CIPEA
CERN Innovation Programme on Environmental Applications

Green computing and HEP: benchmarking the energy efficiency of HEP data processing workloads

Domenico Giordano, Stefan Roiser, Andrea Valassi (CERN IT),
Olivier Mattelaer (UCLouvain)

27 June 2022, CIPEA Innovation Day
https://indico.cern.ch/event/1166768/contributions/4928987/

Large Scale Distributed Data Processing

- LHC experiments currently use a world-wide distributed infrastructure for data processing
    - 170 sites, ~ 1.5 Mio CPU cores, +storage, +network
    - Expecting large demand increase for HL-LHC era

- The global electrical power consumption for information and communication technologies is expected to rise as well
    - With the relative portion for data centers increasing

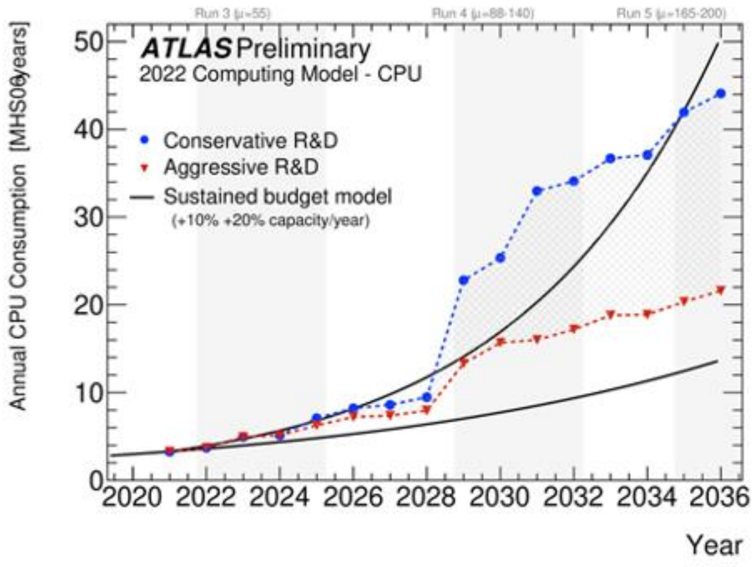

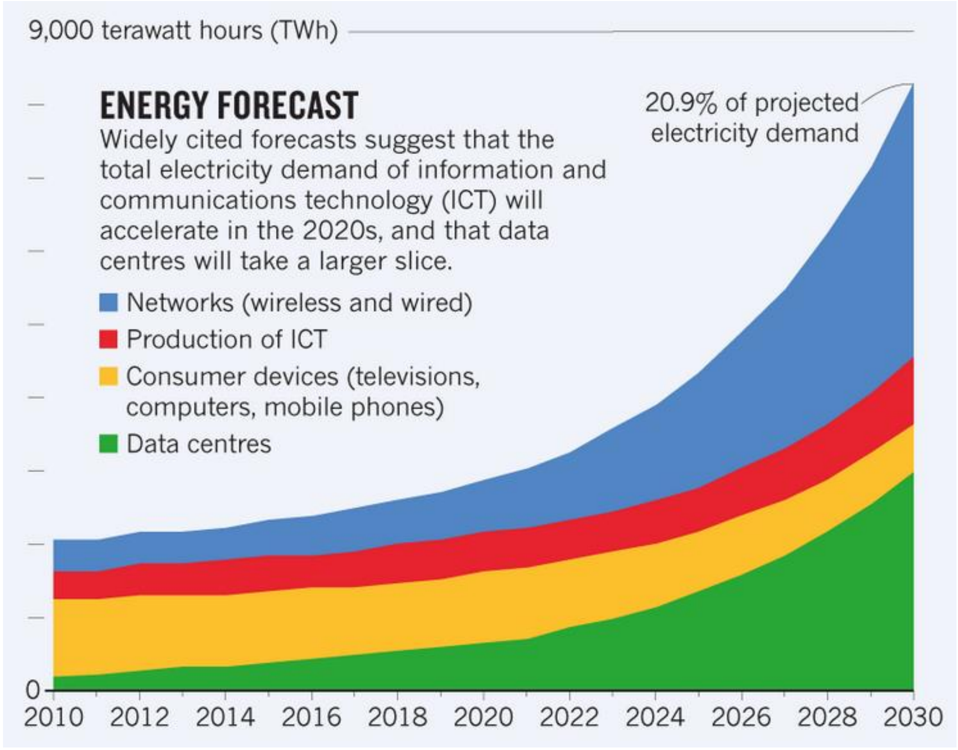

Worst case prediction of global electrical ICT power consumption
Nature 561, 163-166 (2018), https://doi.org/10.1038/d41586-018-06610-y

Comparing data processing on different models of CPUs and hardware accelerators (GPUs, FPGAs…)

- GPUs allow highly parallelized execution of workloads, with O(1000) compute cores
  - Benefit: substantial performance improvements (for software applications which fit this paradigm)
  - Cost: increased power consumption for the GPU
  - Is this a more energy efficient processing model?

