Energy&Carbon Footprint Of CMS WFs

Andrea Piccinelli
On behalf of the CMS collaboration

WLCG Environmental Sustainability Workshop - Dec 12th, 2024

The context

Computing and HEP

➔ Role of computing technologies more and more crucial for scientific advancements

- AI/HPC implementations in particle physics
- AI implementations in Science/Industry (computer vision, LLMs, GenAI)
- Numerical simulations
- …and also at CMS, of course!

More and more energy needed for performing computing activities!

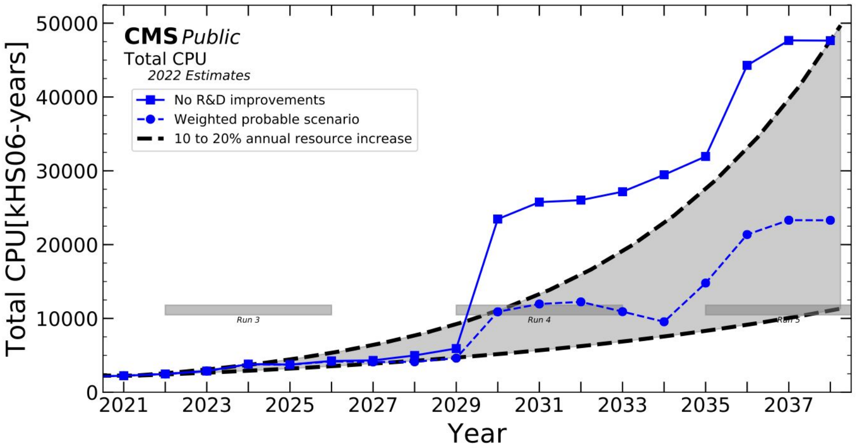

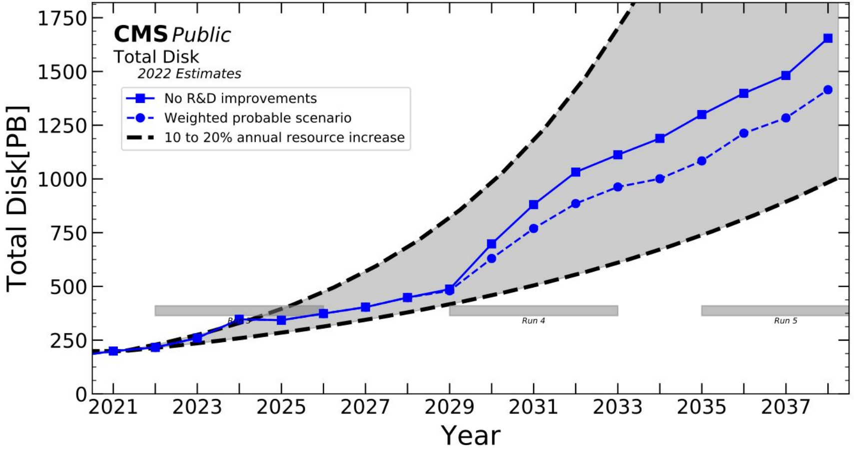

A sustainable effort?

The energy footprint of physics experiment

- ➔ Electricity consumption from data centres and AI expected to double by 2026 (IEA 2024 report)
- ➔ Dennard Scaling: Future generation of processors not expected sustain our needs
    - ◆ Improve resource usage efficiency
- ➔ Increasing awareness of environmental impact of physics experiments

