Energy Efficiency Analysis and Optimization for Present and Future HPC Systems

PASC 2024

Pay Giesselmann
German Climate Compute Center

Motivation:

Saving energy
Mono-topical workload should allow optimization
Recently high energy cost

Agenda

Review energy efficiency of current production system:

- DKRZ energy consumption
- Idle node optimization
- Energy efficient system settings
- Job-specific monitoring

Study possibly heterogeneous future systems:

- ICON CPU & GPU scaling
- ICON kernel benchmarks

German Climate Compute Center (DKRZ)

Levante

- 2982 CPU nodes (AMD Epyc 7763)
- 60 GPU nodes (4x NVIDIA A100 80G)

- Job statistics *(median/N50)*:
  - Runtime: 00:03h / 02:50h
  - Nodes: 1.0 / 50

- Measurement infrastructure
  - Facility: Power & energy
  - Rack (PSU)
  - Node: (lm5066, power) +/- 5% acc., 4Hz sampling, 60s collection
  - Socket: RAPL, nvml

DKRZ Energy Usage

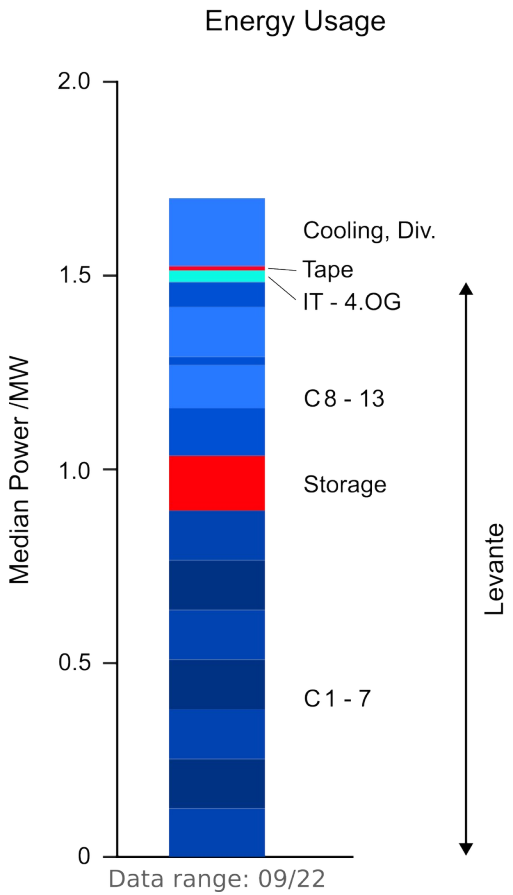

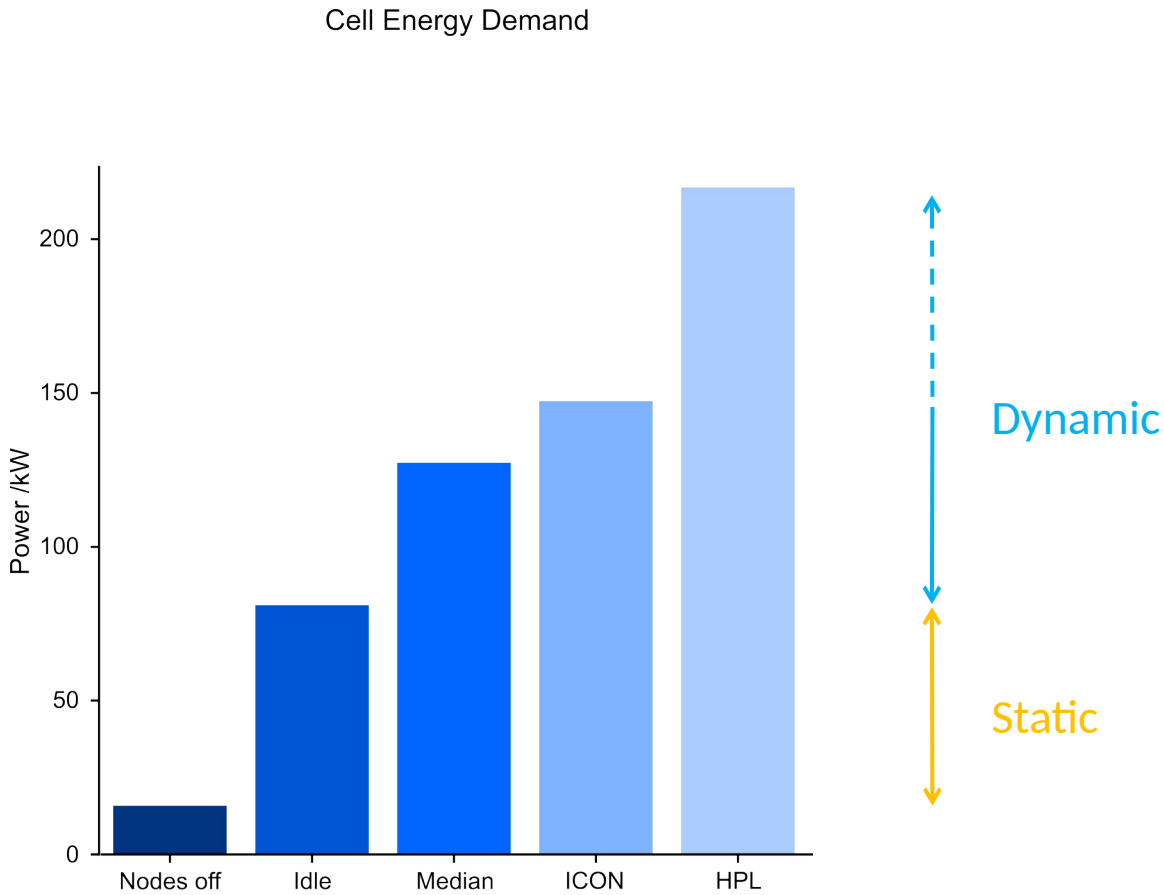

Idle Nodes

- Set CPU governor *powersafe*
- Allow deep sleep states
- Deployed in SLURM epilog

- Switch fabric p-state P0 to P3
  - P0: 250W / node
  - P3: 190W / node
  - 25% less energy

- Idle shutdown needs reliable reboot

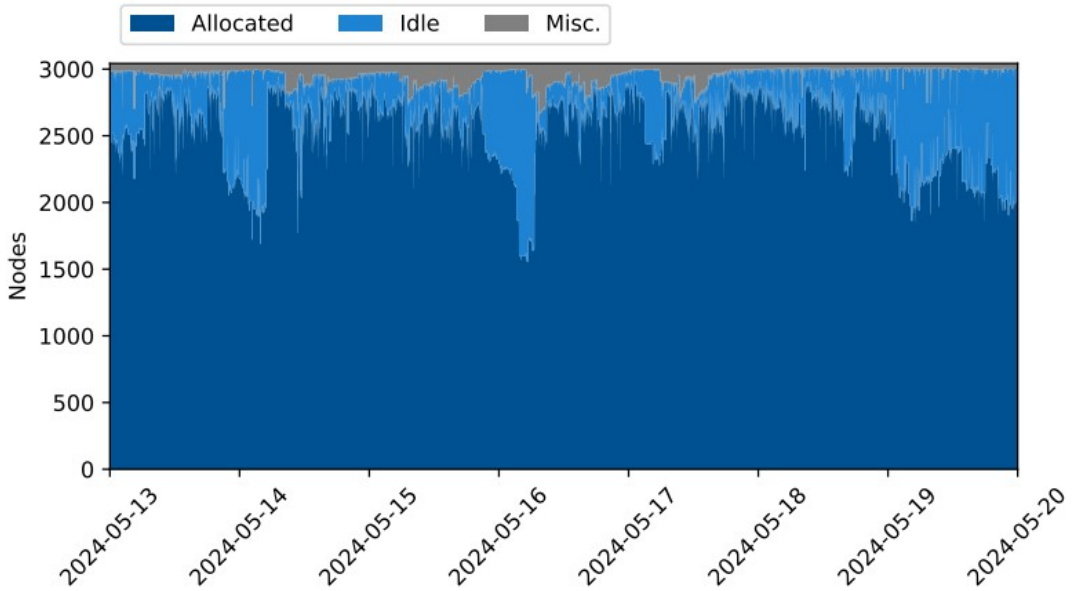

- Levante statistics for 2024:
- 84 % Allocated, 14 % Idle, 2 % Misc.