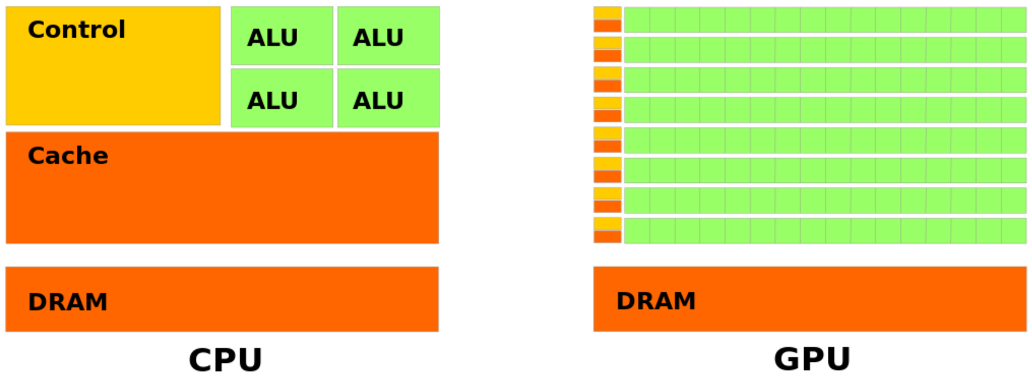

- Different CPU architectures (e.g. x86, ARM) also differ in performance and power needs

General question:
<u>which architecture gives us the best computing performance per Watt?</u>

Goals and objectives of this project

- Measure electrical power efficiency of large-scale scientific data processing workloads in High Energy Physics via a systematic process
- Envisaged outcomes:
  - Benchmark HEP data processing workloads on different architectures (CPU, GPU, CPU+GPU…)
    - Measure both computing performance (events per second) and electrical power consumption (Watts)
    - Benchmark metric: “events processed per second per Watt” (events per Joule) – *physics per Watt!*
  - Develop methodology and infrastructure which can be applied to qualify power efficient “green” computing and make it generally applicable for large scale scientific data processing workloads
- Potential positive impact on CERN computing center, WLCG data centers and the scientific computing community and society in general in terms of reducing electrical power consumption for data processing

HEP data processing workloads

- One important prerequisite of the project is the ability to run the same HEP software workloads on different computing architectures to get the same physics results
  - Leverage on the expertise and tools of the HEPiX benchmarking WG – HEP-score is a compute benchmark based on the HEP workloads most representative of LHC processing on Grid CPUs

- More specific recent examples to compare CPUs and GPUs:
  - “Madgraph4gpu” – a collaboration of CERN IT, UCLouvain, Argonne to port and optimize the Madgraph5_aMC@NLO event generator on GPUs and vectorized CPUs
    - A first alpha release of the software is available to compare CPU and GPU execution
  - “Adept” – a Geant4 detector simulation R&D project hosted by CERN EP to port electromagnetic calorimeter detector response onto GPUs
    - A first proof of concept release is ready for testing on GPUs

Infrastructure support

The project will leverage on already existing infrastructure:

- HEPiX Benchmarking Working Group
  - Infrastructure to execute and analyze containerized HEP workloads in a reproducible way
- CERN IT compute fabrics team
  - Electrical power consumption measurements is already feasible with the current CERN computer center monitoring infrastructure, but must integrated in our infrastructure for our needs
- Other test beds at CERN or elsewhere
  - Via CERN/Openlab connections to more test beds for various other hardware accelerators and platforms can be established

Summary

The “Benchmarking the energy efficiency of HEP data processing workloads” project

- Measures the power efficiency of scientific data processing workloads in High Energy Physics comparing various hardware architectures (CPUs, GPUs…)
- Defines a methodology and a system of metrics to qualify power efficiency for scientific software applications, which can be extended to other scientific domains
- Provides a path to steer scientific data processing towards a green computing paradigm