NET ZERO HPC - NOBLE DREAM OR INEVITABLE GOAL?

Prof Michèle Weiland
Director of Research & Met Office Joint Chair

IPCC's definition of Net Zero

Net zero CO_2 emissions

Net zero *carbon dioxide (CO_2)* emissions are achieved when *anthropogenic* CO_2 emissions are balanced globally by anthropogenic CO_2 removals over a specified period. Net zero CO_2 emissions are also referred to as carbon neutrality. See also *Net zero emissions* and *Net negative emissions.*

IPCC, 2018: Annex I: Glossary [Matthews, J.B.R. (ed.)]. In: *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty* [Masson-Delmotte, V., P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, P.R. Shukla, A. Pirani, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J.B.R. Matthews, Y. Chen, X. Zhou, M.I. Gomis, E. Lonnoy, T. Maycock, M. Tignor, and T. Waterfield (eds.)]. Cambridge University Press, Cambridge, UK and New York, NY, USA, pp. 541-562, doi:10.1017/9781009157940.008.

Every tonne of CO_2 emissions adds to global warming

Global surface temperature increase since 1850–1900 (°C) as a function of cumulative CO_2 emissions ($GtCO_2$)

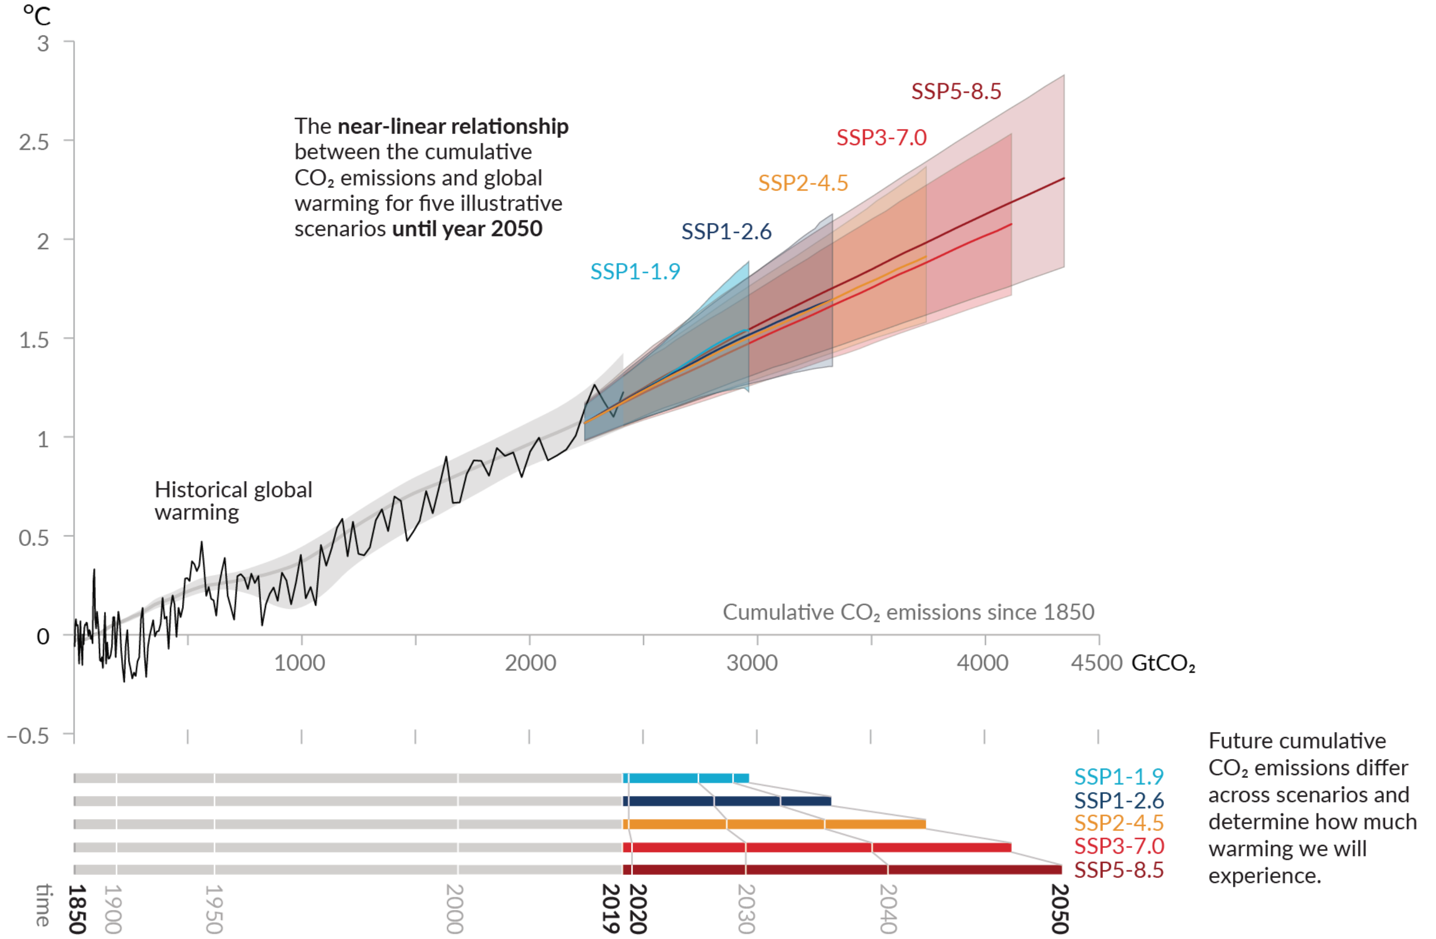

Figure SPM.10 in IPCC, 2021: Summary for Policymakers. In: ***Climate Change 2021: The Physical Science Basis***. *Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)]. Cambridge University Press, Cambridge, UK and New York, NY, USA, pp. 3–32, doi: 10.1017/9781009157896.001

Net Zero Digital Research Infrastructure

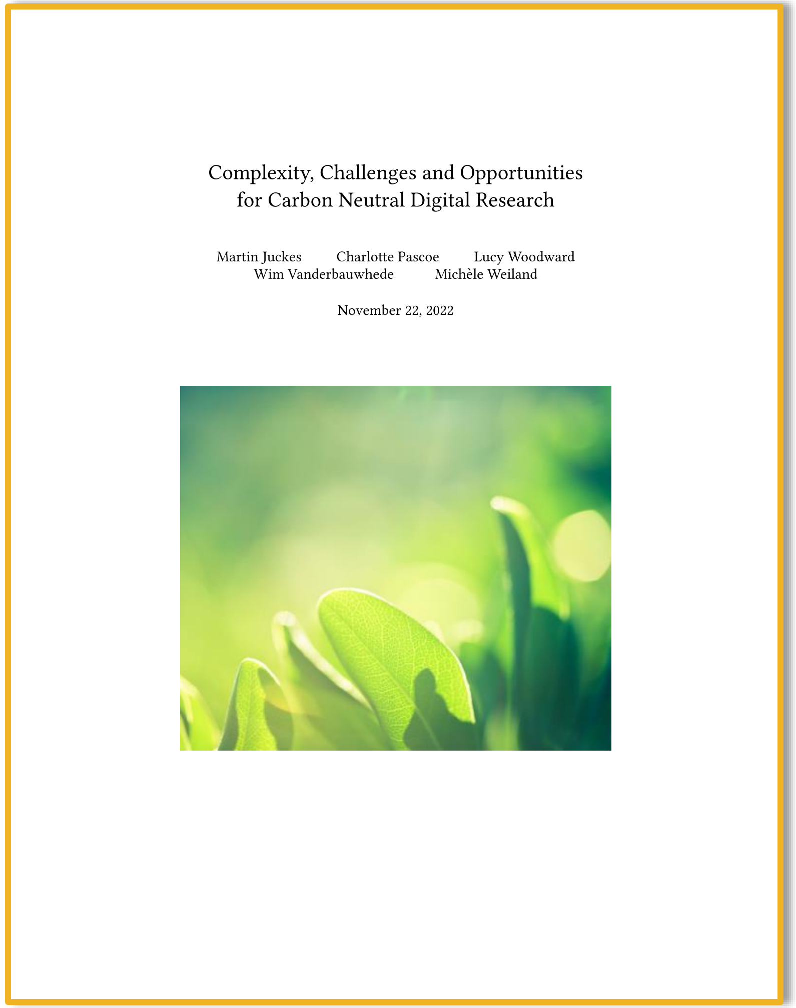
Complexity, Challenges and Opportunities
for Carbon Neutral Digital Research

Martin Juckes Charlotte Pascoe Lucy Woodward
Wim Vanderbauwhede Michèle Weiland

November 22, 2022

“This report is addressed to UKRI Digital Research Infrastructure (DRI) stakeholders in UK Research and Innovation (UKRI), its constituent Research Councils and in the Universities and other institutions which own and operate many digital research facilities which are majority funded by UKRI.

The aim of the UKRI Net Zero Digital Research Infrastructure Scoping Project (hereafter “the project”) is to report on evidence and make recommendations for a roadmap to a Net Zero UKRI DRI by 2040 or sooner. The project, via this interim report and future reports, provides recommendations for reducing and avoiding carbon emissions. It also reviews options for dealing with unavoided emissions through carbon capture, biochar and offsetting.”

https://net-zero-dri.ceda.ac.uk

HPC & CO^2 emissions

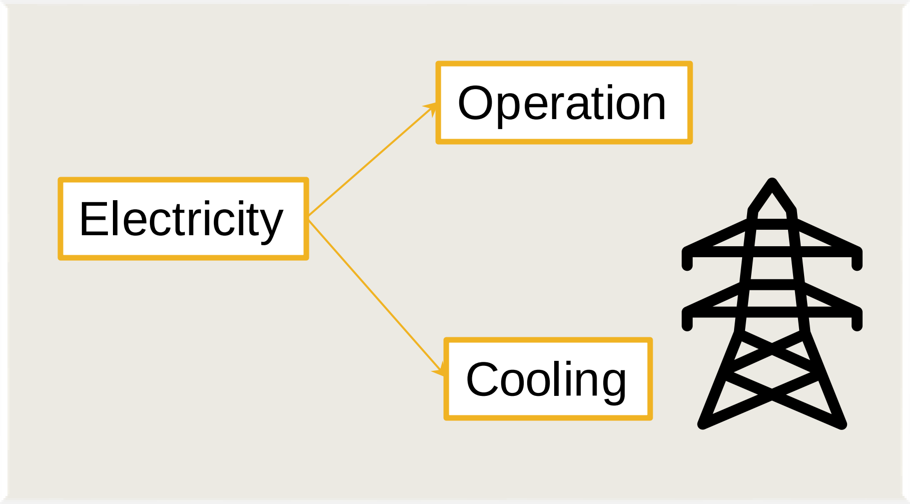

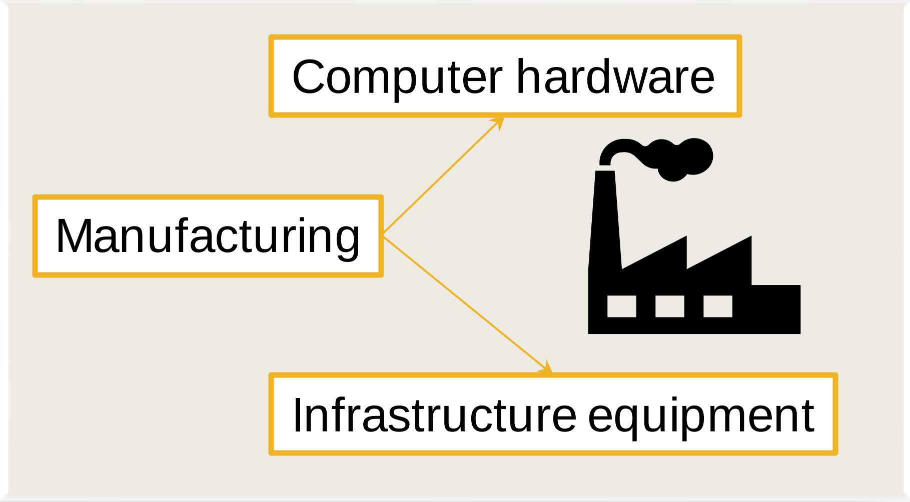

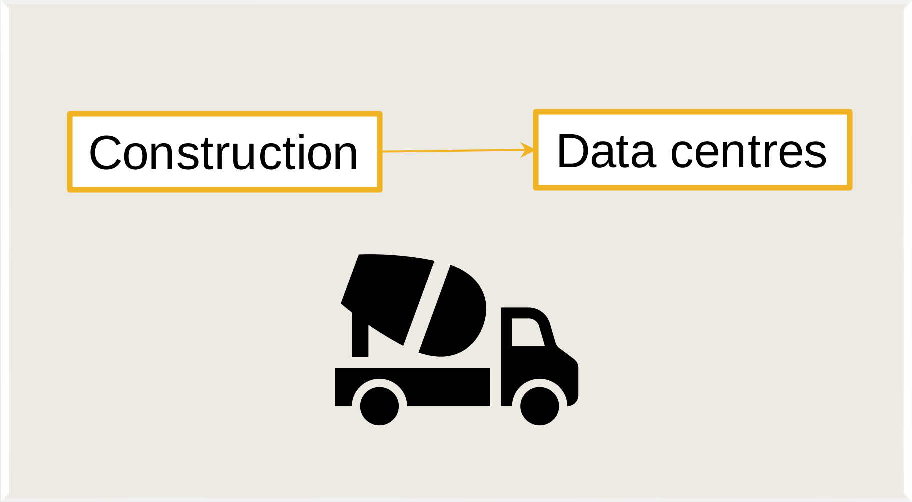

- Power reported for Top10 systems over the past 5 years
- #1 system ranges from 9.7MW to 30MW
- #10 system in Nov’22 has 3rd highest power draw
- Only in 2022 all Top10 systems submitted power
- Accumulative power for Top10 in 2022: 120MW

