ErUM-Data Sustainability workshop 2024

Sustainability at HITS: from energy monitoring to carbon-aware computing

Oleksiy M. Kozlov

Staff scientist, Sustainable Computing Advisor

The Exelixis Lab
Heidelberg Institute for Theoretical Studies (HITS gGmbH)
exelixis-lab.org | h-its.org

November 18, 2024

Energy monitoring: group servers

- Node-level power
  - IPMI → ~every server
- telegraf plugin
  - 30s query interval
- Store to influxDB
- Grafana dashboard

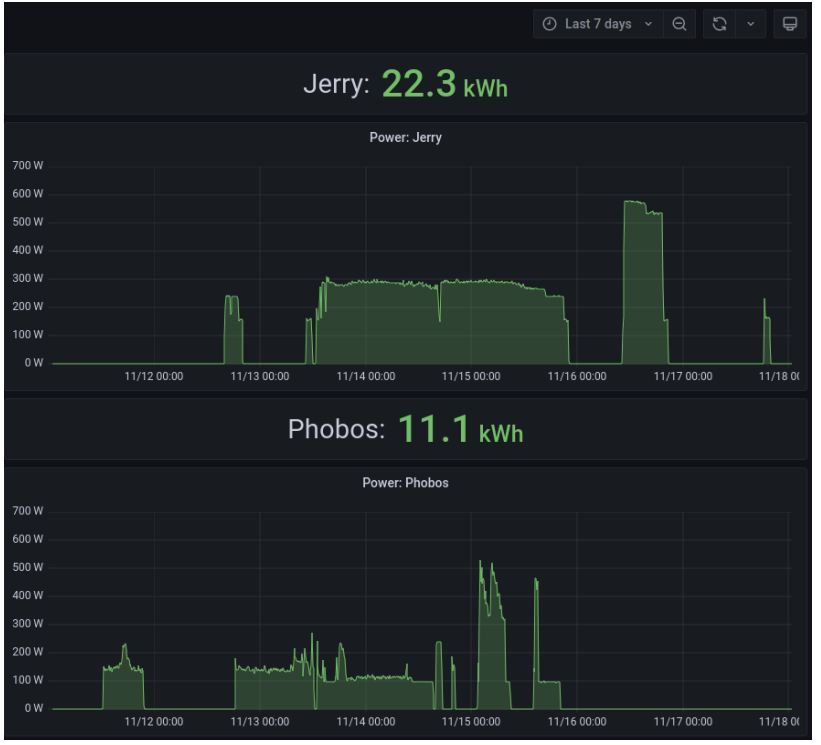

HowTo: https://github.com/amkozlov/ipmi-grafana

PowerOff-on-Idle

- Server idle power: 100 W – 400 W
- Low utilization → wasted energy
- Detect idle → PowerOff
- PowerOn via IPMI

