Sustainability

Introduction

We had discussions on green computing last year, I won't recap..
Focus on a few recents informations I collected from various sources to launch the discussion

L. Duflot

Carbon cost

- DELL server lifetime carbon assessment
- Empreinte fabrication
- Construction is ~50% of the overall carbon footprint over the entire lifetime of a server
  - RAM and SSD are carbon costly
- Caveat: these studies are based on machine config quite different from our usual WN and with shorter lifetime..
- WLCG: different disk server configurations worldwide
  - 12x8TB server in 10+2 RAID6 → 1.9W/TB
  - 14x16TB server in 12+2 RAID6 → 1.2W/TB
  - 36x16TB server in (7+2)xN RAID6 → 0.82W/TB
-

Zach Marshall ATLAS C&S week Oct 2023

Processors

GridPP
UK Computing for Particle Physics

HEPScore Measurements

Machine	CPU	Threads	HEPScore
Test AMD	1x EPYC 7643 HT	96	1359
Std AMD WN	2x EPYC 7513 HT	128	1887
AMD Bergamo	2x EPYC 9754 HT	512	7497
Test ARM	1x ALTRA 80	80	1513
Farm ARM	2x ALTRA 80	160	2619 ??
Ampere Max	1x ALTRA Max	128	2065

x4

Not x2 ??

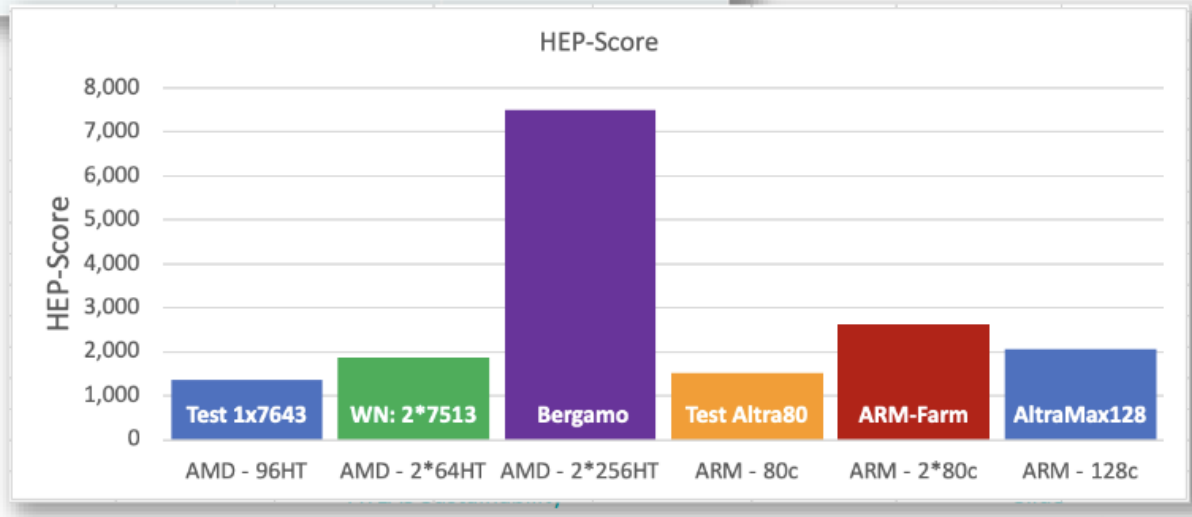

David Britton, University of Glasgow

Processors

GridPP
UK Computing for Particle Physics

HEPScore/Watt

Machine	CPU	Threads	HS/Watt	
Test AMD	1x EPYC 7643 HT	96	3.7	
Std AMD WN	2x EPYC 7513 HT	128	4.2	
AMD Bergamo	2x EPYC 9754 HT	512	6.3	50% improvement from 7513 WN
Test ARM	1x ALTRA 80	80	5.6	
Farm ARM	2x ALTRA 80	160	5.2 ??	
Ampere Max	1x ALTRA Max	128	7.0	25% improvement from 80c

Note added: We use Average Watts to calculate HS/Watt. Other people also use Max Watts, so care needs to be taken comparing numbers.