Thomas Ilsche (TUD): Idle-Power Optimization on HPC Systems, EuroCC, 23.02.24

Allocated Node Settings

- AMD Zen 3 Epyc has limited energy relevant knobs
- Alter system settings TDP and fabric p-state
- Measure runtime impact and energy to solution

Allocated Node Settings

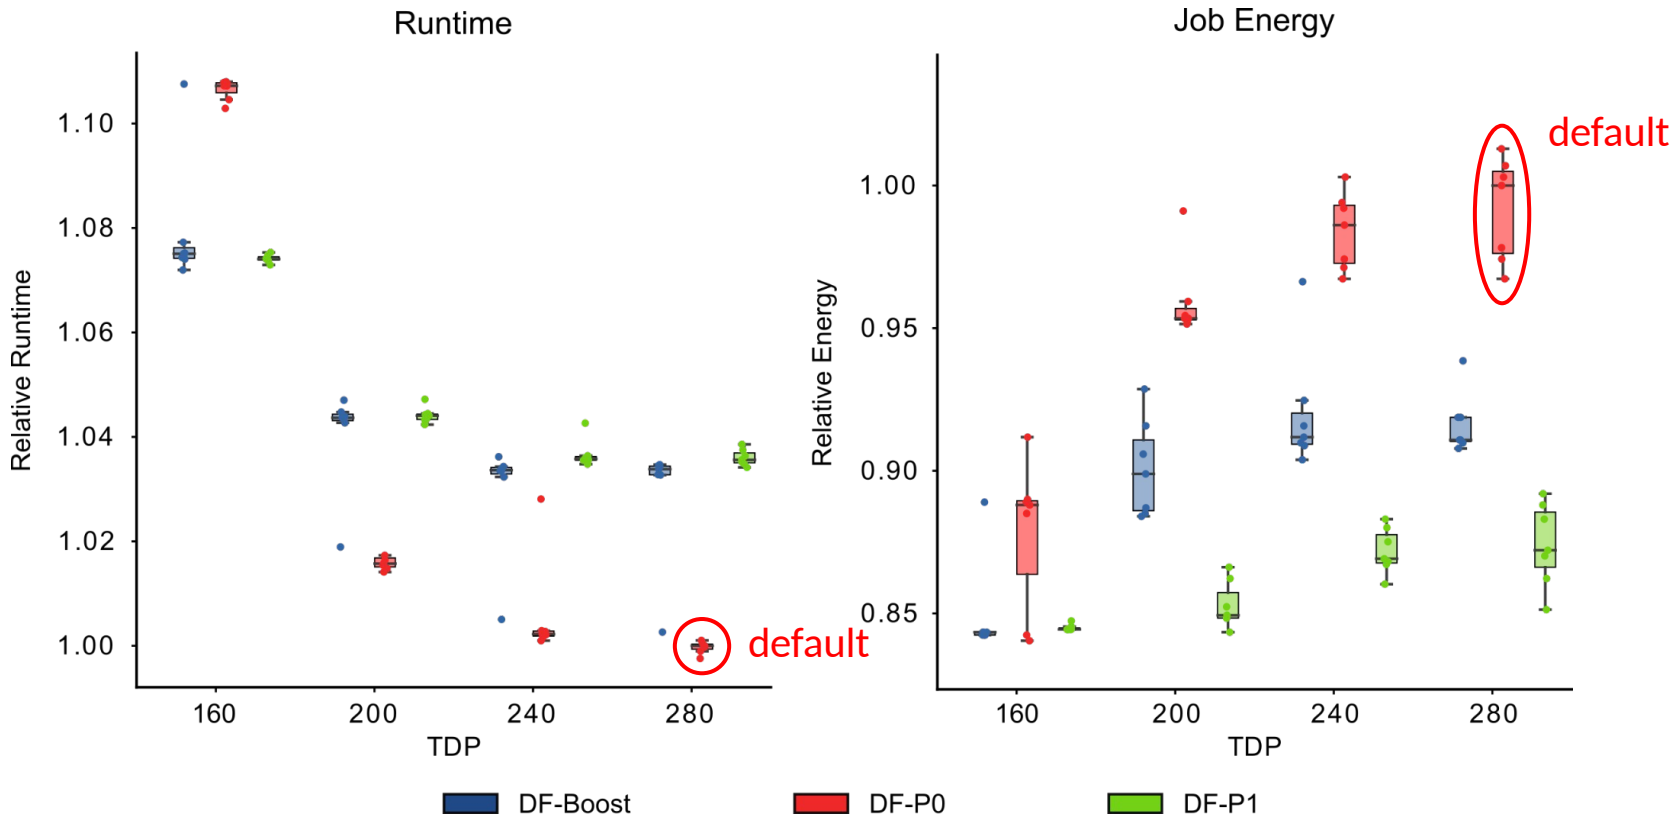

Allocated Node Settings

- AMD Zen 3 Epyc has limited energy relevant knobs
- Alter system settings TDP and fabric p-state
- Measure runtime impact and energy to solution