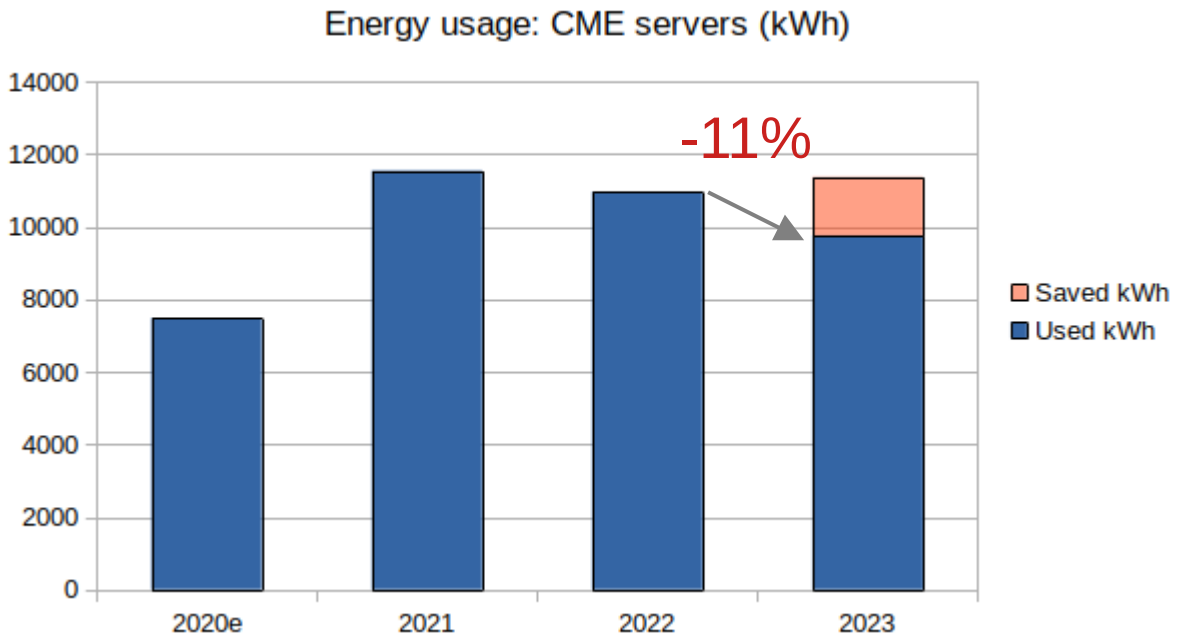

Savings 03.2023 – 11.2024: 3500 kWh (19%)