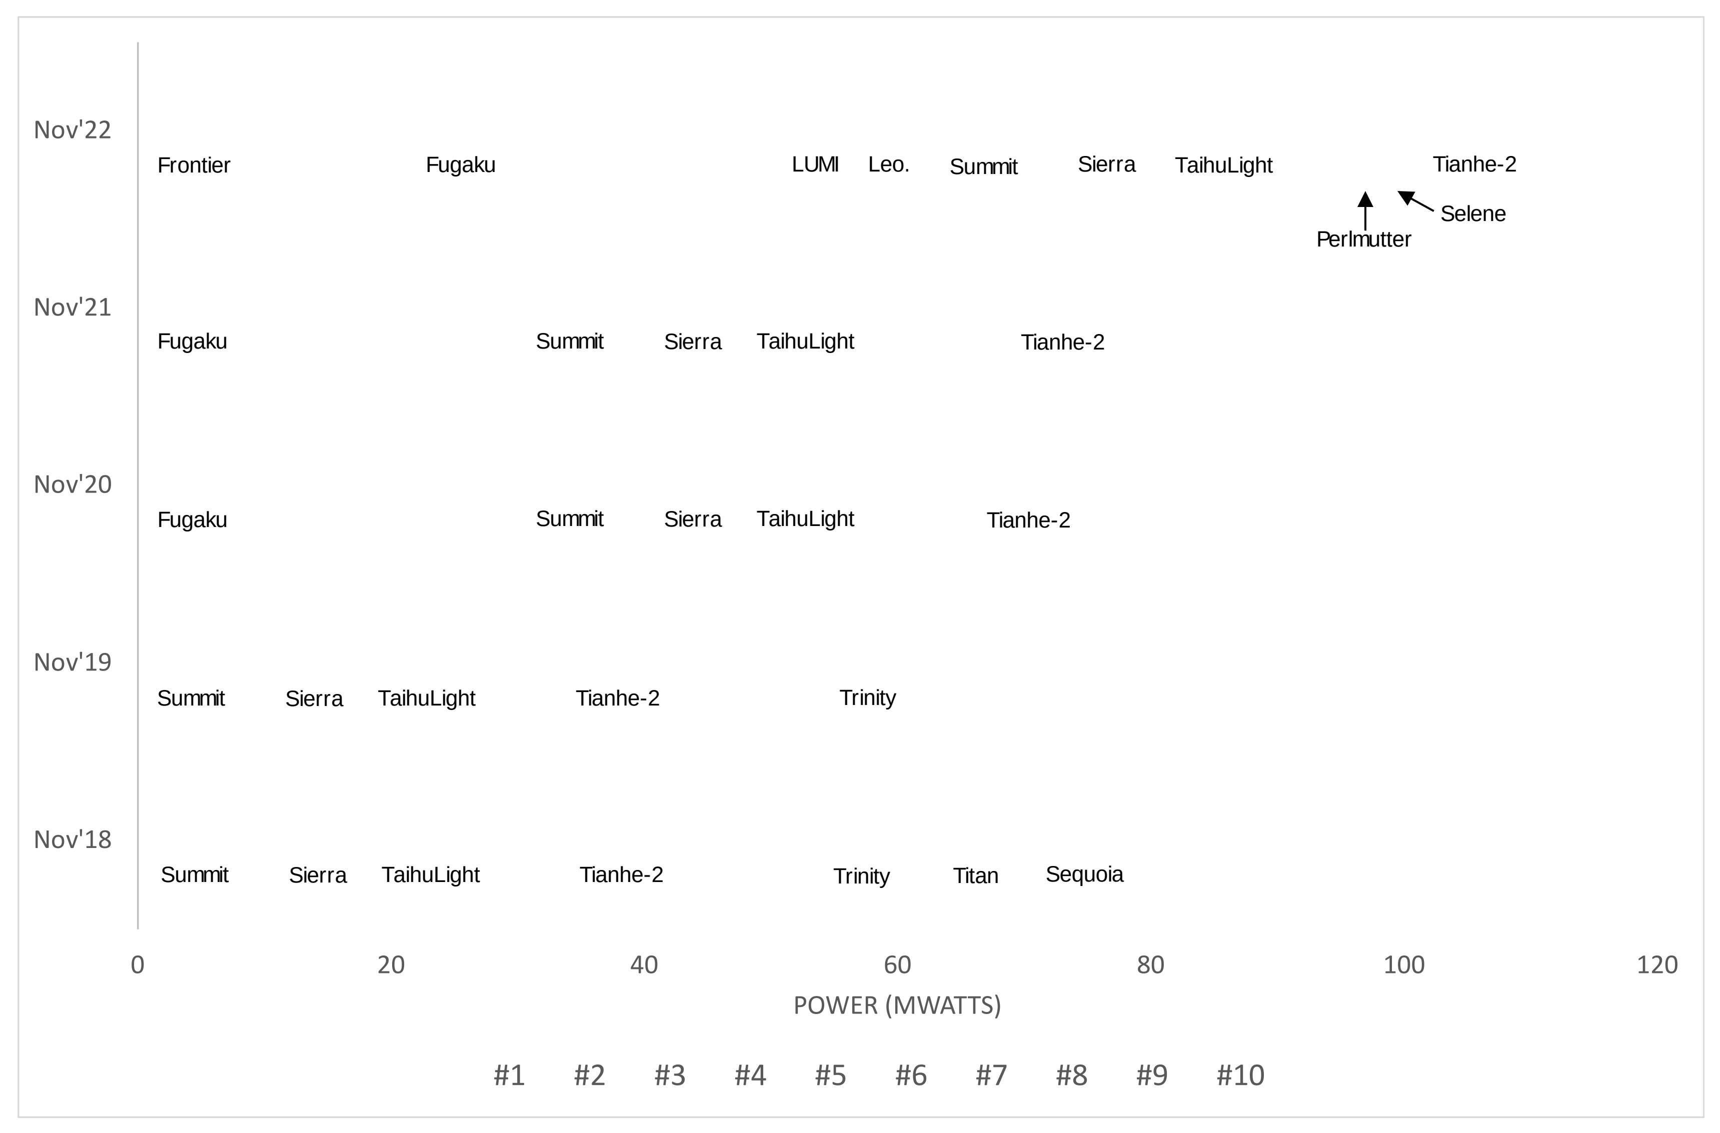

- HPL & HPCG performance normalized to Frontier
- >1 means “better”

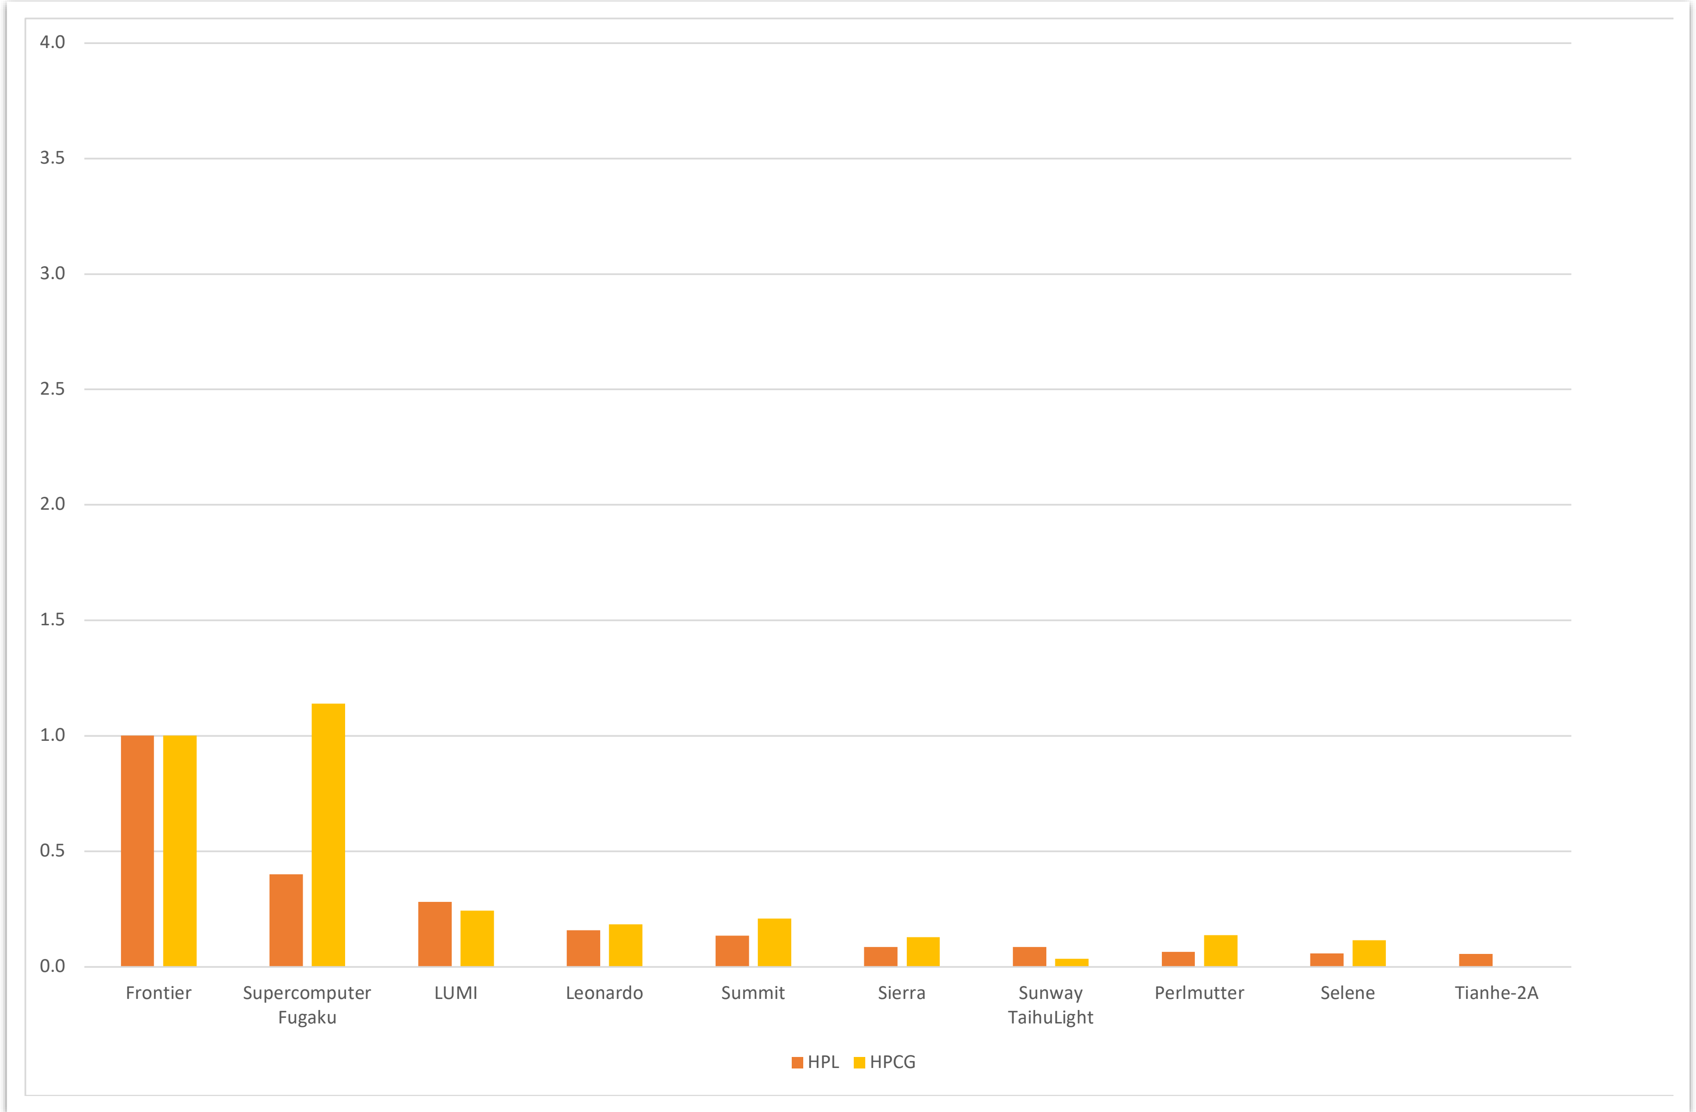

- HPL & HPCG performance normalized to Frontier
- >1 means "better"
- HPL % of peak close to Frontier for all systems
  - 70-75% on average