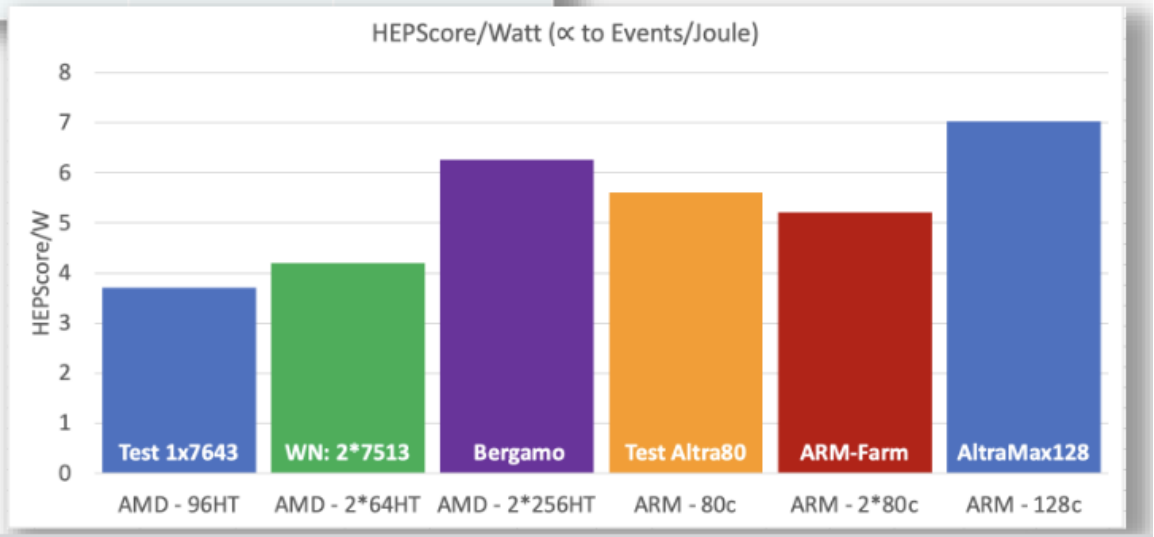

GridPP
UK Computing for Particle Physics

Energy Efficiency vs Clock

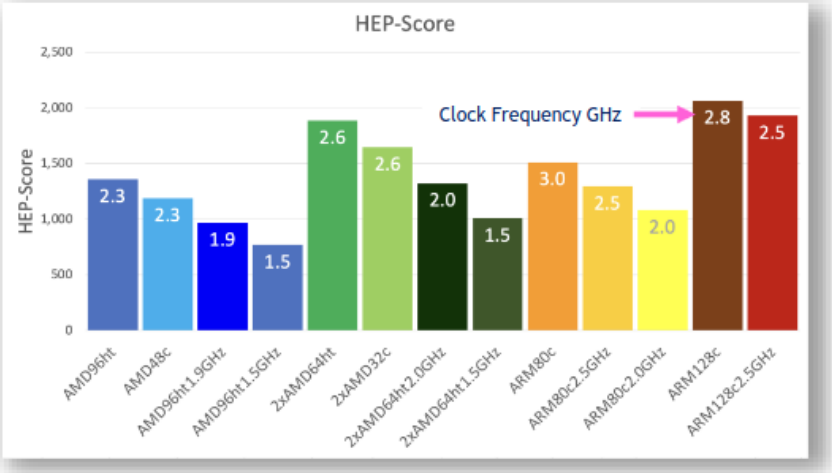

- Running one step below max clock frequency seems to be the sweet-spot for all CPUs tested.
- On ARM128, frequency reduction lowers HEPScore by 6.5% but saves another 9% of the energy.
- No data on Bergamo (bug?)