Scripts & HowTo: https://github.com/amkozlov/idle-sleep

Cluster powercap: Winter 2022/23

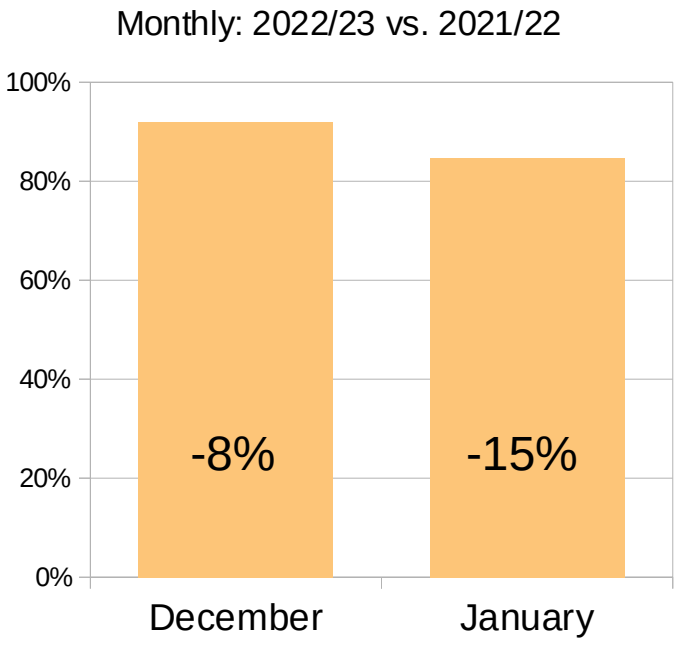

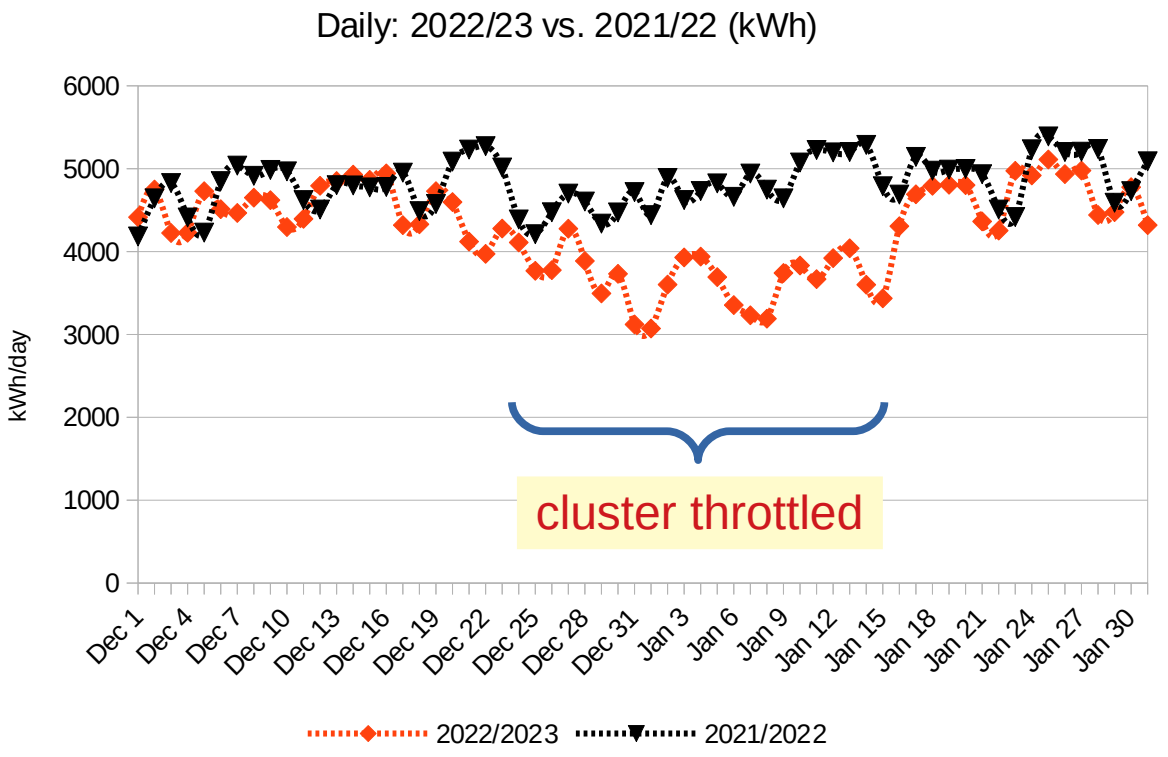

Static powercap: CPU 75%, GPU 70% => visible energy savings at HITS campus level!

Carbon-aware power scaling

“Electricity traffic light”