- Compute: Switch P0 to P1
  - 2.6% longer runtime
  - 5.8% less energy to solution

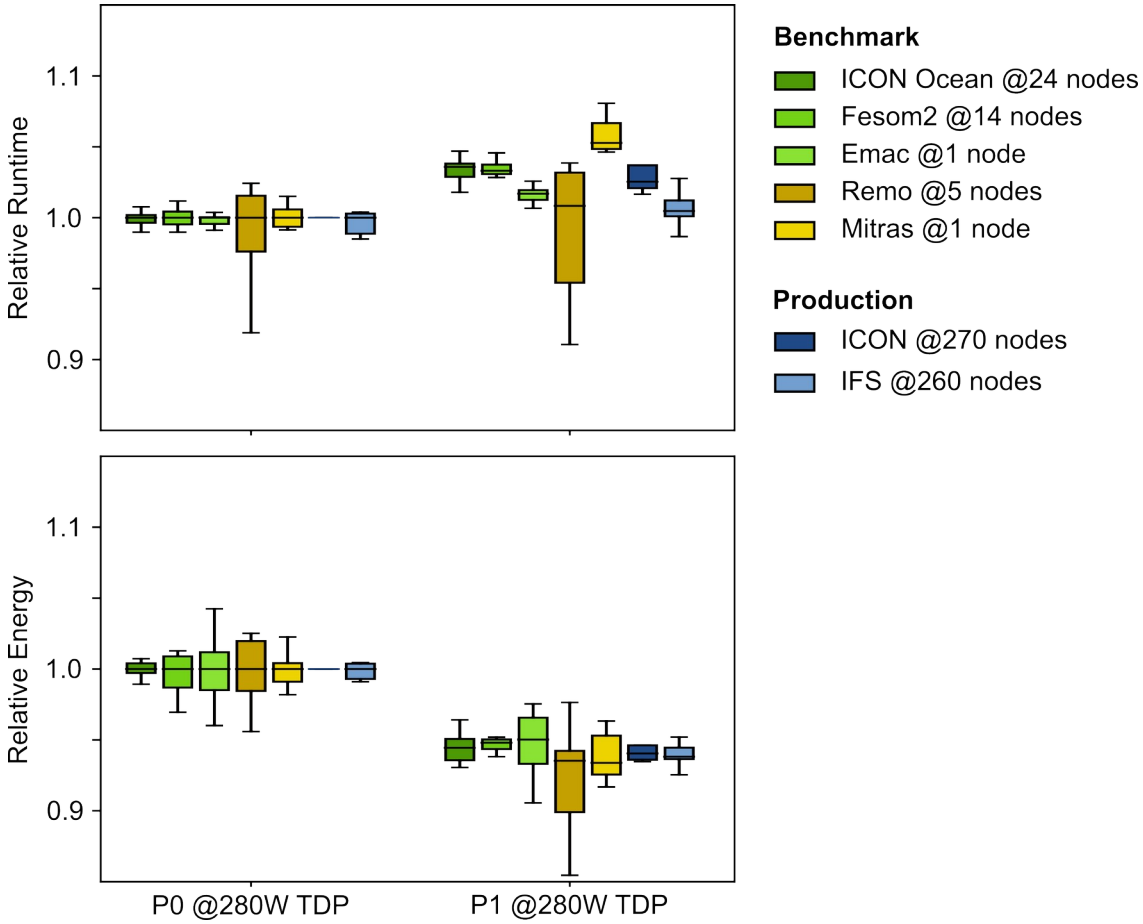

Impact of System Settings

10.11.23: Switch from fabric p-state P0 to P3 on **idle** nodes

Total savings 2023: 30 MWh

Savings 2024 until 30.05: 99 MWh

04.01.24: Switch from fabric p-state P0 to P1 on **allocated** nodes

Savings 2024 until 30.05 (9.2M node-hours): 347 MWh

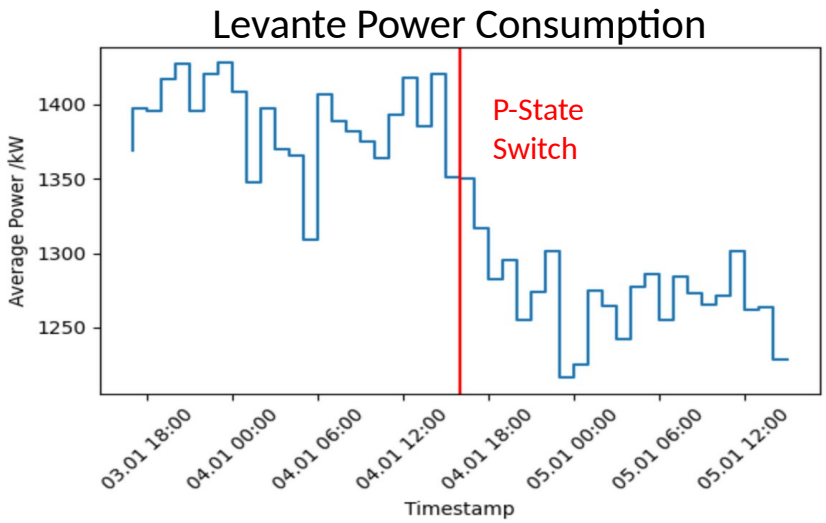

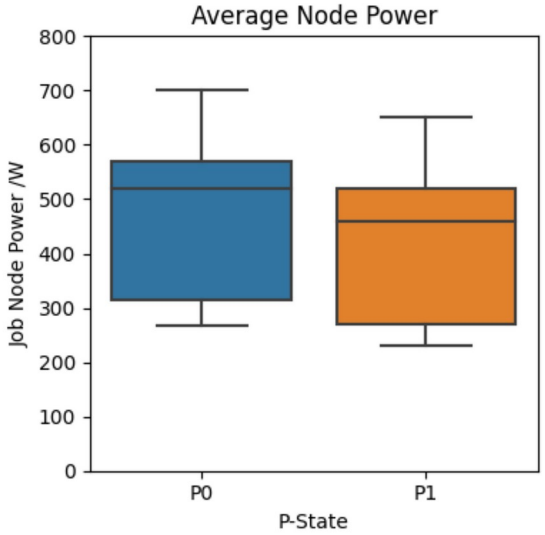

Job-Specific Monitoring

10393026 (levante)
LGM_0.08

8 , 2048
compute

Start: **21/05/2024, 10:51:56**
Duration: **0:18:28** completed
Walltime: **1:55:00**