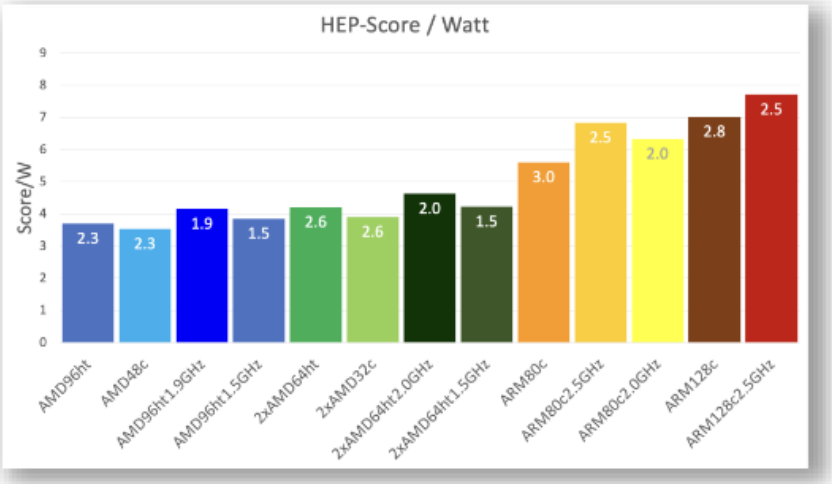

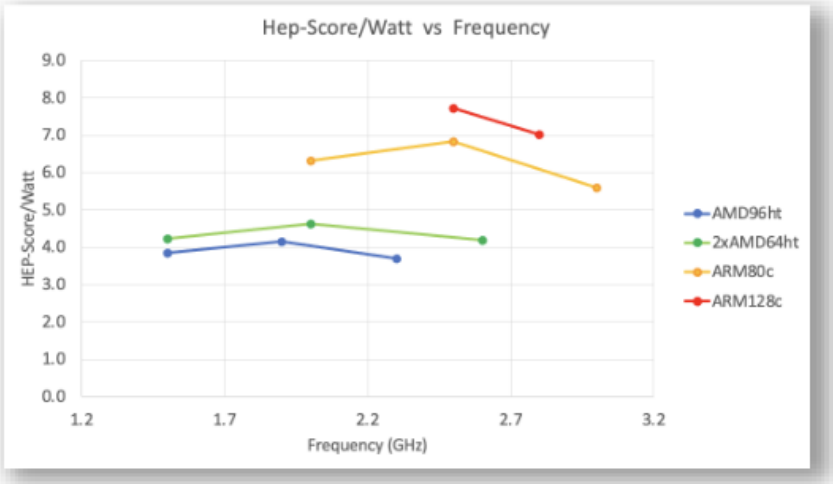

Processors

ORAP51

- ADASTRA at CINES
  - The new machine is a scaled-down copy of LUMI (Finland) or FRONTIER (USA)
  - AMD Genoa + AMD GPU
  - Liquid cooled (1 rack weights 4t!)
  - Same overall performances as the machine it replaces (Intel Haswell) but 1/3 of eletricity
  - Fewers racks: CINES would have enough room and cooling for an exascale machine!
- European processor(s)
  - Rhea architechture ARM out in 2024 (?)
  - Designed for HPC, next design is “server oriented”
  - European design but not fabrication
- RISC-V arch:
  - A credible server proc seems to be years away

PUE

What affects PUE? Anything we can recommend?

- Transition from air to water cooling (LRZ example: 1.4 → 1.1)
- Ambient temperature, weather, building construction (CERN old center: 1.6; PCC: ~1.1; RAL: ~1.2; NERSC 1.06)

Pretend your Berkeley-based data center has a PUE of 1.2, and you could build a new center with a PUE of 1.1. How long does the new data center need to run before you have amortized the carbon from construction and are really "winning"?

(1.2-1.1)/1.2 * 5 MW datacenter at 0.12 mTCO2e/MWh vs. 14,000 m^2 at 8.7 kg CO2e/m2

→ **4 months (super worth it)**

Pretend your CERN-based data center has a PUE of 1.6, and you could build a new center with a PUE of 1.1. How long does the new data center need to run before you have amortized the carbon from construction and are really "winning"?

(1.6-1.1)/1.6 * 4 MW datacenter at 0.075 kg CO2e/kWh vs. 10,000 m^2 at 8.7 kg CO2e/m2

→ **2 months (even more worth it)**

CPU pacing

- Adapt CPU freq to smooth out electricity peaks (more expensive, dirtier)
  - Modulations during the day: we need a constant feed of information and be reactive
- In ATLAS, studied in DE and UK
- Devil's in the details…
  - Jobs will take longer clock time:
    - Adapt batch system so they are not killed in time limit
    - Need work on proxy/tokens renewal since it also has expiration date
  - Accounting to be adapted
  - Pledges will have to anticipate the average CPU power loss

Miscelanous

- ATLAS:
  - Task summary now has an evaluation of the average carbon cost (JEDI)
  - We could have less crash in user analysis….
- Data Center Model being developped by Bristol
-

Liens

- CHEP
  - Tier 2 footprint, IRISCAST, WLCG vs energy
- Forum ORAP51
  - ARM, RICS-V, ADASTRA, HPC roadmap from AMD
- Articles
  - Carbon footprint
  - DELL server lifetime carbon assessment
  - Empreinte fabrication
  - Empreinte carbone de l'heure de calcul
  -