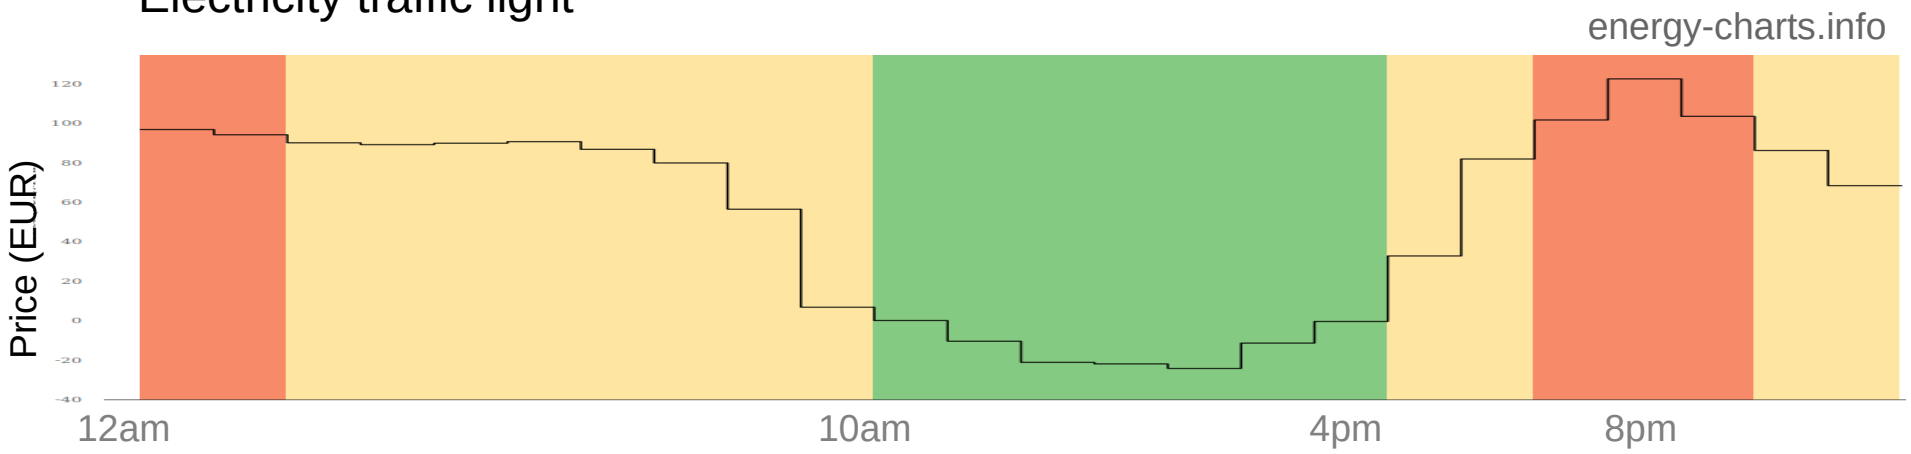

Use less energy in red, more in green phases → EcoFreq

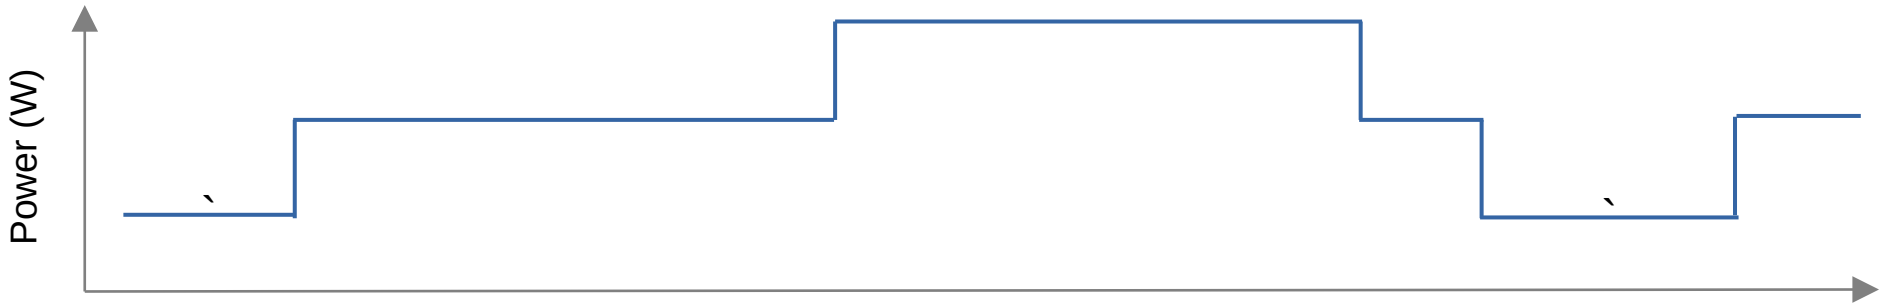

Proof-of-concept: EcoFreq

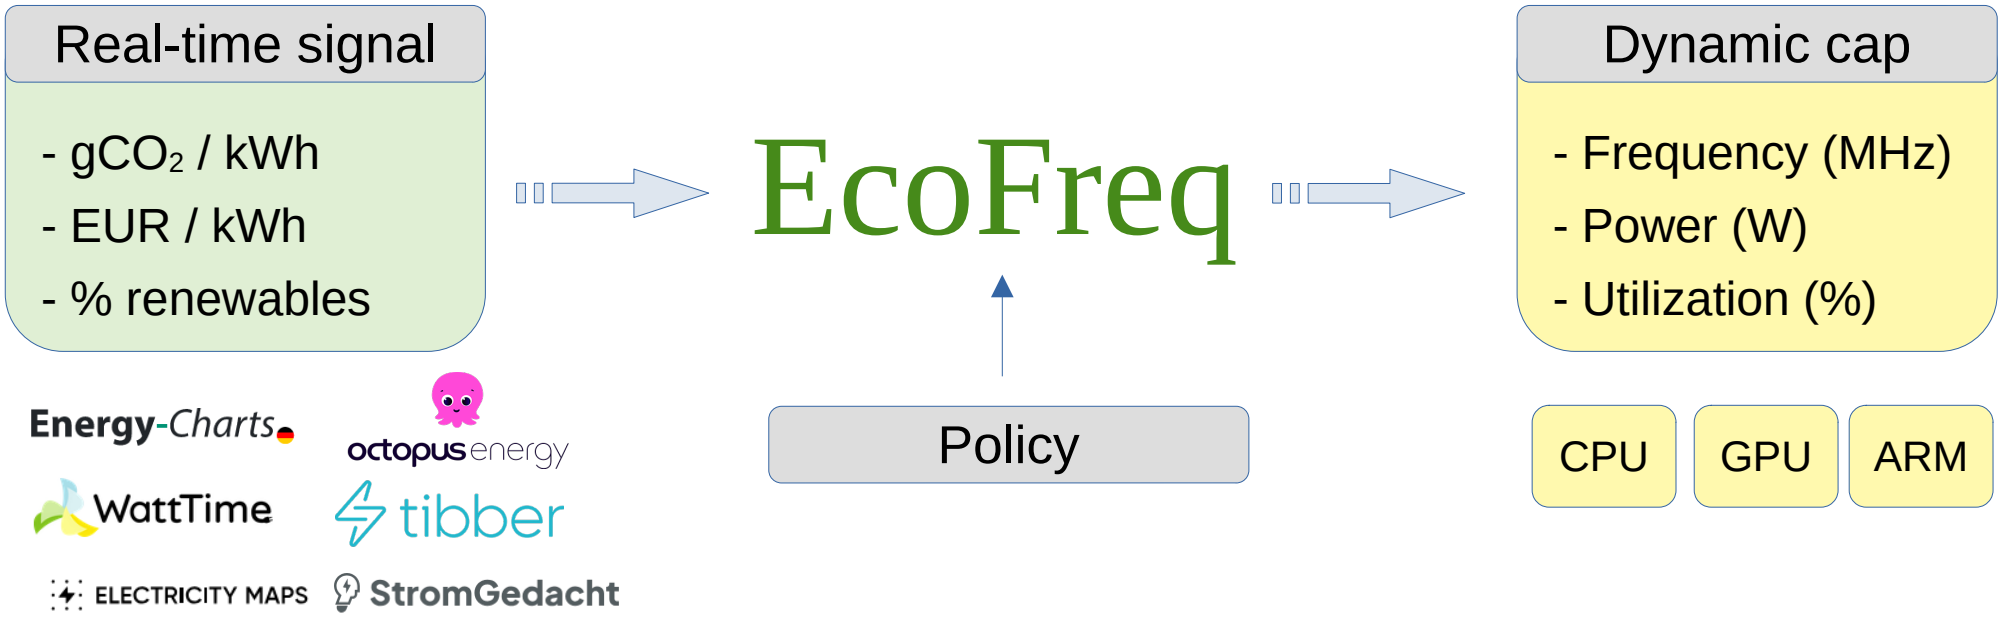

- ✅ Zero workload modification
- ✅ Low latency (seconds)
- ✅ Broad hardware support

EcoFreq: evaluation

- Over-proportional savings
  - 15-18% lower CO_2 & electricity cost
  - @ 10% throughput loss
  - Adjustable via scaling policy
- Presented at ISC'24
  - Paper: https://doi.org/10.23919/ISC.2024.10528928
  - Code: https://github.com/amkozlov/eco-freq

