling

Reducing the CPU frequency could save power, but job will last longer:

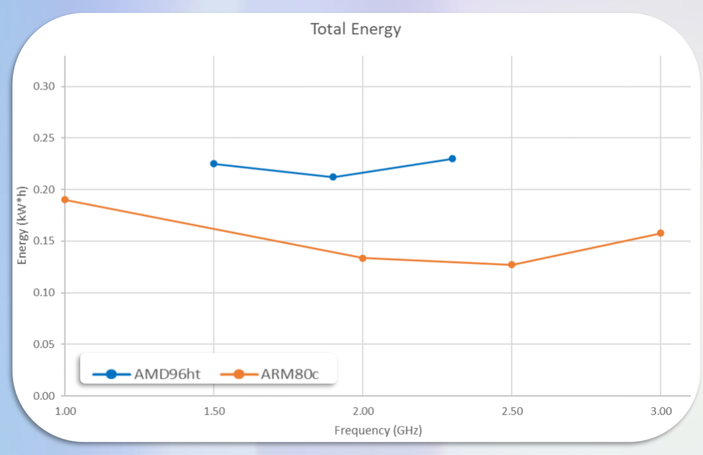



The integrated energy used for a job varies modestly with clock speed but has a minimum value at intermediate frequency settings on both machines.

**Possible beneficial trade-off in throughput for an overall decreased power consumption?**

## Outlook

An energy plug-in has been developed for the HEP-Score suite and will be included in the upcoming release. This will:

- provide a standard benchmarking tool to assess not only performances but also energy efficiency;
- allow Grid sites to make informed decision when purchasing new hardware.

Most LHC experiments have validated **ARM** builds of their software, and some non-LHC experiments are moving too.

| | *ARM build* | *Physics Validated* |
|---|---|---|
| ATLAS | yes | yes |
| CMS | yes | yes |
| ALICE | yes | yes |
| LHCb | yes | almost |
| Belle2 | yes | no |

## Conclusions

Upcoming **ARM** farm at Glasgow site: we are about to get about two thousand **Altra® ARM** cores, donated by Ampere® (USA).

They will become part of the advertised resources of our Tier2 site: a separate batch queue will be created to expose **ARM** resources on the Grid, such that they can be targeted by the experiments.

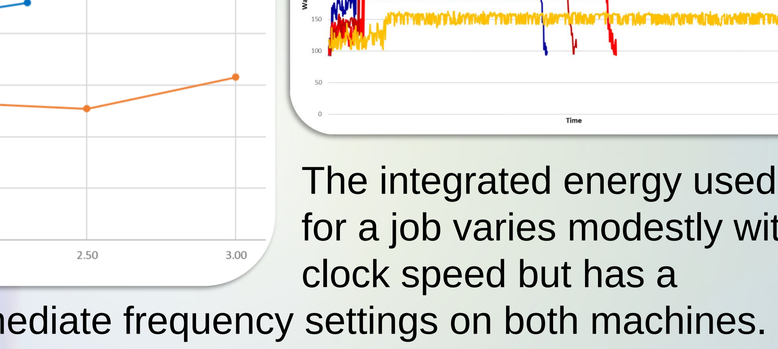


**The availability of untapped ARM resources will further push experiments to develop/build the appropriate software to make the most use of them.**

## References

[1] E.Simili, et al., "Power Efficiency in HEP (x86 vs. ARM)", ACAT2022, https://indico.cern.ch/event/1106990/contributions/4991256/
[2] Ampere® Altra® Multi Core Server Processors, https://amperecomputing.com/processors/ampere-altra
[3] D.Laurie, "IPMItools: Intelligent Platform Management Interface", https://github.com/ipmitool/ipmitool
[4] D.Giordano, et al. "HEPiX Benchmarking Solution for WLCG Computing Resources", Comput Softw Big Sci 5, 28 (2021).