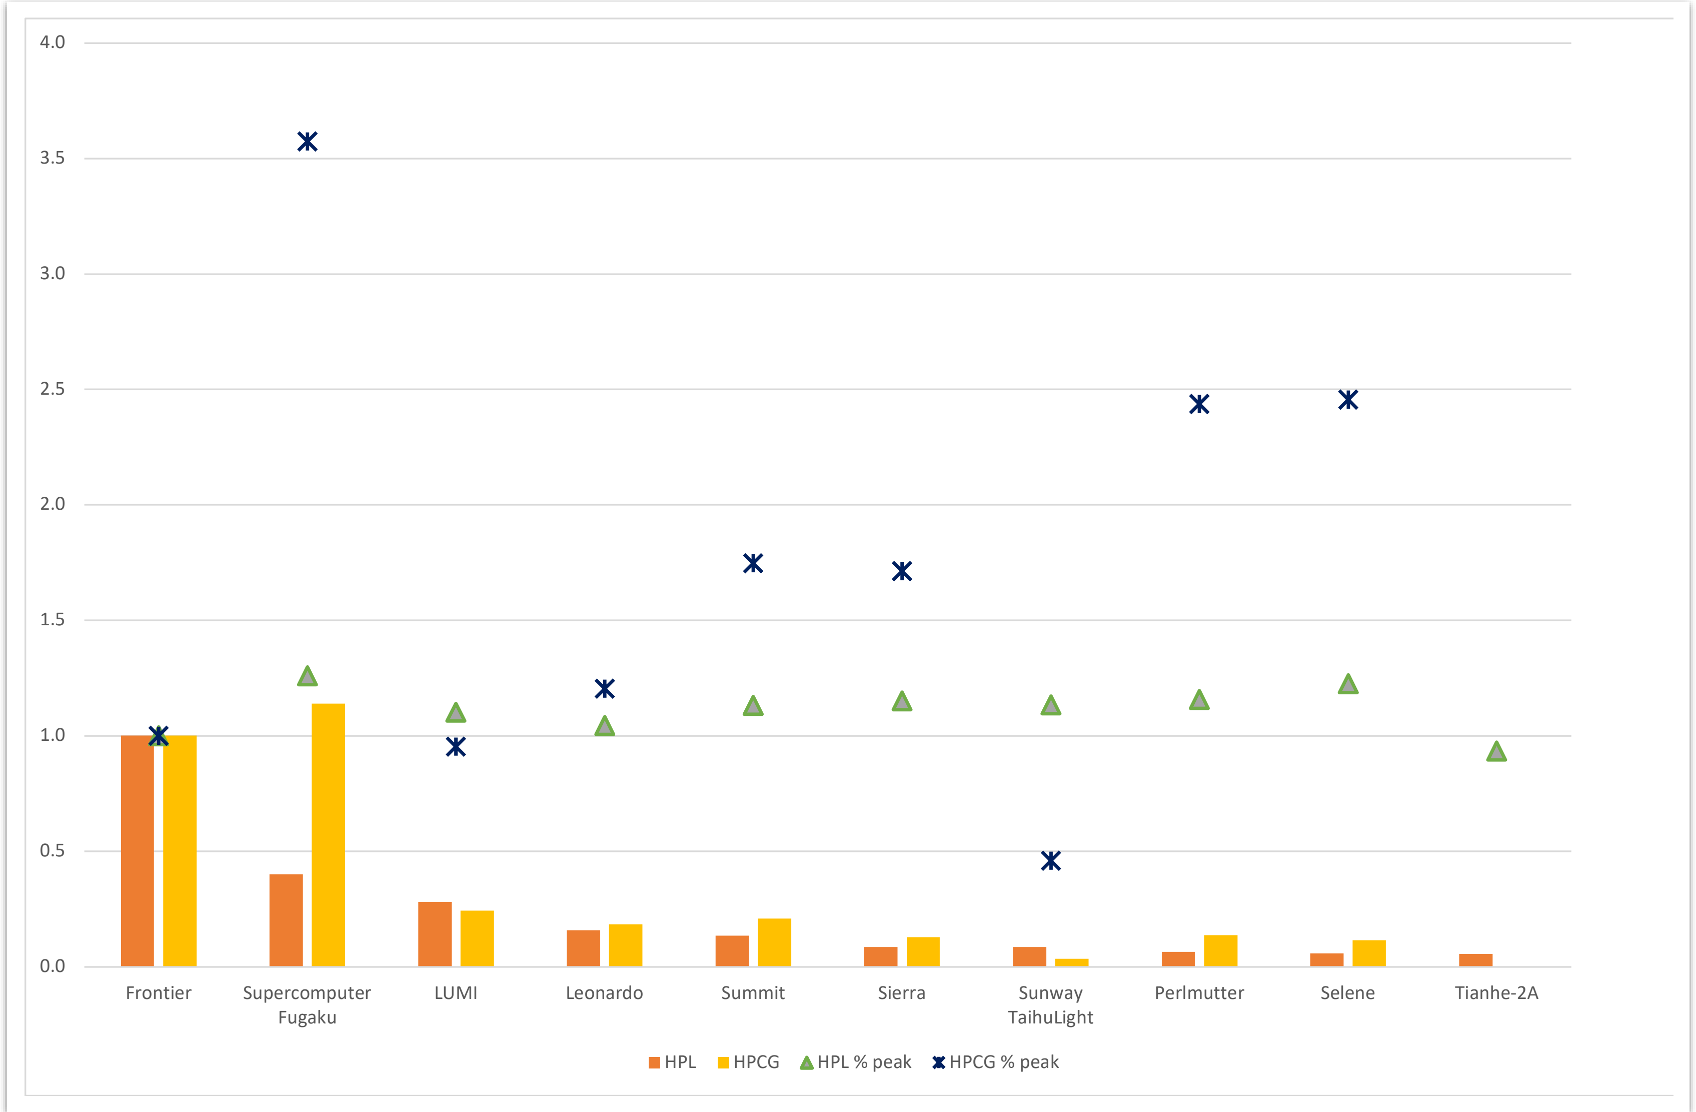

- HPL & HPCG performance normalized to Frontier
- >1 means “better”
- HPL % of peak close to Frontier for all systems
  - 70-75% on average
- Green500 list uses HPL power to evaluate efficiency

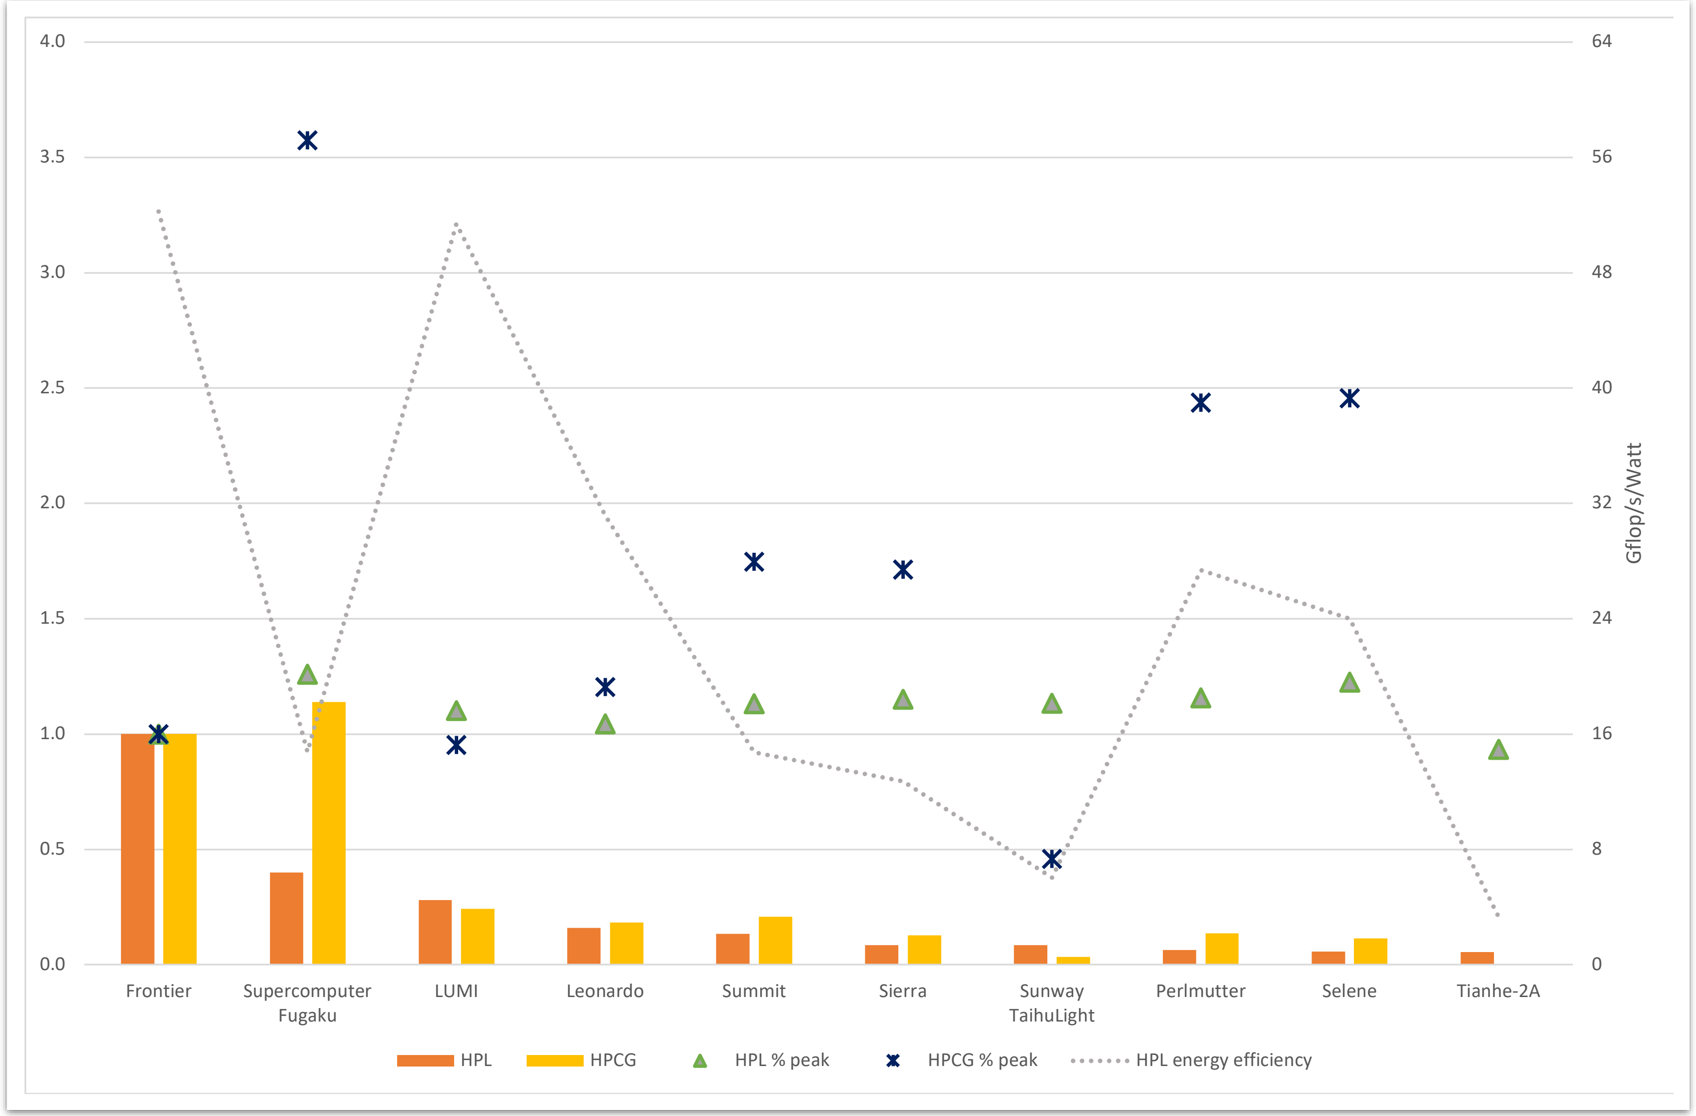

- HPL & HPCG performance normalized to Frontier
- >1 means "better"
- HPL % of peak close to Frontier for all systems
  - 70-75% on average
- Green500 list uses HPL power to evaluate efficiency
- "HPCG" energy efficiency assumes 50% of HPL power