Editorial | Published: 07 March 2023

Physics should acknowledge its environmental impact and act on it

Nature Reviews Physics **5**, 133 (2023) | Cite this article

1640 Accesses | **1** Citations | **2** Altmetric | Metrics

CERN Environment Report 2023

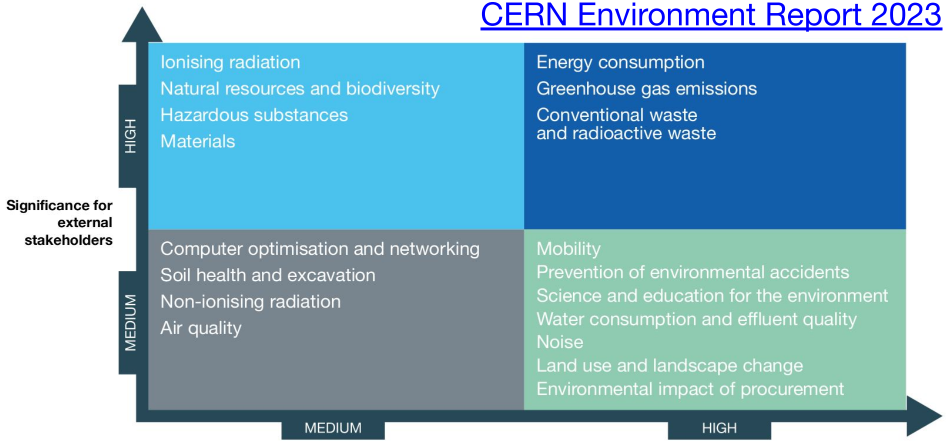

CERN MATERIALITY MATRIX 2022
The topics identified as being of lower significance to all stakeholders are not comprehensively covered in this report but are subject to monitoring by CERN.

Need for a quantitative measurement!

Measuring the carbon footprint with a CMS scenario

➔ Estimating Energy Consumption and Carbon Costs of GENSIM jobs at LHC
  - ◆ Preliminary results from Francesco Minarini, PhD candidate in Physics Supervised by D. Bonacorsi and T. Diotalevi (University and INFN Bologna)
  - ◆ Academia-industry joint effort on NextGEN EU

➔ Goal: Measuring energy consumption and the carbon footprint of a CMSSW GENSIM job
  - ◆ Pheno event generation and Geant4 simulation → First step of CMS simulation

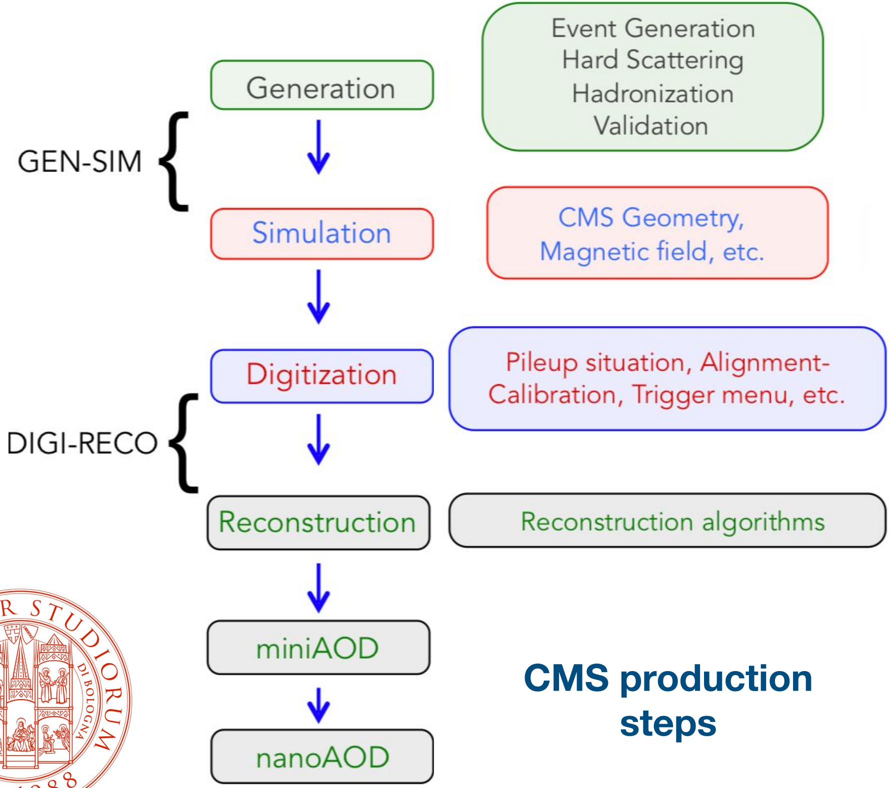

CMS production steps

Theoretical foundations

The energy consumption

$$E_{A\rightarrow X} = T \times (n_c \times P_c \times u_c + n_m \times P_m) \times \mathrm{PUE}$$