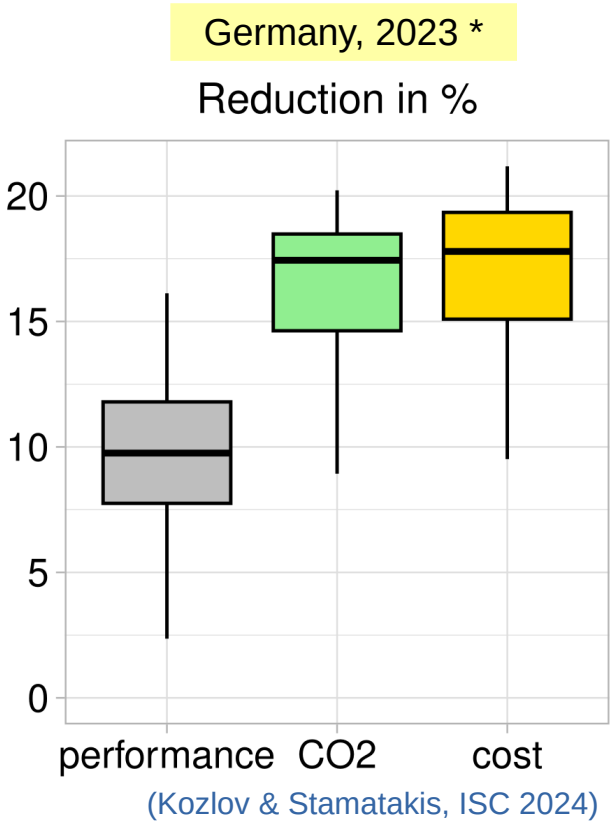

(Kozlov & Stamatakis, ISC 2024)

* electricity day-ahead market data; 3-step scaling policy (100%/80%/60%); 14 HPC workloads, Xeon Cascade Lake

Future / ongoing work

- EcoFreq on the cluster
  - SLURM integration
- Sustainability dashboard
  - Energy, CO2, travel etc.
- Rooftop solar
  - 5-7% of consumption
- Warm water cooling

```
#!/bin/bash
#SBATCH -N 8
#SBATCH -t 24:00:00
#SBATCH -q ecofreq
```

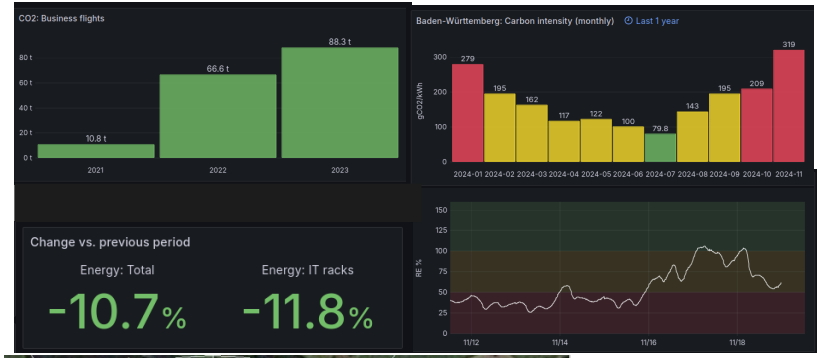

Q & A

Backup

PowerOff-on-Idle

- Idle detection cron job
  - No active sessions + CPU utilization < 0.5 → idle
  - Idle since 1 hour → poweroff
  - Can be temporarily disabled:

```
deimos$ ecosleep disable 12h
Server will not be powered off until: Di 16. Mai 10:13:34 CEST 2023
```

- PowerOn via SSH

```
laptop$ alias | grep ecowake
alias ecowake='ssh kozlovay@XXXX.h-its.org sudo /hits/fast/cme/ecosleep/wakeup.sh'

laptop$ ecowake deimos
kozlovay@XXX.h-its.org's password:
Chassis Power Control: Up/On
waiting for
deimos ..........................................................
Server is back online!
```