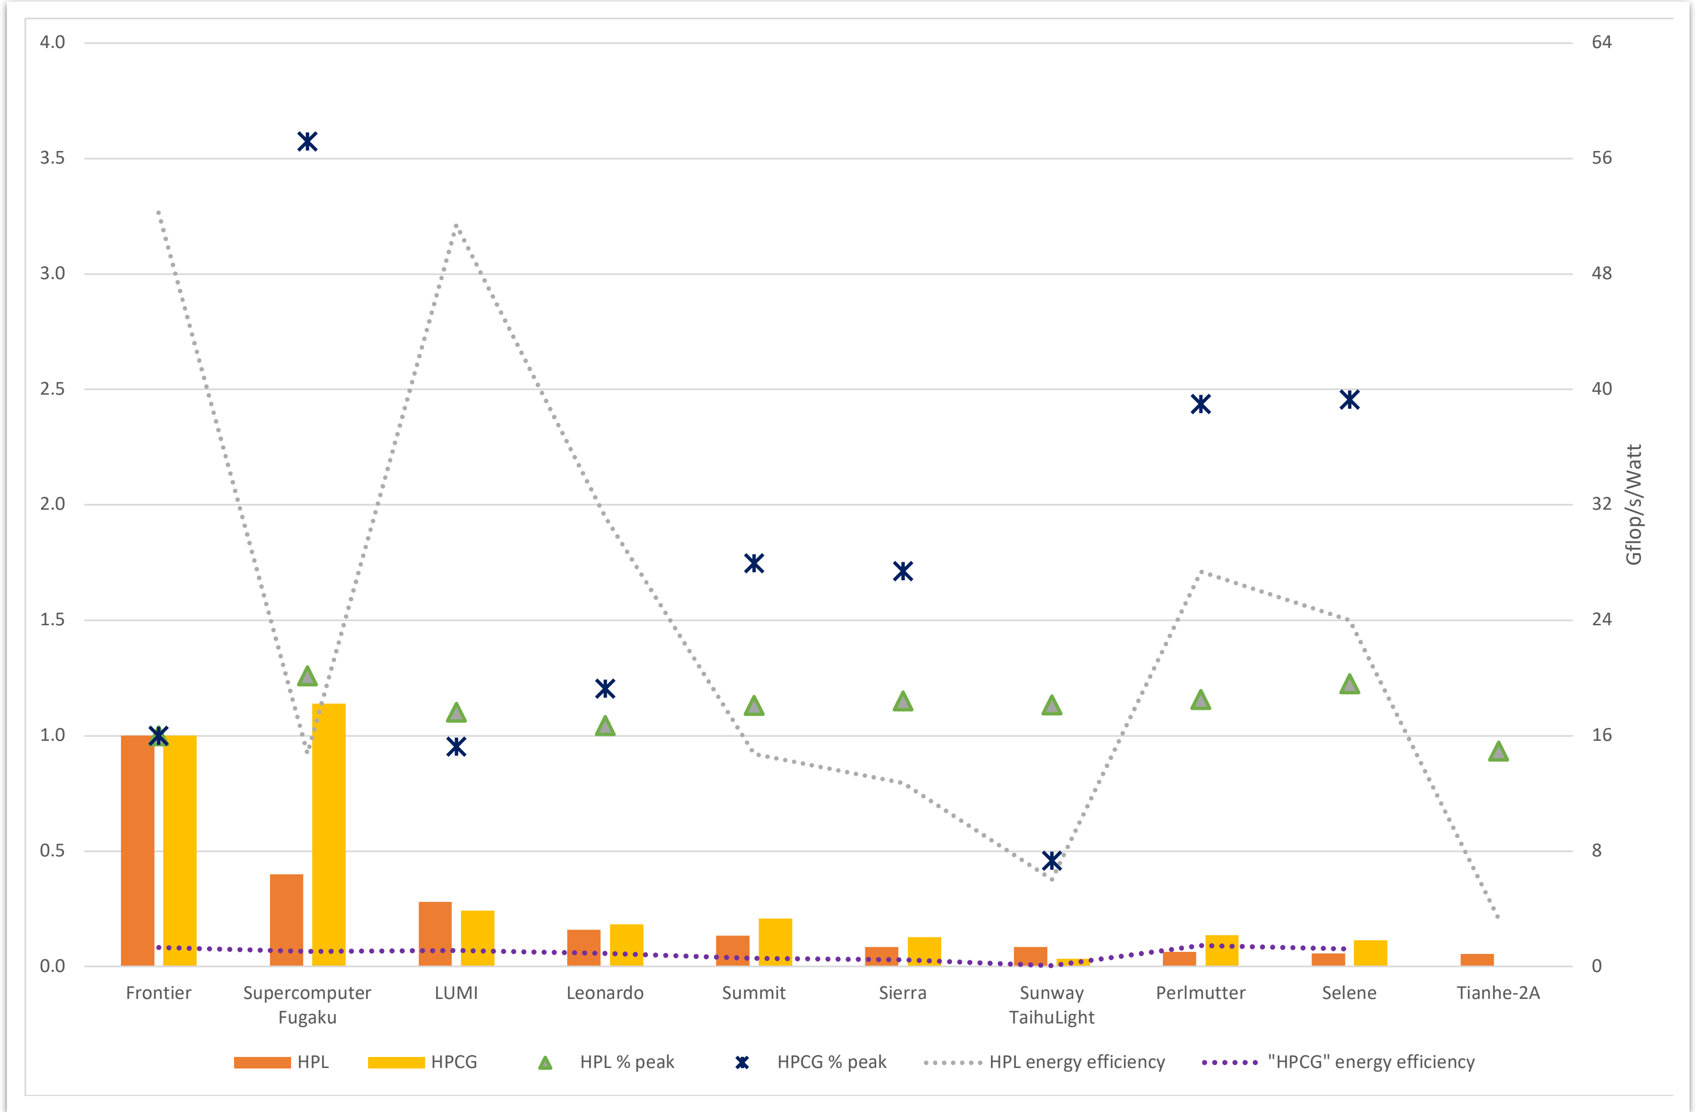

Impact of inefficient I/O configuration

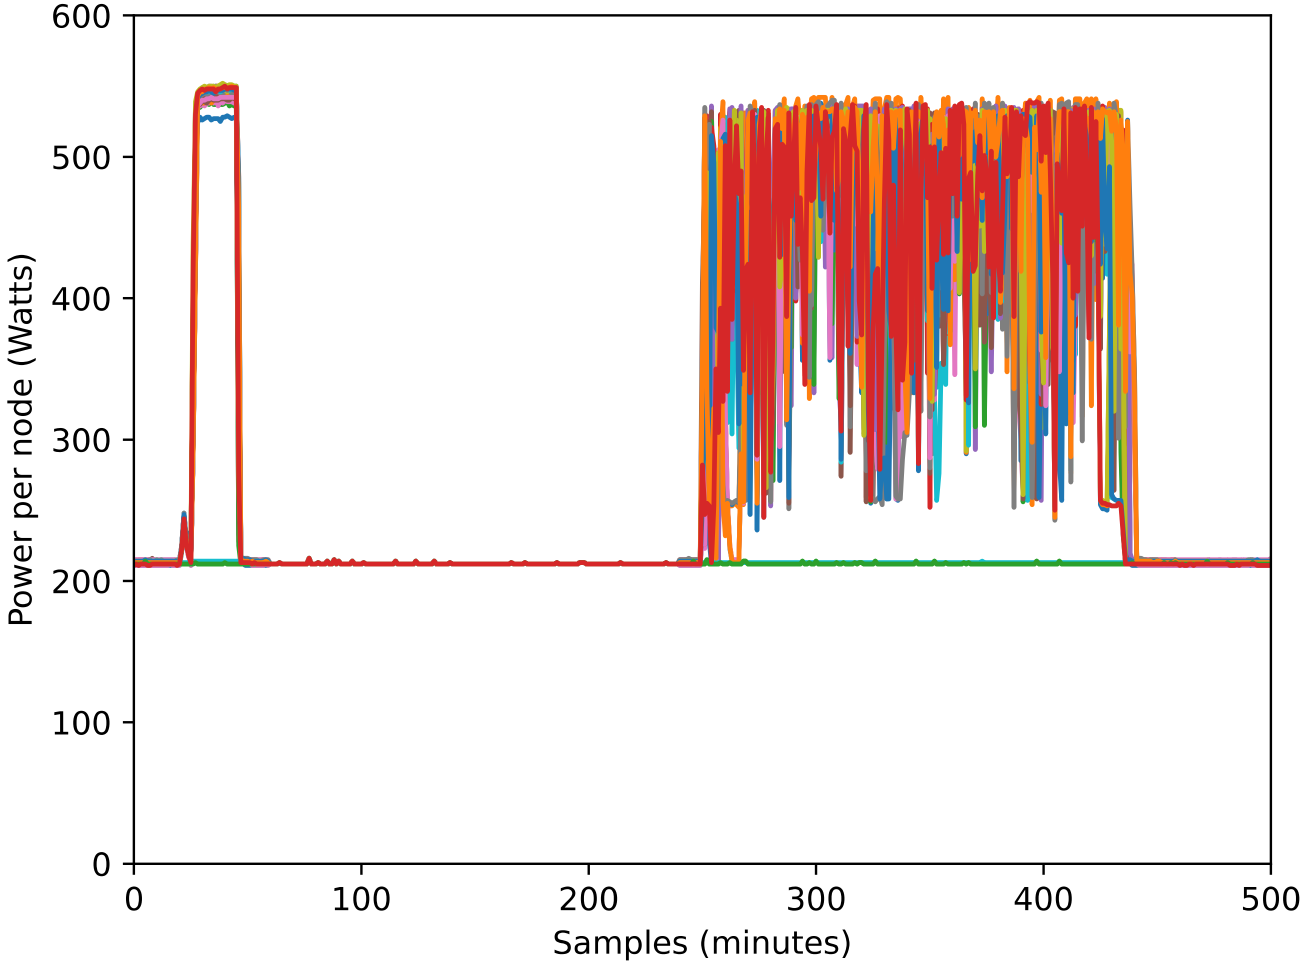

- Node-level power measurement
  - Each line represents power draw for 1 node
  - Full system, 34 nodes in total
  - Idle power draw: 213W
- Two identical aerodynamics simulations with OpenFOAM using 32 nodes
  - On the left: **no** I/O
  - On the right: **excessive** I/O

Impact of inefficient I/O configuration

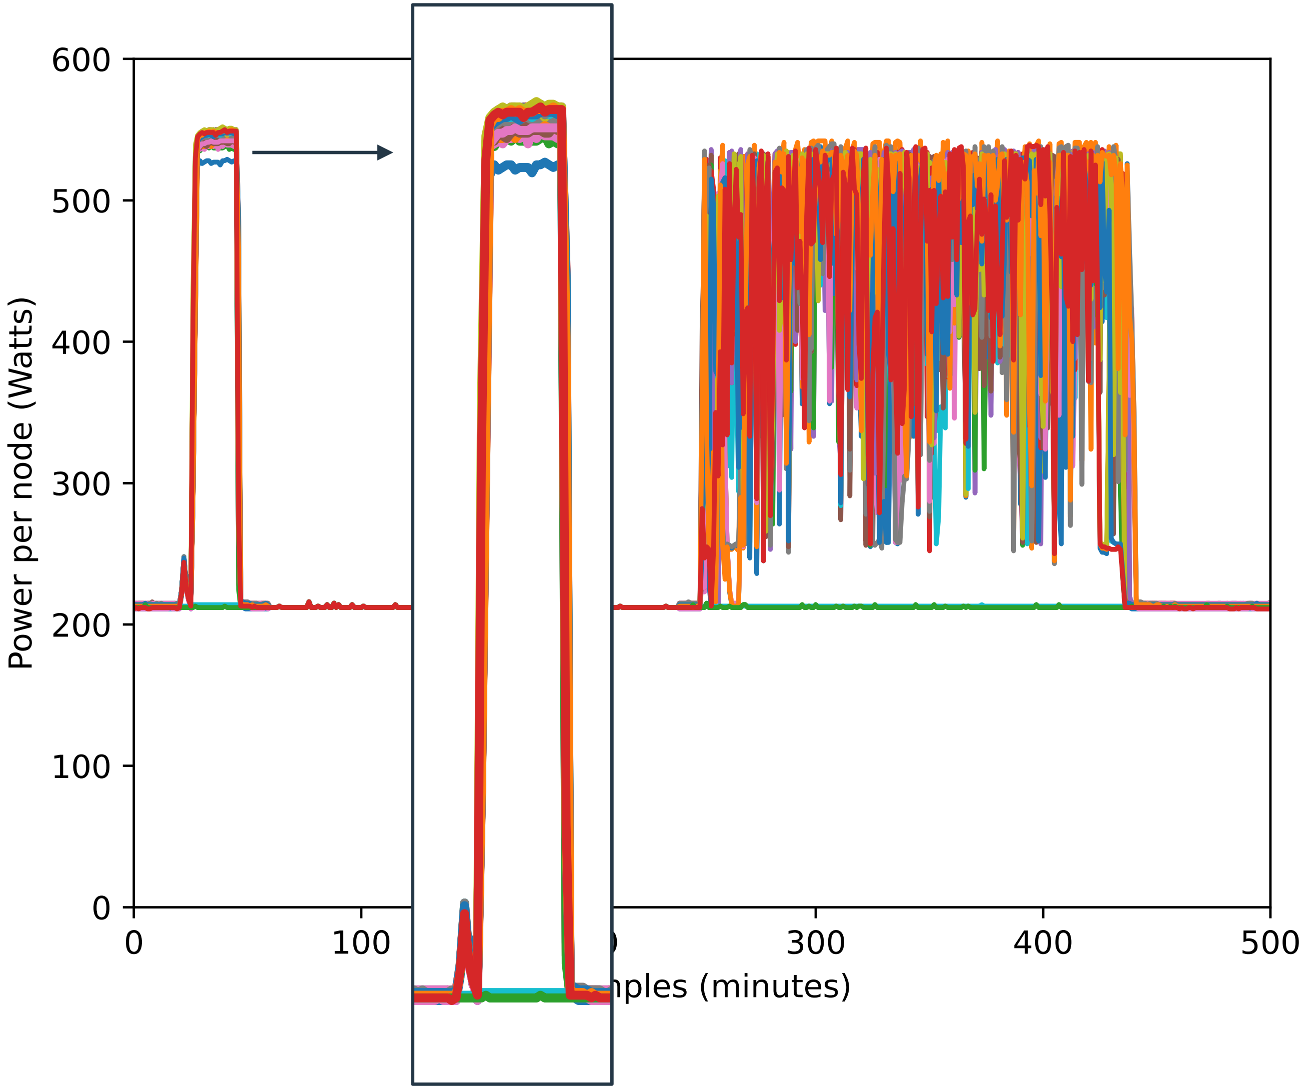

- Node-level power measurement
  - Each line represents power draw for 1 node
  - Full system, 34 nodes in total
  - Idle power draw: 213W
- Two identical aerodynamics simulations with OpenFOAM using 32 nodes
  - On the left: **no** I/O
  - On the right: **excessive** I/O