- ➔ T = Elapsed computing time (h)
- ➔ n_c = number of used cores
- ➔ P_c = power draw normalized to computing cores (kW)
- ➔ u_c = CPU usage factor in [0 (low usage), 1 (high usage)] interval
- ➔ n_m = allocated RAM memory (GB)
- ➔ P_m = power draw of RAM (kW)
- ➔ PUE = Power Usage Efficiency

ADVANCED SCIENCE Influence Series

Research Article | Open Access

Green Algorithms: Quantifying the Carbon Footprint of Computation

Loïc Lannelongue, Jason Grealey, Michael Inouye

First published: 02 May 2021 | https://doi.org/10.1002/advs.202100707 | Citations: 67

UNIVERSITY OF NOTRE DAME

Theoretical foundations

From energy consumption to carbon footprint

$$F = E_{(A \to X)} \times CI$$

CI = Country-wise Carbon Intensity coefficient of electricity production [gCO2/kWh]

CI strictly depends on how energy is produced and exchanged on the energy market
→ Susceptible to geopolitical situation of where the data center is located **Warning**

Measuring the energy consumption

$$E_{A\to X} = T \times (n_c \times P_c \times u_c + n_m \times P_m) \times \mathrm{PUE}$$

Static Data	*Dynamic Data*
P_m	u_c
P_c	n_m
n_c	T
PUE	

Static Data → gathered in an external Config.toml file

Dynamic Data → measurable leveraging some properties of the filesystem (content of `/proc`) of Linux-powered machines

No particular authentication method required to get them

Experimental Setup

➔ Computing machines

◆ Two “Slurm-managed” Intel nodes offered by INFN-CNAF’s student-oriented facility

Node 1 → CPU: 2x E5-2640v2 (16 physical cores), RAM: 128 GB, HT: ON

Node 2 → CPU: 2x E5-2683v3 (28 physical cores), RAM: 128 GB, HT: ON

Monitoring target

➔ CMS containerised HEP Benchmark (open-access)