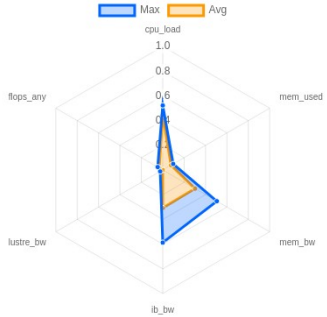

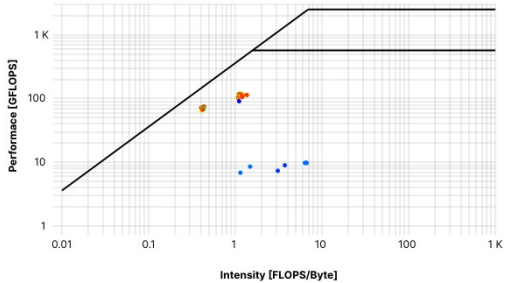

Manage Tags | Metrics

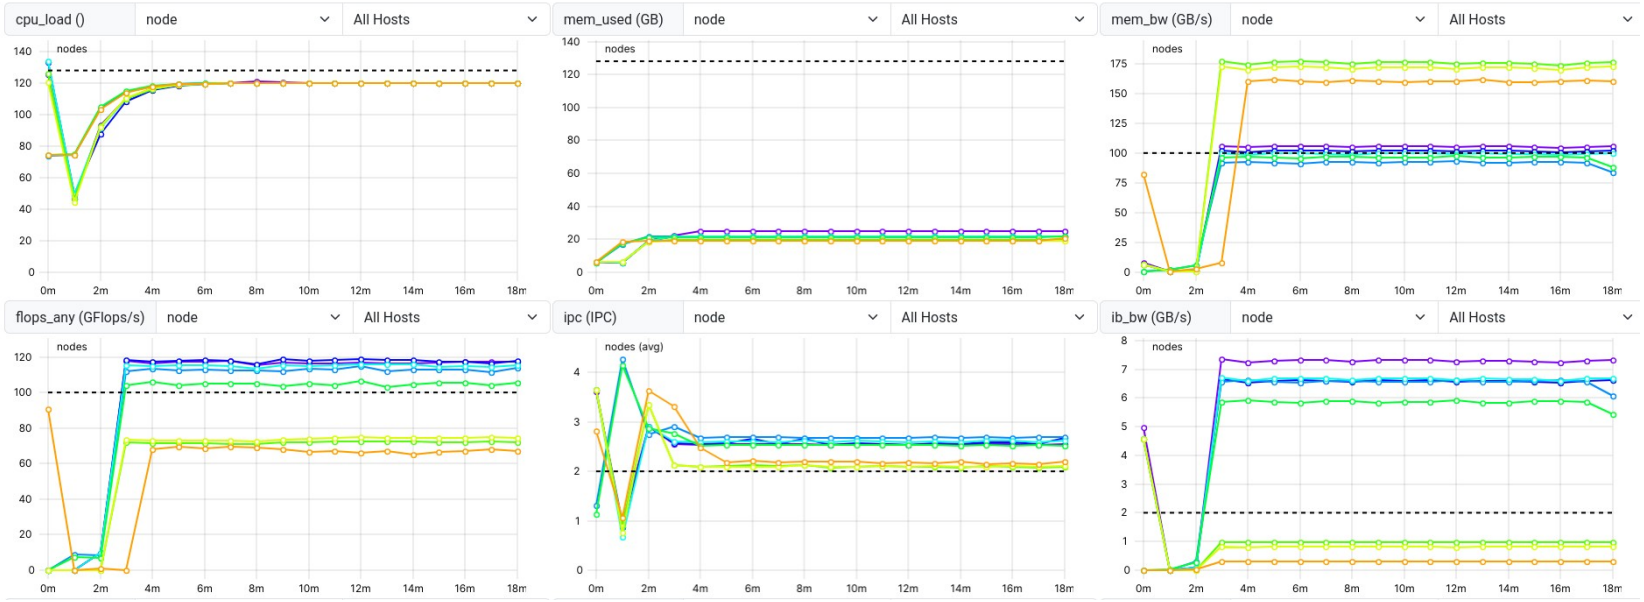

Job-Specific Monitoring

Raising Awareness

- Unknown number of inefficient jobs
- Need efficiency definition
- Want job classification and eventually user notification