Impact of inefficient I/O configuration

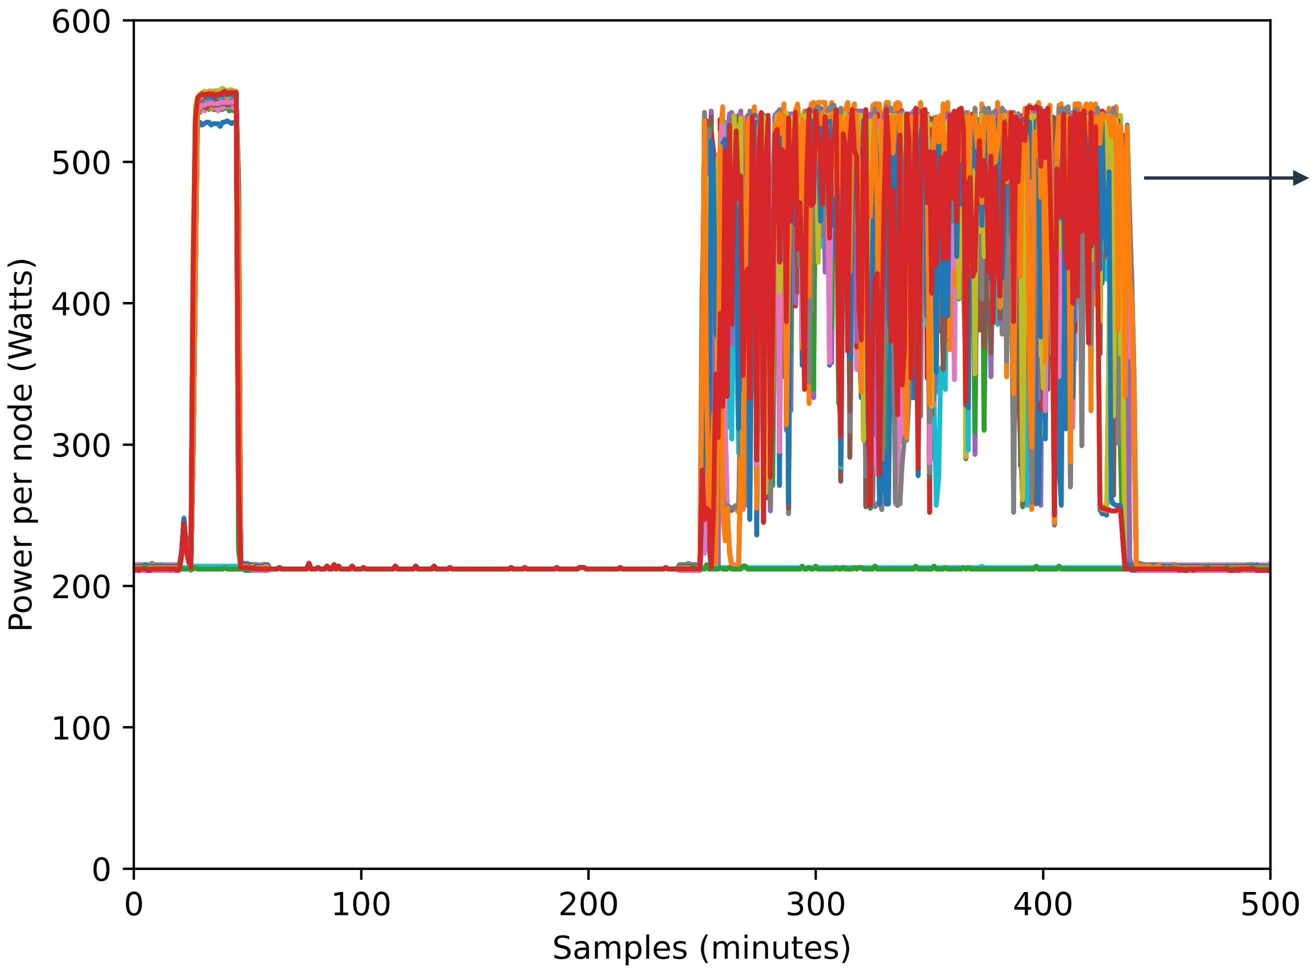

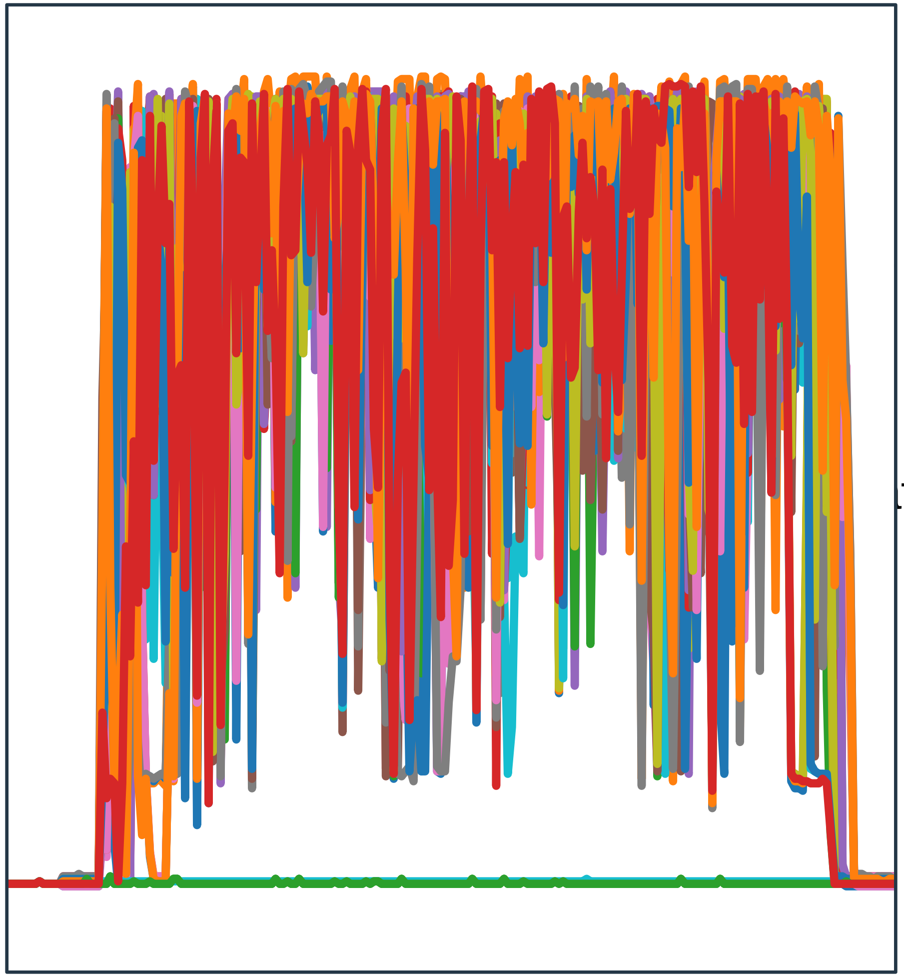

1 node

tions

Impact of inefficient I/O configuration

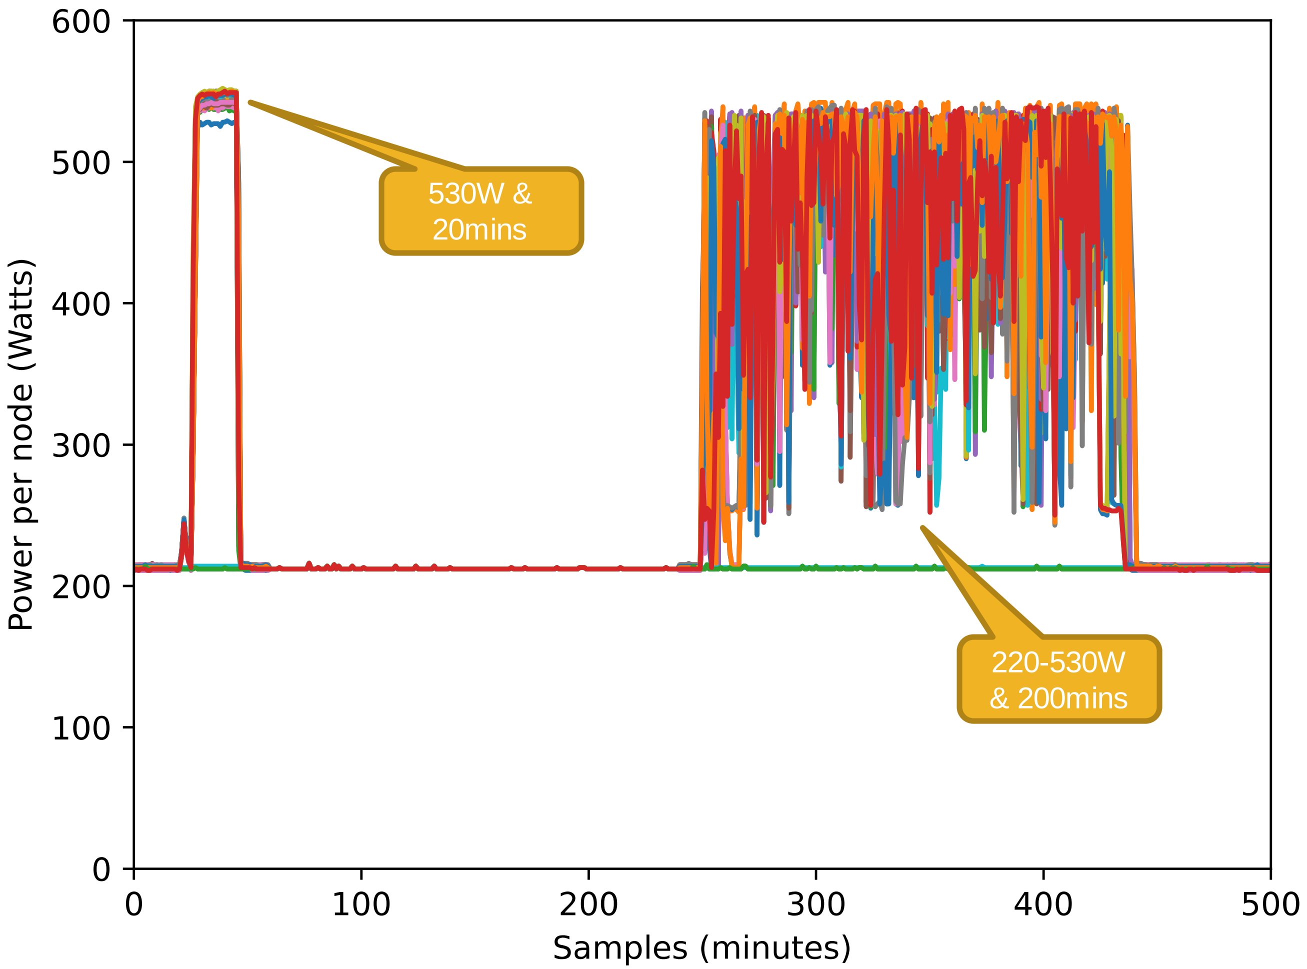

- Node-level power measurement
  - Each line represents power draw for 1 node
  - Full system, 34 nodes in total
  - Idle power draw: 213W
- Two identical aerodynamics simulations with OpenFOAM using 32 nodes
  - On the left: **no** I/O
  - On the right: **excessive** I/O
- Excessive I/O means network contention & frequent stalling

CPU clock frequency

- Frequency ∝ power draw
- Often taken care of by Dynamic Voltage & Frequency Scaling (DVFS)
  - Heat and power management
  - Influenced by governor

- Might not do what you expect…
- Simple test on Cirrus (Intel Broadwell)
  - FIRESTARTER → 2.2GHz
  - HPCG → 2.4GHz
  - OpenFOAM (simpleFoam) → 2.6GHz

High CPU frequency on memory bound application is wasteful

Very compute bound, CPU runs hot!

Very memory bound, CPU runs cool!