◆ CMS GENSIM multithreaded job

SW framework: CMSSW v12_5_0

◆ Top-antitop pair events events

at 14 TeV with Run3 condition

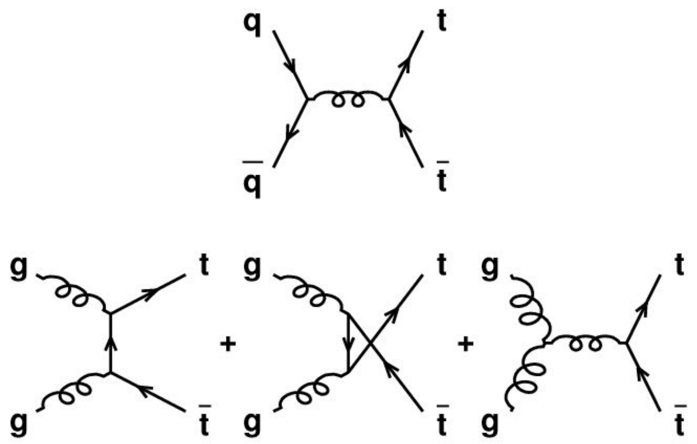

PoC: Monitoring workflow

In parallel, the containerized probe software targets the executable

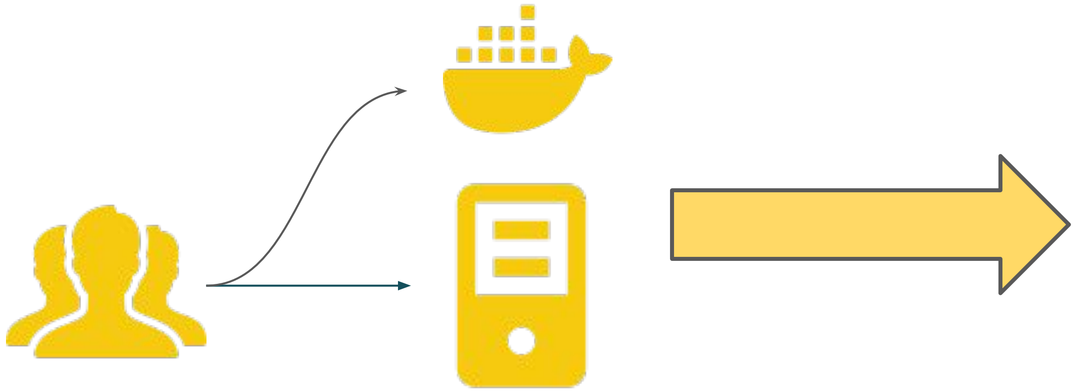

Physics Job submission happens as usual

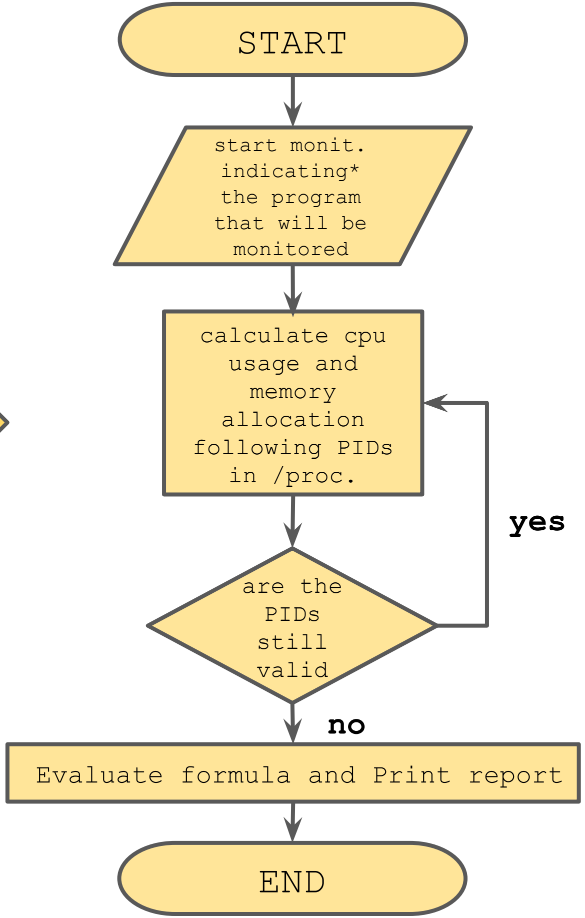

PoC: Monitoring workflow

In paral
softwar

Physics Job su
happens as usu

```
Now evaluating carbon footprint of execution ...

REQUESTED CORES: 30
TDP_PER_CORE: 12
AVG_CPU_USAGE: 0.763323
AVG_MEM_ALLOC (GB): 3.35681
RAM W: 0.375
CORE W: 274.796
MEM W: 1.2588
ABSORBED W: 276.055
ELAPSED T: 0.354481
PUE: 1.65
CARBON INTENSITY: 314
YOUR FOOTPRINT: 50.6993 gCO2e
```

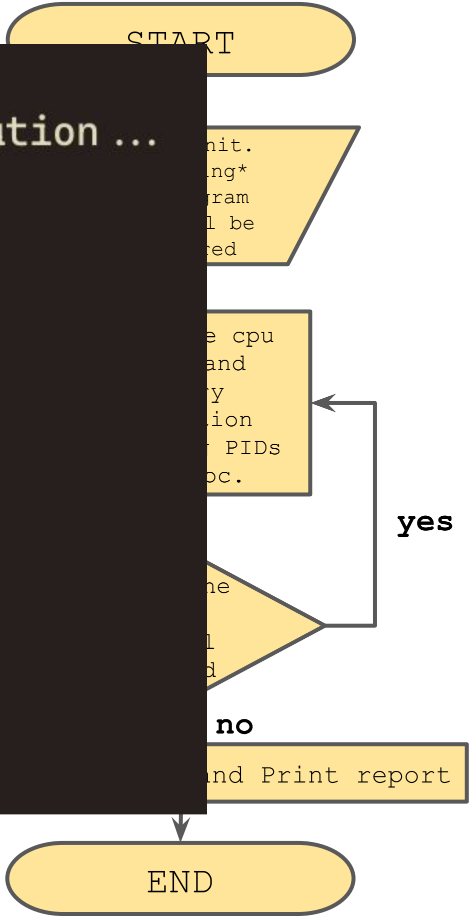

Results

Data are the geometric means of 3 separate measurements on the same node.