long init:

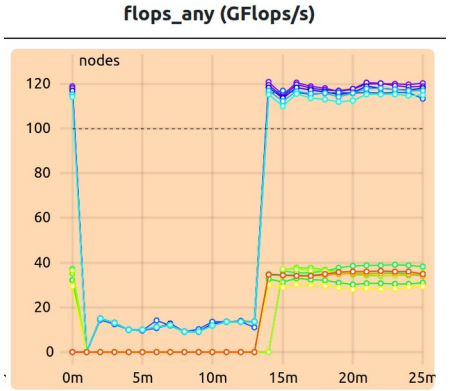

omp missing:

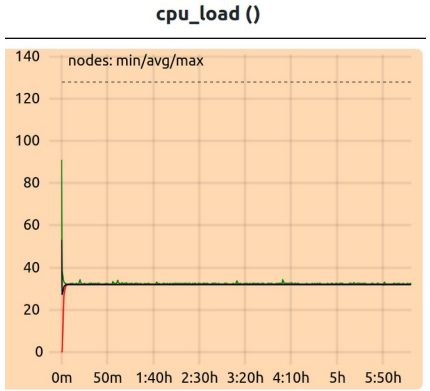

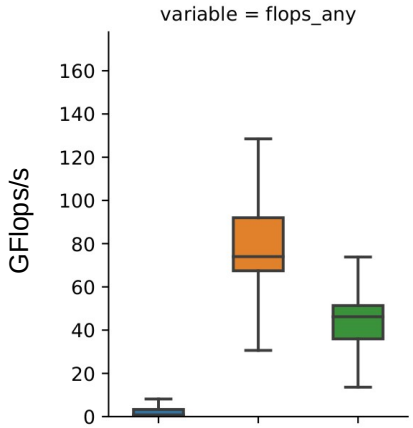

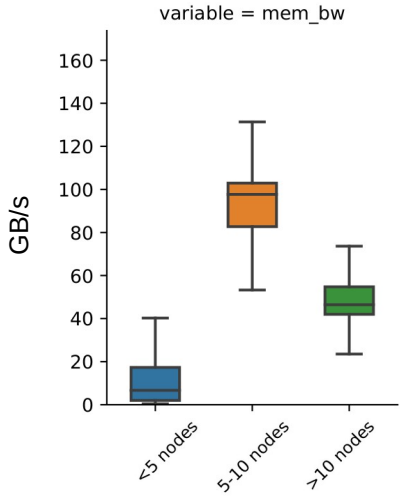

From Present to Future:

Is a Heterogeneous Cluster the Answer to Energy Efficiency?