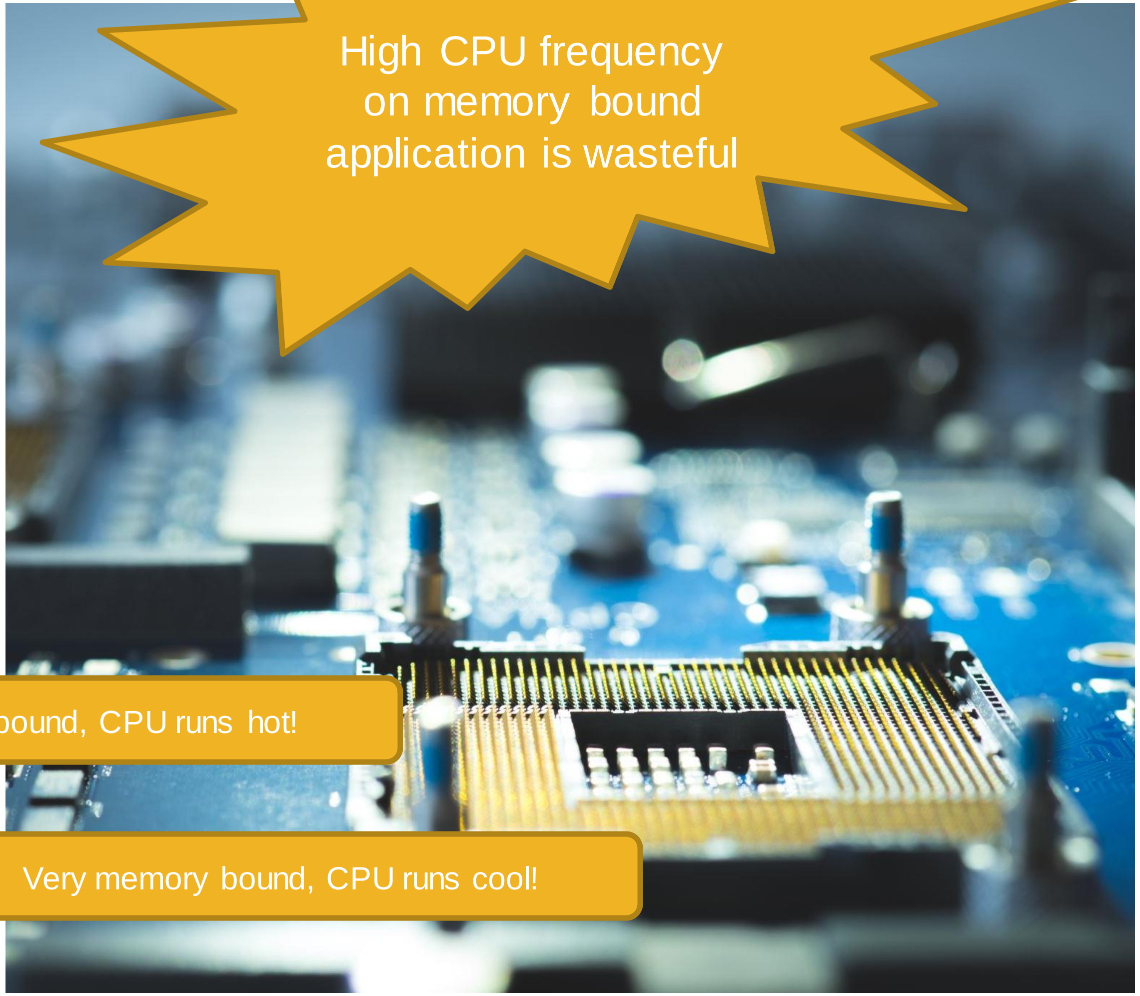

Frequency scaling & STREAM

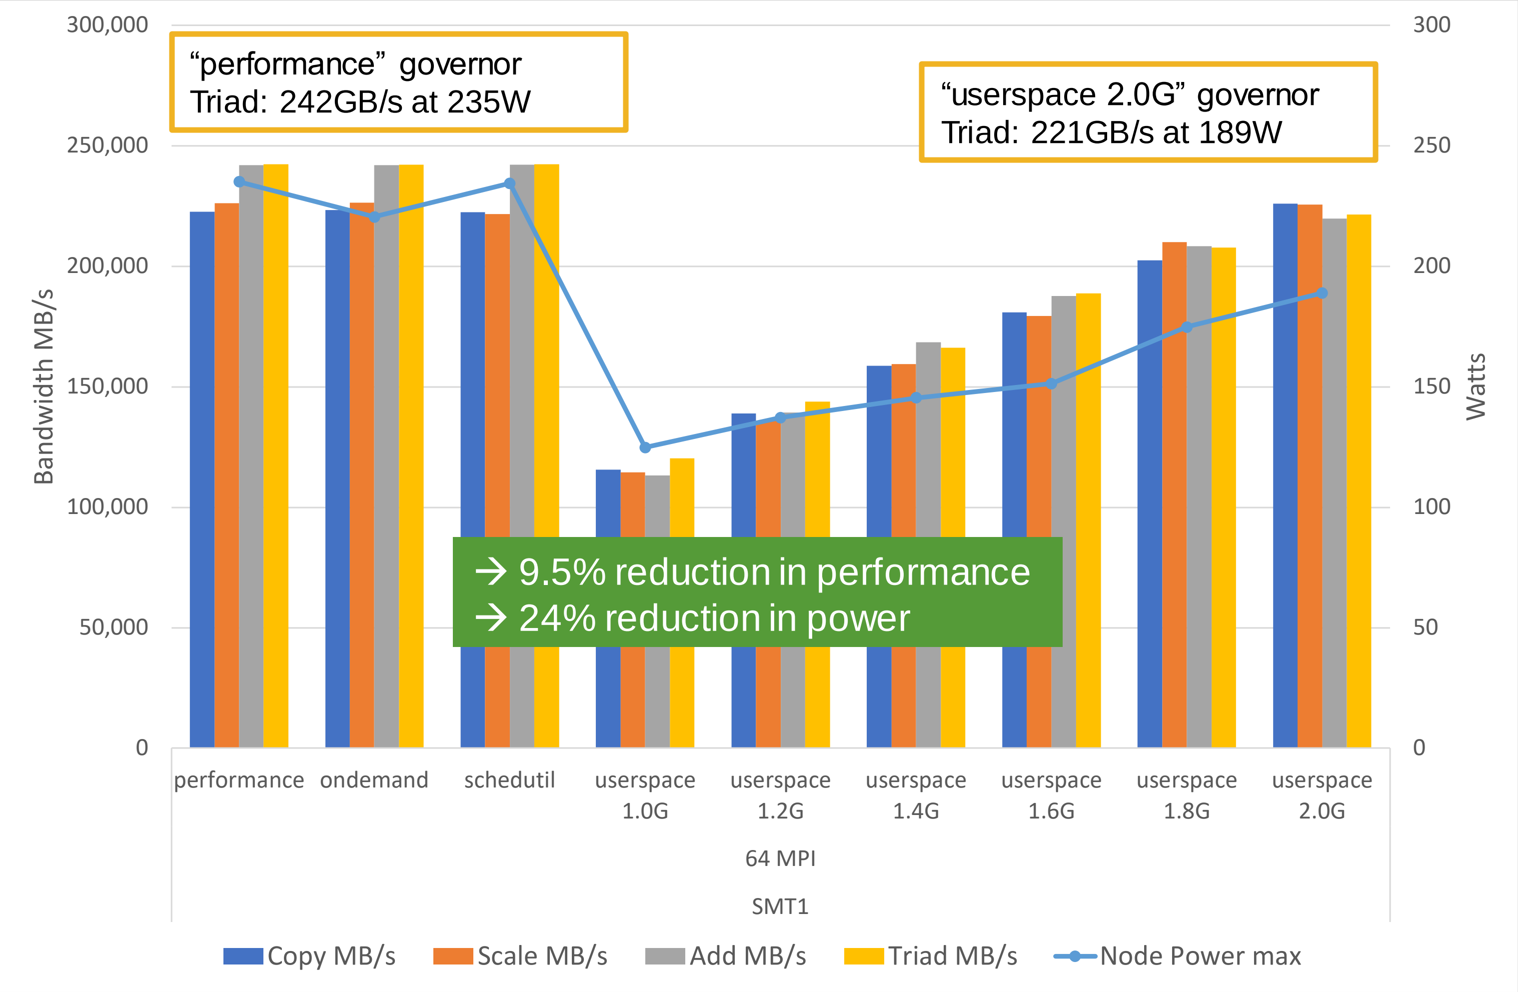

- Marvell ThunderX2 Arm-based test system
- Influencing CPU frequency using different governors
- CPU can boost up to 2.5GHz
  - Default is 2.2GHz

Frequency and memory scaling & STREAM

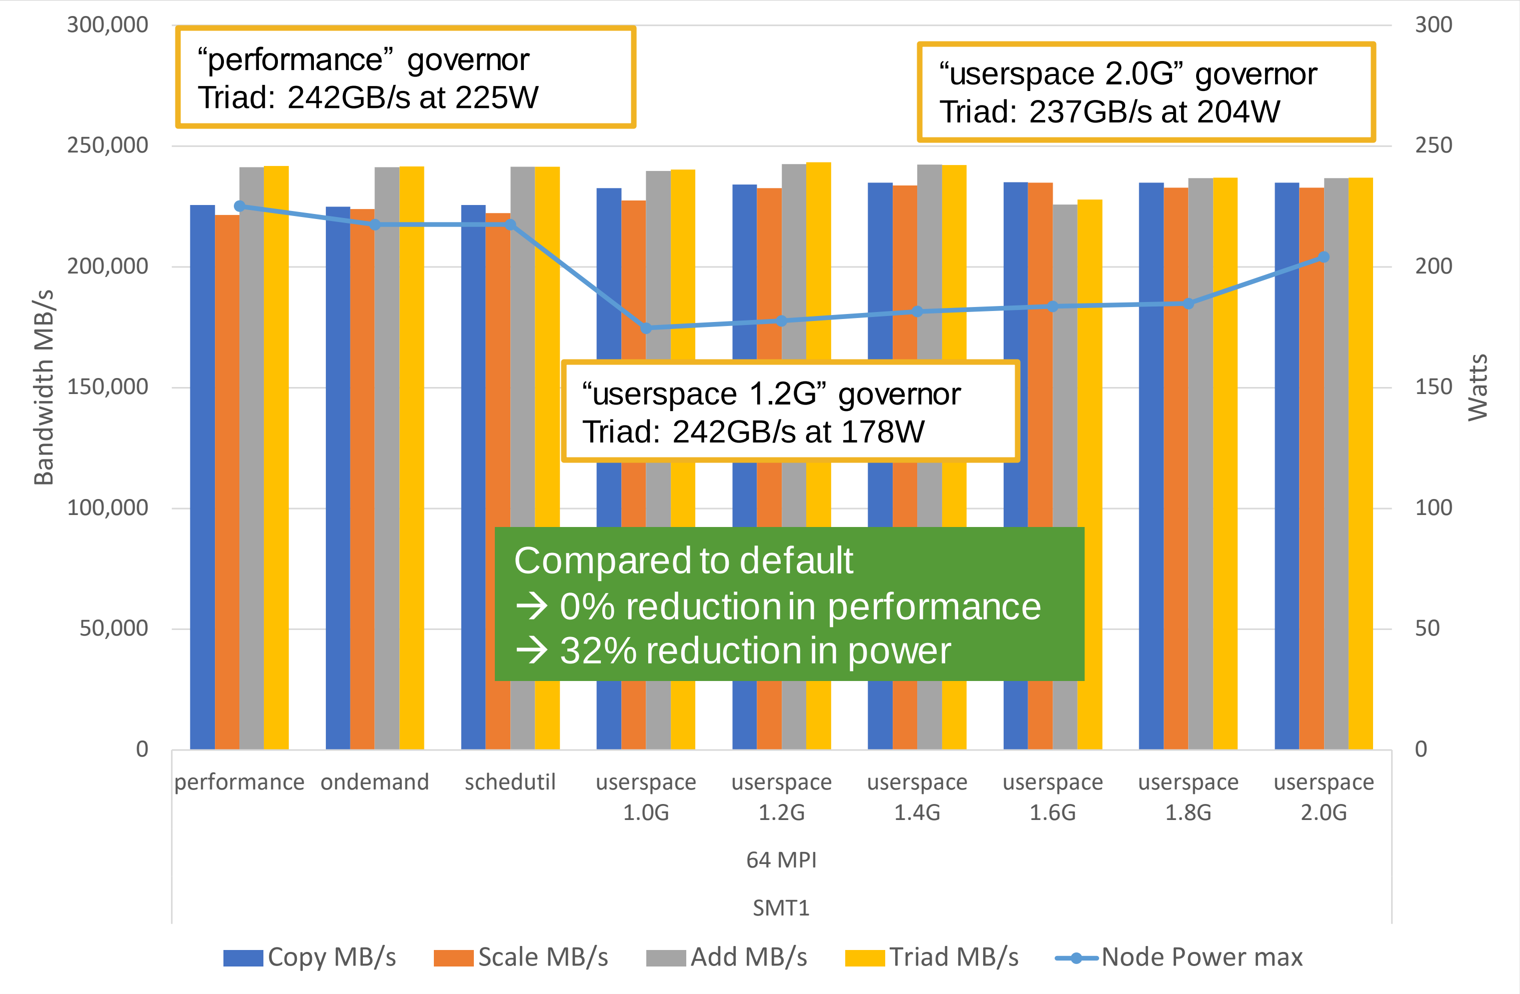

- Marvell ThunderX2 Arm-based test system
- Influencing CPU frequency using different governors
- Influencing memory subsystem frequency by enabling memturbo
  - Increased from 2.2 to 2.3GHz
  - CPU can no longer boost and is limited to 2.2GHz