INFN-CNAF E5-2640v2 node

CPUs	Evts per core	kWh	elap.T
30	333	3.058	5.13
28	357	3.026	5.44
24	416	3.017	6.36
20	500	2.940	7.44
16	625	2.921	9.22
12	833	2.887	12.16
8	1250	2.818	17.80

INFN-CNAF E5-2683v3 node

CPUs	Evts per core	kWh	elap.T
54	185	1.51	1.87
50	200	1.49	2.01
42	238	1.48	2.39
30	333	1.47	3.32
28	357	1.46	3.53
16	625	1.41	5.96
8	1250	1.17	11.11

Results

Metrics scaled in [0,1] to harmonize scales and identify the working point

$$X' = \frac{X - X_{min}}{X_{max} - X_{min}}$$

INFN-CNAF E5-2640v2 node

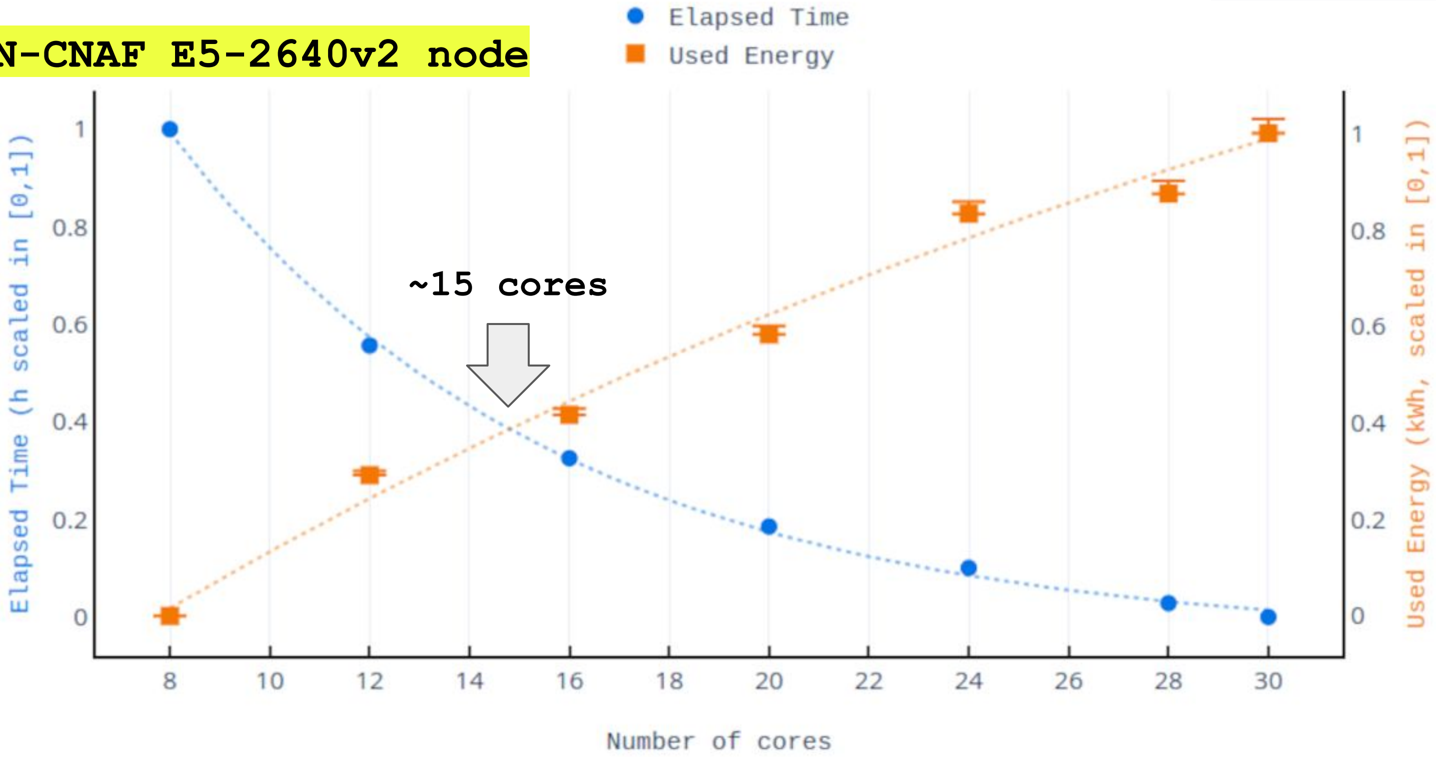

Results

Metrics scaled in [0,1] to harmonize scales and identify the working point

$$X' = \frac{X - X_{\min}}{X_{\max} - X_{\min}}$$

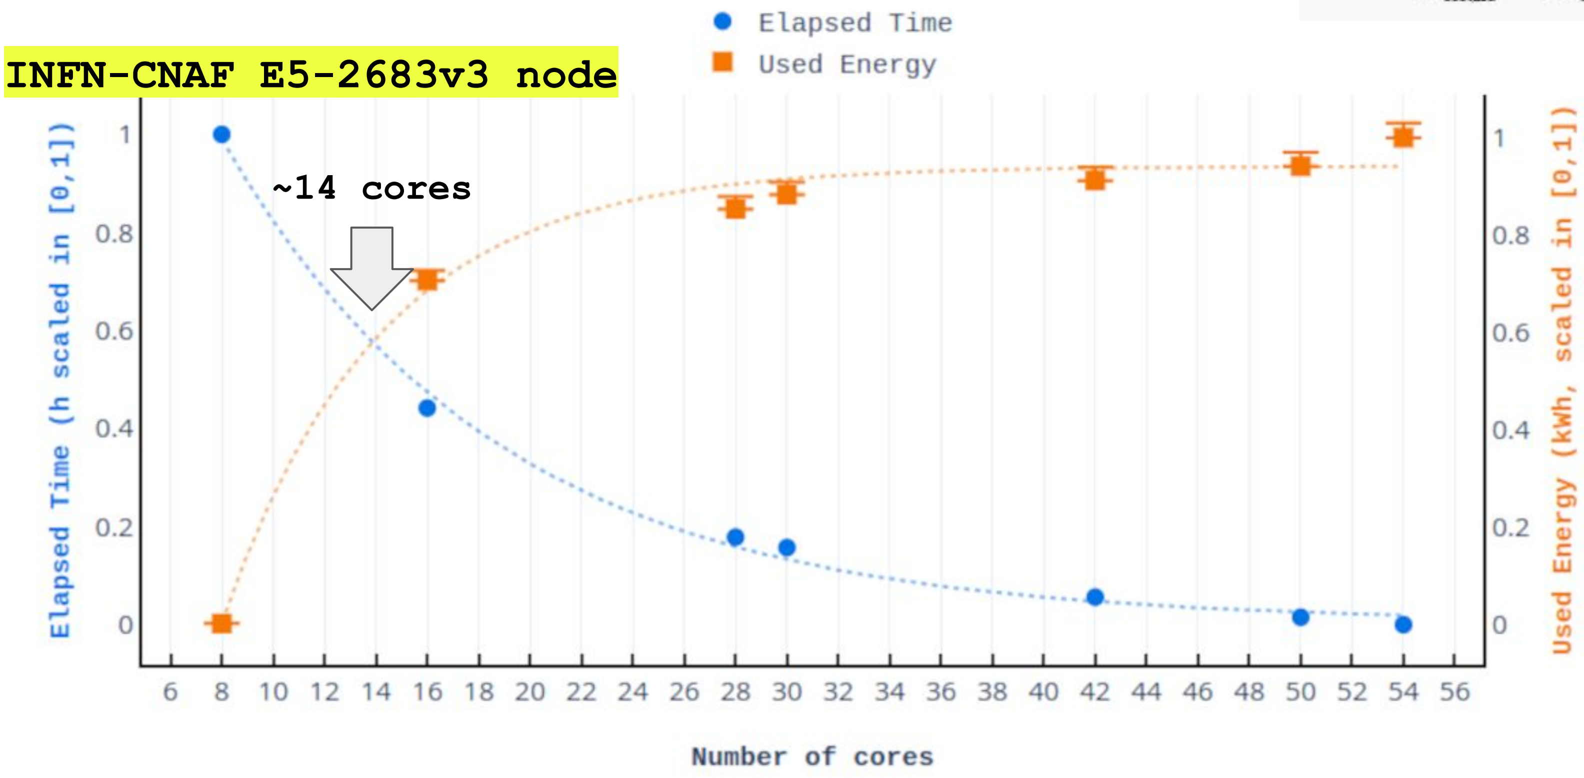

WP validation

INFN-CNAF E5-2640v2 node

Experiment

➔ 1 GENSIM workload
run over 30 cores

➔ 2 GENSIM workloads
run over 14 cores each

Number of events per core tailored to obtain the same final event size

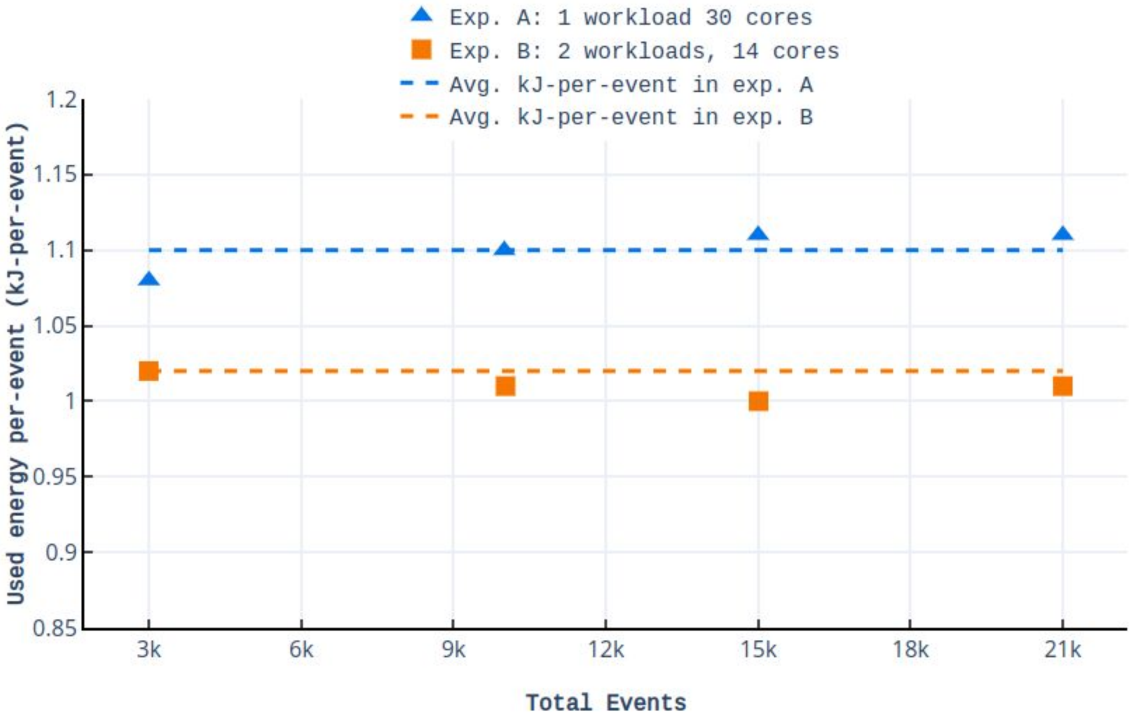

Same event size
with less energy expense!
(at the price of running for longer)

CMS UNIVERSITY OF NOTRE DAME

Conclusions

➔ Milestones
- The developed tool can measure the energy consumption of an arbitrary computing task
- Possible margins for performance improvements energy-wise
- Optimizing the job submission seems a promising task in this perspective

➔ Future outlook
- More detailed estimation of errors and biases
- Structuring the PoC for footprint data collection into a database
- Automatic checkout pipeline of footprint (e.g., power data analyses).