Pay Giesselmann (DKRZ)

GPU Energy Efficiency

Setup:

- ICON R2B8 (10km resolution)
- Strong scaling from 0.25 to 1 SYPD
- CPU and GPU energy to solution

Result:

- ICON on Levante GPU ~3x more energy efficient
- Energy penalty for scaling ICON on GPU is ~3 times larger

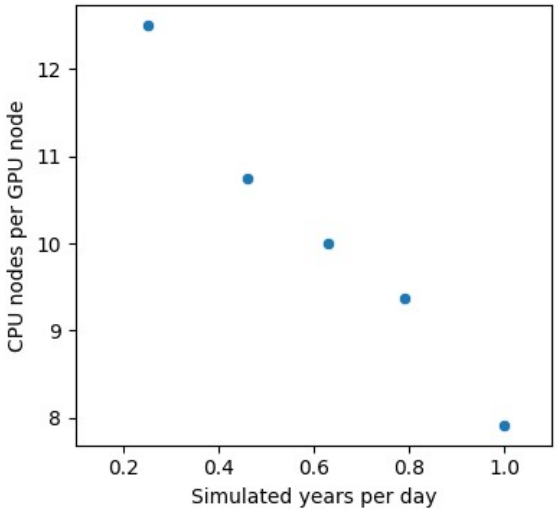

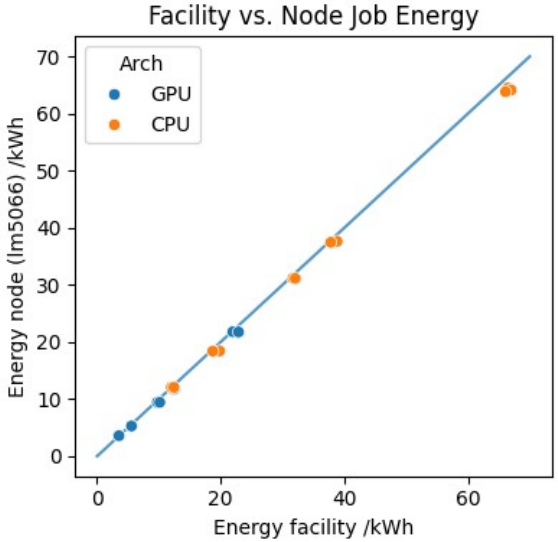

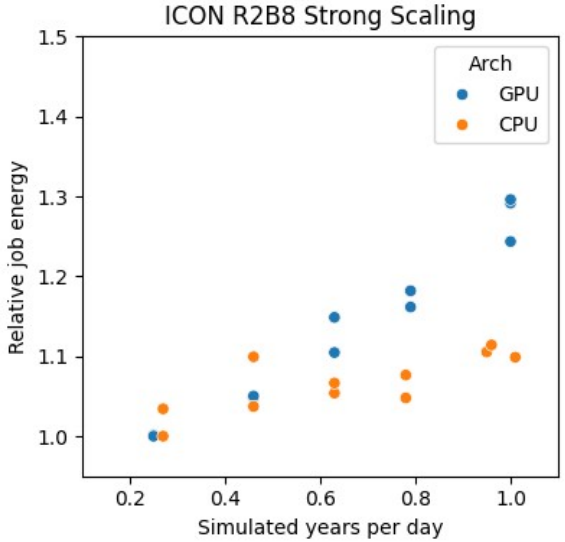

Heterogeneity beyond CPU-GPU?

- **Can we setup an ICON simulation where each component runs on its most energy-efficient architecture?**
- What is the most energy efficient architecture for a component?
- How to build a test-cluster for this?
- Does it save energy?

Hardware:
NVIDIA Grace Grace
NVIDIA Grace Hopper
AMD Genoa
Intel Sapphire Rapid HBM
FUJITSU A64FX
NEC SX-Aurora TSUBASA

Software:

PROXMOX

Spack

slurm workload manager

MQ

ICON Kernel

- Split ICON into kernel
  - e.g. non_hydro, nwp_radiation, …
- Launch kernel in MPI pseudo ranks
- Measure energy & runtime per cell update