EPCC’s path towards Net Zero

Long term investment

- Data centre efficiency requires **long term strategic** investment
  - There is of an (inevitable) upfront cost in emissions
- Oldest ACF machine room from 1970s
- Infrastructure must support new developments in power and cooling

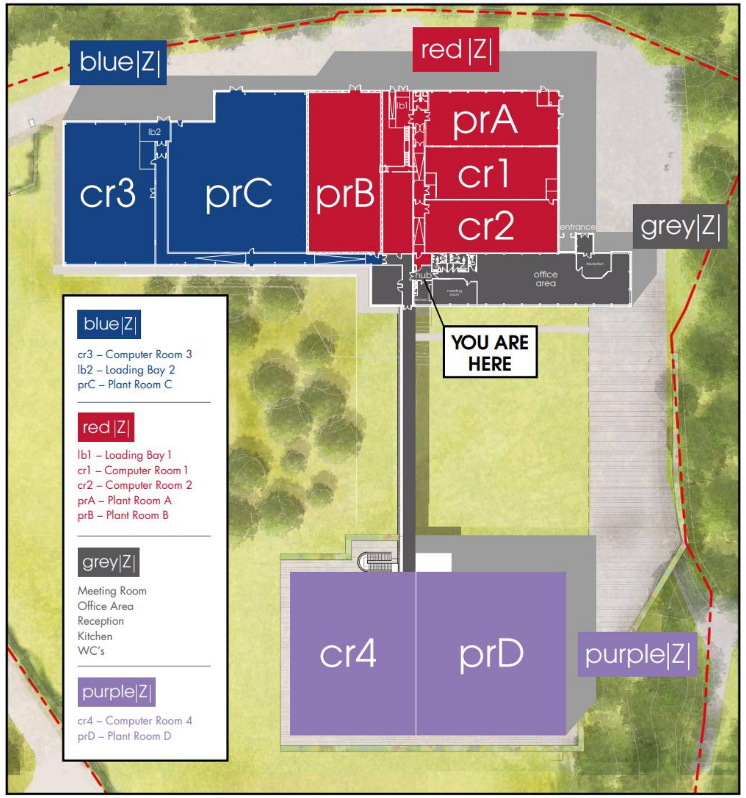

epcc

Renewable energy

Procured using 100% **certified** renewable energy framework agreement

- ✓Renewable Energy Guarantee of Origin (REGO)
- ✓Classed as Net Zero

Both **cost** and **carbon** efficient

So, job done?
Well not quite!

https://www.gov.scot/publications/electricity-framework-agreement/

Weather at the ACF

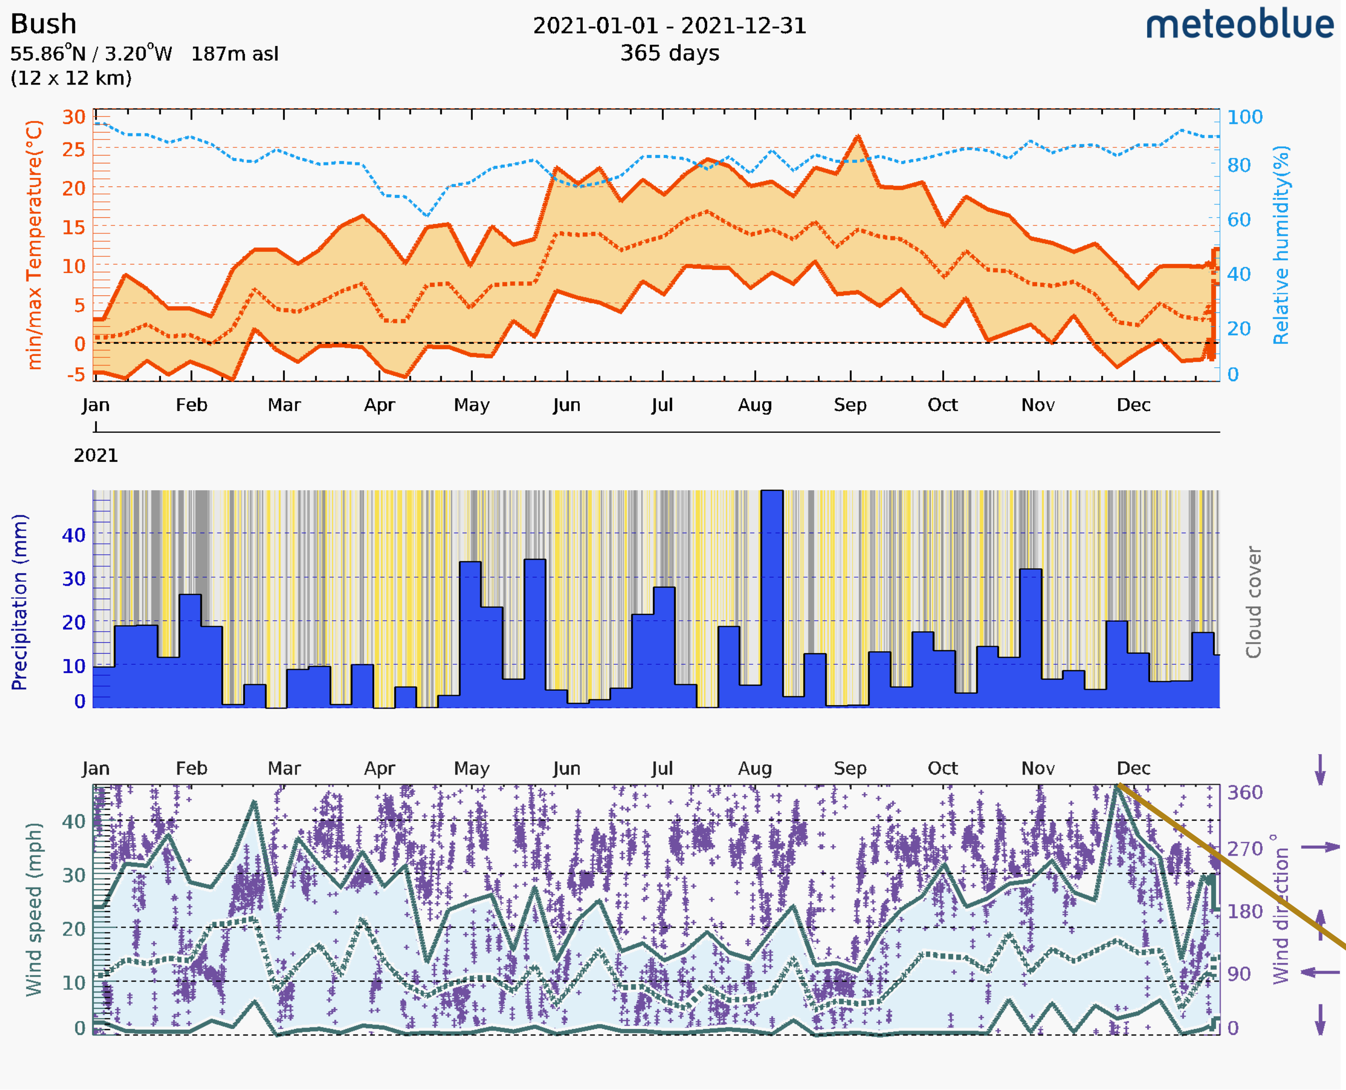

https://www.meteoblue.com/en/weather/historyclimate

Free cooling!

Warm water is pumped to the roof, outside air & fans chill it

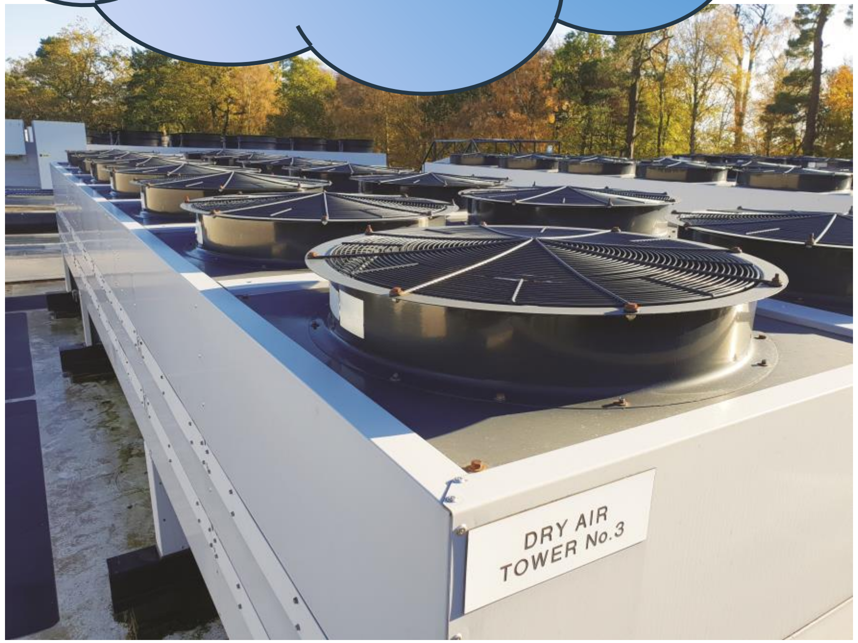

Storm Arwen…

Other forms of cooling we use

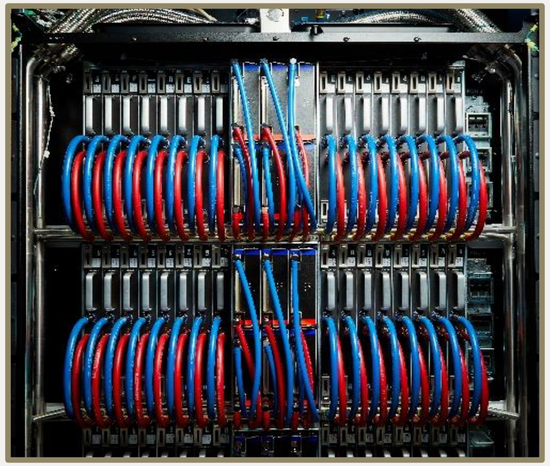

Direct liquid cooling

Adiabatic cooling

Temperature controlled racks

Geothermal battery feasibility project

Problem: We want to be able to reuse our excess heat, but nowhere nearby can use it

Solution: Move the heat to where it is useful

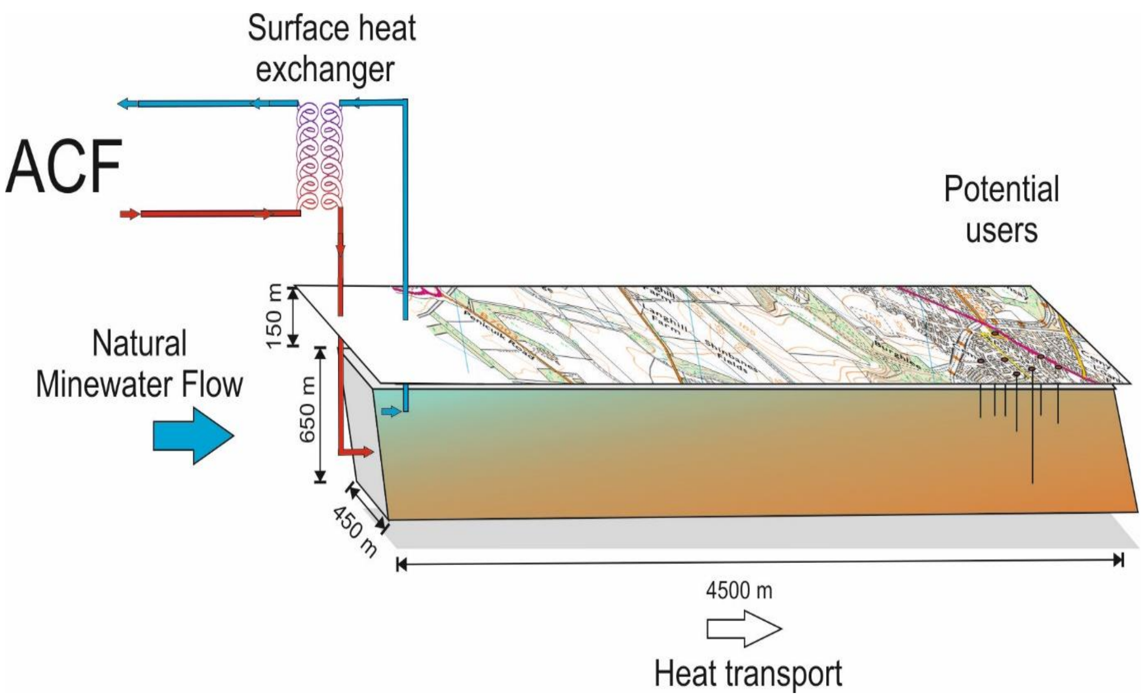

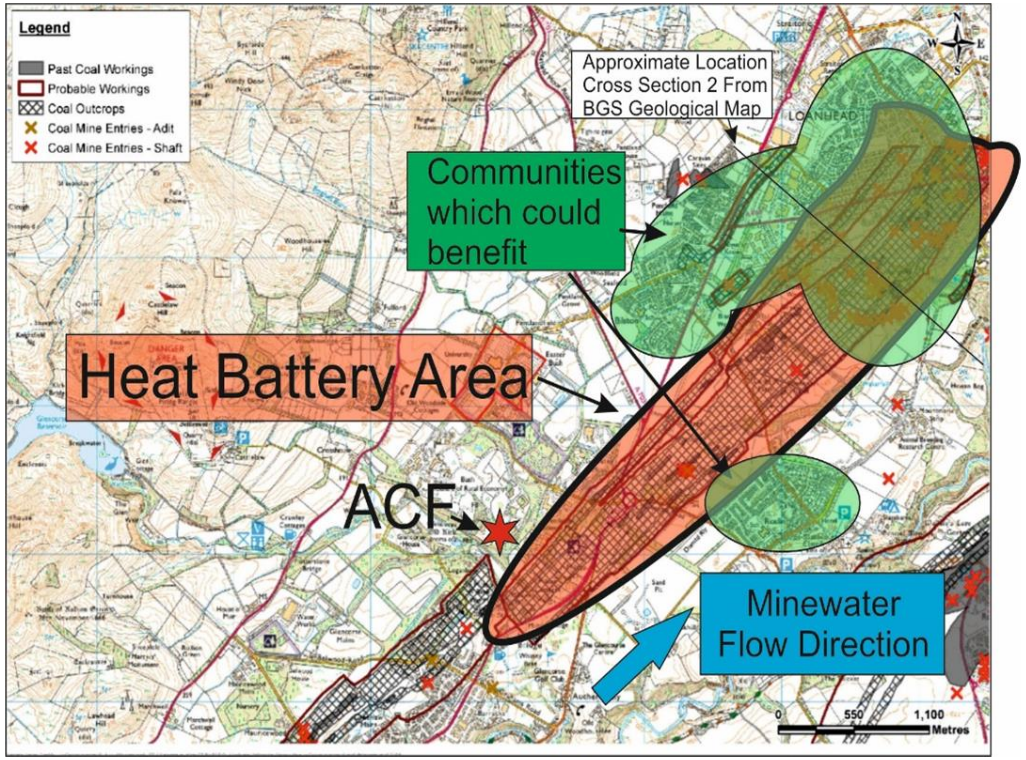

EPYC modes

- **Power** deterministic
  - Allows the highest possible performance
  - CPU will run as fast as it can for given TDP or power input – might vary
- **Performance** deterministic
  - Will deliver the same predictable performance across CPUs
  - Might result in slightly different power consumption

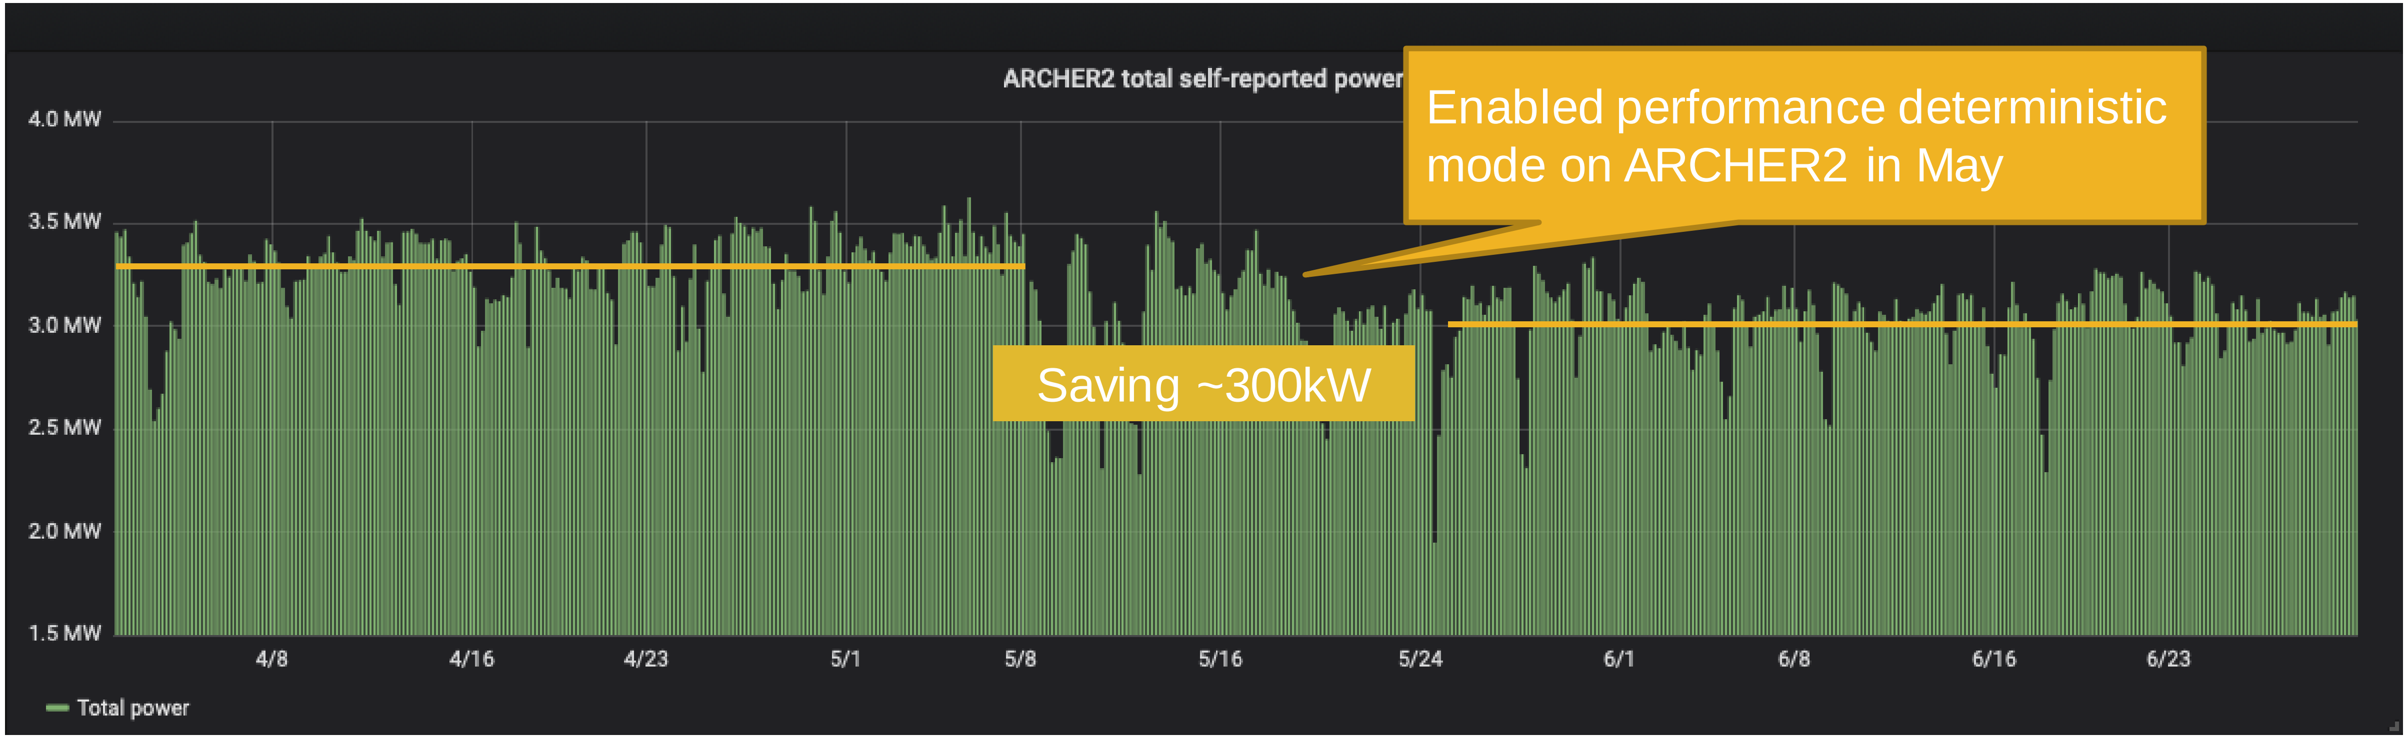

ARCHER2 CPU frequency reduction

Summary of relative energy and performance at 2.00 GHz, compared to 2.25 GHz

Benchmark (single node)	Energy	Performance
VASP (TiO_2)	-20%	-1%
CASTEP (Al Slab)	-13%	-1%
GROMACS (1400k atoms)*	-5%	-15%
OpenSBLI (TGV 512ss)	-20%	-5%
LAMMPS (LJ 8M atoms)**	-4%	-21%
NAMD (STMV 1M atoms)**	-5%	-33%

* Data from Laura Moran, EPCC
** Data from Douglas Shanks, HPE

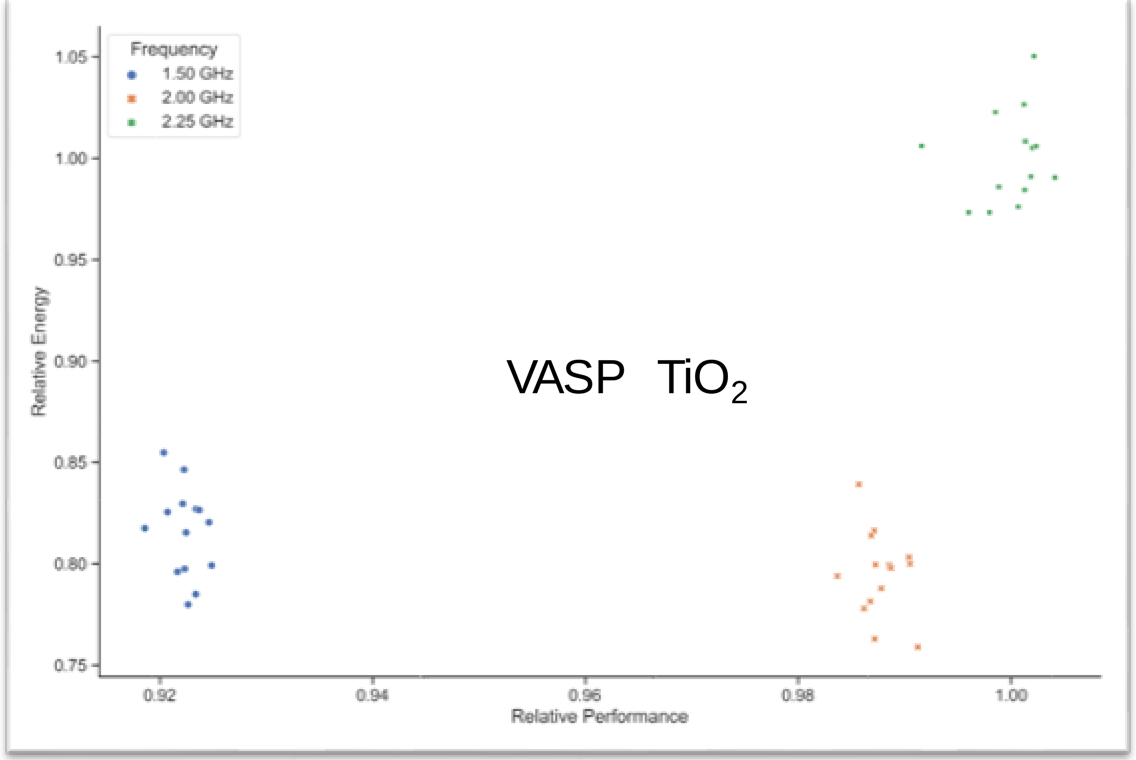

- Investigating reducing default CPU frequency from 2.25GHz to 2.0GHz
- Allow users to override this default, provide higher frequency defaults for some codes

epcc

EPCC's contributions to Net Zero DRI project

ARINZRIT
Identifying socio-technical forces

ENERGETIC
Testing Heterogeneous Architecture

HPC JEEP
Job Efficiency & Energy Usage

IRISCAST
Audit of Carbon Costs

- Taking snapshots of IRIS digital research infrastructure energy consumption over fixed periods of usage
- Developing a carbon model to evaluate the overall carbon usage of IRIS for those periods, taking into account active usage (energy used during the runs) and embodied usage (carbon spent building & providing the DRI)
- HPC-JEEP are presenting in the UKRI Net Zero DRI Project session (Friday morning)

ARCHER2 Net Zero Case Study

- Commissioned by Net Zero DRI project
- Aim is to understand the emissions resulting from the operation of the ARCHER2 service
- Case Study to be published in 2023

- Power, cooling and heat re-use
- Minimising emissions
- System software changes for energy reduction
- Energy utilisation at the application level
- Influencing user behaviour

Education

Efficient software is important → the **developers**' responsibility

Efficient *use* of software is equally important → the **users**' responsibility

Deployment of software is important → the **system providers**' responsibility

- Education is key – enable developers/users/system providers to understand implications of their choices
  - As a University department and national service provider, one of our core missions
  - EPCC's MSc in HPC has included lectures on power & energy efficiency for several years

Final thoughts

HPC systems are **scientific instruments** that are used to find solutions to many of the problems humanity faces

→ to discover new vaccines

→ to design new renewable energy solutions

→ and even to model the climate, in order to more accurately predict climate change and its impact

Significantly reducing scientific throughput is a false economy

Net Zero HPC must be achieved in a fair manner while maintaining, or indeed increasing, the amount of science we do

Net Zero HPC is not just a "noble dream"…

- Very much an inescapable reality
- Provided a perspective **from EPCC's point of view**
  - All centres/sites will be different, but striving for operationally efficiency must be at the heart
- Renewable energy provision is key
  - Challenge is that (global) demand is continuously growing