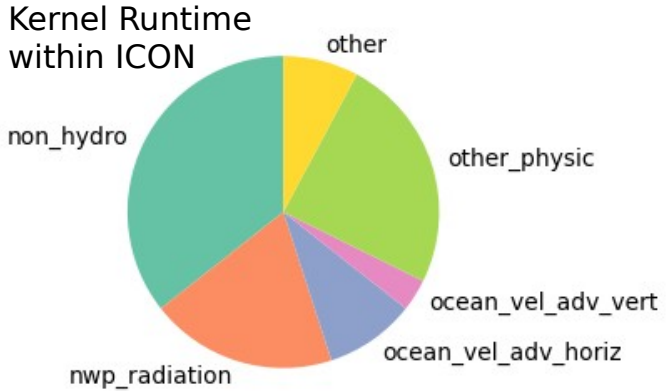

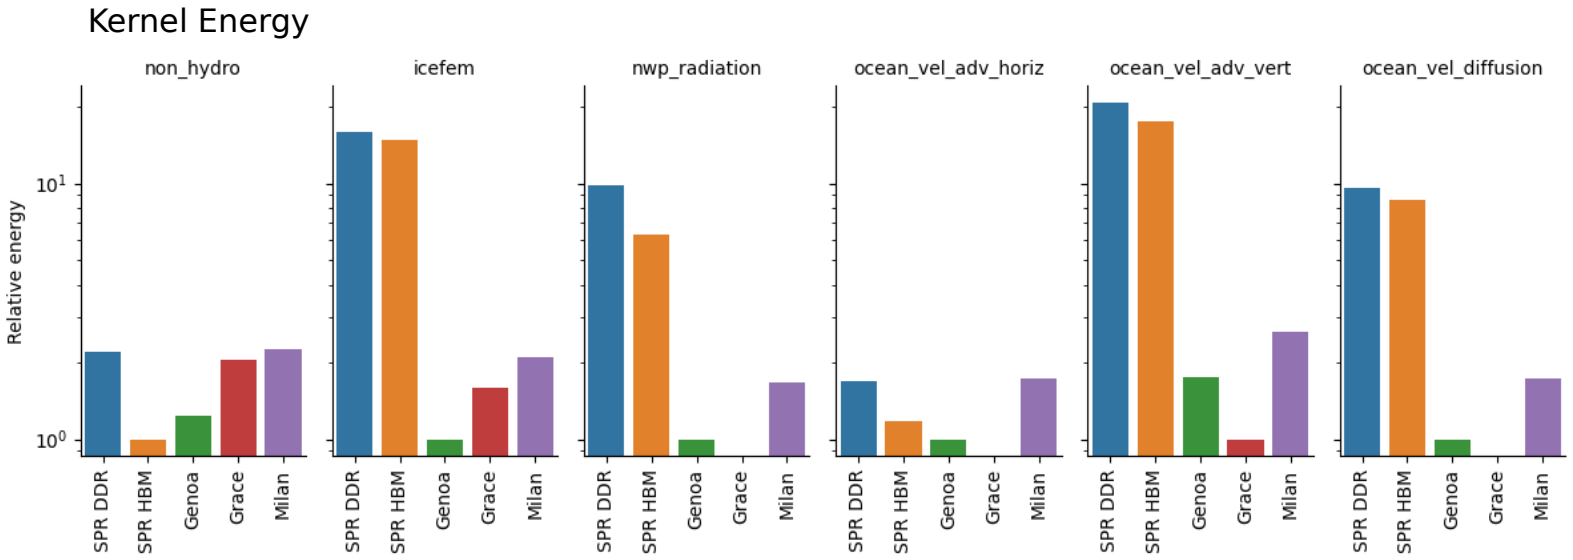

Heterogeneous ICON

- Challenges:
  - Kernels are somewhat synthetic
  - Need coupling to become component (beyond our scope)
  - Load balancing
- Available components: Atmosphere, Ocean, HAMOCC
- Possible configuration:
  - Atmosphere: Intel SPR HBM / NVIDIA Hopper
  - Ocean, HAMOCC: AMD Genoa

Summary

- Unraveling datacenter energy sinks
- Optimizing idle node settings
- Optimizing compute node settings: static & dynamic
- Job specific performance metrics

- CPU & GPU scaling
- Distributing ICON to its most energy efficient hardware(s)

Acknowledgments

DKRZ:
Applications Dept.
Hendryk Bockelmann
Jan Frederik Engels
Julius Plehn

Claudia Frauen
Dominik Zobel

Systems Dept.
Carsten Beyer

Eviden:
Stephan Jauré

EEHPC Partners:
FAU | HLRS | RWTH | HPE

EECliPs Partners:
TUD | ParTec AG | Eviden

Federal Ministry
of Education
and Research

Conflict of interest: AMD & Intel sponsored parts of test-cluster

Backup

Monitoring Stack